Volume 19, Number 4 October 2023

IJCB

INTERNATIONAL JOURNAL OF CENTRAL BANKING

IJCB INTERNATIONAL JOURNAL OF CENTRAL BANKING Volume 19, Number 4 October 2023

International Journal of Central Banking

Volume 19, Number 4 October 2023

The *International Journal of Central Banking* is published bimonthly (ISSN: 1815-4654). Online access to the publication is available free of charge at **www.ijcb.org**.

International Journal of Central Banking
Printing & Fulfillment K1-120
Federal Reserve Board
Washington, DC 20551

Phone: 202-452-3425
Fax: 202-728-5886
E-mail: editor@ijcb.org

ISSN: 1815-4654

International Journal of Central Banking

Macroprudential Policy and Aggregate Demand*

André Teixeira[a] and Zoë Venter[b]
[a]ISEG, REM/UECE, Universidade de Lisboa
[b]Católica Lisbon School of Business and Economics

This paper assesses the impact of macroprudential policy (MaPP) on aggregate demand in the EU between 2000 and 2019. Using a difference-in-differences approach, we find that MaPP reduces household consumption and increases firm investment. These effects are relatively mild in the short run but become more pronounced in the long run. Our findings point to a weaker macroeconomic impact than suggested in previous studies.

JEL Codes: E21, E22, E52, E58, O47.

1. Introduction

Macroprudential policy (MaPP) is back in fashion, and rightly so. Few economists today would dispute that MaPP is a powerful weapon in the arsenal of crisis economics. But what do we know about its effects on aggregate demand? How does it affect consumption and investment? Answering these questions is crucial to assess the overall impact of MaPP.

So far, the existing literature has focused on the effects of MaPP on output (e.g., Lim et al. 2011; Cerutti, Claessens, and Laeven 2017). Not surprisingly, these studies find an inverse relationship between the adoption of MaPP and economic growth. Their story

*We would like to thank António Afonso, Steven Ongena (the editor), and two anonymous reviewers of the IJCB for their fruitful comments. We are also grateful to Brantly Callaway and Pedro Sant'Anna for their suggestions on the implementation of the difference-in-differences. The opinions expressed herein are those of the authors. Author contact: Teixeira (corresponding author): ISEG, REM/UECE, Universidade de Lisboa, R. Miguel Lúpi 20, 1249-078, Lisbon, Portugal. This work is supported by Fundação para a Ciência e a Tecnologia. E-mail: ateixeira@iseg.ulisboa.pt. Venter: Católica Lisbon School of Business and Economics, Palma de Cima, 1649-023, Lisbon, Portugal. E-mail: zventer@ucp.pt.

is intuitively simple but it has important implications: MaPP constrains credit with searing consequences on growth. This poses the question: should we evaluate the effectiveness of MaPP through the lens of output growth? After all, the goal of MaPP is to tame credit and slow down growth. Perhaps a more pertinent question to ask is whether MaPP affects households and firms in a similar way.

In this paper, we argue that the effects of MaPP on consumption and investment depend—directly or indirectly—on the financial constraints imposed on households and firms. If MaPP tightens borrowing constraints for everyone in the same way, both households and firms will be forced to save more and borrow less. This should lead to a decline in both consumption and investment. However, if MaPP makes access to credit more difficult for households than for firms, consumption is likely to fall but investment should remain stable or even increase if banks expand credit to the corporate sector. The converse could also be true: if MaPP makes access to credit more difficult for firms than for households, investment is likely to plummet and consumption should remain stable or even increase if banks shift their lending to households. In these last two scenarios, MaPP may have a profound effect on consumption and investment. By how much is an empirical issue, which we address via a novel difference-in-differences (DiD) approach.

The purpose of this paper is to extend the existing literature in three directions. First, we isolate the effects of MaPP on spending components of aggregate demand, particularly on consumption and investment. A shortcoming of previous papers is that they do not explain how MaPP influences private spending and imperils growth. Relative to these papers, we directly link the adoption of MaPP to movements in household consumption and firm investment.

Second, we depart from traditional regressions and time-series models to establish causality from MaPP to aggregate demand. This is important because MaPP is usually implemented in response to contemporaneous events. By using the first wide-scale staggered DiD in a policy setting, we are able to estimate the effects of MaPP in a setting with multiple countries and variation in treatment timing.

Third, we distinguish between the short- and long-run effects of MaPP. Separating out the two is an empirically difficult matter, but we are able to estimate a single interpretable treatment effect parameter that accounts for the dynamic effects of MaPP. This allows us

to examine the effects of MaPP over time and determine whether they are more pronounced in the short or long run. This distinction between short- and long-run effects has received surprisingly little attention in the literature, but it is crucial to our understanding of the overall impact of MaPP.

As we shall see, our results indicate that MaPP has asymmetric effects on consumption and investment. In particular, we find that households in countries that implement MaPP increase their savings rate by 1.87–3.63 percentage points. This corresponds to a sharp increase of one quarter in savings. Furthermore, we find that MaPP boosts firm investment by a whopping 5.05–6.63 percentage points over time. These results are statistically significant and stand up to several robustness checks.

The rest of the paper is organized as follows. Section 2 reviews the existing literature. Section 3 provides a detailed explanation of the staggered DiD. Section 4 discusses the empirical results and investigates the robustness of our findings. Section 5 concludes.

2. Literature Review

A vast literature examines the effects of MaPP on financial stability. Most of these papers suggest that MaPP curtails lending (Lim et al. 2011; Dell'Ariccia et al. 2012) and reduces excessive leverage (Claessens, Ghosh, and Mihet 2013). Moreover, MaPP lessens the probability of a crisis (e.g., Galac and Kraft 2011), especially in housing markets (Crowe et al. 2011; Kuttner and Shim 2013). This early literature provides compelling evidence that MaPP is an effective tool to manage financial cycles and reduce systemic risk.

However, recent research finds that MaPP may have deleterious effects on growth. Most notably, Angelini, Neri, and Panetta (2014) show that banking regulation decreases the steady-state level of output. A few general equilibrium models also show that MaPP can be used to correct externalities in aggregate demand (e.g., Farhi and Werning 2016). These results seem to be consistent with conventional theory on the relationship between credit and spending. If MaPP restricts access to credit, it may force constrained households to reduce consumption (e.g., Hall 2011). In a scenario of rapid deleveraging, MaPP may even increase precautionary savings, which

is likely to depress aggregate demand even further (e.g., Eggertsson and Krugman 2012; Guierrieri and Lorenzoni 2017).

This theoretical work finds empirical support in studies that use regressions and time-series models to estimate the effects of MaPP on growth (e.g., Lim et al. 2011; Cerutti, Claessens, and Laeven 2017; Akinci and Olmstead-Rumsey 2018). This stream of research amassed a remarkable body of evidence on a negative relationship between MaPP and output growth. These studies generally conclude that MaPP should be tightened in boom periods and loosened in bust periods. Yet, they offer little explanation on the transmission channels of MaPP, i.e., the way in which MaPP is supposed to have affected output. This is key to our understanding of the causal effects of MaPP on growth.

Until now, only a few empirical papers have used causal techniques to identify the impact of MaPP on growth. Behncke (2020) uses a simple DiD to estimate the effects of MaPP on lending using data from 25 banks in Switzerland. Her findings show that MaPP constrains lending with no unintended consequences on credit risk. But most papers point to important redistributive effects of MaPP. For example, Defusco, Johnson, and Mondragon (2020) use a DiD strategy to exploit a policy-induced discontinuity in the debt-service-to-income ratio in the United States. They report a substantial increase in borrowing costs following the adoption of MaPP, with ill effects on the distribution of leverage in the mortgage market. Interestingly, Acharya et al. (2020) also find a similar reallocation of mortgage credit after running a DiD model on loan-level data from Ireland.

A major drawback of traditional DiD strategies is that they restrict the analysis to micro-level data from a single country. This is perhaps the simplest way to estimate the treatment effects of MaPP on credit. Although this DiD approach sheds light on the impact of MaPP on credit, it leaves many questions unanswered. For instance, how does MaPP affect spending across countries? How does the length of exposure to MaPP influence consumption and investment over time? By construction, traditional DiD approaches are unable to answer these questions because countries implement MaPP in different time periods.

Overall, there has been, in both theory and empirical work, an obvious push for generality on the effects of MaPP. Yet, it remains

unclear how MaPP affects households and firms. The existing work is premised on the assumption that MaPP affects all agents in the same way. But this is unlikely because MaPP imposes different financial constraints on households and firms. A more interesting way to assess the overall impact of MaPP is to disentangle its causal effects on households and firms. It is to these matters that we turn next.

3. Methodology

3.1 Method

Our methodology is based on the DiD approach with staggered treatment adoption proposed by Callaway and Sant'Anna (2021). Similarly to a standard DiD, this method allows for a causal interpretation and it circumvents the restrictive assumptions of regressions and time-series models. But unlike a standard DiD, it enables us to estimate the average treatment effects of MaPP in a setting with multiple countries and variation in treatment timing.

To do so, let us start with some notation. We consider τ periods where $t = 1, \ldots, \tau$ and D_t is a binary variable that equals 1 when a country implements MaPP in quarter t and 0 otherwise. We then define G_g equal to 1 when a country is first treated in quarter g and 0 otherwise. Lastly, we assign C equal to 1 to the countries that never implement MaPP in our sample (i.e., "never treated") and 0 otherwise. This implies that each country in our sample will have exactly one G_g or C equal to 1.

The generalized propensity score $p_g(X)$ is then defined as the probability that a country is treated conditional on having covariates X and belonging to group g or the control group, i.e., $p_g(X) = P(G_g = 1|X, G_g + C = 1)$. The observed outcome in each period t is estimated as follows:

$$Y_t = D_t Y_t(1) + (1 - D_t) Y_t(0), \tag{1}$$

where $Y_t(1)$ and $Y_t(0)$ are the potential outcomes in time t with and without treatment, respectively.

In contrast to a standard DiD, our main causal parameter of interest is a group-time average treatment effect ($ATT(g,t)$). Simply

put, the $ATT(g,t)$ gives us the average treatment effect experienced by group g in time t, with "group" being defined as the first period of implementation of MaPP, as below:

$$ATT(g,t) = E[Y_t(1) - Y_t(0) \,|G_g = 1]. \tag{2}$$

In our panel data setup, under the assumptions of parallel trends, irreversibility of treatment, and covariate overlap and for $2 \leq g \leq t \leq \tau$ the $ATT(g,t)$ for group g in period t can be non-parametrically identified and estimated as below:[1]

$$ATT(g,t) = E\left[\left(\frac{G_g}{E\lfloor G_g \rfloor} - \frac{\frac{p_g(X)C}{1-p_g(X)}}{E\left[\frac{p_g(X)C}{1-p_g(X)}\right]}\right)(Y_t - Y_{g-1})\right]. \tag{3}$$

Equation (3) allows us to assess how the effect of MaPP varies by group and time. It is worth noting that the $ATT(g,t)$ weights up observations from the control group that share similar characteristics to those in each treated group. This reweighting procedure ensures that the covariates of the treated group and the control group remain balanced.

Next, we aggregate the $ATT(g,t)$ across g and t to interpret the overall effects of MaPP. Given that many, if not most, treated groups will comprise a single country, the easiest way to obtain an "overall" $ATT(g,t)$ is to use a simple average, as follows:

$$\frac{2}{\tau(\tau-1)} \sum_{g=2}^{\tau} \sum_{t=2}^{\tau} \mathbf{1}\{g \leq t\} ATT(g,t). \tag{4}$$

[1]Callaway and Sant'Anna (2021) show that Equation (3) enables us to identify the treatment effects under the assumptions of parallel trends, irreversibility of treatment, and covariate overlap. The first assumption was tested using the Cramér-von-Mises (CvM) test that fails to reject the parallel trends assumption (Appendix C). The second assumption states that a country that adopts MaPP is forever treated, which is consistent with the behavior of countries in our sample that rarely reverse MaPP. The last assumption simply requires a control group for every treatment period.

Alternatively, we can compute a weighted average of each $ATT(g,t)$ by putting more weight on the $ATT(g,t)$ of groups that are exposed to MaPP for longer, as below:

$$\frac{1}{k}\sum_{g=2}^{\tau}\sum_{t=2}^{\tau}\mathbf{1}\{g \leq t\}ATT(g,t)\,P\left(G=g\right), \tag{5}$$

where $k = \sum_{g=2}^{\tau}\sum_{t=2}^{\tau}\mathbf{1}\{g \leq t\}\,P(G=g)$ so that the weights on the $ATT(g,t)$ sum to 1.

In our baseline model, the results are computed using the doubly robust method,[2] no covariates, and a "not yet treated" control group. Statistical significance is assessed using clustered bootstrapped standard errors at the country level, which also account for the autocorrelation of the data. Of course, making inference based on several $ATT(g,t)$ can be troublesome. In the following subsections, we explain how the choice and timing of MaPP can bias the estimates of the overall $ATT(g,t)$. We then describe in detail how we estimate the treatment effect parameters to circumvent these issues. This can be done by computing group-time treatment effects and dynamic effects.

3.1.1 Group-Time Treatment Effects

The adoption of MaPP is a choice of each country. Therefore, countries that implement MaPP earlier may also experience the effects of being treated earlier. A caveat of combining the $ATT(g,t)$ across g and t using a simple average is that we may overweight the effect of early-treated groups with more observations in post-treatment periods. To get around this issue, we compute the $ATT(g,t)$ specific to each treated group and we average them across all post-treatment periods:

$$\tilde{\theta}_S(g) = \frac{1}{\tau - g + 1}\sum_{t=2}^{\tau}\mathbf{1}\{g \leq t\}ATT(g,t). \tag{6}$$

[2]The ATT uses OLS regression to compute the difference between the treated and control groups for each observation; these differences are then weighted according to the probability of each observation occurring.

Equation (6) is the time-averaged treatment effect for countries in group g. In simple terms, it is an average of each available $ATT(g,t)$ in a particular group g across time. The "overall" ATT, θ_S, can then be estimated by aggregating the group-specific treatment effects across groups, as below:

$$\theta_S = \sum_{g=2}^{\tau} \tilde{\theta}_S(g) P(G = g). \tag{7}$$

Equation (7) is our main measure of the overall impact of MaPP on aggregate demand. Although it may seem similar to Equation (5), there is an important difference in the weights. While Equation (5) assigns more weight to groups with a higher number of post-treatment periods, the weights in Equation (7) depend only on group size. In this way, Equation (7) does not overweight the effects of earlier-treated groups and provides an unbiased estimate of the effects of MaPP on each treated group g.

3.1.2 Dynamic Treatment Effects

The effects of MaPP on aggregate demand may also depend on the length of exposure to these policies. One may expect larger effects of MaPP to occur over longer horizons when households and firms have had the time to adjust their behavior. However, a caveat of parameter (5) is that it does not explicitly consider a country's length of exposure to MaPP. To account for this, we begin by averaging the group-time $ATT(g,t)$ into treatment effects at different lengths of exposure to treatment, as follows:

$$\tilde{\theta}_D(e) = \sum_{g=2}^{\tau} \sum_{t=2}^{\tau} \mathbf{1}\{t - g + 1 = e\} ATT(g,t) P(G = g | t - g + 1 = e), \tag{8}$$

where e is the length of exposure to MaPP.

A length of exposure equal to 0 estimates the average effect of MaPP across groups in the quarter of implementation of MaPP. To make the point most clearly, suppose that $e = 1$. Then, Equation (8) estimates a value for the $ATT(g,t)$ based on group size for $g = t = 0$. This will be the estimate of the $ATT(g,t)$ in the first

quarter after MaPP adoption. When $e = 2$, Equation (8) estimates a different value for the $ATT(g,t)$ based on group size for all groups where $t - g = 1$. This will be the estimate of the $ATT(g,t)$ for all the countries exposed to MaPP for two quarters. The $ATT(g,t)$ is computed iteratively in this way for $e = 0, \ldots, 40$.

The θ_D captures the dynamic evolution of treatment effects by averaging $\tilde{\theta}_D(e)$ over all possible values of e, as below:

$$\theta_D = \frac{1}{\tau - 1} \sum_{e=1}^{\tau-1} \tilde{\theta}_D(e). \tag{9}$$

Equation (9) is our main estimate of the dynamic effects of MaPP. Once again, the crucial difference between θ_D, θ_S, and Equation (5) is in the weights: θ_D puts more emphasis on $ATT(g,t)$ when g is significantly less than t (i.e., when e is large). This allows for groups with a longer exposure to MaPP to be weighted more when there is a relatively small number of groups with long periods of exposure. This parameter is particularly suitable to measure how the treatment effects of MaPP evolve over time.

3.2 Data

Our empirical setting uses quarterly data on 21 European countries spanning the period 2000:Q1 to 2019:Q4.[3] In our baseline model, the main variables of interest are the household savings rate and firm investment rate. These two measures are often compared, and the data are readily available from Eurostat. To provide robustness checks, we also consider other proxies like household consumption to GDP and non-financial corporations (NFCs) gross fixed capital formation (GFCF) to GDP. Summary statistics are presented in Table 1.

We assign to the treated group every country that implements prudential rules to reduce banks' exposure to household and firm

[3]Our initial data set comprises the 27 member states of the European Union plus the United Kingdom. We exclude Cyprus, Malta, Lithuania, and Luxembourg due to severe swings in savings and investment. Additionally, Bulgaria, Greece, and Romania are removed because data are missing in the pre-treatment or post-treatment periods.

Table 1. Summary Statistics

Sample	Obs.	Mean	Std. Dev.	Min.	Max.
HH Savings Rate: Full Sample					
HH Savings Rate	1200	0.11	0.04	0.00	0.20
Dependency Ratio	1200	0.26	0.04	0.15	0.36
GDP per Capita	1200	7.34	3.11	1.17	18.58
Firm Investment Rate: Full Sample					
Firm Investment Rate	1200	0.24	0.06	0.09	1.25
Dependency Ratio	1200	0.26	0.04	0.15	0.36
GDP per Capita	1200	7.37	3.07	1.05	18.58
HH Savings Rate: Control Group (Full)					
HH Savings Rate	760	0.12	0.04	0.00	0.20
Dependency Ratio	760	0.26	0.04	0.18	0.36
GDP per Capita	760	6.49	2.53	1.17	11.23
HH Savings Rate: Not Yet Treated Countries					
HH Savings Rate	360	0.09	0.03	0.00	0.18
Dependency Ratio	360	0.24	0.03	0.18	0.34
GDP per Capita	360	4.94	2.66	1.17	9.32
Firm Investment Rate: Control Group (Full)					
Firm Investment Rate	818	0.24	0.05	0.12	0.41
Dependency Ratio	818	0.26	0.04	0.18	0.36
GDP per Capita	818	6.51	2.54	1.05	11.23
Firm Investment Rate: Not Yet Treated Countries					
Firm Investment Rate	418	0.26	0.05	0.12	0.41
Dependency Ratio	418	0.25	0.04	0.18	0.34
GDP per Capita	418	5.03	2.56	1.05	9.24

Note: The table reports summary statistics of pre-treatment variables for control and treatment groups. The full control group includes never treated and not yet treated countries. The statistics are further disaggregated into a subset of not yet treated countries that have not yet adopted MaPP during the sample period.

risks. This includes the loan-to-value (LTV) ratio, debt-service-to-income (DSTI) ratio, and loan restrictions.[4] The data were collected from the International Monetary Fund (IMF) iMaPP database (Alam et al. 2019) and updated with information from the European Central Bank's Macroprudential Bulletins.

Most countries in our sample end up implementing MaPP at some point in time. This may raise concerns about the size and heterogeneity of our control group. For example, our control group may comprise only low-income or high-income economies with very different characteristics to an average country. As such, we force the control group in the baseline model to include countries that have "not yet" implemented MaPP. This increases the size of the control group at the expense of treatment heterogeneity. The remaining models use alternative specifications for the control group, the estimation method, and the aggregation method.

Although EU countries are fairly homogeneous, there could be covariate-specific trends in aggregate demand across groups. In particular, the literature on the secular drivers of savings suggests that demographics and inequality could influence private spending.[5] To account for these factors, we run alternative specifications of our model including the dependency ratio and GDP per capita. Detailed descriptions of every variable are available in Appendix A.

4. Results

4.1 Baseline Results

Table 2 presents the estimates for the impact of MaPP on household savings. The bulk of the results indicate that MaPP leads to a surge in savings. In the baseline model, the group-time treatment effect of MaPP increases savings by 1.94 percentage points. This impact is surprisingly consistent across models, ranging from 1.87 to 2.41 percentage points. To be clear, this estimate for savings is not a small number. The average savings rate in our sample is only 11.33 percent,

[4] In a similar spirit to Lim et al. (2011), we focus on loan-targeted MaPP.

[5] Other potential drivers of savings may include government debt and unemployment. However, these factors are not explicitly included as control variables in the DiD because they also influence the choice of MaPP and could render our results invalid. In any case, GDP per capita should partially capture these effects.

Table 2. Impact of MaPP on Household Savings Rate, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0194**	0.0232**	0.0241**	0.0187**
Standard Error	0.0053	0.0060	0.0062	0.0059
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0212**	0.0119	0.0302**	0.0363**
Standard Error	0.0080	0.0085	0.0153	0.0180
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0224	0.0293*	0.0332*	0.0064
Standard Error	0.0171	0.0167	0.0197	0.0221

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on savings across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on savings over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

which means that MaPP pushes savings up by approximately one quarter.

The vast majority of group-time treatment effects are statistically significant, and they stand up to robustness checks that control for demographics and income. They also hold when we restrict the control group to "never treated" countries. When we control for demography and income, the group-time treatment effects are based on the assumption that only countries with similar dependency ratios and GDP per capita would follow a similar trend in savings in the absence of MaPP. These conditional results indicate that MaPP leads to a rise in savings of 2.12–2.41 percentage points. Altogether, both unconditional and conditional results suggest that households increase their savings over and above what they would have in the absence of MaPP.

An interesting question is whether the impact of MaPP on savings is more profound in the short or in the long run. This can be assessed by examining the dynamic effects of MaPP. Our results show that the impact on savings gets stronger as countries are exposed to MaPP for longer. When we consider the length of exposure, the dynamic impact of MaPP on savings ranges from 2.93 to 3.63 percentage points. This can be visually inspected in Figure 1, which depicts the dynamic impact of MaPP on savings under the assumption of unconditional parallel trends. The dynamic effect on savings remains positive across time and becomes stronger as the length of exposure to MaPP increases, especially after year 3.

The uncanny finding that MaPP has a lower impact in the short run is not entirely new. A couple of papers using regressions to estimate the effects of MaPP over a four-year window report a severe contraction in credit and output around year 3 (e.g., Borio and Shim 2007; Richter, Schularick, and Shim 2019). Our estimates now provide a potential explanation for the plunging output: households reduce their consumption and increase savings around year 3.

Table 3 presents our estimates for the impact of MaPP on firm investment. The group-time treatment effects of MaPP on firm investment are somewhat more modest. Most of our models report a positive and significant impact on investment, but a few of the estimates are not statistically significant. This is not unexpected because firms often plan their investments in advance. It may be the case that the effects of MaPP are mostly dynamic.

Figure 1. Dynamic Impact of MaPP on Household Savings Rate, Full Sample, 2000–19

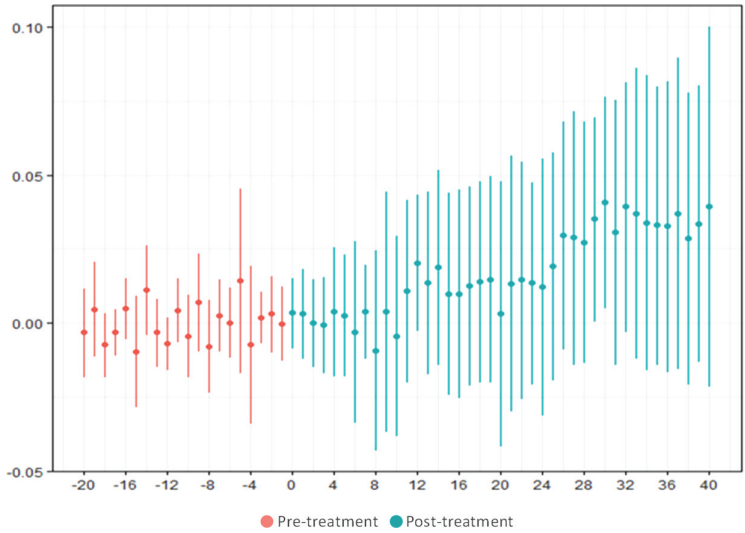

Note: The red dots are the pre-treatment pseudo group-time average treatment effects, which are plotted to pre-test the parallel trends assumption. The blue dots are the post-treatment group-time average treatment effects, which measure the average effect of adopting MaPP in quarter nth for all groups that implement MaPP in that quarter. The x-axis is the length of exposure to MaPP. A length of exposure equal to 0 corresponds to the average effect of implementing MaPP across groups in the first quarter after the adoption of MaPP; equal to –1 corresponds to the quarter before groups implement MaPP; and equal to 1 corresponds to the first quarter after initial adoption.

Not surprisingly, the results for the dynamic effects on investment are far more enlightening. Our estimates show that MaPP increases firm investment by 5.05–6.63 percentage points over time. This corresponds to an increase of more than one quarter in firm investment, since the average investment rate in our sample is 23.94 percent. As before, Figure 2 displays the dynamic effects of MaPP on investment. This figure forcibly shows that the short-run effects on investment are relatively mild but they build up significantly with time. The firm investment picks up around year 4, which suggests that firms are slower to adjust to MaPP than households. The dynamic effects on investment peak at 6.63 percentage points when we control for differences in income across countries.

Table 3. Impact of MaPP on Firm Investment Rate, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0074**	0.0006	0.0062	0.0068**
Standard Error	0.0030	0.0034	0.0050	0.0021
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0021	0.0165**	0.0591**	0.0505*
Standard Error	0.0031	0.0053	0.0260	0.0274
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0534**	0.0575**	0.0525**	0.0663**
Standard Error	0.0231	0.0239	0.0238	0.0262

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on investment across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on investment over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Figure 2. Dynamic Impact of MaPP on Firm Investment Rate, Full Sample, 2000–19

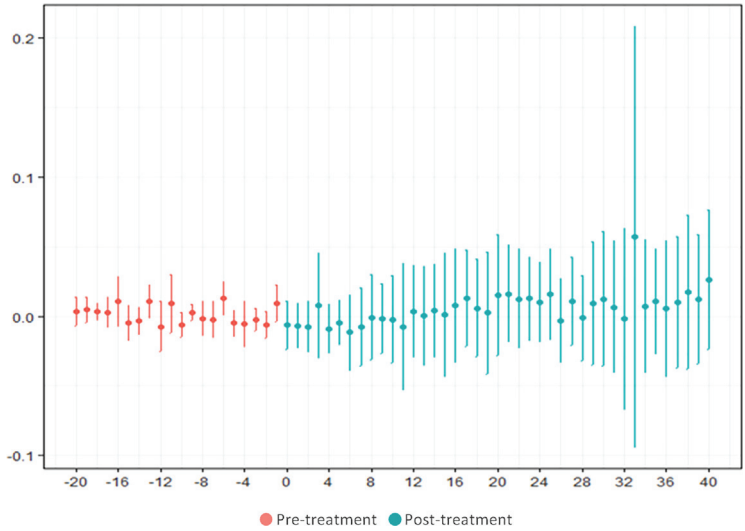

Note: The red dots are the pre-treatment pseudo group-time average treatment effects, which are plotted to pre-test the parallel trends assumption. The blue dots are the post-treatment group-time average treatment effects, which measure the average effect of adopting MaPP in quarter nth for all groups that implement MaPP in that quarter. The x-axis is the length of exposure to MaPP. A length of exposure equal to 0 corresponds to the average effect of implementing MaPP across groups in the first quarter after the adoption of MaPP; equal to –1 corresponds to the quarter before groups implement MaPP; and equal to 1 corresponds to the first quarter after initial adoption.

In summary, our results suggest that MaPP increases firm investment at the expense of household consumption. The immediate effects of MaPP are relatively modest, but they pick up in the long run. If we make a crude comparison between our estimates for savings and investment, firm investment increases twice as much as the decrease in household consumption. These results point to a weaker macroeconomic cost of MaPP than reported in previous papers.

4.2 Individual Policy Tools

In this section, we disaggregate the effects of each MaPP tool on aggregate demand. We hope to cast light on the tools that have

the greatest impact on private spending. To study individual policy choices, we ensure that the treated countries have not yet implemented another MaPP at the time of treatment. We also examine only countries that implement one household- or firm-targeted policy to isolate the impact of this policy choice.

Tables 4–9 provide the estimations of the impact of LTV ratios, loan restrictions, and DSTI ratios on household savings and firm investment. In Tables 4 and 7, we see that the implementation of LTV ratios results in a 3.39 and 3.83 percentage point increase in savings and investment, respectively. When looking at the dynamic effects, we find that implementing an LTV ratio pushes savings and investment up by 4.25 and 8.17 percentage points over time, respectively. The dynamic impact is statistically significant in all cases. In the case of group-time treatment effects, we lose some statistical significance for savings when we condition on GDP per capita. Still, the impact remains positive, which is consistent with our main findings.

Tables 5 and 8 provide the estimates for the impact of loan restrictions on savings and investment, respectively. In short, we find that loan restrictions have little impact on savings and investment. A possible explanation for this result is that loan restrictions are more likely to affect emerging economies than advanced economies. This is because loan restrictions usually target foreign-currency lending, certain types of liabilities, and excessive leverage (e.g., Cerutti, Claessens, and Laeven 2017).

Tables 6 and 9 show the impact of the DSTI ratio on savings and investment, respectively. The results indicate that the adoption of a DSTI ratio spurs savings by 7.64 percentage points. Over time, the dynamic effects of the DSTI ratio result in a marked increase in savings of 4.74–8.74 percentage points. Of course, one should interpret these numbers with caution because inference is based on small treated groups. But if our results are more than chance, the DSTI ratio has serious consequences for households. It literally pushes savings up and sends consumption sharply downward. Interestingly, the impact of the DSTI ratio on investment is nearly zero in our baseline model. If we condition on the dependency ratio and the GDP per capita, we obtain mixed results but the impact on investment is always relatively meager. The DSTI ratio, then, has a greater impact on consumption than on investment.

Table 4. Impact of MaPP on Household Savings Rate, LTV, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0339**	0.0413**	0.0153	0.0331**
Standard Error	0.0065	0.0050	0.0120	0.0072
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0411**	0.0079	0.0425**	0.0528**
Standard Error	0.0068	0.0163	0.0217	0.0189
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0086	0.0418*	0.0529**	0.0474**
Standard Error	0.0243	0.0221	0.0196	0.0183

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on savings across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on savings over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Table 5. Impact of MaPP on Household Savings Rate, Loan Restrictions, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0064	0.0033	0.0243	0.0066
Standard Error	0.0090	0.0178	0.0396	0.0090
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0035	0.0260	0.0055	0.0002
Standard Error	0.0173	0.0360	0.0106	0.0239
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0264	0.0056	0.0004	0.0281
Standard Error	0.0512	0.0108	0.0247	0.0477

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on savings across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on savings over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Table 6. Impact of MaPP on Household Savings Rate, DSTI, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0764**	0.0849**	0.0874**	0.0764**
Standard Error	0.0076	0.0169	0.0122	0.0071
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0849**	0.0874**	0.0764**	0.0849**
Standard Error	0.0225	0.0202	0.0076	0.0169
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0874**	0.0764**	0.0849**	0.0474**
Standard Error	0.0122	0.0071	0.0225	0.0202

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on savings across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on savings over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Table 7. Impact of MaPP on Firm Investment Rate, LTV, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0383**	0.0350**	0.0376**	0.0386**
Standard Error	0.0034	0.0041	0.0041	0.0037
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0351**	0.0385**	0.0817**	0.0787**
Standard Error	0.0051	0.0035	0.0243	0.0219
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0828**	0.0820*	0.0787**	0.0835**
Standard Error	0.0281	0.0232	0.0217	0.0285

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on investment across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on investment over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Table 8. Impact of MaPP on Firm Investment Rate, Loan Restrictions, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0045	0.0054	0.0026	0.0045
Standard Error	0.0039	0.0048	0.0034	0.0042
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0054	0.0026	0.0045	0.0054
Standard Error	0.0037	0.0035	0.0039	0.0048
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0026	0.0045	0.0054	0.0026
Standard Error	0.0034	0.0042	0.0037	0.0035

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on investment across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on investment over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Table 9. Impact of MaPP on Firm Investment Rate, DSTI, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0059	–0.0189*	0.0394**	0.0060
Standard Error	0.0052	0.0101	0.0028	0.0046
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	–0.0192*	0.0389**	0.0088	–0.0146
Standard Error	0.0113	0.0029	0.0093	0.0162
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0299**	0.0090	–0.0148	0.0294**
Standard Error	0.0131	0.0107	0.0167	0.0138

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on investment across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on investment over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Our bottom-line result is that measures that directly restrict access to credit—mainly the LTV and the DSTI ratio—have a stronger impact on the spending components of aggregate demand. But it is worth noting that their impact is strikingly different: while the LTV ratio affects both household consumption and firm investment, the DSTI ratio has a greater effect on household consumption.

4.3 Robustness Checks

Our attempt to establish robustness takes two tacks. First, we test if our results are robust to alternative proxies for the dependent variables. In doing so, we provide reassuring evidence on the validity of our results. Second, we check if our results hold when we restrict the sample to include only countries that never loosen or remove MaPP. This addresses the main limitation of the staggered DiD, which assumes that countries that adopt MaPP will never reverse these policies.

4.3.1 Alternative Dependent Variables

A potential concern with our analysis is that we only look at the effects of MaPP using a single measure of household savings and firm investment. To explicitly address this issue, we rerun the DiD on alternative proxies of spending, particularly on household consumption to GDP and NFC GFCF to GDP.

Tables 10 and 11 provide the estimations of the DiD using these alternative proxies. We find that household consumption is 0.67 percentage point lower than what it would have been in the absence of MaPP. The dynamic effects on consumption raise this number to 2.20–2.83 percentage points over time. These estimates are in the same ballpark as the ones obtained earlier. Some of these results have less statistical significance, but they all point to a negative impact of MaPP on consumption in the short and long run. This ties in with our previous finding that household savings rise sharply in response to MaPP.

The impact on GFCF is also similar to before. Once again, the group-time average treatment effects on investment are close to zero, but we do find strong dynamic effects. As the length of exposure to MaPP increases, the impact on GFCF becomes more pronounced.

Table 10. Impact of MaPP on Household Consumption to GDP, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	–0.0067*	–0.0084**	–0.0104**	–0.0050
Standard Error	0.0036	0.0033	0.0037	0.0033
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	–0.0053*	–0.0037	–0.0247**	–0.0283**
Standard Error	0.0030	0.0069	0.0108	0.0117
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	–0.0156	–0.0220**	–0.0232**	–0.0037
Standard Error	0.0137	0.0109	0.0110	0.0169

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on consumption to GDP across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on consumption to GDP over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Table 11. Impact of MaPP on NFC GFCF to GDP, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0036	0.0003	0.0038	0.0039
Standard Error	0.0025	0.0022	0.0038	0.0027
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0020	0.0103**	0.0434**	0.0395**
Standard Error	0.0022	0.0035	0.0138	0.0176
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0428**	0.0435**	0.0417**	0.0498**
Standard Error	0.0143	0.0175	0.0141	0.0167

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on NFC GFCF to GDP across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on NFC GFCF to GDP over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

We estimate that the impact on investment can be as high as 3.95–4.98 percentage points. Perhaps more interestingly, we again find that investment only picks up three years after the adoption of MaPP. These results suggest that banks may need some time to adjust to new regulation or that agents may not be as forward-looking as previously thought. The underlying causes of the surge in investment are hard to pinpoint in our analysis, but one thing is clear: MaPP has a lower impact on investment in the short run than in the long run. This result is statistically significant and holds across all our model specifications. This is interesting because MaPP is often implemented in response to a shock. But perhaps this is already too late.

4.3.2 Restricted Sample

A caveat of our staggered DiD is the assumption that once a country is treated, it will remain treated throughout the sample period. This could bias the results if countries reverse MaPP at some point in time. To account for this possibility, we rerun our models using a restricted sample that includes only countries that never loosen or remove MaPP.[6]

Tables 12 and 13 show the estimations of the DiD using the restricted sample. In general, our results continue to hold throughout. They show that, on average, households save between 1.96 and 2.74 percentage points more than they would have in the absence of MaPP. The dynamic impact suggests an increase of up to 4 percentage points over time, particularly after three years of exposure to MaPP, as shown in Figure 3. This matches all our previous findings.

Turning to investment, one gratifying result is that our estimates become more statistically significant in the restricted sample. The group-time average treatment effect is now statistically significant in half of our models, with an estimated impact of 0.62–1.14 percentage points. The dynamic effects of MaPP indicate that investment increases by 6.78–8.22 percentage points, which is only slightly higher than the impact found in the main results. When we plot the

[6]Only Denmark, the Netherlands, and Poland remove or loosen a prudential rule during our sample period. These countries are excluded from this restricted sample.

Table 12. Impact of MaPP on Household Savings Rate, Restricted Sample, 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0204**	0.0274**	0.0238*	0.0196**
Standard Error	0.0058	0.0065	0.0072	0.0063
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0270**	0.0096	0.0400**	0.0530**
Standard Error	0.0078	0.0068	0.0184	0.0145
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0313**	0.0389**	0.0510**	0.0106
Standard Error	0.0160	0.0183	0.0152	0.0198

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on savings across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on savings over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Table 13. Impact of MaPP on Firm Investment Rate, Restricted Sample 2000–19

Model	**(1)**	**(2)**	**(3)**	**(4)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Not Yet Treated	Not Yet Treated	Never Treated
Estimation Method	Doubly Robust	Regression	Regression	Doubly Robust
Aggregation Method	Group	Group	Group	Group
Covariates	—	Dependency Ratio	GDP per Capita	—
ATT	0.0112**	0.0091**	0.0062*	0.0096**
Standard Error	0.0026	0.0032	0.0032	0.0015
Model	**(5)**	**(6)**	**(7)**	**(8)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Never Treated	Never Treated	Not Yet Treated	Not Yet Treated
Estimation Method	Regression	Regression	Doubly Robust	Regression
Aggregation Method	Group	Group	Dynamic	Dynamic
Covariates	Dependency Ratio	GDP per Capita	—	Dependency Ratio
ATT	0.0089**	0.0114*	0.0822**	0.0770**
Standard Error	0.0021	0.0037	0.0237	0.0196
Model	**(9)**	**(10)**	**(11)**	**(12)**
Period	2000–19	2000–19	2000–19	2000–19
Treatment Group	Not Yet Treated	Never Treated	Never Treated	Never Treated
Estimation Method	Regression	Doubly Robust	Regression	Regression
Aggregation Method	Dynamic	Dynamic	Dynamic	Dynamic
Covariates	GDP per Capita	—	Dependency Ratio	GDP per Capita
ATT	0.0678**	0.0776**	0.0748**	0.0793**
Standard Error	0.0251	0.0237	0.0216	0.0257

Note: The table reports the aggregated group treatment effect ($ATT(g,t)$) parameters estimated as in Equation (7) to evaluate the impact of MaPP on investment across groups. The aggregated dynamic treatment effect parameters estimated as in Equation (9) are also reported to examine the impact of MaPP on investment over time. $ATT(g,t)$ is the average treatment effect experienced by group g in time t. Statistical significance is assessed using clustered bootstrapped standard errors at the country level that also account for the autocorrelation of the data. ** represents statistical significance at a 5 percent level and * represents statistical significance at a 10 percent level.

Figure 3. Dynamic Impact of MaPP on Household Savings Rate, Restricted Sample, 2000–19

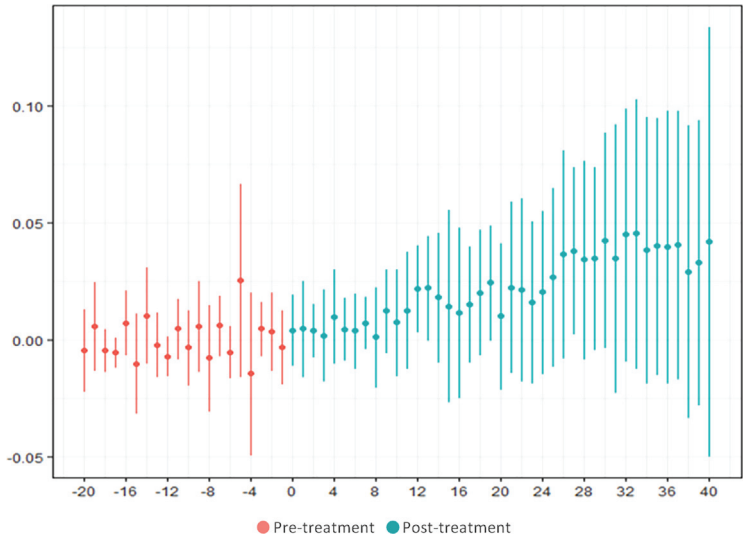

Note: The red dots are the pre-treatment pseudo group-time average treatment effects, which are plotted to pre-test the parallel trends assumption. The blue dots are the post-treatment group-time average treatment effects, which measure the average effect of adopting MaPP in quarter nth for all groups that implement MaPP in that quarter. The x-axis is the length of exposure to MaPP. A length of exposure equal to 0 corresponds to the average effect of implementing MaPP across groups in the first quarter after the adoption of MaPP; equal to –1 corresponds to the quarter before groups implement MaPP; and equal to 1 corresponds to the first quarter after initial adoption.

dynamic effects in Figure 4, we can see that GFCF starts to rise significantly after year 4, which is also in line with the pattern of the firm investment rate. This is reassuring despite the fact that we use a different sample.

5. Conclusions

In this paper, we have investigated the causal effects of MaPP on aggregate demand. Using a novel DiD approach with staggered adoption, we find that MaPP reduces household consumption in the short and long run, while increasing firm investment in the long run. These

Figure 4. Dynamic Impact of MaPP on Firm Investment Rate, Restricted Sample, 2000–19

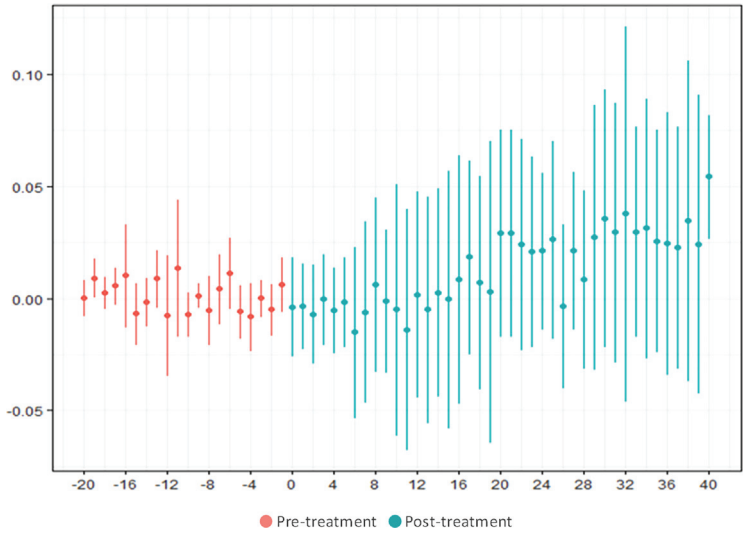

Note: The red dots are the pre-treatment pseudo group-time average treatment effects, which are plotted to pre-test the parallel trends assumption. The blue dots are the post-treatment group-time average treatment effects, which measure the average effect of adopting MaPP in quarter nth for all groups that implement MaPP in that quarter. The x-axis is the length of exposure to MaPP. A length of exposure equal to 0 corresponds to the average effect of implementing MaPP across groups in the first quarter after the adoption of MaPP; equal to –1 corresponds to the quarter before groups implement MaPP; and equal to 1 corresponds to the first quarter after initial adoption.

results clarify a point that is too often overlooked in the literature: consumption and investment are important transmission channels through which MaPP affects growth.

But why does MaPP hinder consumption and facilitate investment? We believe the answer lies in the type of financial constraint imposed by MaPP. When we look at individual MaPP tools, we find that LTV and DSTI ratios have a deleterious impact on household consumption even though the LTV ratio has a positive effect on firm investment. A potential explanation for this result is that LTV ratios target mainly home loans. This is likely to foster financial stability, which may lead to a surge in investment at the cost of lower

consumption. Finally, we find little evidence that loan restrictions affect aggregate demand, at least, in advanced economies.

Some limitations of our model point to potential research opportunities. First, the staggered DiD assumes that a country becomes forever treated after implementing MaPP. An unintended consequence is that we cannot fully capture the effects of loosening, tightening, or removing MaPP. We address this caveat by rerunning our model on a restricted sample of countries that never loosen or remove MaPP. But it would also be interesting to assess how changes in the overall macroprudential stance affect consumption and investment.

Second, the staggered DiD cannot completely disentangle the effects of MaPP from other country-specific macroeconomic events. This is because MaPP is usually adopted in response to contemporaneous events. That said, our aggregation of average treatment effects into a single causal parameter reduces estimation uncertainty. Furthermore, our results are robust to a number of different dependent variables, samples, and controls. Future work could investigate how the interaction of MaPP with other policies influences aggregate demand.

Lastly, our results provide suggestive evidence that some MaPP tools have a disproportionally high impact on aggregate demand. Yet, we are unable to fully disaggregate the effects of individual tools because most of the countries in our sample implement DSTI ratios in conjunction with LTV ratios or loan restrictions. Understanding how the design of MaPP influences aggregate demand remains a potentially fruitful area for research.

In spite of these caveats, our results offer useful policy guidance. An important finding is that MaPP has a weaker macroeconomic cost than previously suggested in the literature. If left unattended, MaPP can have pernicious effects on consumption; but if properly managed, MaPP can also lead to higher investment over time. The overall macroeconomic impact, then, depends on a country's policy objectives. If private consumption is in a free fall, MaPP may aggravate the consequences for households, particularly if countries implement LTV and DSTI ratios. But if private consumption is relatively stable, then MaPP can be an effective tool to restore financial stability and boost investment in the long run.

Another important finding is that the effects of MaPP on aggregate demand seem to only gain traction after three years. This

finding is interesting because MaPP is usually tightened in response to a crisis, but our results suggest that this is already too late. Indeed, MaPP may only send demand downward at the height of the crisis. Instead, our results support the view that policymakers should continuously adjust MaPP in much the same way as monetary policy. But given that MaPP has the ability to drive spending, policymakers should be cautious about using it liberally.

Appendix A. Data Sources and Details

Table A.1. Data Sources and Details

Variable	Type	Source	Details
NFC GFCF to GDP	Dependent	Eurostat	GFCF consists of resident producers' acquisitions, less disposals of fixed assets plus certain additions to the value of non-produced assets realized by productive activity, such as improvements to land. Seasonally and calendar adjusted. Quarterly data.
Firm Investment Rate	Dependent	Eurostat	Gross fixed capital formation (P51) divided by gross value-added (B1G) of NFC. Seasonally and calendar adjusted. Quarterly data.
Household (HH) Savings Rate	Dependent	Eurostat	Gross saving (B8G) divided by gross disposable income adjusted for changes in pension entitlements (B6G + D8net). Seasonally and calendar adjusted. Quarterly data.
HH Consumption to GDP	Dependent	Eurostat	Private consumption expenditure consists of expenditure incurred for the direct satisfaction of individual or collective needs by private households or non-profit institutions serving households. Seasonally and calendar adjusted. Quarterly data.
Dependency Ratio	Control	Eurostat	Ratio between the number of persons aged 65 and over (age when they are generally economically inactive) and the number of persons aged between 15 and 64. This indicator is published annually, and it was assumed constant for all quarters within the year.
GDP per Capita	Control	ECB MPDB	Gross domestic product at market prices. Million euros. Seasonally and calendar adjusted. Divided by total population. Total population is published annually, and it was assumed constant for all quarters within the year. Quarterly data.
HH Loan Restrictions	Policy	IMF iMaPP	Household loan restrictions include mainly loan limits and may be conditioned on loan characteristics like the maturity, the size, the type of interest rate, and the LTV ratio. Index cumulated to a quarterly frequency.
NFC Loan Restrictions	Policy	IMF iMaPP	Firm loan restrictions include mainly loan limits and may be conditioned on loan characteristics like the maturity, the size, the type of interest rate, and the LTV ratio. Index cumulated to a quarterly frequency.
LTV Ratio	Policy	IMF iMaPP	Limits to the loan-to-value ratios, including those mostly targeted at housing loans, but also those targeted at automobile loans, and commercial real estate loans. Index cumulated to a quarterly frequency.
DSTI Ratio	Policy	IMF iMaPP	Limits to the debt-service-to-income ratio and the loan-to-income ratio, which restrict the size of debt services or debt relative to income. They include those targeted at housing loans, consumer loans, and commercial real estate loans. Index cumulated to a quarterly frequency.

Appendix B. Unit-Root Tests

Table B.1. Unit-Root Tests

	Observations	Statistic	P-Value
Maddala and Wu (1999) Test			
HH Savings Rate	1200	193.751	0.000
HH Consumption to GDP	1600	129.887	0.000
Firm Investment Rate	1200	236.652	0.000
NFC GFCF to GDP	1280	212.983	0.000
Dependency Ratio	1600	26.644	0.948
HH GDP per Capita	1600	9.896	0.000
Pesaran (2007) Test			
HH Savings Rate	1200	–6.821	0.000
HH Consumption to GDP	1600	–7.454	0.000
Firm Investment Rate	1200	–8.984	0.000
NFC GFCF to GDP	1280	–8.542	0.000
Dependency Ratio	1600	–0.502	0.308
HH GDP per Capita	1600	–2.335	0.000

Note: The table presents the first-generation Maddala and Wu (1999) test and second-generation Pesaran (2007) test for panel unit roots results based on H_o: All panels contain unit roots and H_a: At least one panel is stationary. The results of an inverse chi-squared test are presented above with both the test statistic and the p-value being displayed. The presence of a unit root is always rejected because the p-value is less than 0.1 except for the case of the dependency ratio.

Appendix C. Cramer-von-Mises Tests

Table C.1. Cramer-von-Mises Tests

	HH Savings Rate		
Covariates	**No Covariates**	**Dependency Ratio**	**GDP per Capita**
CvM Test Statistic	0.0673	0.0451	0.0560
CvM Critical Value	0.3253	0.3501	0.4912
CvM P-Value	0.8400	0.9760	0.9210
	Firm Investment Rate		
Covariates	**No Covariates**	**Dependency Ratio**	**GDP per Capita**
CvM Test Statistic	0.1372	0.1397	0.1586
CvM Critical Value	0.4527	0.7462	1.0857
CvM P-Value	0.8190	0.9640	0.9460

Note: The table presents the CvM test for the presence of (un)conditional parallel pre-trends based on H_o: (Un)conditional parallel pre-trends hold and H_a: (Un)conditional parallel pre-trends do not hold. The results of the Wald-type test are presented above with both the test statistic and the p-value being displayed. Note that we always fail to reject the presence of parallel trends, as the p-value is greater than 0.10.

Appendix D. Control and Treated Groups in the DiD

Table D.1. Control and Treated Groups in the DiD

Variable	**Treated Group**	**"Never" Treated Control Group**
HH Savings Rate	Czech Republic Denmark Finland Hungary Ireland The Netherlands Poland Portugal Sweden United Kingdom	Austria Belgium France Germany Italy

(*continued*)

Table D.1. (Continued)

Variable	Treated Group	"Never" Treated Control Group
HH Consumption to GDP	Croatia Czech Republic Denmark Estonia Finland Hungary Latvia The Netherlands Poland Portugal Slovakia Slovenia Sweden United Kingdom	Austria Belgium France Germany Italy
Firm Investment Rate	Czech Republic Denmark Estonia Finland France The Netherlands Poland Portugal Sweden	Austria Belgium Germany Italy United Kingdom
NFC GFCF to GDP	Czech Republic Denmark Estonia Finland France Hungary The Netherlands Poland Portugal Sweden	Austria Belgium Germany Italy United Kingdom

Note: List of countries in the control and treated groups for the DiD estimations on household consumption and firm investment. A country is assigned to the treatment group if it implements MaPP at some point in time in the sample period. A country is assigned to the "never treated" control group if it "never" implements MaPP in the sample period. An important point to note is that the control group in our main models will also include countries that have "not yet" implemented MaPP at the time of implementation of MaPP for every group g.

Appendix E. MaPP Adoption

Table E.1. MaPP Adoption

Country	Date of Implementation	Policy Implemented
Croatia	2006:Q4	LTV
Czech Republic	2015:Q2	LTV, Household Loan Restrictions
Denmark	2003:Q2	Household Loan Restrictions, NFC Loan Restrictions
Estonia	2015:Q1	LTV, Household Loan Restrictions, DSTI
Finland	2010:Q1	LTV
France	2018:Q3	NFC Loan Restrictions
Hungary	2010:Q1	LTV, Household Loan Restrictions, DSTI
Ireland	2001:Q4	LTV
Latvia	2007:Q1	NFC Loan Restrictions
Netherlands	2007:Q1	DSTI
Poland	2006:Q4	Household Loan Restrictions
Portugal	2018:Q3	LTV, Household Loan Restrictions, DSTI
Slovakia	2014:Q4	LTV
Slovenia	2016:Q3	LTV, DSTI
Sweden	2004:Q3	LTV
United Kingdom	2009:Q1	Household Loan Restrictions

Note: Date of first implementation of MaPP for every country in our sample and brief description of the policy.

References

Acharya, V. V., K. Bergant, M. Crosignani, T. Eisert, and F. McCann. 2020. "The Anatomy of the Transmission of Macroprudential Policies." IMF Working Paper No. 20/58.

Akinci, O., and J. Olmstead-Rumsey. 2018. "How Effective are Macroprudential Policies? An Empirical Investigation." *Journal of Financial Intermediation* 33 (January): 33–57.

Alam, Z., A. Alter, J. Eiseman, G. Gelos, H. Kang, M. Narita, E. Nier, and N. Wang. 2019. "Digging Deeper — Evidence on the Effects of Macroprudential Policies from a New Database." IMF Working Paper No. 19/66.

Angelini, P., S. Neri, and F. Panetta. 2014. "The Interaction between Capital Requirements and Monetary Policy." *Journal of Money, Credit and Banking* 46 (6): 1073–1112.

Behncke, S. 2020. "Effects of Macroprudential Policies on Bank Lending and Credit Risks." Working Paper No. 2020-06, Swiss National Bank.

Borio, C., and I. Shim. 2007. "What Can (Macro-)Prudential Policy do to Support Monetary Policy?" BIS Working Paper No. 242.

Callaway, B., and P. H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.

Cerutti, E., S. Claessens, and L. Laeven. 2017. "The Use and Effectiveness of Macroprudential Policies: New Evidence." *Journal of Financial Stability* 28 (February): 203–24.

Claessens, S., S. R. Ghosh, and R. Mihet. 2013. "Macro-prudential Policies to Mitigate Financial System Vulnerabilities." *Journal of International Money and Finance* 39 (December): 153–85.

Crowe, C. W., D. Igan, G. Dell'Ariccia, and P. Rabanal. 2011. "How to Deal with Real Estate Booms." IMF Staff Discussion Note No. 11/02.

Defusco, A. A., S. Johnson, and J. Mondragon. 2020. "Regulating Household Leverage." *Review of Economic Studies* 87 (2): 914–58.

Dell' Ariccia, G., D. Igan, L. Laeven, and H. Tong. 2012. "Policies for Macrofinancial Stability: How to Deal with Credit Booms." IMF Staff Discussion Note No. 12/06.

Eggertsson, G., and P. Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics* 127 (3): 1469–1513.

Farhi, E., and I. Werning. 2016. "A Theory of Macroprudential Policies in the Presence of Nominal Rigidities." *Econometrica* 84 (5): 1645–1704.

Galac, T., and E. Kraft. 2011. "Macroprudential Regulation of Credit Booms and Busts — The Case of Croatia." World Bank Policy Research Working Paper No. 5772.

Guerrieri, V., and G. Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *Quarterly Journal of Economics* 132 (3): 1427–67.

Hall, R. E. 2011. "The Long Slump." *American Economic Review* 101 (2): 431–69.

Kuttner, K. N., and I. Shim. 2013. "Can Non-interest Rate Policies Stabilise Housing Markets? Evidence from a Panel of 57 Economies." BIS Working Paper No. 433 (November).

Lim, C., F. Columba, A. Costa, P. Kongsamut, A. Otani, M. Saiyid, T. Wezel, and X. Wu. 2011. "Macroprudential Policy: What Instruments and How to Use Them? Lessons from Country Experiences." IMF Working Paper No. 11/238.

Maddala, G. S., and S. Wu. 1999. "A Comparative Study of Unit Root Tests with Panel Data and a New Simple Test." *Oxford Bulletin of Economics and Statistics* 61 (S1): 631–52.

Pesaran, M. H. 2007. "A Simple Panel Unit Root Test in the Presence of Cross-section Dependence." *Journal of Applied Econometrics* 22 (2): 265–312.

Richter, B., M. Schularick, and I. Shim. 2019. "The Costs of Macroprudential Policy." *Journal of International Economics* 118 (May): 263–82.

Unanticipated and Backward-Looking: Deflations and the Behavior of Inflation Expectations*

Ryan Banerjee and Aaron Mehrotra
Bank for International Settlements

We analyze the behavior of inflation expectations during periods of deflation using a large cross-country data set of professional forecasters' expectations. We find that deflation episodes are largely unanticipated by forecasters and feature much larger forecast errors than periods of high inflation rates. Furthermore, inflation expectations become more backward-looking during deflations, especially among forecasters that are pessimistic about growth. The results have implications for a risk-management approach to monetary policy when an economy faces very low inflation.

JEL Codes: E31, E58.

1. Introduction

The widespread shift to an unusually low inflation environment after the Great Financial Crisis (GFC) of 2008–09, and Japan's earlier experience, brought policymakers' concerns about deflation back to the fore. A downward drift in inflation expectations and deflation

*The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank for International Settlements. Matthias Loerch and Alessandro Barbera provided excellent research assistance. We are grateful to co-editor Òscar Jordà and two anonymous referees for helpful comments that very much improved the paper. We thank participants at the 49th Money, Macro and Finance Conference, the 2018 Annual Conferences of the Royal Economic Society, the Western Economic Association, the Central Bank Research Association, and seminars held at the BIS and the Swiss National Bank. We also thank Marlene Amstad, Pradyumna Dash, Bill English, Andy Filardo, Petra Gerlach, Marco Lombardi, Gianni Lombardo, Roland Meeks, Elmar Mertens, Dubravko Mihaljek, Benoit Mojon, Jouchi Nakajima, Maritta Paloviita, Alon Raviv, Katja Schmidt, and James Yetman for helpful comments and suggestions.
Author e-mails: ryan.banerjee@bis.org; aaron.mehrotra@bis.org.

risks were an important rationale for the introduction of unconventional monetary policy in the euro area (e.g., Draghi 2015). And as central banks globally were again easing policy over the course of 2019, observers were making reference to deflation risks as one factor (e.g., Summers 2019). Earlier, avoiding deflation was a significant driver behind the accommodative U.S. monetary policy stance of the early 2000s, in line with a "robust optimal" approach to monetary policy to avoid very bad outcomes (e.g., Giannoni 2007).

Perhaps surprisingly, given concerns about deflation, empirical research has so far not focused on the behavior of inflation expectations during periods of falling prices.[1] Yet, expectations matter for wage and price setting and therefore for any second-round effects from falling prices. The behavior of expectations is thus likely to affect both the depth and duration of deflation (see, e.g., Fuhrer 2017, Nishizaki, Sekine, and Ueno 2014, and Williams 2009). Benhabib, Schmitt-Grohe, and Uribe (2002) show how non-fundamental revisions to expectations can contribute to the occurrence of liquidity traps. Some policy proposals for escaping a liquidity trap aim to generate expectations of higher inflation so that the real ex ante interest rate falls (e.g., Krugman 1998; Svensson 2001). Relatedly, Mertens and Williams (2018) argue that inflation expectations can become anchored at a level below the central bank's inflation target if the economy is expected to spend a significant amount of time at the zero lower bound.

In this paper, we analyze the behavior of inflation expectations during deflations. To this aim, we use surveys of professional forecasters from Consensus Economics from 1990 onwards for 42 advanced and emerging market economies. The global nature of the data set is highly relevant, as deflations have not been limited to advanced economies—30 of the 47 deflationary episodes identified in our data set occur in emerging market economies (EMEs), mostly in Asia and Central and Eastern Europe.

We first document that deflation episodes are largely unanticipated by forecasters, even at short forecasting horizons. Considering all one-year-ahead inflation forecasts made before a deflation episode yields a large median overprediction of close to 2 percentage points.

[1] Our companion paper, Banerjee and Mehrotra (2021), which focuses on inflation forecast disagreement, is an exception.

This contrasts with a median forecast error—an underprediction—of around 0.6 percentage point prior to episodes of high inflation, and close-to-zero forecast errors during times of moderate positive inflation rates. Thus, our results concur with the evidence in Ahearne et al. (2002) regarding the unanticipated nature of Japan's deflation in the 1990s, but apply for a much larger set of economies and time periods.

We then formally examine the effect of deflations on inflation expectations, using data for a large panel of forecasters. We start by examining the degree of backward-lookingness of one-year-ahead inflation expectations during deflations. This short forecast horizon is related to the frequency with which most prices and wages are adjusted and may therefore provide a timely measure of the risk that expectations could become unmoored. Our panel regressions, based on a similar specification as in Ehrmann (2015), suggest that deflations lead to increased backward-lookingness of expectations, compared with other time periods. The dependence of inflation expectations on past inflation outcomes rises by close to 50 percent during deflations compared with other periods. Notably, while estimations that do not control for deflations suggest that periods of low positive inflation rates—those with inflation rates below 1 percent—also appear to be associated with greater backward-lookingness, this finding is not robust to controlling for deflation episodes explicitly in the estimation. Thus, our results suggest that it becomes significantly more difficult to get inflation back to target when an economy is already experiencing deflation. We also show that an alternative methodology yields results that are highly similar in nature: impulse responses from local projections indicate that short-term inflation expectations become more sensitive to actual inflation in a deflationary environment.

Next, we analyze the behavior of expectations during deflations using five-year-ahead forecasts. We apply these longer-term measures of expectations to examine the anchoring of inflation expectations in two ways. First, using the same approach as for short-term expectations, we analyze whether deflations render the five-year-ahead inflation expectations backward-looking. Second, applying regressions similar to Buono and Formai (2018) and Lyziak and Paloviita (2017), we analyze whether long-term measures of inflation expectations become more sensitive to changes in

short-term expectations. Notably, both tests suggest that inflation expectations become unanchored during periods when prices are falling. This stands in stark contrast to periods of high inflation, as well as to those of low positive inflation (below 1 percent), for which we do not find evidence of unanchoring.

Our results also highlight significant heterogeneity in the expectations behavior across individual forecasters during deflations. Using forecaster-level data on short-term inflation expectations, we show that significant backward-lookingness in inflation expectations during deflations obtains especially in the sample of forecasters who are pessimistic about the short-term growth prospects of the economy. By contrast, there is less backward-lookingness in expectations for the sample of more optimistic forecasters. Moreover, during periods of moderately positive inflation rates, we find the degree of backward-lookingness to be highly similar between optimistic and pessimistic forecasters.

Our results—in particular, the changing behavior of expectations once the economy hits deflation—have implications for a risk management approach of monetary policy (see, e.g., Greenspan 2004). As it becomes increasingly difficult to return inflation to target once prices are falling, it may be prudent to preemptively react to deflation risks. In addition, optimal policy under adaptive expectations implies a stronger policy response than under rational expectations (e.g., Gaspar, Smets, and Vestin 2006).

Our paper is related to different strands of literature. It adds to the vast and expanding literature that uses surveys to analyze the behavior of inflation expectations (see, e.g., Faust and Wright 2013, Kozicki and Tinsley 2012, Mehrotra and Yetman 2018, and the references therein). Similar to our paper, various studies have analyzed the relationship between surveys of inflation expectations and inflation outcomes to examine the anchoring of expectations (e.g., Ehrmann 2015, Levin, Natalucci, and Piger 2004, Lyziak and Palovita 2017). However, to our knowledge, no empirical study has previously focused explicitly on the behavior of inflation expectations during deflations in a large sample of countries. While the empirical approach in our paper is similar to Ehrmann (2015), we make additional contributions by highlighting the role of heterogeneity across forecasters in explaining increased backward-lookingness of forecasters; analyze the behavior of both short-term and longer-term

(five-year-ahead) forecasts; and focus the analysis on deflation periods which have so far been under-researched in the literature.

The paper also contributes to empirical research on the behavior of inflation expectations during the post-GFC period. Coibion and Gorodnichenko (2015) suggest that rising inflation expectations can account for the missing disinflation during the Great Recession. Ehrmann (2015) analyzes the anchoring of expectations when inflation is below the central bank's target, while International Monetary Fund (IMF) (2016) evaluates changes in anchoring both over time and conditional on monetary policy. The Bank of Japan (2016) evaluates the behavior of various measures of inflation expectations after the introduction of quantitative and qualitative monetary easing (QQE) in April 2013. Natoli and Sigalotti (2018) propose novel techniques based on the distribution of financial market data to examine the anchoring of inflation expectations during the post-GFC period. And Kenny and Dovern (2017) use data from surveys of professional forecasters to analyze how the distribution of long-run inflation expectations has changed in the euro area after the GFC.

The paper also relates to studies on the macroeconomic record of deflations. While the experience from the Great Depression suggests that deflations have highly adverse implications for economic activity, the more benign deflations during the latter part of the 19th century provide less evidence of negative feedback loops between deflation and output (see Bordo and Filardo 2005 and Borio et al. 2015). On the other hand, Eichengreen, Park, and Shin (2017), using producer price inflation, find stronger evidence of damaging effects of deflation on growth. Burdekin and Siklos (2010) and Smith (2006) provide overviews of the empirical evidence relating to deflation, and Fisher (1933) introduces the phenomenon of debt deflation.

Finally, the paper relates to empirical studies that consider the importance of inflation dynamics in the risk-management approach for monetary policy. Kilian and Manganelli (2008) provide evidence that the Federal Reserve under Greenspan's chairmanship was responding to changes in the balance of risks in inflation and output growth, rather than changes in the conditional means of these variables. Similarly, Gnabo and Moccero (2015) find that risks to the inflation outlook were an important consideration for U.S. monetary policy in the Greenspan era, while Andrade, Ghysels,

and Idier (2015) document that the Fed's policy rate is affected by right-tail inflation risks, measured from distributions of survey expectations.

This paper is structured as follows. The next section describes the data and discusses the identification of deflation episodes in our sample. Section 3 presents stylized facts about the behavior of expectations around deflations. This is followed in Section 4 by a formal investigation of how deflation affects both short- and longer-term inflation expectations. Section 5 analyzes possible reasons behind the forecast behavior and Section 6 discusses policy implications. Finally, concluding comments are provided in Section 7.

2. Data and Identification of Deflationary Episodes

We use surveys of professional forecasters from Consensus Economics. These data are available for a relatively long history and are collected in a comparable fashion across a large number of countries, both advanced and emerging. Having a global data set is essential for the analysis, given the large number of deflation episodes in emerging market economies. Regarding the favorable forecasting performance of subjective expectations, Faust and Wright (2013) find that survey measures of inflation expectations tend to improve the forecasts that come from a large number of different forecasting models.

At the same time, we acknowledge the limitation of using surveys of professional forecasters, as the expectations of firms and households—the price and wage setters—would arguably be more economically relevant. Moreover, the inflation expectations behavior of firms and households has been shown to deviate systematically from that of professional forecasters (see, e.g., Coibion et al. 2020). However, while there are recent studies examining firms' and households' expectations behavior in individual economies (e.g., Coibion, Gorodnichenko, and Kumar 2018; D'Acunto et al. 2021), comparable data are not available for a large number of economies, in particular EMEs.

Each month, Consensus Economics polls a panel of experts from public and private economic institutions, mostly investment banks and research institutions, about their predictions for the main macroeconomic variables for the current and next calendar year.

Given that the fixed-event nature of the forecasts—expectations of inflation during a calendar year—implies a changing forecast horizon between different months, we transform the fixed-event forecasts to one-year-ahead forecasts by computing a weighted average of current and next-year forecasts. This approach has been widely used in the literature (see, e.g., Dovern, Fritsche, and Slacalek 2012; Siklos 2013). With h as the forecast horizon, the 12-month-ahead forecast is computed as

$$\hat{\pi}_{t+12|t} = \frac{h}{12}\hat{\pi}_{t+h|t} + \frac{12-h}{12}\hat{\pi}_{t+12+h|t}, \tag{1}$$

where $\frac{h}{12}$ and $\frac{12-h}{12}$ denote the weights, i.e., the shares of current and next-year forecasts in the forecast period.

For some countries in the sample, longer-term survey forecasts are also available. The longer-term surveys are conducted and published twice a year. The data are not available at the forecaster level, as only the mean forecasts are published. We use these forecasts in parts of the analysis.

Our data cover 42 economies, 12 advanced and 30 emerging. The length of the data set depends on the availability of inflation forecast data. For advanced economies, the data start earliest in 1990, yielding a maximum of 319 monthly observations per country (see Table A.1 in the appendix for details). For emerging markets, the starting dates vary by region. For most countries in emerging Asia, the data start in late 1994; for Latin America in 2001; and for Central and Eastern Europe in 2007. The number of forecasters varies both across countries and within the same country across time. The average number of forecasters in the country-specific samples varies from 8 in Lithuania to 30 in the United Kingdom; in the full sample, the average number of forecasters per country is 15.

Our inflation data are for headline consumer price inflation (CPI, year on year).[2] While developments in core inflation would also be interesting, expectations data are not widely available for this measure. Moreover, for some EMEs where volatile components such as food comprise a large share of the consumption basket, developments

[2]A partial exception is the United Kingdom, where inflation refers to retail price (RPIX) inflation until 2004 and CPI inflation thereafter.

in core inflation may be less relevant than those in headline inflation (eg., Anand, Prasad, and Zhang 2015).

We focus on deflation episodes characterized by negative headline inflation rates (year on year) for at least six consecutive months. Furthermore, a country is regarded as exiting the deflation episode only in the third consecutive month of positive inflation rates that follow deflation. This classification ensures that very short bouts of negative inflation rates do not count as individual deflation episodes. Moreover, it avoids longer deflation periods being classified as several shorter ones, if they are interrupted only by one or two months of positive inflation rates.

The 47 deflation episodes identified in our sample are shown in Figure 1. Three periods with greater occurrence of deflations stand out. First, various Asian economies experienced deflation around the time of the Asian financial crisis: Hong Kong SAR, mainland China, Chinese Taipei, Singapore, and Thailand. Japan also experienced a long spell of deflation as its domestic banking crisis occurred. The second, more global, bout of deflations took place during the GFC. The third relatively widespread period of falling prices occurred in 2014–15, affecting many European countries but also some emerging economies in Asia. Over time, deflations were increasingly associated with near-zero interest rates (blue lines in Figure 1).

Overall, the deflation periods are relatively widely dispersed across countries, as in 16 countries deflation episodes occurred only once. Seventeen deflation episodes took place in advanced economies and 30 in emerging markets, while 11 occurred in countries that were part (or later became part) of the euro area. Hong Kong SAR experienced the lowest inflation outcome within a single deflation episode in the sample (−6.1 percent). On average the minimum inflation outcome across all the different deflation episodes was −1.7 percent. Table A.2 in the appendix shows details of the identified deflations, including their length and the minimum inflation outcomes and levels of expectations during these time periods.

In order to compare the behavior of expectations during deflations with other inflation environments, we construct two other dummy variables. First, we consider periods of high inflation, defining an economy to have high inflation if the headline inflation rate is above 4 percent for at least six consecutive months. Furthermore,

Figure 1. Deflation Episodes

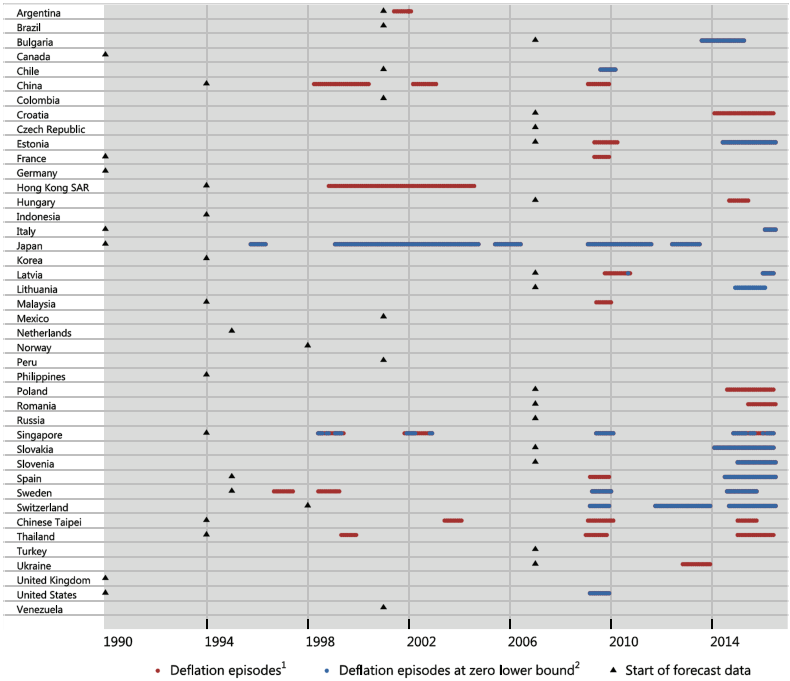

Note: [1]Negative headline consumer price inflation (CPI, year on year) for at least six consecutive months. A deflation episode ends if subsequently at least three consecutive months of positive inflation rates occur. [2]Deflation episodes occurring when policy rates are at or below 0.5 percent. If policy rate data are not available, money market interest rate data are used.

the country only exits a high-inflation episode in the third consecutive month of inflation rates below 4 percent. The high-inflation episodes using this definition are documented in Figure A.1 in the appendix. For many EMEs—in particular, in Latin America but also in other emerging economy regions—such periods cover large parts of the sample.

Moreover, in order to examine whether there are differences between low (mostly positive) inflation and deflation, we use a 1 percent inflation threshold to similarly classify episodes of low inflation. Finally, we omit as outliers all observations with inflation rates exceeding 10 percent and/or policy rates above 100 percent.

3. Descriptive Evidence of Expectations around Deflations

We find that deflations are largely unanticipated by professional forecasters. The first row of Table 1 shows median forecast errors associated with different inflation environments. It considers all forecasts (median expectation across forecasters for country c) made one year before a particular inflation environment prevailed. Thus, for deflations, the forecasts may have been made one year before a deflation episode started or, in the case of longer deflations, the economy may already have been experiencing deflation when the forecast was made. The median forecast error associated with deflations, defined as realization less forecast, is close to 2 percentage points (-1.9 percent; an overprediction). This stands in contrast with high inflations (threshold of 4 percent; see Section 2) where the median forecast error is only 0.6 percentage point (an underprediction). For low inflations which, with the threshold of 1 percent, are a superset of the deflation episodes, the forecast error is -1.3 percentage points. Thus, forecasters seem to be able to better anticipate future inflation trends during periods of higher inflation. For all other periods, which correspond to moderate positive inflation rates, the forecast error is close to zero.

Table 1 also documents forecast errors by types of deflation. To do this, we divide the deflations into those that last longer than the sample average; those that are deeper than sample average; and those where prices fall at a faster rate than sample average (where speed is defined as deepness of deflation divided by its duration). Forecast errors are particularly large for deflations that saw greater price declines (median of -3.1 percentage points; second row of Table 1) and those that featured faster price declines (-4.2 percentage points). By contrast, longer deflations tend to have lower forecast errors, possibly because forecasters have more time to adjust their expectations downwards (see also row 3 of Table 1).

We then consider similar evidence for longer-term, five-year-ahead, expectations. These surveys, conducted and published twice a year, are available during 28 of the 47 deflations identified in Figure 1 (see also Table A.2). Not surprisingly, deflations are completely unanticipated by this measure, with a median overprediction of 2.5 percentage points for forecasts made five years prior to all months

Table 1. Short-Term Forecasts for Inflation and Growth before and during Different Inflation Environments

	One-Year-Ahead Inflation Forecast Errors Prior to Different Inflation Environments			
	High Inflation	**Low Inflation**	**Deflation**	**Other**
Median Forecast Error, % Pts.	0.628	–1.318	–1.916	–0.113
	[0.518, 0.726]	[–1.381, –1.254]	[–2.059, –1.795]	[–0.150, –0.084]
Obs.	1,642	1,868	792	4,108
	One-Year-Ahead Inflation Forecast Errors Prior to Different Types of Deflations			
	Longer Deflations	**Deflations with Greater Price Declines**	**Deflations with Faster Price Declines**	
Median Forecast Error, % Pts.	–1.567	–3.130	–4.150	
	[–1.699, –1.440]	[–3.371, –2.738]	[–4.429, –3.855]	
Obs.	457	303	165	
	One-Year-Ahead Inflation Forecasts Made during Different Environments			
	High Inflation	**Low Inflation**	**Deflation**	**Other**
Median Forecast, %	5.100	0.917	0.480	2.400
	[5.000, 5.225]	[0.894, 0.953]	[0.418, 0.515]	[2.367, 2.424]
Obs.	1,848	1,898	800	4,317

(*continued*)

Table 1. (Continued)

	One-Year-Ahead Inflation Forecasts Made during Different Types of Deflations			
	Longer Deflations	**Deflations with Greater Price Declines**	**Deflations with Faster Price Declines**	
Median Forecast, %	0.331	0.368	0.661	
	[0.250, 0.434]	[0.286, 0.500]	[0.569, 0.780]	
Obs.	457	311	173	
	One-Year-Ahead GDP Growth Forecasts Made during Different Environments			
	High Inflation	**Low Inflation**	**Deflation**	**Other**
Median Forecast, %	4.417	1.838	1.385	2.808
	[4.267, 4.543]	[1.780, 1.900]	[1.257, 1.493]	[2.775, 2.850]
Obs.	1,847	1,898	800	4,317

Note: 90 percent confidence intervals are in brackets. Forecast error is defined as realization minus forecast. Forecasts are computed as the median expectation across forecasters at each point in time, for country c. The table shows the forecasts and forecast errors as medians over the months associated with the different inflation environments.

an economy spends in deflation (first row of Table 2). Again, deflations with faster price declines feature the highest absolute forecast errors, with overpredictions of 3 percentage points (second row of Table 2).

The difficulty in anticipating deflations remains if we consider short-term forecasts at the tails of the forecast distribution. Even for the "pessimistic" forecasts, i.e., those at the 10th percentile of the forecast distribution, the median one-year-ahead forecast error associated with deflations amounts to −1.5 percentage points. For the "optimistic" forecasts, i.e., those at the 90th percentile, these short-term forecasts are off the mark by −2.4 percentage points.

Consistent with evidence of the unanticipated nature of deflations, only 3 percent of the median one-year-ahead forecasts in the data set are negative. This compares with a higher share of negative inflation outcomes in the sample, at 11 percent. However, in cases where the one-year-ahead forecasts do fall into negative territory, expectations of future GDP growth also adjust downwards—the median one-year-ahead GDP growth forecast is at 0.9 percent when forecasters expect negative inflation rates, compared with a median forecast of 2.9 percent during other periods. This suggests that professional forecasters associate deflations with adverse growth outcomes.

Once deflation actually hits, professional forecasters adjust their short-term expectations downwards. The median one-year-ahead inflation forecast formed during months when an economy is already experiencing deflation amounts to 0.5 percent (third row of Table 1). This contrasts with median forecasts of 0.9 percent made during low-inflation episodes, 5.1 percent during high-inflation episodes, and 2.4 percent during all other times. Inflation expectations are lower for longer deflations, as well as those that feature greater price declines (0.3–0.4 percent; fourth row).

The downward adjustment in inflation forecasts is also reflected in GDP growth forecasts. The median one-year-ahead GDP growth forecast during deflations amounts to 1.4 percent, compared with 1.8 percent during low-inflation episodes, 4.4 percent during high-inflation episodes, and 2.8 percent during all other times (see fifth row of Table 1).

In addition to short-term expectations, the five-year-ahead expectations also adjust downward when the price level is declining.

Table 2. Five-Year-Ahead Forecasts for Inflation before and during Different Inflation Environments

	Five-Year-Ahead Inflation Forecast Errors Prior to Different Inflation Environments			
	High Inflation	**Low Inflation**	**Deflation**	**Other**
Median Forecast Error, % Pts.	1.665	–1.855	–2.516	–0.090
	[1.427, 2.080]	[–1.957, –1.735]	[–2.697, –2.300]	[–0.209, –0.007]
Obs.	153	149	53	334
	Five-Year-Ahead Inflation Forecast Errors Prior to Different Deflation Environments			
	Longer Deflations	**Deflations with Greater Price Declines**	**Deflations with Faster Price Declines**	
Median Forecast Error, % Pts.	–2.574	–2.690	–3.049	
	[–2.889, –2.294]	[–3.691, –2.293]	[–3.759, –2.683]	
Obs.	29	14	13	

(*continued*)

Table 2. (Continued)

	Five-Year-Ahead Inflation Forecasts Made during Different Environments			
	High Inflation	**Low Inflation**	**Deflation**	**Other**
Median Forecast, %	3.9	1.9	1.85	2.3
	[3.6, 4.2]	[1.9, 2]	[1.7, 2.1]	[2.3, 2.4]
Obs.	192	204	76	489
	Five-Year-Ahead Inflation Forecasts Made during Different Deflation Environments			
	Longer Deflations	**Deflations with Greater Price Declines**	**Deflations with Faster Price Declines**	
Median Forecast, %	1.7	2.3	2.6	
	[1.332, 1.868]	[1.7, 3]	[2.2, 2.844]	
Obs.	47	25	15	

Note: 90 percent confidence intervals are in brackets. Forecast error is defined as realization minus forecast. Forecasts are the mean expectation across forecasters (as reported by Consensus) at each point in time, for country c. The table shows the forecasts and forecast errors as medians over the months associated with the different inflation environments.

The long-term expectations move down by close to 0.5 percentage point during deflations, compared with periods of moderate positive inflation (Table 2, third row). This downward movement is notable, given the well-documented inertia in long-run survey measures of inflation.[3] Yet, the five-year-forecasts in our sample never reach negative territory, and they only dip below 1 percent during 1.3 percent of the observations in our sample—all corresponding to Japan's deflation and occurring at various dates between 2001 and 2012.

Notably, the short-term inflation and GDP growth forecasts become significantly more volatile once the economy is in deflation: the standard deviation (normalized by the mean) of the one-year-ahead median inflation forecast is 3.8 during deflations, compared with 0.8 during all periods. For GDP growth, the corresponding figures are 1.3 and 0.7, respectively. The increase in volatility potentially reflects greater uncertainty about future macroeconomic conditions during deflations.

4. Estimation Results

Beyond the previous informal evidence, do we observe changes in the behavior of survey forecasts when an economy is experiencing deflation? In what follows, we formally analyze the degree to which expectations depend on past inflation and compare it across different inflation environments. A similar approach to evaluate backward-lookingness has been used in various studies; see, e.g., Blanchard (2016), Ehrmann (2015), Levin, Natalucci, and Piger (2004), and Lyziak and Paloviita (2017).

4.1 Short-Term Forecasts during Deflations

We commence with one-year-ahead inflation expectations. While these are short term in nature and likely affected by current transitory shocks affecting the economy, they could provide timely measures of potential changes in the degree of anchoring. Changes

[3]For example, in the case of Canada, the five-year-ahead forecast for inflation remained unchanged at 2.0 from October 2005 to October 2008, and then again from October 2009 to October 2012.

in short-term inflation expectations may well spill over to those at longer horizons, as Buono and Formai (2018) and Lyziak and Paloviita (2017) document for the euro area, potentially affecting the credibility of monetary policy. Moreover, for the setting of monetary policy, the most relevant horizon is arguably related to the frequency with which most prices and wages are adjusted, and hence has an important impact on inflation dynamics. Indeed, Fuhrer (2017) finds that short-term expectations play a quantitatively important role for actual inflation outcomes in estimated Phillips curves for Japan and the United States.

We consider a panel fixed effects regression of the type

$$E_{i,t}\left(\pi_{c,t+12}\right) = \alpha_i + \beta_1 \pi_{c,t-1} + \beta_2 D^{infl}_{c,t-1} + \beta_3 D^{infl}_{c,t-1} \pi_{c,t-1} + \varepsilon_{i,t}, \tag{2}$$

where $E_{i,t}\left(\pi_{c,t+12}\right)$ denotes the one-year-ahead inflation expectation by forecaster i for country c, formed in period t. $\pi_{c,t-1}$ is lagged inflation. $D^{infl}_{c,t-1}$ is a dummy variable capturing a particular inflation environment; we define individual dummy variables for episodes of deflation, low inflation, or high inflation (see Section 2). These dummy variables are interacted with lagged inflation to infer whether there is evidence of increased backward-lookingness during the different inflation environments, compared with other periods.[4] α_i denote forecaster fixed effects. All data sources are given in Table A.3 in the appendix.

The dummy variables and actual inflation interacted with them are lagged so as to match the timing of the information set based on which expectations are formed. Forecasters will generally have access to some information about the previous month's inflation outcome (and whether the economy was in deflation/high inflation) at the time they submit the expectations to Consensus Economics but not for the current month. As an example, consider Germany, where forecasters would receive information about the January 2021 CPI inflation outcome in the flash inflation release by Destatis on January 28, 2021. They would then submit their inflation expectations for

[4]Such analysis is related to research that examines potential non-linearities in the Phillips curve; see, e.g., Doser et al. (2017).

Table 3. One-Year-Ahead Inflation Expectations and Deflation

Variable	(1)	(2)	(3)	(4)
$\pi_{c,t-1}$	0.525***	0.496***	0.478***	0.479***
	(0.00939)	(0.0102)	(0.0142)	(0.0138)
$D^{defl}_{c,t-1}$		0.0458		0.0395
		(0.107)		(0.110)
$D^{High\ Infl}_{c,t-1}$		–0.0403	–0.0938	–0.101
		(0.0758)	(0.0801)	(0.0803)
$D^{defl}_{c,t-1}\pi_{c,t-1}$		0.239**		0.372***
		(0.101)		(0.0949)
$D^{High\ Infl}_{c,t-1}\pi_{c,t-1}$		0.0198	0.0374*	0.0375*
		(0.0170)	(0.0192)	(0.0194)
$D^{Low\ Infl}_{c,t-1}$			–0.165***	–0.0698**
			(0.0308)	(0.0296)
$D^{Low\ Infl}_{c,t-1}\pi_{c,t-1}$			0.107***	–0.119***
			(0.0404)	(0.0201)
Obs.	124,252	124,252	124,252	124,252
R-squared	0.766	0.767	0.767	0.767

Note: Dependent variable is the expectation by an individual forecaster for one-year-ahead inflation. Robust standard errors clustered by forecaster in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. All models include forecaster fixed effects.

the current and next calendar year to the February Consensus Economics publication on the survey date of February 8, 2021, which is before the February flash release.

Equation (2) is estimated by ordinary least squares using forecaster-level data, yielding over 124,000 observations for the full sample. We use heteroskedasticity-consistent standard errors clustered by forecaster, allowing the residuals to be correlated within the same forecaster over time.

The estimates in Table 3 show that deflations render inflation expectations more backward-looking. Column 1 indicates that one-year-ahead inflation expectations are generally backward-looking in our sample—the coefficient on the lagged inflation term is 0.525 with a standard error of 0.009. This is consistent with the short-term nature of the one-year-ahead expectations. Column 2 adds dummy variables for deflation and high-inflation periods, together

with interaction terms of lagged inflation with deflation and high-inflation dummy variables, respectively. The estimates in column 2 suggest that deflations are associated with more backward-looking inflation expectations, as shown by an economically and statistically significant coefficient estimate on the interaction term between deflation and lagged inflation. The sum of the backward-looking terms amounts to $0.496 + 0.239 = 0.735$—this implies that the dependence of expectations on past inflation increases by around 50 percent during deflations, compared with normal periods. By contrast, a similar increase in backward-lookingness is not observed for high-inflation periods.

We also analyze whether longer, deeper, or faster deflations (see Section 3) have different implications for inflation expectations. To do this, we include the same interaction variable between lagged inflation and the deflation dummy as in the baseline model, together with an additional interaction variable between lagged inflation and a deflation dummy for longer/deeper/faster deflations, respectively. Table A.4 in the appendix shows that the additional interaction variables for faster and deeper deflations enter with positive and statistically significant coefficients, suggesting that such deflations feature greater backward-lookingness. However, the additional interaction variable for longer deflations enters with a negative and statistically significant coefficient. This could arise, as forecasters have time to re-adjust their expectations during longer deflations. Indeed, Table 1 showed that forecast errors associated with longer deflations are lower than those that occur when price declines are of shorter duration. In all models, the baseline interaction term between all deflations and lagged inflation remains positive and statistically significant.

An important issue is whether the results are indeed driven by deflations, or whether they reflect the effect of low inflation on inflation expectations more broadly. For example, Ehrmann (2015) finds that expectations become more backward-looking when inflation is below the central bank's target. As our sample includes a wide range of economies with various different monetary policy regimes, we use our low-inflation dummy variables to investigate this (using a positive 1 percent inflation threshold; see Section 2). Including the interaction variable associated with low inflations in column 3 of Table 3 indeed suggests that low inflations are associated with more backward-looking expectations, with a coefficient estimate roughly

half of that observed for deflations in column 2. However, this finding is not robust to the inclusion of the deflation dummy variables in the estimation—column 4 shows that when the different inflation environments are considered jointly, deflations alone appear to render inflation expectations more backward-looking in a way that is economically significant.[5]

One may argue that a 4 percent inflation threshold is too low for economies with a history of higher inflation and/or higher inflation targets and this may result in biased estimates of backward-lookingness in periods when inflation is neither high nor low. To address this concern, Table A.5 in the appendix presents results with a country-specific high-inflation threshold, defined as inflation above the mean of the country-specific sample. In these regressions, the results for backward-lookingness during deflations are similar to the baseline estimates. However, the interaction term between high inflation and lagged inflation is more statistically significant. At the same time, the high-inflation dummy variable gets a statistically significant negative sign, suggesting some "mean reversion" in expectations when an economy is experiencing higher inflation rates.[6]

Overall, the finding of greater backward-lookingness in expectations concurs with other evidence. For instance, the Bank of Japan (2016) finds a significant adaptive component of inflation expectations in the Japanese economy. The study attributes this feature in expectations formation to the prolonged deflation such that expectations have not been anchored at 2 percent. Similarly, Banca d'Italia (2017) reports that wage contracts in Italy have been increasingly linked to past rather than forecast inflation, which could make it more difficult to return inflation to target.

4.2 Evidence from Local Projection Regressions

Similar results related to the relationship between inflation outcomes and expectations are obtained with an alternative methodology. In

[5]This finding is robust to considering a dummy variable for low inflations that excludes periods when the price level is falling. The results are available upon request.

[6]The results for deflations are also robust to considering 5 percent or 6 percent as alternative high-inflation thresholds. At the same time, the high-inflation dummy and the associated interaction variables lose significance when 6 percent high-inflation thresholds are considered.

particular, impulse responses computed from local projection regressions suggest that expectations respond more strongly to changes in inflation outcomes in negative inflation environments.

To show this, we consider local projection regressions of the type[7]

$$\pi^e_{c,t+h} - \pi^e_{c,t-1} = \alpha_c + \beta_1 \Delta\pi_{c,t} + \beta_2 \pi^e_{c,t-1} + e_{c,t}, \tag{3}$$

where $\pi^e_{c,t+h}$ is the one-year-ahead inflation expectation (median across forecasters for country c, formed in period $t+h$); $\pi^e_{c,t-1}$ is the initial level of one-year-ahead inflation expectation (formed in period $t-1$); $\Delta\pi_{c,t}$ denotes the change in actual inflation between period t and $t-1$; and $e_{c,t}$ is the error term. We consider impulse responses between zero and six months from the inflation increase, given that the shortest deflation episode in our sample lasts for six months.

Figure 2, upper row, shows that the response of one-year-ahead expectations to 1 percentage point increases in inflation is higher in a deflationary environment. A 1 percentage point change in actual inflation leads, after one month, to a 0.55 percentage point change in inflation expectations in the same direction during deflations, whereas the change in expectations is 0.3 point during other periods. The peak effect on expectations is also higher during deflations, although the 90 percent confidence bands are considerably wide, reflecting the much smaller estimation samples when deflationary periods are considered. The results suggest that survey expectations respond more strongly to inflation outcomes in the deflation subsamples than during other periods.

We also apply an alternative way of identifying the deflationary episodes by considering the level of expected inflation. We consider periods where, instead of actual inflation, the median one-year-ahead inflation expectation is below zero, resulting in a still smaller subsample than in the case where actual inflation is negative. The lower row of Figure 2 displays the impulse responses, highlighting that expectations respond significantly more to movements in actual inflation during periods when forecasters expect negative inflation. When one month has passed from a change in inflation by 1 percentage point, one-year-ahead inflation expectations change by

[7]See Jordà (2005).

Figure 2. Impulse Responses to 1 Percentage Point Increases in Inflation Outcomes

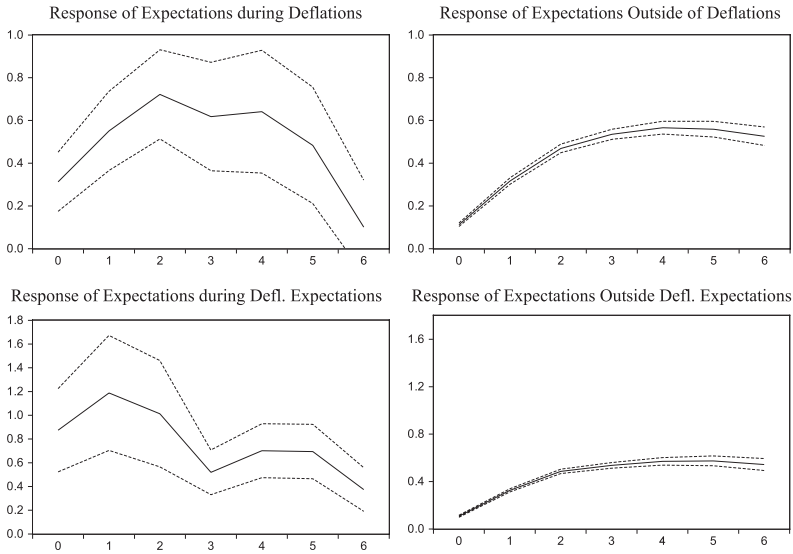

Note: Based on local projection regressions; in percentage points. The upper row shows the responses of one-year-ahead inflation expectations to 1 percentage point increases in actual inflation during (left) and outside (right) deflations. The lower row shows the responses of inflation expectations during periods when one-year-ahead expectations are below (left) or above (right) zero. The x-axis shows the number of months that have passed from the inflation increase. Dashed lines are 90 percent confidence intervals.

1.2 percentage point in the same direction when forecasters expect deflation, and by 0.3 percentage point during other periods. The corresponding peak responses over six months are 1.2 and 0.6 percentage points, respectively, again suggesting greater sensitivity of expectations to movements in actual inflation during deflationary periods.

4.3 Evidence from Longer-Term Forecasts

We then investigate the anchoring of expectations during deflations using five-year-ahead forecasts. As our first test of anchoring, we examine whether deflations result in more backward-looking longer-term expectations. To do this, we replace one-year-ahead forecasts in

Table 4. Five-Year-Ahead Inflation Expectations and Deflation

Variable	(1)	(2)	(3)	(4)
$\pi_{c,t-1}$	0.0820**	0.0521***	0.0585**	0.0582**
	(0.0336)	(0.0158)	(0.0245)	(0.0245)
$D^{defl}_{c,t-1}$		–0.0221		–0.0588
		(0.0976)		(0.103)
$D^{High\ Infl}_{c,t-1}$		–0.293	–0.272	–0.279
		(0.328)	(0.314)	(0.316)
$D^{defl}_{c,t-1}\pi_{c,t-1}$		0.0796		0.164**
		(0.0474)		(0.0678)
$D^{High\ Infl}_{c,t-1}\pi_{c,t-1}$		0.0705	0.0637	0.0646
		(0.0758)	(0.0725)	(0.0728)
$D^{Low\ Infl}_{c,t-1}$			–0.0168	0.0495
			(0.0644)	(0.0671)
$D^{Low\ Infl}_{c,t-1}\pi_{c,t-1}$			0.0352	–0.0903
			(0.0481)	(0.0648)
Obs.	880	880	880	880
R-squared	0.871	0.872	0.872	0.872

Note: Dependent variable is the mean expectation across forecasters for five-year-ahead inflation. Robust standard errors clustered by country in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. All models include country fixed effects.

Equation (2) (forecaster-level data) with their five-year-ahead equivalent (country-level data). As only the mean long-run expectation across forecasters is available at each point in time, the forecaster fixed effects are replaced by country fixed effects. The results are shown in Table 4.

We find that deflations are associated with more backward-looking expectations behavior even when the long-term forecasts are considered. First, we observe some backward-lookingness for the five-year-ahead expectations when the entire sample period is included (column 1 in Table 4). Not surprisingly, the coefficient estimate is much smaller than in the case of one-year-ahead expectations (0.082 in Table 4 versus 0.525 in Table 3). However, the increase in the dependence of longer-term expectations on past inflation is very large for deflations—in column 4, where all the different inflation environments are included, the coefficient estimate on the interaction

term between deflation and lagged inflation is around three times as high as that on lagged inflation alone. Thus, we observe changes in expectations behavior even for the (relatively inertial) longer-term survey expectations, such that it becomes more difficult to get inflation back to target.

A second test of anchoring of inflation expectations is the dependence of long-term expectations on short-term expectations. If long-run expectations are well anchored, they should not respond to movements in short-term inflation expectations. For this test, we use both five-year and one-year-ahead forecasts, the former as the dependent variable in Equation (2) and the latter replacing lagged inflation in the right-hand side of the same equation. Given that only the mean five-year-ahead forecasts are published, we also use mean forecasts for the one-year-ahead expectations.

The results, shown in Table 5, provide further evidence about the unanchoring of survey expectations during deflations. Column 1 suggests that long-run survey expectations are related to short-run survey expectations in the whole sample, with a statistically significant coefficient estimate of 0.219. Adding interaction variables for high inflation and deflation periods with short-term inflation expectations in column 2, we find that the coefficient on the interaction variable associated with deflations is economically and statistically highly significant. The estimate of 0.240 in column 2 suggests that deflations are associated with greater dependence of long-run survey expectations on short-run expectations, with an effect that is 40 percent higher in magnitude than in normal periods. The significance of deflations remains when low inflations are included in the estimation (column 4). Finally, no similar evidence of unanchoring is obtained for the high inflation periods, concurring with the evidence in Table 4.

5. Reasons for Forecast Behavior

Why do expectations become more backward-looking during deflations? In this section we examine a number of potential explanations, including beliefs about the macroeconomic situation during deflations, a more persistent inflation process during deflationary episodes, and constraints on monetary policy.

Table 5. Sensitivity of Long-Run Inflation Expectations to Short-Term Expectations

Variable	(1)	(2)	(3)	(4)
$\pi^e_{c,t-1}$	0.219***	0.172***	0.203***	0.206***
	(0.0557)	(0.0328)	(0.0460)	(0.0476)
$D^{defl}_{c,t-1}$		–0.0402		–0.0822
		(0.111)		(0.0766)
$D^{High\ Infl}_{c,t-1}$		–0.751*	–0.676	–0.666
		(0.402)	(0.431)	(0.437)
$D^{defl}_{c,t-1}\pi^e_{c,t-1}$		0.240**		0.247**
		(0.105)		(0.104)
$D^{High\ Infl}_{c,t-1}\pi^e_{c,t-1}$		0.175	0.149	0.147
		(0.112)	(0.123)	(0.125)
$D^{Low\ Infl}_{c,t-1}$			0.0516	0.148
			(0.153)	(0.127)
$D^{Low\ Infl}_{c,t-1}\pi^e_{c,t-1}$			0.0678	–0.0442
			(0.0959)	(0.0787)
Obs.	879	879	879	879
R-squared	0.889	0.896	0.895	0.896

Note: Dependent variable is the five-year-ahead forecast (mean across forecasters). π^e denotes the one-year-ahead inflation forecast. Robust standard errors clustered by country in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. All models include country fixed effects.

It appears plausible that, in an unusual economic environment, forecasters are increasingly getting their cues from actual inflation outcomes. However, such behavior does not appear to be uniform across forecasters. Indeed, diving deeper into the forecaster-level data, we observe significant heterogeneity in forecast behavior during deflations.[8] In particular, we find that significant backward-lookingness in inflation expectations during deflations obtains especially in the sample of forecasters who are more pessimistic about the short-term growth prospects of the economy. Table 6 splits the sample of forecasters into those whose GDP forecasts are below and

[8]This analysis is only pursued for the one-year-ahead expectations, as only the mean forecasts are published for the longer-term (e.g., five-year-ahead) expectations.

Table 6. Heterogeneity among Forecasters Based on GDP Growth Expectations

	Below-Median GDP Forecast		Above-Median GDP Forecast	
Variable	**(1)**	**(2)**	**(3)**	**(4)**
$\pi_{c,t-1}$	0.486***	0.471***	0.493***	0.470***
	(0.0150)	(0.0194)	(0.0101)	(0.0133)
$D^{defl}_{c,t-1}$	0.154	0.143	–0.0469	–0.0463
	(0.174)	(0.179)	(0.0827)	(0.0842)
$D^{High\ Infl}_{c,t-1}$	0.0814	0.0306	–0.114	–0.188**
	(0.106)	(0.111)	(0.0911)	(0.0956)
$D^{defl}_{c,t-1}\pi_{c,t-1}$	0.356**	0.486***	0.129	0.260***
	(0.169)	(0.157)	(0.0845)	(0.0815)
$D^{High\ Infl}_{c,t-1}\pi_{c,t-1}$	–0.00249	0.0119	0.0348*	0.0571**
	(0.0225)	(0.0255)	(0.0201)	(0.0224)
$D^{Low\ Infl}_{c,t-1}$		–0.0543		–0.0941***
		(0.0392)		(0.0312)
$D^{Low\ Infl}_{c,t-1}\pi_{c,t-1}$		–0.118***		–0.111***
		(0.0290)		(0.0220)
Obs.	56,927	56,927	59,751	59,751
R-squared	0.736	0.736	0.833	0.833

Note: Dependent variable is the expectation by an individual forecaster for one-year-ahead inflation. Columns 1 and 2 show inflation expectations behavior for forecasters who predict below-median GDP growth; columns 3 and 4 show results for forecasters who predict above-median GDP growth. Robust standard errors clustered by country in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. All models include forecaster fixed effects.

above the median across forecasters, respectively (computed separately for each country). In the more parsimonious specification, increased backward-lookingness obtains only for those forecasters who have lower GDP growth expectations (column 1). For those with a more optimistic outlook on activity (column 3), inflation expectations do not become more backward-looking in a statistically significant way during deflations. By contrast, during periods of moderate positive inflation, the degree of backward-lookingness is highly similar among pessimists and optimists.

Similar evidence obtains in a model that also includes the dummy variable on low inflation. In this case, while also the more optimistic

forecasters seem to display greater backward-lookingness, the coefficient estimate on the interaction variable between deflation and lagged inflation is higher (0.49 versus 0.26) for the more pessimistic forecasters (columns 2 and 4 in Table 6). Thus, anchoring of near-term inflation expectations appears to weaken especially for forecasters who expect lower GDP growth. Again, for periods of moderate inflation, backward-lookingness is similar for both pessimists and optimists.

The point inflation forecasts for the two groups of forecasters are also different during deflations. For the median pessimistic forecaster, the one-year-ahead inflation expectation formed during deflations is 0.28 percent, compared with 0.58 percent for an optimistic forecaster. By contrast, both the relative and absolute differences between forecasters are much smaller outside of deflations, with one-year-ahead inflation forecasts at 2.47 percent for pessimists and 2.56 percent for optimists.

Such evidence on forecast heterogeneity, particularly during adverse time periods, is consistent with other literature. Andrade et al. (2019) document, using forecast data from the Federal Reserve's forward guidance in 2011–12, that the same announcement on maintaining future interest rates low led to highly contrasting interpretations across professional forecasters. The optimists interpreted forward guidance as a promise of a future accommodative monetary policy stance and therefore better macroeconomic conditions, whereas the pessimists regarded it as signaling worse future macroeconomic outcomes. Similarly, Banerjee and Mehrotra (2021) show that deflations are associated with an economically and statistically significant rise in forecast disagreement across professional inflation forecasters. By contrast, disagreement about future inflation appears lowest when inflation rates are close to 2 percent—a common inflation target. Bachmann, Berg, and Sims (2015) report that the sign of the relationship between readiness to purchase durable goods and inflation expectations differs across households, with a positive relationship documented only for a small fraction of households that are "good" forecasters. Moreover, the prevailing negative relationship between durables spending attitudes and inflation expectations becomes stronger at the zero lower bound.

Relatedly, evidence about the macroeconomic implications of deflations is far from conclusive, potentially contributing to the

divergent views across forecasters. Borio et al. (2015) document that, the Great Depression aside, the historical link between output growth and consumer price deflation is weak, and Bordo and Filardo (2005) highlight the mostly benign deflation record of the 19th century. By contrast, Eichengreen, Park, and Shin (2017), using producer price inflation, find stronger evidence of damaging effects of deflation on growth. Davis (2015) finds that expectations of deflation led to lower durable consumption in the United States during the Great Depression; similar evidence is provided for Japan's deflation experience in Hori and Shimizutani (2005).

Besides pessimistic perceptions about the economic implications of deflations, there are other possible factors driving greater forecast persistence. One possibility is that inflation outcomes themselves could be more persistent during deflations. Then, greater backward-lookingness could merely reflect this change in the inflation process rather than any other change in the behavior of expectations per se. However, we do not find evidence of increased inflation persistence during deflations: estimating a simple AR(1) model for inflation with country-level panel data for the deflation episodes leads to an autoregressive coefficient at the first lag of 0.787 (standard error of 0.050).[9] Considering the entire sample instead, the coefficient is 0.943 (standard error of 0.010). Thus, inflation actually appears somewhat less persistent during deflations.

Another factor potentially affecting the behavior of expectations during deflations is the possibility that conventional interest rate policy is constrained. Forecast dynamics and the macroeconomic implications of deflations may be expected to be more adverse if monetary policy is perceived to lack the tools to return inflation to target. In order to examine to what extent interest rates close to zero are affecting the results, we exclude periods with policy interest rates at levels of 0.5 percent or below from the estimations that use near-term inflation expectations. The results are shown in Table A.6 in the appendix.

We find that greater backward-lookingness in expectations during deflations does not hinge on the zero interest rate floor. When periods of near-zero interest rates are excluded, the point estimate

[9]These models include country fixed effects, and the standard errors are clustered by country.

on the interaction variable between lagged inflation and the deflation dummy actually increases in absolute terms. The sum of the backward-looking terms during deflations, when zero lower bound (ZLB) periods are excluded, is close to one (column 4 of Table A.6 in the appendix). This suggests that the fall in expectations during deflations does not arise due to perceived lack of potency of monetary policy.

6. Implications for Policy from More Backward-Looking Expectations during Deflation

Higher backward-lookingness in expectations also has implications for optimal monetary policy. Consider, for instance, the differences for optimal policy between rational expectations and expectations formed through adaptive learning. In the latter case, the private sector forms its inflation expectations based on the past behavior of inflation, updating their beliefs as the available data change. Gaspar, Smets, and Vestin (2006) show that under this environment, optimal policy responds in a persistent manner to cost-push shocks, allowing to stabilize inflation expectations in this way.[10] The higher is the perceived inflation persistence by the private sector, the stronger and more persistent the optimal response. Thus, if expectations are increasingly adaptive as the inflation rate falls, a stronger policy response would be called for when negative cost-push shocks push inflation down from already low levels. In Evans, Guse, and Honkapohja (2008) that features the ZLB, the possibility of a deflationary spiral arises around the lower inflation equilibrium under learning, requiring a strong fiscal and monetary response to guarantee a lower bound on inflation.

The fact that significant backward-lookingness in inflation expectations during deflations is driven by forecasters who are also pessimistic about the short-term growth prospects may have implications for monetary and fiscal coordination in such states. In particular, during deflationary episodes, there may be a role for fiscal

[10]The authors show that this decreases inflation persistence and volatility at little cost in terms of output volatility. See also Gaspar, Smets, and Vestin (2011), who note that departures from rational expectations raise the potential for instability, increasing the importance of managing inflation expectations.

stimulus to boost short-term growth and thus minimize pessimistic beliefs about near-term growth prospects (Nakata and Schmidt 2022). This in turn may help anchor inflation expectations (Evans and Honkapohja 2010).

7. Conclusion

In this paper, we have analyzed the behavior of inflation expectations during periods of deflation, using a large cross-country data set of Consensus forecasts. The global nature of the data is essential—out of the 47 deflation episodes identified in our sample, 30 occur in emerging market economies. We document that deflations are largely unanticipated by forecasters, even at short forecasting horizons. Thus, our results concur with the evidence in Ahearne et al. (2002) for Japan's deflation in the 1990s, but apply for a much larger sample of economies and time periods. In the econometric analysis, we first use one-year-ahead forecasts that could be relevant for actual wage and price setting, but also for providing timely signals that expectations risk becoming unmoored. Using these short-term forecasts, we show that the dependence of inflation expectations on past inflation outcomes rises by close to 50 percent during deflations, compared with other periods, implying that it becomes significantly more difficult to get inflation back to target when an economy is already experiencing deflation.

We additionally document that increased backward-lookingness in inflation expectations is driven by those forecasters who are more pessimistic about the short-term growth prospects. By contrast, inflation expectations of forecasters who are more optimistic about growth during deflations remain better anchored. In this way, deflations differ from other time periods that feature highly similar backward-lookingness across optimists and pessimists.

We also use longer-term, five-year-ahead, inflation expectations. We find that deflations are associated with greater backward-lookingness also in terms of the long-term expectations, as well as with greater dependence of long-term on short-term expectations. Both pieces of evidence are notable given the generally inertial nature of long-term survey expectations. They also differ from the results for other inflation environments and indicate that inflation expectations become unanchored during deflations.

Our results, in particular the changing behavior of expectations once the economy hits deflation, are supportive of a risk-management approach to monetary policy (see, e.g., Greenspan 2004). As it becomes increasingly difficult to return inflation to target once prices are falling and expectations have become more backward-looking, it may be prudent to preemptively react to deflation risks. Previous evidence for the United States suggests that the Federal Reserve may have already been responding to inflation in this manner (e.g., Kilian and Manganelli 2008). The global loosening of monetary policy in 2019 as inflation remained absent but trade tensions loomed large is also suggestive of a risk-management approach being applied in other economies.

Appendix

Figure A.1. High-Inflation Episodes

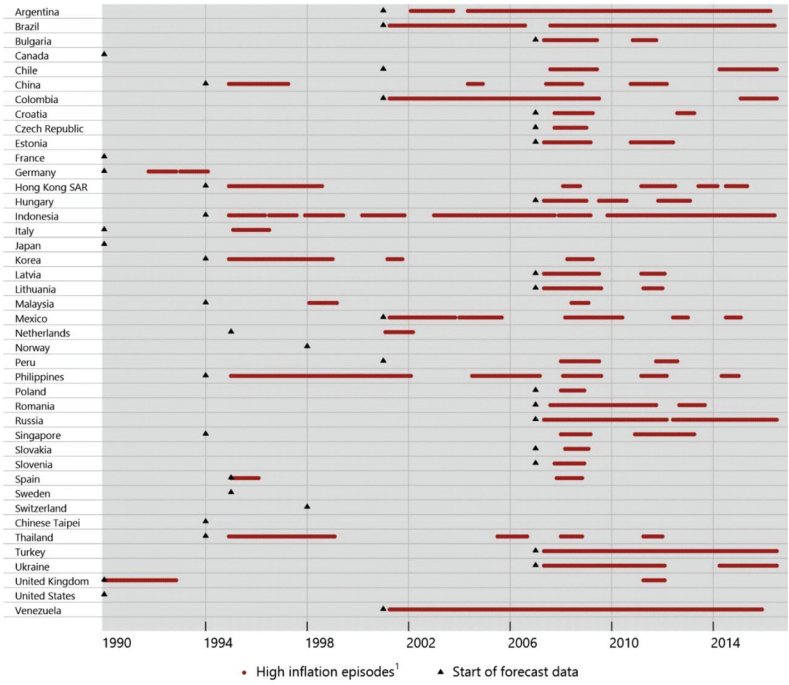

Note: [1]High-inflation periods are defined as at least six consecutive months of above 4 percent inflation rates. A high-inflation period ends if at least three consecutive periods with below 4 percent inflation rates occur.

Table A.1. Data Coverage

Economy	Start	End	Number of Months	Number of Forecasters			Total Observations
				Minimum	Maximum	Average	
Argentina	Apr. 01	Apr. 16	181	7	25	15.1	2,734
Bulgaria	May 07	Jun. 16	110	8	16	11.9	1,309
Brazil	Apr. 01	Jun. 16	183	13	27	18.0	3,303
Canada	Jan. 90	Jun. 16	318	11	20	15.2	4,847
Switzerland	Jun. 98	Jul. 16	218	6	17	13.3	2,891
Chile	Apr. 01	Jul. 16	184	11	23	17.0	3,122
China	Dec. 94	Jun. 16	259	11	24	17.9	4,624
Colombia	Apr. 01	Jul. 16	184	7	16	12.2	2,251
Czech Republic	May 07	Jun. 16	110	13	22	17.0	1,869
Germany	Jan. 90	Jul. 16	319	12	32	26.5	8,457
Estonia	May 07	Jul. 16	111	6	12	9.2	1,023
Spain	Jan. 95	Jul. 16	259	7	19	13.9	3,607
France	Jan. 90	Jul. 16	319	6	26	17.8	5,692
UK (CPI)	Jan. 05	Jun. 16	138	19	29	24.5	3,382
UK (RPI)	Jan. 90	Dec. 04	180	18	39	29.5	5,308
Hong Kong SAR	Dec. 94	Jun. 16	259	9	21	15.4	3,984
Croatia	May 07	Jun. 16	110	7	17	11.3	1,238
Hungary	May 07	Jun. 16	110	12	25	16.5	1,810
Indonesia	Dec. 94	Jul. 16	260	7	20	13.9	3,601
Italy	Jan. 90	Jul. 16	319	6	21	13.8	4,406

(*continued*)

Table A.1. (Continued)

Economy	Start	End	Number of Months	Number of Forecasters			Total Observations
				Minimum	Maximum	Average	
Japan	Jan. 90	Jun. 16	318	5	25	17.2	5,481
Korea	Dec. 94	Jul. 16	260	9	22	15.2	3,949
Lithuania	May 07	Jun. 16	110	6	12	8.0	883
Latvia	May 07	Jun. 16	110	7	14	9.5	1,043
Mexico	Apr. 01	Jul. 16	184	9	27	19.3	3,549
Malaysia	Dec. 94	Jun. 16	259	10	23	15.1	3,916
Netherlands	Jan. 95	Jul. 16	259	7	14	9.4	2,430
Norway	Jun. 98	Jun. 16	217	5	12	9.2	2,007
Peru	Apr. 01	Jul. 16	184	7	17	11.8	2,177
Philippines	Dec. 94	Jul. 16	260	12	21	15.4	1,526
Poland	May 07	Jun. 16	110	12	25	17.3	1,901
Romania	May 07	Jul. 16	111	7	19	12.4	1,374
Russia	May 07	Jul. 16	111	10	24	16.3	1,807
Sweden	Jan. 95	Jun. 16	258	6	18	13.4	3,470
Singapore	Dec. 94	Jun. 16	259	8	19	14.7	3,800
Slovenia	May 07	Jul. 16	111	7	16	10.6	1,179
Slovakia	May 07	Jun. 16	110	8	15	11.4	1,249
Thailand	Dec. 94	Jul. 16	260	8	22	13.9	3,623
Turkey	May 07	Jul. 16	111	9	21	15.2	1,683
Chinese Taipei	Dec. 94	Jul. 16	260	9	22	14.3	3,720
Ukraine	May 07	Jul. 16	111	8	19	14.4	1,599
United States	Jan. 90	Jun. 16	318	16	33	25.8	8,203
Venezuela	Apr. 01	Dec. 15	177	7	19	13.6	2,402

Table A.2. Deflation Episodes

Economy	Start	End	Length (Months)	Minimum Current-Year Forecast	Minimum Next-Year Forecast	Minimum Five-Year-Ahead Forecast	Minimum Inflation Outcome
Argentina	Jun. 01	Feb. 02	9	–1.6	–0.9	1.5	–1.7
Chile	Aug. 09	Mar. 10	8	–1.3	2.7	3	–2.3
Bulgaria	Aug. 13	Apr. 15	21	–1.3	1.1	NA	–2.6
Estonia	May 09	Apr. 10	12	–0.8	–0.2	NA	–2.2
Estonia	Jun. 14	Jul. 16	26	–0.4	1.3	NA	–1.3
Croatia	Feb. 14	Jun. 16	29	–0.7	1.0	NA	–1.9
Hungary	Sep. 14	Jun. 15	10	–0.1	1.5	2.6	–1.4
Lithuania	Dec. 14	Feb. 16	15	–0.7	1.4	NA	–1.8
Latvia	Oct. 09	Oct. 10	13	–3.0	–3.1	NA	–4.2
Latvia	Jan. 16	Jun. 16	6	–0.1	1.7	NA	–0.8
Poland	Aug. 14	Jun. 16	23	–0.9	0.7	2.2	–1.3
Romania	Jun. 15	Jul. 16	14	–1.2	–0.2	NA	–3.5
Slovenia	Jan. 15	Jul. 16	19	–0.5	0.8	NA	–0.8
Slovakia	Feb. 14	Jun. 16	29	–0.3	0.9	NA	–0.8
Ukraine	Nov. 12	Dec. 13	14	–0.3	3.3	NA	–0.8
Switzerland	Mar. 09	Dec. 09	10	–0.5	0.6	1.4	–1.2
Switzerland	Oct. 11	Dec. 13	27	–0.6	–0.1	1.3	–1.1
Switzerland	Sep. 14	Jul. 16	23	–1.3	–0.3	1.3	–1.4
Spain	Mar. 09	Dec. 09	10	–0.3	1.3	2.2	–1.4
Spain	Jul. 14	Jul. 16	25	–1.1	0.5	1.8	–1.3
France	May 09	Dec. 09	8	0.1	1.1	2	–0.7
Italy	Feb. 16	Jul. 16	6	–0.1	0.9	1.7	–0.5
Japan	Oct. 95	May 96	8	0.0	0.2	NA	–0.7
Japan	Feb. 99	Oct. 04	69	–1.0	–0.9	0.3	–1.6
Japan	Jun. 05	Jun. 06	13	–0.2	0.1	1.3	–1.0

(*continued*)

Table A.2. (Continued)

Economy	Start	End	Length (Months)	Minimum Current-Year Forecast	Minimum Next-Year Forecast	Minimum Five-Year-Ahead Forecast	Minimum Inflation Outcome
Japan	Feb. 09	Aug. 11	31	–1.3	–1.0	0.7	–2.5
Japan	Jun. 12	Jul. 13	14	–0.2	–0.1	0.3	–0.9
Sweden	Sep. 96	Jun. 97	10	0.8	1.3	NA	–0.4
Sweden	Jun. 98	Apr. 99	11	0.3	0.6	NA	–1.2
Sweden	Apr. 09	Jan. 10	10	–0.4	0.8	1.9	–1.9
Sweden	Aug. 14	Oct. 15	15	–0.2	0.7	2.2	–0.4
United States	Mar. 09	Dec. 09	10	–1.0	1.6	2.2	–2.1
China	Apr. 98	Jun. 00	27	–1.3	1.2	3.9	–2.2
China	Mar. 02	Feb. 03	12	–0.8	0.3	1.7	–1.3
China	Feb. 09	Dec. 09	11	–0.7	1.4	3	–1.8
Hong Kong SAR	Nov. 98	Aug. 04	70	–3.6	–1.5	1.4	–6.1
Malaysia	Jun. 09	Jan. 10	8	0.6	2.0	3.3	–2.5
Singapore	Jun. 98	Jun. 99	13	–0.6	–0.5	NA	–1.5
Singapore	Nov. 01	Dec. 02	14	–0.4	1.1	1.9	–1.1
Singapore	Jun. 09	Feb. 10	9	–0.2	1.4	2.2	–0.9
Singapore	Nov. 14	Jun. 16	20	–0.5	0.5	1.9	–1.6
Thailand	May 99	Dec. 99	8	0.3	2.2	NA	–1.2
Thailand	Jan. 09	Nov. 09	11	–1.0	2.1	2.7	–4.4
Thailand	Jan. 15	Jun. 16	18	–0.8	1.4	2.5	–1.3
Chinese Taipei	Jun. 03	Feb. 04	9	–0.3	0.5	NA	–1.0
Chinese Taipei	Feb. 09	Feb. 10	13	–1.2	0.7	NA	–2.3
Chinese Taipei	Jan. 15	Oct. 15	10	–0.4	1.1	NA	–0.9

Note: See the definition of deflation episodes in Figure 1. The length of deflation episodes in this table includes the two consecutive months of positive inflation rates that potentially follow deflation. The minimum forecast refers to the lowest median forecast across forecasters (for five-year-ahead forecasts, minimum of mean forecasts; the latter as reported by Consensus). NA indicates that the forecast is not available.

Table A.3. Data Sources and Definitions

Variable	Source	Data Transformations
CPI Inflation	National Data	
Inflation Expectations	Consensus Economics	
GDP Growth Expectations	Consensus Economics	
Policy Interest Rate	Datastream; National Data	Where policy rates are not available, money market interest rates are used.

Table A.4. Degree of Backward-Lookingness, by Type of Deflation

Variable	(1)	(2)	(3)
$D^{defl}_{c,t-1}\pi_{c,t-1}$	0.664***		
	(0.194)		
$D^{longer\ defl}_{c,t-1}\pi_{c,t-1}$	–0.596***		
	(0.197)		
$D^{defl}_{c,t-1}\pi_{c,t-1}$		0.090***	
		(0.027)	
$D^{deeper\ defl}_{c,t-1}\pi_{c,t-1}$		0.408***	
		(0.139)	
$D^{defl}_{c,t-1}\pi_{c,t-1}$			0.136***
			(0.034)
$D^{faster\ defl}_{c,t-1}\pi_{c,t-1}$			0.699***
			(0.246)
Obs.	124,252	124,252	124,252
R-squared	0.770	0.768	0.771

Note: Dependent variable is the expectation by an individual forecaster for one-year-ahead inflation. Only the coefficients capturing backward-lookingness are displayed. Models are estimated separately with dummy variables for longer/deeper/faster deflations, respectively. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. All models include forecaster fixed effects.

Table A.5. One-Year-Ahead Inflation Expectations, Using a Country-Specific High-Inflation Threshold

Variable	(1)	(2)	(3)	(4)
$\pi_{c,t-1}$	0.525***	0.485***	0.466***	0.467***
	(0.00939)	(0.0260)	(0.0328)	(0.0316)
$D^{defl}_{c,t-1}$		0.0193		0.0321
		(0.0933)		(0.106)
$D^{High\ Infl}_{c,t-1}$		–0.118***	–0.196***	–0.163***
		(0.0288)	(0.0394)	(0.0343)
$D^{defl}_{c,t-1}\pi_{c,t-1}$		0.249**		0.356***
		(0.116)		(0.0976)
$D^{High\ Infl}_{c,t-1}\pi_{c,t-1}$		0.0424**	0.0613**	0.0557**
		(0.0184)	(0.0249)	(0.0229)
$D^{Low\ Infl}_{c,t-1}$			–0.208***	–0.106**
			(0.0385)	(0.0429)
$D^{Low\ Infl}_{c,t-1}\pi_{c,t-1}$			0.127**	–0.0928***
			(0.0584)	(0.0326)
Obs.	124,252	124,252	124,252	124,252
R-squared	0.766	0.767	0.767	0.768

Note: Dependent variable is the expectation by an individual forecaster for one-year-ahead inflation. The high-inflation dummy is constructed using a country-specific inflation threshold (average inflation in the country over the sample). Robust standard errors clustered by country in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. All models include forecaster fixed effects.

Table A.6. One-Year-Ahead Inflation Expectations, Excluding the ZLB

Variable	(1)	(2)	(3)	(4)
$\pi_{c,t-1}$	0.531***	0.492***	0.485***	0.486***
	(0.0112)	(0.0124)	(0.0166)	(0.0159)
$D^{defl}_{c,t-1}$		0.210		0.166
		(0.177)		(0.182)
$D^{High\ Infl}_{c,t-1}$		–0.0546	–0.0740	–0.0785
		(0.0832)	(0.0857)	(0.0870)
$D^{defl}_{c,t-1}\pi_{c,t-1}$		0.343**		0.542***
		(0.136)		(0.127)
$D^{High\ Infl}_{c,t-1}\pi_{c,t-1}$		0.359*	0.0413**	0.0413*
		(0.0190)	(0.0209)	(0.0212)
$D^{Low\ Infl}_{c,t-1}$			–0.116***	0.00658
			(0.0384)	(0.0354)
$D^{Low\ Infl}_{c,t-1}\pi_{c,t-1}$			0.116**	–0.196***
			(0.0523)	(0.0242)
Obs.	101,442	101,442	101,442	101,442
R-squared	0.745	0.746	0.745	0.747

Note: Dependent variable is the expectation by an individual forecaster for one-year-ahead inflation. Robust standard errors clustered by country in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent level, respectively. All models include forecaster fixed effects.

References

Ahearne, A., J. Gagnon, J. Haltmaier, and S. Kamin. 2002. "Preventing Deflation: Lessons from Japan's Experience in the 1990s." International Finance Discussion Paper No. 729, Board of Governors of the Federal Reserve System.

Anand, R., E. Prasad, and B. Zhang. 2015. "What Measure of Inflation Should a Developing Country Central Bank Target?" *Journal of Monetary Economics* 74 (September): 102–16.

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon. 2019. "Forward Guidance and Heterogeneous Beliefs." *American Economic Journal: Macroeconomics* 11 (3): 1–29.

Andrade, P., E. Ghysels, and J. Idier. 2015. "Tails of Inflation Forecasts and Tales of Monetary Policy." Mimeo.

Bachmann, R., T. Berg, and E. Sims. 2015. "Inflation Expectations and Readiness to Spend: Cross-sectional Evidence." *American Economic Journal: Economic Policy* 7 (1): 1–35.

Banca d'Italia. 2017. "Private Sector Contract Renewals in 2016." *Economic Bulletin* 1/2017: Box, p. 32.

Banerjee, R., and A. Mehrotra. 2021. "Disagreeing during Deflations." *Journal of Money, Credit and Banking* 53 (7): 1867–85.

Bank of Japan. 2016. "Development in Inflation Expectations over the Three Years Since the Introduction of Quantitative and Qualitative Monetary Easing (QQE)." *Bank of Japan Review* 2016-E-13 (October).

Benhabib, J., S. Schmitt-Grohe, and M. Uribe. 2002. "Avoiding Liquidity Traps." *Journal of Political Economy* 110 (3): 535–63.

Blanchard, O. 2016. "The Phillips Curve: Back to the '60s?" *American Economic Review* 106 (5): 31–34.

Bordo, M., and A. Filardo. 2005. "Deflation and Monetary Policy in a Historical Perspective: Remembering the Past or Being Condemned to Repeat It?" *Economic Policy* 20 (441): 799–844.

Borio, C., M. Erdem, A. Filardo, and B. Hofmann. 2015. "The Costs of Deflations: A Historical Perspective." *BIS Quarterly Review* (March): 31–54.

Buono, I., and S. Formai. 2018. "New Evidence on the Evolution of the Anchoring of Inflation Expectations." *Journal of Macroeconomics* 57 (September): 39–54.

Burdekin, R., and P. Siklos, eds. 2010. *Deflation. Current and Historical Perspectives*. Cambridge University Press.

Coibion, O., and Y. Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7 (1): 197–232.

Coibion, O., Y. Gorodnichenko, and S. Kumar. 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review* 108 (9): 2671–2713.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte. 2020. "Inflation Expectations as a Policy Tool?" *Journal of International Economics* 124 (May): Article 103297.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber. 2021. "Exposure to Grocery Prices and Inflation Expectations." *Journal of Political Economy* 129 (5): 1615–39.

Davis, J. 2015. "The Asymmetric Effects of Deflation on Consumption Spending: Evidence from the Great Depression." *Economics Letters* 130 (May): 105–8.

Doser, A., R. Nunes, N. Rao, and V. Sheremirov. 2017. "Inflation Expectations and Nonlinearities in the Phillips Curve." Working Paper No. 17-11, Federal Reserve Bank of Boston.

Dovern, J., U. Fritsche, and J. Slacalek. 2012. "Disagreement among Forecasters in G7 Countries." *Review of Economics and Statistics* 94 (4): 1081–96.

Draghi, M. 2015. "Introductory Statement to the Press Conference." Frankfurt am Main, January 22.

Ehrmann, M. 2015. "Targeting Inflation from Below: How Do Inflation Expectations Behave?" *International Journal of Central Banking* 11 (S1): 213–49.

Eichengreen, B., D. Park, and K. Shin. 2017. "Should the Dangers of Deflation Be Dismissed?" *Journal of Macroeconomics* 52 (June): 287–307.

Evans, G., E. Guse, and S. Honkapohja. 2008. "Liquidity Traps, Learning and Stagnation." *European Economic Review* 52 (8): 1438–63.

Evans, G., and S. Honkapohja. 2010. "Expectations, Deflation and Macroeconomic Policy." In *Twenty Years of Inflation Targeting*, ed. D. Cobham, O. Eitrheim, S. Gerlach, and J. Qvigstad. Cambridge University Press.

Faust, J., and J. Wright. 2013. "Forecasting Inflation." *Handbook of Economic Forecasting,* Vol. 2A, ed. G. Elliott and A. Timmermann, 2–56. Amsterdam: Elsevier.

Fisher, I. 1933. "The Debt-Deflation Theory of Great Depressions." *Econometrica* 1 (4): 337–57.

Fuhrer, J. 2017. "Japanese and U.S. Inflation Dynamics in the 21st Century." IMES Discussion Paper No. 2017-E-5.

Gaspar, V., F. Smets, and D. Vestin. 2006. "Adaptive Learning, Persistence, and Optimal Monetary Policy." *Journal of the European Economic Association* 4 (2–3): 376–85.

———. 2011. "Inflation Expectations, Adaptive Learning and Optimal Monetary Policy." *Handbook of Monetary Economics*, Vol. 3B, ed. B. M. Friedman and M. Woodford, 1055–95 (chapter 19). North-Holland.

Giannoni, M. 2007. "Robust Optimal Monetary Policy in a Forward-looking Model with Parameter and Shock Uncertainty." *Journal of Applied Econometrics* 22 (1): 179–213.

Gnabo, J.-Y., and D. Moccero. 2015. "Risk Management, Nonlinearity and Aggressiveness in Monetary Policy: The Case of the US Fed." *Journal of Banking and Finance* 55 (June): 281–94.

Greenspan, A. 2004. "Risk and Uncertainty in Monetary Policy." Remarks at the meetings of the American Economic Association, San Diego, CA, January 3.

Hori, M., and S. Shimizutani. 2005. "Price Expectations and Consumption under Deflation: Evidence from Japanese Household Survey Data." *International Economics and Economic Policy* 2 (2–3): 127–51.

International Monetary Fund. 2016. "Global Disinflation in an Era of Constrained Monetary Policy." In *World Economic Outlook: Subdued Demand: Symptoms and Remedies*, 121–70 (chapter 3). Washington, DC: International Monetary Fund.

Jordà, Ò. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.

Kenny, G., and J. Dovern. 2017. "The Long-term Distribution of Expected Inflation in the Euro Area: What Has Changed since the Great Recession?" ECB Working Paper No. 1999.

Kilian, L., and S. Manganelli. 2008. "The Central Banker as a Risk Manager: Estimating the Federal Reserve's Preferences under Greenspan." *Journal of Money, Credit and Banking* 40 (6): 1103–29.

Kozicki, S., and P. A. Tinsley. 2012. "Effective Use of Survey Inflation in Estimating the Evolution of Expected Inflation." *Journal of Money, Credit and Banking* 44 (1): 145–69.

Krugman, P. 1998. "It's Baaack! Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity* (2).

Levin, A., F. Natalucci, and J. Piger. 2004. "The Macroeconomic Effects of Inflation Targeting." *Review* (Federal Reserve Bank of St. Louis) 86 (4): 51–80.

Lyziak, T., and M. Paloviita. 2017. "Anchoring of Inflation Expectations in the Euro Area: Recent Evidence Based on Survey Data." *European Journal of Political Economy* 46 (January): 52–73.

Mehrotra, A., and J. Yetman. 2018. "Decaying Expectations: What Inflation Forecasts Tell Us about the Anchoring of Inflation

Expectations." *International Journal of Central Banking* 14 (5): 55–101.

Mertens, T., and J. Williams. 2018. "What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices." Working Paper No. 2018-3, Federal Reserve Bank of San Francisco.

Nakata, T., and S. Schmidt. 2022. "Expectations-Driven Liquidity Traps: Implications for Monetary and Fiscal Policy." *American Economic Journal: Macroeconomics* 14 (4): 68–103.

Natoli, F., and L. Sigalotti. 2018. "Tail Co-movement in Inflation Expectations as an Indicator of Anchoring." *International Journal of Central Banking* 14 (1): 35–71.

Nishizaki, K., T. Sekine, and Y. Ueno. 2014. "Chronic Deflation in Japan." *Asian Economic Policy Review* 9 (1): 20–39.

Siklos, P. 2013. "Sources of Disagreement in Inflation Forecasts: An International Empirical Investigation." *Journal of International Economics* 90 (1): 218–31.

Smith, G. 2006. "The Spectre of Deflation: A Review of Empirical Evidence." *Canadian Journal of Economics* 39 (4): 1041–72.

Summers, L. 2019. "Global Economy is at Risk from a Monetary Policy Black Hole." *Financial Times*, October 11.

Svensson, L. 2001. "The Zero Bound in an Open Economy: A Foolproof Way of Escaping from a Liquidity Trap." *Monetary and Economic Studies* 19 (S-1, February): 277–312.

Williams, J. C. 2009. "The Risk of Deflation." FRBSF Economic Letter No. 2009-12, Federal Reserve Bank of San Francisco.

Central Clearing and Systemic Liquidity Risk*

Thomas B. King,[a] Travis D. Nesmith,[b] Anna Paulson,[a] and Todd Prono[b]
[a]Federal Reserve Bank of Chicago
[b]Board of Governors of the Federal Reserve System

By stepping between bilateral counterparties, central counterparties (CCPs) transform credit exposure, thereby improving financial stability. But large CCPs are concentrated and interconnected with major global banks. Moreover, although they mitigate credit risk, CCPs create liquidity risks, because they require participants to provide cash. Such requirements increase with market volatility; consequently, CCP liquidity needs are inherently procyclical. This procyclicality makes it more challenging to assess CCPs' resilience in the rare event that one or more large financial institutions default. Liquidity-focused macroprudential stress tests could help to assess and manage this systemic liquidity risk.

JEL Codes: G23, G21, G28, E58, N22.

*We thank Ivana Ruffini for her comments and for providing substantial insights into the data. We also thank Kelsey Burr, Marta Chaffee, Roy Cheruvelil, Josh Gallin, Richard Haynes, Mike Kiley, Beth Klee, David Li, Jennifer Lucier, David Murphy, Sam Schulhofer-Wohl, and Jonathan Rose for helpful comments and feedback. We also wish to thank the three referees whose careful review and comments improved the paper. Most of all, we want to recognize the debt we owe our former colleague Robert Cox; he not only influenced this paper with his comments, but more generally influenced our understanding of the topic. David Kelley, Tim Larach, and John Spence provided excellent research assistance. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff, the Federal Reserve Bank of Chicago, the Board of Governors, or the Federal Reserve System. Author contact: (Thomas King and Anna Paulson) Federal Reserve Bank of Chicago, 230 S. LaSalle St., Chicago IL 60604; (Travis D. Nesmith and Todd Prono) Board of Governors of the Federal Reserve System, 20th & C Sts. NW, Washington DC 20551; E-mails: thomas.king@chi.frb.org; travis.d.nesmith@frb.gov; anna.paulson@chi.frb.org; todd.a.prono@frb.gov.

1. Introduction

Central counterparties (CCPs) play a critical role in the financial system. By inserting itself between the two participants in a securities or derivatives transaction, a central counterparty guarantees payments that could otherwise be jeopardized by the default of either participant. CCPs' importance has grown significantly since the 2008 financial crisis, in part as a result of new regulatory mandates requiring central clearing of over-the-counter (OTC) derivatives. By the first half of 2018, over 75 percent of the notional value of interest rate swaps (IRS) and credit default swaps (CDS)—the two largest categories of OTC derivatives affected by clearing mandates—was centrally cleared.[1] The risks managed and posed by CCPs are a subject of intense interest for financial regulators and policymakers, particularly as certain large CCPs have been designated to be systemically important.

We review the central clearing landscape and highlight some financial stability issues, focusing on the large CCPs that are the most critical for the U.S. financial system. A robust literature discusses CCPs, and we synthesize some of this work in our overview. We also make three novel contributions. First and most importantly, we highlight a potential risk of central clearing: CCPs may place significant liquidity strains on the banking system precisely in the moments when the banking system is least able to bear such strains. This risk stemming from CCPs' procyclical need for liquidity was arguably under-appreciated when an earlier version of this paper circulated, but became more widely recognized following the severe market volatility resulting from the COVID-19 pandemic.[2] The increased attention has focused on the procyclicality of margin—for example, see Futures Industry Association (FIA) (2020) and ISDA Clearing Member Committee (2021) for industry responses and Financial Stability Board (FSB) (2021) and BCBS-CPMI-IOSCO (2022) for the global regulatory response—which is only part of the potential risk. CCPs are connected to and dependent

[1]Clearing percentages calculated from Bank for International Settlements (BIS) data at https://www.bis.org/statistics/derstats.htm.

[2]King et al. (2020) was published just prior to the market volatility in early 2020.

upon globally systemically important banks (G-SIBs) through a variety of channels in addition to direct clearing relationships. These connections include banks' obligations to post margin to CCPs; their required contributions to CCP default funds; and their provision of lines of credit to CCPs. The liquid resources that CCPs demand through each of these and other channels are inherently *procyclical* with respect to market conditions. The materiality of this risk may have been under-appreciated partly due to a focus in the literature on OTC derivatives clearing in isolation. Despite the impact of COVID-19 on market volatility, the full scope of the potential liquidity risk has not received as much attention because there were no defaults requiring CCPs to use liquidity resources beyond margin. Second, given this procyclical liquidity risk, we discuss the potential for macroprudential liquidity stress tests to improve the preparedness of CCPs, G-SIBs, and other actors by looking at potential liquidity needs across the system. Third, we exploit "quantitative disclosure" data on CCPs in combination with market volatility data. The quarterly disclosure data, which most large CCPs began reporting publicly in 2015, present a rich source of potential information for researchers and policymakers.[3] Although the length of the data limits analysis, it is long enough to at least indicate relationships between clearing and market volatility.

At the outset, it is important to recognize that in most circumstances CCPs improve financial stability relative to a world in which trades are bilateral.[4] There are a variety of ways in which they do so, including by insulating members from each others' defaults; by simplifying and reducing members' gross exposures through the netting of positions; by centralizing risk management within a small number of relatively transparent entities; by pooling financial resources to address extreme tail risks; and by increasing the transparency and predictability of market operations. De Bandt and Hartmann (2019) create a typology of systemic risk and identify the first mechanism that can lead to widespread instability within the financial

[3]See Alfranseder et al. (2018) for further discussion of this data and its usefulness for monitoring CCPs. We access the data through Clarus CCPView.

[4]For example, Kiff et al. (2010) argue that well-run CCPs reduce systemic risk, in principle, and FSB (2017b, p. 1) concludes that reforms have mitigated systemic risk in OTC derivatives largely because of increased central clearing.

system as *contagion*—the potential transfer of distress "horizontally" from one financial intermediary or market to another.[5] CCPs are exactly targeted at mitigating such contagion, by serving as a "fire break" to contain defaults. Indeed, several recent studies—Duffie and Zhu (2011) and Amini, Cont, and Minca (2016), among others—find that CCPs reduce contagion stemming from potential counterparty defaults. Arguably, CCPs performed well during the COVID-19 crisis (CCP12, 2020).

However, De Bandt and Hartmann (2019) also identify *interconnectedness* as a potential driver of financial instability. Large CCPs are, virtually by definition, highly concentrated and interconnected. It is thus worth exploring channels through which CCPs could amplify systemic risk through interconnectedness, even as they dampen contagion risk.[6] Some previous papers (e.g., Faruqui, Huang, and Takáts 2018) have noted the interconnectedness risk associated with CCPs being unable to meet their obligations following member default. We highlight a different problem: the ability of CCPs to fulfill their obligations in stressful periods involves their accessing contingent liquidity from members; such procyclical liquidity risk can be a source of systemic risk that could limit how effective CCPs are at reducing systemic risk overall. Our discussion recalls Bernanke (1990), who, in the wake of the 1987 stock market crash, expressed concerns both about the limits to CCPs' ability to address systematic risks on their own and about the potential for CCPs to be a source of risk.

The potential for CCPs to produce procyclical liquidity strains was recently illustrated by the large increases in margin requirements during the extraordinary market volatility in early 2020 brought on by the COVID-19 pandemic; BCBS-CPMI-IOSCO (2022) reports that daily variation margin calls, which cover already realized losses on cleared portfolios, increased around $115 billion from early 2020

[5]De Bandt and Hartmann (2019) divide systemic risk into horizontal risks that remain in the financial system and vertical risks where financial system risks create risks in the real economy.

[6]FSB (2017a, 2018a) details the high interconnectedness between the largest CCPs and 11 to 20 large financial institutions. Barker et al. (2017) argue that banks need to model their interconnected exposures to CCPs and that the modeling is extremely complex.

to the peak in aggregate. At the same time, initial margin requirements, which cover CCPs against potential losses if a clearing member defaults, increased $300 billion in aggregate. These requirements were met, but the potential stress could have been worse, as it was attenuated by the strong liquidity position of banks. Critically, there were no major concurrent defaults, like Lehman during the 2009 financial crisis.

Defaults can create additional liquidity draws, as was seen in the September 2018 default of a large clearing member at the Swedish CCP, Nasdaq Clearing. The default originated due to the member's inability to meet the CCPs' liquidity demands, which were in turn generated by extremely large market movements (in this case, in the spread between Nordic and German Power futures). Losses stemming from the default quickly ate through much of the CCP's prefunded buffer, and surviving members found themselves required to replenish €107 million of that buffer within a few days.[7] While some elements of this event were idiosyncratic—and it fortunately occurred amid otherwise benign market conditions—the difficulties in liquidating the portfolio and the resulting impact on other clearing members show the potential for liquidity problems at a CCP to amplify systemic risk. The implication is that the observed liquidity strains at CCPs during COVID-19 did not fully reflect the systemic risk that could have materialized if there had been a significant default.

Procyclicality of margin requirements has been studied by Murphy, Vasios, and Vause (2014) and Glasserman and Wu (2018), among others. Regulatory bodies have also highlighted margin procyclicality as a concern prior to March 2020 (e.g., CPMI-IOSCO 2012, 2016). But this issue has not been explicitly tied to systemic risk, and we are not aware of any studies that focus on the impact of this procyclicality on bank condition.[8]

[7] Based on materials downloaded from NASDAQ Nordic's website on August 4, 2020.

[8] Duffie (2014) studies how collateral demands shift with the introduction of mandatory margin and central clearing. Similarly, Heller and Vause (2012) show that the current set of rules can place significant liquidity burdens on clearing members, potentially contributing to their failure, and Gibson and Murawski (2013) study banks' trading behavior and welfare under different margin regimes.

Moreover, we emphasize the additional demands that CCPs may place on banks during stressful periods, beyond margin requirements. Most large banks, particularly G-SIBs, have multiple types of obligations toward multiple CCPs simultaneously. Some of these obligations are more likely to come due or to be larger during times of heightened market volatility. These obligations have received little attention in the literature even though, during severe crises, they could be substantial.[9] Liquidity obligations to CCPs could concentrate at banks during an already stressful period and could strain a bank's ability to manage their liquidity, either by making it challenging to maintain regulatory liquidity requirements or in extremis by exceeding a bank's funding capacity.[10] In either case, CCPs' liquidity demands could increase systemic risk, either by decreasing the liquidity available to support other market functions or by increasing the likelihood that payments are missed or even that institutions default.

To be clear, we do not argue that the potential for procyclical liquidity demand outweighs the systemic benefits provided by CCPs, or that the world would be better off if CCPs did not make such demands. Both prudence and regulatory directives lead CCPs to build stronger defenses during periods of greater realized and anticipated stress, and some amount of procyclicality is an unavoidable consequence of the dynamic risk management required of CCPs. Nevertheless, CCPs are most critical precisely when systemic stress occurs, heightening concerns over whether CCPs' dynamic risk management sufficiently anticipates systemic stress. Although individual CCPs regularly stress-test the sufficiency of their own liquid resources, that "micro" approach is not designed to estimate the potential liquidity impacts across multiple CCPs at a time of

However, the analyses in these papers are static; they do not consider how liquidity demands vary over time with financial market conditions. Sidanius and Zikes (2012) estimate that collateral demand from margin requirements (both in and out of CCPs) would be approximately twice as large under "stressed" market conditions as it is under "normal" conditions.

[9]Maruyama and Cerezetti (2019) also study the inherent procyclicality of CCP risk management, including some of the additional demands that banks could face, but do not focus as much on financial stability.

[10]Under extreme market stress, regulatory requirements may be relaxed, increasing the amount of availability liquidity, but not completely forestalling the possibility that liquidity demands could exceed capacity.

financial stress.[11] This observation motivates the call for macroprudential liquidity stress tests that look across multiple institutions simultaneously with a focus on the systemic impact. (See CPMI-IOSCO 2017 and Anderson, Cerezetti, and Manning 2020.) Such macroprudential stress tests could provide insights into potential aggregate liquidity demands in response to an extreme but plausible market event. At the end of this paper, we discuss some proposals for incorporating CCP liquidity into a macroprudential stress-testing framework. Insights from such tests could be crucial in more accurately assessing whether CCPs and the broader financial system could weather any procyclical demands for liquidity emanating from central clearing. In addition, most bank stress-testing implementations (e.g., those surveyed in Borio, Drehmann, and Tsatsaronis 2014 and Schuermann 2014) have not incorporated liquidity exposures to CCPs. After CCP macroprudential stress testing matures, integrating bank and CCP macroprudential stress testing should be considered.

Our discussion of the procyclicality of CCP demands fits into a broader literature showing how synchronized liquidity needs in the financial sector can be destabilizing and lead to systemic risk.[12] Adrian and Shin (2008, 2010) show that intermediaries broadly decrease leverage during periods of financial stress, and they argue that this behavior is indicative of procyclical liquidity in the financial system as a whole. Allen, Carletti, and Gale (2009), Acharya, Shin, and Yorulmazer (2011), and Heider, Hoerova, and Holthausen (2015) present models in which bank demand for liquidity is inefficiently high during crises, contributing to instability of the interbank market. Ashcraft, McAndrews, and Skeie (2011), Acharya and Merrouche (2012), and Berrospide (2021) document that banks did

[11]In addition to CCPs' own liquidity stress tests, the U.S. Commodity Futures Trading Commission (CFTC) and the European Securities and Markets Authority (ESMA) both have conducted liquidity stress testing using their own scenarios applied jointly to sets of CCPs, but the objective has remained focused on the micro objective of assessing individual CCPs' resilience (see CFTC 2017 and ESMA 2018a, respectively).

[12]The contingent obligations of banks to CCPs also have implications for system-wide liquidity measurement. For example, though potentially large, these commitments do not generally factor into measures of aggregate liquidity conditions that have been proposed in the academic literature (Berger and Bouwman 2009, 2017; Bai, Krishnamurthy, and Weymuller 2018).

indeed hoard liquidity during the 2008 financial crisis.[13] These shifts in the demand for safe and liquid assets can cause financial institutions to fire-sale securities (Shleifer and Vishny 2011, Greenwood, Landier, and Thesmar 2015) and reduce lending (Cornett et al. 2011, Iyer et al. 2014). The increasing contingent obligations of banks to CCPs constitute an additional source of liquidity pressure—likely correlated with both asset and liability draws during stressful periods—adding to the funding pressures faced by banks during stress.[14]

Section 2 of this paper provides a brief overview of the U.S. CCP landscape and discusses how CCPs fit into the broader financial system. Section 3 reviews the basics of CCP operation and risk management. Section 4 discusses various ways that CCP risk management can extract liquidity from the rest of the financial system during times of financial stress. Section 5 describes the role of liquidity stress tests at both the CCP and the macro level. Section 6 concludes.

2. The U.S. CCP Landscape

As cataloged by Kroszner (1999, 2000), CCPs developed in the United States to address counterparty risks in response to financial crises. Indeed, the essential function of a CCP is transforming counterparty exposure. A CCP steps into each trade between its clearing members. By design, a CCP is a very restricted counterparty, as it

[13]Cornett et al. (2011) and Ippolito et al. (2016) emphasize that bank exposure to liquidity risk during stressful periods can come both from creditor demands and from off-balance-sheet commitments, such as lines of credit. While the ability of banks to provide liquidity on both sides of their balance sheets has traditionally been seen as a stabilizing force (Kashyap, Rajan, and Stein 2002), this is only true to the extent that liquidity outflows are imperfectly correlated. Ivashina and Scharfstein (2010) show that draws on lines of credit occurred during the crisis at precisely the time that banks also faced difficulty rolling over their liabilities.

[14]Similar effects extend outside of the commercial banking sector. In particular, liquidity shortages may affect broker-dealers' willingness to provide securities financing to clients. Breach and King (2018) demonstrate this relationship empirically. In Brunnermeier (2009), increases in the collateral demanded by dealers for securities financing leads to asset-price volatility that feeds back into funding conditions in a destabilizing spiral. To the extent that dealers experience or anticipate extraordinary liquidity demands from CCPs, they may respond by making funding conditions more restrictive, increasing the likelihood of such a spiral.

does not take any positions on its own behalf. By being the "buyer to every seller and seller to every buyer," a CCP always maintains a perfectly matched book and therefore takes no direct market risk in the markets it clears. However, counterparty risk is concentrated at the CCP, and the CCP is exposed to contingent market risk upon the default of a clearing member, because it acquires the defaulting member's positions. A CCP then has to take steps to return itself to a matched book and flat market-risk position. (We discuss CCP risk management in the next section.)

CCPs fared well, and served their role as a buffer against defaults, during the 2008 financial crisis. For example, from Valukas (2010), the Chicago Mercantile Exchange (CME) closed out the cleared derivatives portfolio of Lehman Brothers, which had a net value around $21 billion in May 2008, within a few days; Lehman's margin was sufficient to cover auction-related losses amounting to $1.2 billion. Similarly, at LCH SwapClear, Lehman's interest rate swap portfolio, with a notional value of $9 trillion spread over 66,390 trades across five currencies, was closed out by early October (LCH.Clearnet 2008). As discussed in Fleming and Sarkar (2014) and Wiggins and Metrick (2019), central clearing did face some challenges and consequent criticism. Yet, despite this stress, as shown in Figure 1, the percentage of cleared IRS trades (the red line) jumped from less than 25 percent in 2007 to nearly 50 percent by 2009, as traders sought the safety of CCPs during the crisis. In 2009, the G-20 leaders committed to clearing all standardized OTC derivatives contracts through CCPs. The goal was to address risks in the bilateral OTC derivatives' markets, including large concentrations of counterparty exposures, inconsistent risk management, a scarcity of prefunded resources to cover realized losses, a lack of transparency, and adverse feedback loops (e.g., margin spirals). Subsequently, as shown in the figure, clearing of IRS rose to 75 percent.

In the years since the 2008 crisis, a variety of factors have led to further increases in clearing volumes for both exchange-traded and OTC derivatives. The most direct impetus was the establishment of central clearing mandates for OTC derivatives noted above.[15]

[15]Culp (2010) reviews the regulatory history of OTC derivatives' clearing through the Dodd-Frank legislation. FSB (2018b) analyzed the incentives to centrally clear created by various reforms beyond just clearing mandates.

Figure 1. Interest Rate Swaps Outstanding

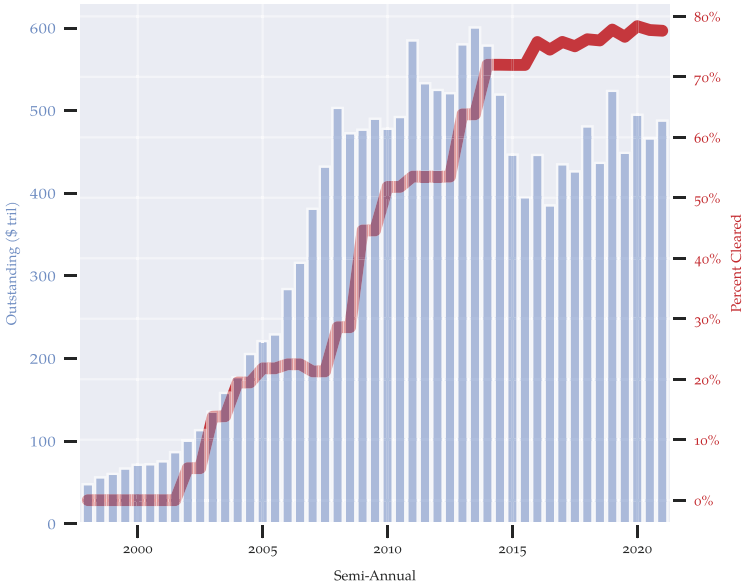

Note: The blue bars, associated with the left-hand scale, represent IRS notional outstanding; the red line, associated with the right-hand scale, shows the percentage centrally cleared.
Source: Bank for International Settlements.

In Figure 1, for example, the jump in cleared interest rate swaps from about 50 percent in 2012 to near 75 percent by 2014 is coincident with mandates coming into force. Similarly, CDS clearing, which began at ICC (formerly ICE Trust) in March 2009 and at ICE Clear Europe (ICEU) in July 2009, steadily climbed to nearly 40 percent by 2017 according to BIS data (not shown in the figure).[16] This increase came despite delays in implementing clearing mandates. For both IRS and CDS, over 80 percent of new dollar-denominated trades are now centrally cleared. Some of the growth in central clearing likely reflects an increased appreciation for the risk-mitigating function of CCPs following the crisis. It also likely reflects market moves away from more exotic and bespoke trades

[16]CME also started clearing CDS in 2009, but terminated the service in early 2018.

to more standardized ones. The increased clearing rate means that clearing volumes have held roughly steady even as notional amounts fell by a third in the last few years. More recent support for central clearing has been provided by bilateral margin requirements; FSB (2018b) found that OTC derivative clearing dramatically accelerated in terms of notional cleared upon the implementation of bilateral margin requirements, even for trades where central clearing was not required. These developments helped to contribute to a more resilient financial system that was able to weather the COVID-19 market stress (FSB 2021).

In the United States six CCPs clear the most important financial markets; these are listed in Table 1. Five of these CCPs are designated as systemically important in the United States. The sixth, London Clearing House SwapClear (LCH), is a U.K. CCP but clears a substantial amount of U.S. dollar-denominated interest rate swaps. These six CCPs can be classified as clearing securities and derivatives, respectively. The securities CCPs include the Fixed Income Clearing Corporation (FICC), which operates two separate clearing services: one for U.S. Treasuries and repurchase agreements (repos), and one for mortgage-backed securities (MBS). The National Securities Clearing Corporation (NSCC) also clears securities; it is the primary clearer for U.S. equity markets, and also clears corporate and municipal bonds. These securities CCPs are all part of the Depository Trust Clearing Corporation. Derivatives CCPs include: the Options Clearing Corporation (OCC), which is the primary clearer for U.S. equity options; the Chicago Mercantile Exchange Inc. (CME), which operates two separate clearing services, one for futures and options (Base), and one for interest rate swaps and swaptions; ICE Clear Credit (ICC), which clears credit default swap indexes and single names; and, LCH SwapClear, which also clears IRS. Derivatives CCPs can further be divided by whether the cleared derivatives are exchange traded or OTC. OCC and CME's Base service clear exchange-traded derivatives. The remainder—CME's IRS service, ICC, and LCH SwapClear—clear OTC derivatives.[17]

[17]Under the Dodd-Frank legislation, the Securities and Exchange Commission (SEC) is the supervisory agency for FICC, NSCC, and OCC while the CFTC is the supervisory agency for CME and ICC. Because ICC is registered with

Table 1. Six Major CCPs Clearing U.S. Markets

CCP	Main Products	Approx. Prefunded Resources ($ Bil.)		Max. Daily Margin Call Since 2015 ($ Bil. Est.)	
		2019	2020:Q1	Pre-2020	2020:Q1
	Securities				
FICC		39.7	66.6	16.2	34.2
GSD	U.S. Treasuries and Repos				
MBSD	Mortgage-Backed Securities				
NSCC	U.S. Equities, Corps., and Munis	12.5	36.7	8.6	24.5
	Derivatives				
OCC	U.S. Equity Options and Futures	57.2	103.0	17.3	31.8
CME		139.4	239.8	16.0	21.3
Base	Commodity and Financial Futures and Options				
IRS	Interest Rate Swaps and Swaptions				
ICC	Credit Default Swaps	36.7	53.3	2.4	9.1
LCH SwapClear	Interest Rate Swaps	170.7	206.9	24.9	32.9

Note: Maximum daily margin call estimated by adding the peak VM and IM calls during each quarter.
Source: Quantitative disclosures via Clarus CCPView. Data through 2021:Q3.

Measuring the size of CCPs is not entirely straightforward, but Table 1 provides some context. The first numeric column presents approximate total prefunded resources in 2019; although substantial, the resources also jumped significantly—for NSCC, nearly tripling—in the first quarter of 2020, as shown in the second numeric column. The last two columns present the maximum daily call for margin since 2015 prior to 2020 and during the first quarter of 2020, respectively. Again, the pre-2020 peaks were large but were dwarfed by the calls during the COVID-19 market stress. Note that the measures of size used here likely understate the economic impact of securities clearing houses, because they do not capture the total cash that flows through securities CCPs, which is substantial due to the settlement demands for cash. Despite measurement challenges, clearly, all the CCPs have significant resources available, but also all have made significant calls for additional funds from the market, both before the market shock in early 2020 and even more so during it. This pattern is consistent with the analysis in BCBS-CPMI-IOSCO (2022).

3. CCP Function and Risk Management

The concentration of counterparty exposure at a CCP means that the success, or failure, of its risk management controls could have profound implications for the markets it clears. Success means defaults are not amplified with potential systemic implications; failure, although designed to be an extremely remote possibility, potentially creates contagion effects between clearing members that would not otherwise have been exposed to the default. CCPs manage their credit and liquidity risks differently than banks, partly reflecting that CCPs do not actively trade financial contracts, but rather manage the risks generated by participants' trading activity.[18]

---

the SEC as a clearing agency, the SEC also has authority over all of its clearing services. The Federal Reserve has authority to participate in the designated supervisor's examinations and reviews of material changes at the designated CCPs. LCH SwapClear is registered with, and supervised by, the CFTC for the clearing of U.S. dollar-denominated positions; the Federal Reserve also participates in international oversight of LCH SwapClear under the auspices of the Bank of England.

[18]In the event of a clearing member default, a CCP may trade in order to hedge and liquidate the defaulter's portfolio, but usually needs to second traders from

Figure 2. Variation and Initial Margin

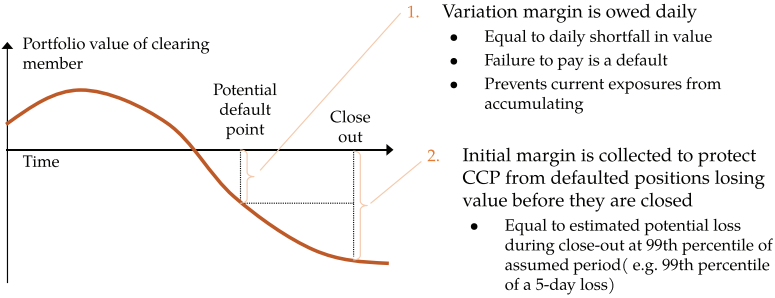

CCP risk management generally includes membership eligibility requirements, the netting of exposures, margin requirements, mutualized financial resources (usually in a default or guaranty fund), and default-management procedures. See Murphy (2013) for a fuller discussion. CCPs may also have additional tools to employ in recovery scenarios, such as assessment powers, variation margin haircutting, contract tear-ups, or loss allocation powers (CPMI-IOSCO 2014, Section 4). In addition to the management of credit risk, CCPs also manage liquidity risk, and have a variety of tools to address liquidity needs, such as access to collateral markets and pre-arranged repo lines. With an eye towards the subsequent discussion of procyclicality, we review two of the most important elements of CCP risk management: margin and mutualized financial resources, which represent prefunded resources held by the CCP. Both resources are critical to managing and mitigating a CCP's contingent risk exposures and are two major drivers of procyclicality in CCPs' resource demands.

Margin is a critical component of CCP risk management. It can be divided into variation margin (VM) and initial margin (IM); Figure 2 illustrates how VM extinguishes a CCP's current exposure and IM covers potential future exposure. We discuss each type of margin in turn.

---

its clearing members to do so. For further discussion of the differences between banks and CCPs, particularly the different roles played by capital and collateral in their respective risk-management approaches, see Manning and Hughes (2016) and Cox and Steigerwald (2018).

Variation margin covers the realized change in a cleared position. Regulations generally require CCPs to collect and pay out in cash the daily change in value that each member's portfolio experiences as VM.[19] In the stylized example in Figure 2, the change in value of the portfolio at the end of the day is negative. The clearing member must pay this amount to the CCP. The total value of the payments of VM exactly equals the amount owed to clearing members whose portfolios gained in value. By marking every portfolio to market daily and exchanging cash to cover the changes, the CCP resets its current exposure to zero every day and prevents exposures from accumulating.[20] The requirement to pay VM, however, also creates a point of failure; if a clearing member fails to make a VM payment for itself or its clients in the time required, the CCP can, and likely will, declare the clearing member to be in default.[21] As emphasized by Maruyama and Cerezetti (2019), VM is the main driver of CCP margin calls. The requirement to meet VM effectively tests each clearing member's performance at least daily, which in some sense transforms the credit exposure it faced with its original counterparty to a liquidity exposure to the CCP (Cont 2017). Though not without its own risks (as we discuss in detail below), the liquidity exposure has the virtue of transparency, since it depends directly on observable market movements, rather than on difficult to observe counterparty actions.

Initial margin is designed to cover the potential future exposure of a clearing member's portfolio from the time of the last VM payment until the portfolio could be liquidated in a default. Because the loss that could be realized on a portfolio is uncertain, IM targets a high quantile—at least the 99th percentile—of market moves over the specified close-out period (also sometimes called the margin period of risk). In Figure 2, IM covers both the VM that was not

[19] There are exceptions. In particular, FICC's clearing structure for mortgage-backed securities collects VM, but does not pass it through. OCC collects little VM, because the covered options positions it clears are generally hedged by the underlying securities.

[20] In addition, VM is sometimes exchanged intra-daily to more actively limit the buildup of current exposures that might occur more rapidly, for example, in more volatile markets.

[21] Armakolla and Laurent (2017) show the importance of members' abilities to meet their obligations for the resilience of a CCP.

paid during a default plus additional losses on the position that could be realized during the close-out. At their core, CCPs' IM estimates for a portfolio bear similarities to standard market risk calculations, like value-at-risk or expected shortfall. However, regulatory requirements also specify that IM covers other exposures that are more difficult to quantify, like estimates of portfolio transaction costs, jump-to-default exposures, and concentration risks. Heckinger, Cox, and Marshall (2017) review both historical and current IM practices. Looking across certain derivatives CCPs for futures and swaps, as shown in Figure 3, the size of IM requirements more than doubled from less than $194 billion at the end of 2013 to over $425 billion in August 2019; that steady increase was matched by roughly an additional $200 billion in March 2020 alone, and at the end of 2021 IM is more than triple that held in 2013 at these CCPs.[22] This growth in IM implies growth in the amount of risk CCPs are managing, although not in direct lockstep.

To cover potential losses beyond margin in the event of a default, CCPs require clearing members to provide mutualized financial resources sufficient to meet a specific coverage target in a wide range of extreme but plausible stress scenarios. The mutualized resources are sized so that the CCP has sufficient financial resources to cover the default of either any single or any pair of clearing members in stressed market conditions.[23] At most CCPs, mutualized resources are maintained in a separate fund, called a default or guaranty fund. The U.S. securities CCPs instead mutualize all IM. In either case, the adequacy of these resources is tested daily by the CCP through its stress-testing program, which is required to apply a variety of extreme but plausible scenarios: both historically observed market stresses and hypothetical scenarios.

To cover any losses from a default, a CCP would use prefunded resources in a prescribed sequence, often called the CCP "waterfall." The defaulter's margin absorbs losses first. If the defaulter's margin

[22]The data include requirements from more than the systemically important CCPs. Besides ICEU, which clears energy futures and options in addition to CDS, and other parts of LCH Ltd. besides SwapClear, requirements from ICE US, which clears a variety of futures and options, and LCH SA, which is the European sited counterpart to LCH Ltd. in the United Kingdom are included.

[23]Specifically, in the United States, FICC, NSCC, and OCC are required to cover the single largest default loss, while CME and ICC are required to cover the largest two default losses, as is LCH SwapClear.

Figure 3. IM for IRS, Futures, and CDS

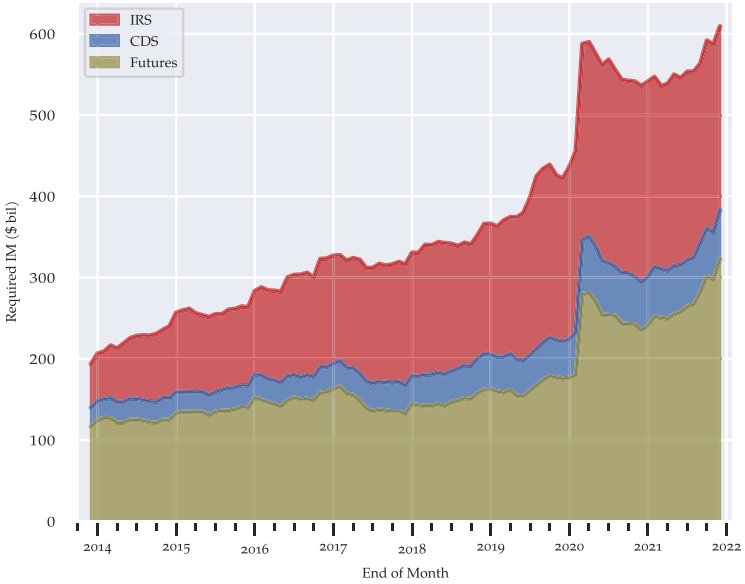

Note: Total requirements held by clearinghouses from clearing members, including add-ons. IRS data include CME and LCH Ltd. Futures data include CME, ICEU, and ICE US. CDS data include CME, ICC, ICEU, and LCH SA. Data are month-end through December 2021.
Source: CFTC.

is exhausted, any additional resources provided by the defaulting party, such as its contributions to a default fund, are applied. If the defaulter's resources are exhausted, before mutualizing losses to surviving clearing members, CCPs typically apply a portion of their own capital—known as "skin in the game"—to cover remaining losses. Mutualized resources would be applied to any remaining losses. CCPs are required to have explicit rules and procedures that address how potentially uncovered credit losses would be allocated. They may also have additional prescribed assessment powers to seek further resources from non-defaulting members.

The strong and consistent risk-management practices adopted by CCPs have been enhanced by stronger regulatory expectations promulgated in the wake of the 2008 financial crisis. This strengthening was facilitated by the development of internationally agreed

risk-management standards set out in the Principles for Financial Market Infrastructures (PFMI).[24] The principles in the PFMI set several expectations for CCPs, including effectively managing all dimensions of the CCP's credit and liquidity risks, employing a robust margining system, and maintaining a minimum level of financial resources to cover potential losses and honor payment obligations, both in extreme market conditions. The principles are designed to ensure that CCPs will halt contagion among their members and mitigate systemic risk across the interconnections.

4. Potential Procyclical Resource Demand

Any reduction in systemic risk afforded by central clearing depends on CCPs performing as designed in stressed markets. As noted above, CCPs have grown both in size and in the scope of the products they clear. Furthermore, individual systemically important banks tend to participate in multiple CCPs, in order to access different markets across different jurisdictions. For example, the recent analysis in FSB (2018a) shows that out of a set of 26 major global CCPs, the 11 largest clearing members in terms of aggregate resources are members of at least 16 of the 26, with a median participation of 22. Many of these same large financial institutions also provide other services to CCPs, such as settlement and investment services and lines of credit. This interconnectedness, together with CCPs' critical functions supporting markets, naturally raises the need to monitor their impact on financial stability.

CCPs' risk-management controls can affect the other financial institutions expected to meet one or more resource demands, particularly if CCP resource demands come at times of heightened financial stress. These demands include (i) variation margin; (ii) initial margin; (iii) settlement requirements; (iv) default fund contributions and assessments; (v) lines of credit and other liquidity arrangements; and (vi) capital and liquidity to absorb the positions of members and their clients in the event of default. In this section, we review each of these potential resource demands in turn.

From a systemic risk perspective, there is a trade-off between rigorously managing the risk of positions cleared through CCPs and

[24]See CPMI-IOSCO (2012).

minimizing interconnected liquidity strains on the system. The failure of a large CCP would be a devastating event for the global financial system; it would likely be accompanied by massive disruptions in financial markets and the severe distress or failure of many other systemically important institutions. From a financial stability perspective, it is important for CCPs to maintain financial resources sized to cover potential losses in a wide range of stress scenarios. One hypothetical way to achieve this would be for CCPs to prefund all of these financial resources ex ante by holding very large quantities of safe and liquid assets at all times. Such an approach would tie up liquidity and collateral in accounts that, most of the time, would be well in excess of any plausible loss that CCPs might experience. Moreover, the requirement to keep large amounts of assets parked at CCPs continuously would impose costs on clearing members; those costs could well induce banks and their clients to move activity into uncleared positions (eliminating the stability benefits of central clearing) or to exit positions altogether, which could result in less hedging activity and a reduction in overall market liquidity. Rather than holding an extremely large buffer of financial resources all the time, CCPs hold smaller (though still significant) amounts of liquid resources during calm periods and increase these resources during times of high financial market volatility, when their market and counterparty risks increase. The increase must be met by a web of banks and other financial institutions, which implies an increase in the risk associated with the CCPs' interconnectedness.

Importantly, the most common resource draws that come from CCPs, margin calls, are not unique to clearing relationships. Prior to implementation of the Dodd-Frank Act and similar regulatory changes abroad, many bank arrangements with non-bank counterparties included provisions for IM, and for most derivative exposures there was some periodic exchange of VM. However, these arrangements were somewhat ad hoc and flexible. With central clearing, margin requirements are universal for cleared contracts and, to a large degree, standardized due to consistent regulatory expectations.[25]

[25]Margin requirements, both for VM and IM, have also been established for bilateral trades (BCBS-IOSCO 2015); for analysis of their impact on the incentives to centrally clear, see FSB (2018b).

Figure 4. Percentage of Client IM Requirements at Five Largest Parent Firms

Note: Futures data include CME, ICEU, and ICE US. Swaps data include both IRS and CDS requirements. IRS data include CME and LCH Ltd. CDS data include CME, ICC, ICEU, and LCH SA. Data are month-end through December 2021. Identity of top five firms can vary from month to month.
Source: CFTC.

Nevertheless, because clearing takes place through a small number of CCPs, the models and rules used by any one CCP can have significant implications for a large number of market participants. In addition, much of the activity at CCPs is concentrated in a few large members. The concentration is illustrated in Figure 4, which shows that the five largest clearing members also account for the majority of IM requirements facing clients for dollar-denominated clearing of both futures and swaps; in fact, at one point the five largest clearing members accounted for nearly three-quarters of clients' swap IM requirements, before declining to around 60 percent. The concentration associated with CCPs and their large members implies that margin calls are likely to come in a more coordinated manner in a world with central clearing, reflecting stronger interconnectedness. Furthermore, as discussed below, CCPs also impose other types

Figure 5. Daily Peak VM Paid versus Realized Volatility

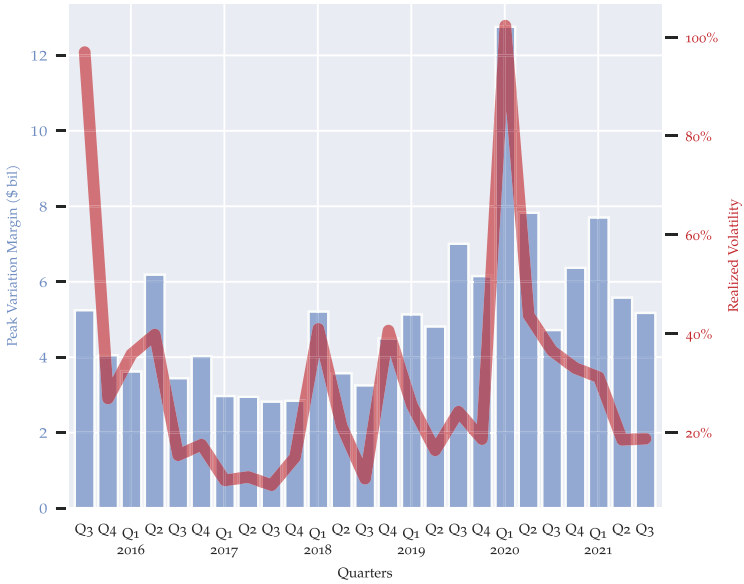

Note: Peak VM, represented by the blue bars associated with the left-hand scale, is the industry average of the maximum paid by/to each CCP in the quarter. Realized volitility is annualized.
Source: Quantitative disclosures accessed via Clarus CCPView and Heber et al. (2009). Data through 2021:Q3.

of obligations on their clearing members—obligations that have no counterparts in bilateral trading—and these are more likely to be triggered in the same states of the world in which margin calls are large.

4.1 Variation Margin

As noted earlier, VM is the amount that a clearing member must post to the CCP to cover marked-to-market changes in the value of its portfolio. Maruyama and Cerezetti (2019) illustrate how VM is the margin component that is most sensitive to market volatility. Because changes are greatest on days when market prices move most, the amount of VM paid to CCPs necessarily increases during times of high realized volatility in financial markets. Figure 5

illustrates this relationship by showing the peak amount of VM collected in each quarter by the CCPs through which U.S. banks clear most of their trades. This margin is plotted against the peak in the realized volatility of the stock market in each quarter. (Although the data do not tell us for certain, it is likely that the peaks in the two series occur on the same days of each quarter.) This comparison is crude because some CCPs do not clear products that are directly linked to equity markets. Nevertheless, the tendency of volatility in most markets to move together makes the comparison informative, and the pattern is clear. Not surprisingly, the peak in both is during the first quarter of 2020.

To meet VM calls, clearing members must make payments to the CCP in a short amount of time, often in as little as one hour, and these payments generally are made in cash. Furthermore, although clearing members who clear on behalf of clients pass through VM calls on client positions, this pass-through sometimes does not occur until the following day, and the clearing member is itself responsible for ensuring that the proper amount of margin is posted to cover their clients' losses. Clearing members and their clients expect to make (or receive) VM daily; however, because they reflect changes in asset prices, the size of VM calls are essentially unpredictable. To address this risk, clearing members maintain reserves of cash or have access to funding markets that can be drawn on with very short notice.

Even in times when financial institutions are not under severe pressure, VM calls due to spikes in volatility can be burdensome. Events surrounding the "Brexit" referendum in the United Kingdom in June 2016 provide an example. The referendum, the result of which surprised many, caused a spate of volatility in financial markets, particularly those involving the British pound, the euro, and associated fixed-income products. LCH, which clears many such instruments, called for large amounts of VM in the wake of this episode. Figure 6 shows that on the peak day during the second quarter of 2016, VM payments to LCH totaled $16 billion, in contrast to the daily average of about $3 billion. These calls came simultaneously with smaller VM calls from other CCPs; Commodity Futures Trading Commission (CFTC) (2016) found that LCH together with four other CCPs called for $27 billion over the two days following the

Figure 6. VM at LCH SwapClear

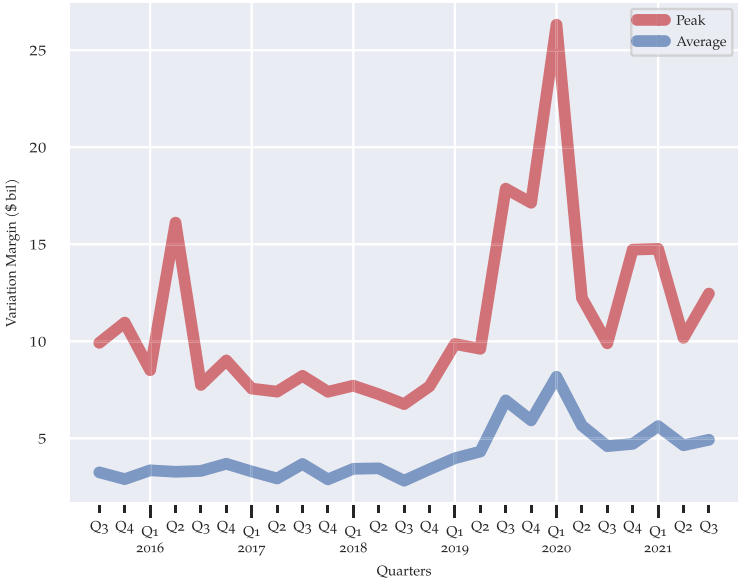

Note: Peak VM, the higher red line, is the maximum paid by/to LCH SwapClear in the quarter. The blue line shows the corresponding quarterly average.
Source: Quantitative disclosures accessed via Clarus CCPView. Data through 2021:Q3.

referendum.[26] The peak demand caused by Brexit was much higher than the elevated average demand during COVID-19, even though much larger calls were made at LCH in the later period.

Under most circumstances, VM payments to CCPs do not reduce aggregate liquid resources available to market participants, because CCPs run matched books, so that every marked-to-market loss they face is accompanied by an equal marked-to-market gain. Thus, VM ought to be passed through, dollar for dollar, from one set of clearing

[26]LCH was also slower to make payments to members with position gains than it was to collect from members with losses, creating further liquidity pressure; some market participants seemed to find the obligation to post this margin onerous, despite the fact that the financial system itself did not appear to be under a particularly high level of pressure (Madigan, Wood, and Becker 2016).

members (those with net losses) to another (those with net gains).[27] Even so, the transfer of funds can cause stress for those with unanticipated liquidity outflows. Furthermore, particularly in times of stress, disruptions and delays in payment chains can occur, potentially preventing the liquid assets from flowing to their recipients in a timely manner. CCPs and settlement banks themselves may deliberately slow the transfer of liquidity out of a fear that their own liquid resources may prove inadequate. The stock-market crash of 1987, for example, was accompanied by a number of disruptions in the payments system, including the reluctance of settlement banks to provide large amounts of intraday credit. As a result, VM payments to clearing members at CME and OCC were delayed, adding to the stress of market participants on an already volatile day (Presidential Task Force on Market Mechanisms 1988). During the COVID-19 stress, some clients did report needing to conduct repo transactions and assets sales to meet margin calls (BCBS-CPMI-IOSCO 2021), which may have added additional strain on funding markets.

Finally, extreme circumstances may also prevent clearing members from receiving the full VM due to them. In particular, if a CCP's resources should be exhausted by clearing member defaults, one option some CCPs have to manage liquidity is to "haircut" the VM it pays out. In this case, the CCP would absorb, at least temporarily, part of the VM it received on loss positions. As another option in such situations, some CCPs may pay out VM in securities collateral, rather than cash, putting the burden of liquidity transformation onto members.

4.2 Initial Margin

Initial margin is posted to CCPs to account for possible deterioration in the value of clearing member positions that might occur between the time of a default and the time the defaulting member's positions can be liquidated. Like VM, IM generally increases during times of financial-market volatility. As discussed in Murphy, Vasios, and Vause (2014), the requirements to cover a quantile of market moves require IM models to be risk sensitive, and therefore procyclical, to at least some degree. However, while the pass-through

[27] There are some exceptions for intraday variation margin calls.

Figure 7. CME Futures IM Requirements and VIX Index

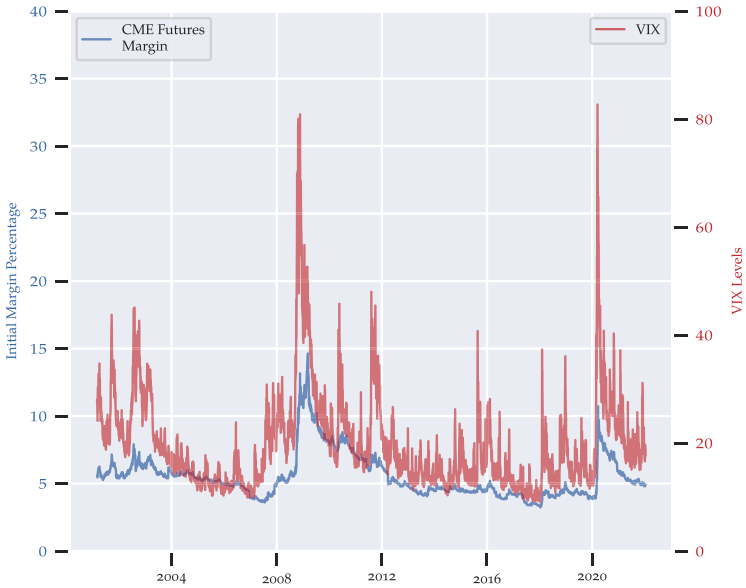

Note: The blue line, associated with the left-hand scale, shows CME's IM requirements on an S&P 500 futures contract as a percentage of the contract value. The red line, associated with the right-hand scale, is the Chicago Board Options Exchange's Volatility Index (VIX), which reflects market expectations of 30-day forward-looking implied volatility. It is calculated from S&P 500 index options.
Source: Daily data 3/1/2001–1/12/2022 from Chicago Board Options Exchange, CME, and Haver Analytics.

of volatility to VM is largely mechanical and out of the CCP's control, the factors determining IM are more complex and depend on the specific practices the CCP follows. These practices are largely spelled out in CCP rule books, but they also involve elements of discretion.

As a baseline, CCPs maintain models of the value that positions could lose with some confidence: typically, 99 percent or 99.5 percent over a specified liquidation period (typically a few days). These tail losses are almost always greater when market volatility is expected to be higher, and consequently IM tends to increase in periods of high realized and implied volatility. As an example, Figure 7 plots the amount of IM that CME has required on an S&P 500 futures

Figure 8. CME: Futures IM Requirement versus VIX Index

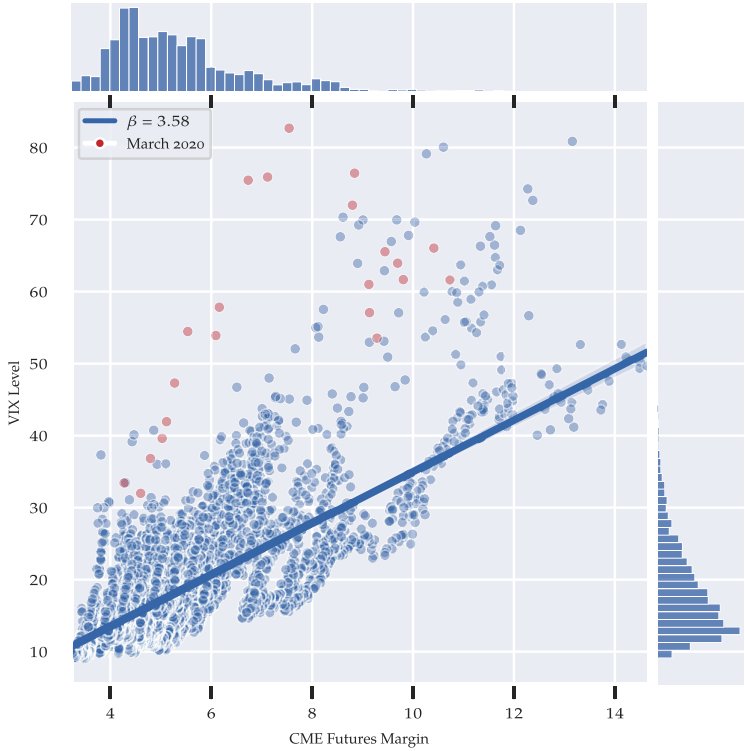

Note: The joint plot shows CME's IM requirements on an S&P 500 futures contract as a percentage of the contract value versus VIX; the strong positive relationship is reflected in the upward-sloping regression line; this regression, which is highly significant as shown by the barely visible 99th percent bootstrapped confidence interval, is estimated robustly to reduce the influence of outliers that would increase the slope further. The marginal distributions, plotted above and to the right as histograms, show that both series have a long positive tail, although presumably margin as a percent is capped. The red points highlight values from March 2020 during the COVID-19 induced market shock.
Source: Daily data 3/1/2001–1/12/2022 from Chicago Board Options Exchange, CME, and Haver Analytics.

contract, as a percentage of the contract value, with the VIX index of implied volatility on the S&P. It is clear that these two series move closely together particularly during large spikes in implied volatility. The strength of this relationship is further shown in Figure 8, which

Figure 9. OCC Quarterly Peak IM Calls

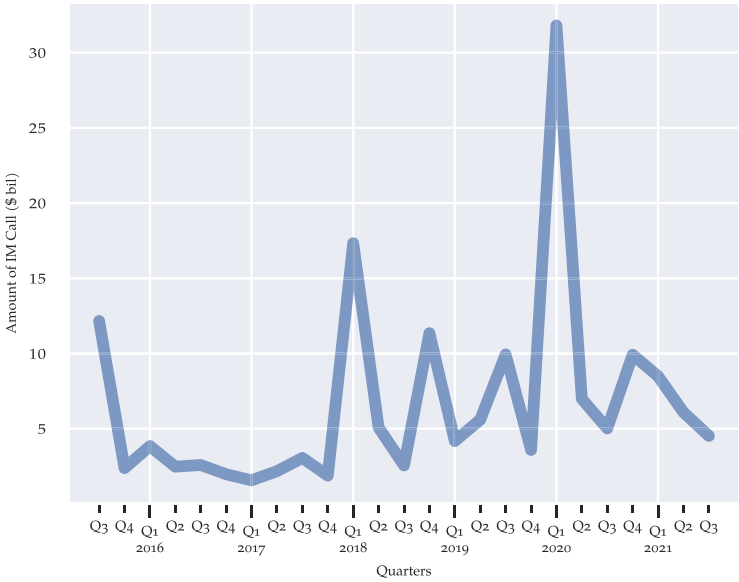

Source: Quantitative disclosures accessed via Clarus CCPView. Data through 2021:Q3.

plots the same data against each other; the significant positive relationship can be seen by the slope of the regression line.[28] The red dots show the behavior during March 2020 as the IM level increased as VIX spiked, but continued to do so even as VIX fell back.

As another example of the connection between volatility and IM, Figure 9 shows peak IM calls at the OCC, which primarily clears equity options; the jumps in equity volatility that occurred in February 2018, December 2018, and March 2020 are clearly identified as causing corresponding jumps in IM during those quarters.

Although it is related to volatility, the total amount of IM required of a clearing member is, in most cases, more complicated than the direct relationships suggested by Figures 7 to 9.

[28]The line is calculated by an M-type estimator with Huber's T norm (Huber 1981) to make the regression robust against the large positive outliers.

Figure 10. Quarterly Changes in Required IM and VIX

Note: The blue bars, associated with the left-hand scale, represent quarterly changes in the sum of the largest calls in required IM across the seven CCPs; the red line, associated with the right-hand scale, shows the change in the maximum VIX level in each quarter.
Source: Quantitative disclosures accessed via Clarus CCPView, and Chicago Board Options Exchange. Data through 2021:Q3.

The net positions that clearing members and their clients maintain with CCPs consist of large arrays of financial instruments that are exposed to different types of risk. CCP margin models attempt to account for the co-movement of different assets and employ a variety of modeling techniques to capture the tails of the joint distributions of returns. In addition, the overall level of margin charged often includes certain "add-ons," which tend to be less sensitive (to varying degrees) to market volatility (CPMI-IOSCO 2016). And, of course, the total amount of IM required tends to rise and fall with the size of participants' positions. Figure 10 shows quarterly changes in the total quantity of IM required by the major U.S. CCPs, plotted against contemporaneous changes in the VIX. The positive correlation is clear, despite the introduction of anti-procyclicality tools

designed to mute such correlation. European Securities and Markets Authority (ESMA) (2018b) codifies some anti-procyclicality tools, which have also been applied to U.S. CCPs due to their global activity. Examples of tools include a buffer that allows for a proportion of margin to be temporarily exhausted following a significant increase in margin requirements, which lessens the size of any IM calls that result from sudden spikes in volatility, and a floor that limits the amount that margin requirements can decline during tranquil periods, thereby muting increases that result when tranquility gives way to turbulence. Despite the prevalent use of these tools, BCBS-CPMI-IOSCO (2022) finds that market volatility was the major driver of IM requirements during the severe COVID-19 stress.

Another reason that CCP IM requirements may not be direct functions of market volatility is that CCPs may attach margin surcharges to particular clearing members to compensate for concentration or counterparty risk. Though counterparty-risk surcharges are rare, they can result in sudden large increases in the amount of IM margin required from financial institutions that are already in distress. For example, as discussed by Heckinger (2014), amid concerns about the risk associated with its repo positions, the futures commission merchant MF Global was downgraded by Moody's and Standard & Poors on the week of October 24, 2011. Over the following few days, LCH (through which these repo positions were cleared) imposed counterparty-risk surcharges that approximately doubled the IM that MF Global was required to hold. The resulting margin calls, which totaled over $500 million, exceeded MF Global's resources, which led directly to its failure. This sort of example illustrates how margin requirements can contain a procyclical element, even if they are not explicitly tied to asset price volatility.[29] Such unusual charges also received significant attention during the "meme

[29]Standards require IM models to incorporate anti-procyclicality measures (CPMI-IOSCO 2012), but ultimately IM will respond to changing risk exposures. Houllier and Murphy (2017), Wong and Pei (2017), and Glasserman and Wu (2018) study the problem of reducing margin procyclicality. O'Neill and Vause (2018) study how a time-varying "macroprudential buffer" added to IM can address fire-sale externalities. Raykov (2012) examines the trade-offs inherent in reducing the procyclicality of IM and finds that reducing it does not necessarily reduce systemic risk.

stock" episode in U.S. equity markets in January 2021 (Mourselas 2021).

Finally, it is not only the amount of IM required that is likely to rise during times of market stress. According to the quantitative disclosure data, approximately half of IM is in the form of securities, rather than cash, and those securities are subject to haircuts. Lewandowska and Glaser (2017) find that one major CCP does not significantly increase its haircuts during times of market stress. Nevertheless, in principle, rising haircuts and declines in market value that could be associated with episodes of market volatility could lead to changes in haircuts and would require the posting of additional collateral even if the margin requirement were unchanged.

4.3 Settlement Requirements

In discussing how VM and IM can vary with market volatility, there is a tendency to focus on the change in market prices. However, market volatility usually is also associated with increased trading activity. As a result, portfolio valuations can change because both prices and positions fluctuate. Importantly, increases in trading activity directly create liquidity needs at securities CCPs in order to settle the trades. Because the cash needed to settle securities corresponds directly to the full value of the securities trades, the resulting liquidity needs can be simultaneously material and procyclical. In terms of materiality, from the latest quantitative disclosures, the sum of the peak payment obligations in the prior 12 months was just under $194 billion. For derivatives, most trades do not require up-front payments to settle, so the effect there is at most minor.[30] From a systemic view, the liquidity needs generated by the settlement of centrally cleared securities need to be combined with the liquidity needs from VM and IM in assessing the resiliency of the clearing system overall.

[30]Derivatives contracts can have settlement requirements, such as if an option is exercised or a future matures, but such requirements likely are not strongly procyclical. CDS potentially can generate procyclical payments due to the defaults of reference entities, but such payments are difficult to predict or regularly observe.

Figure 11. Change in Average Default Fund Assessments

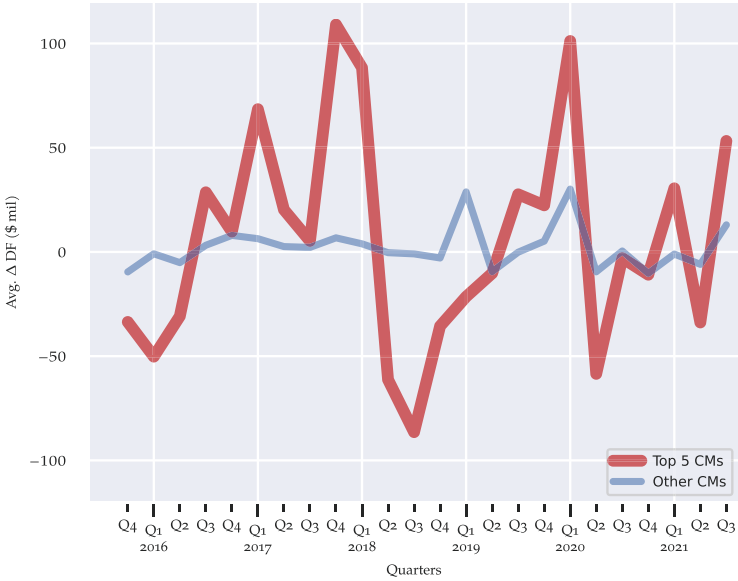

Note: The thicker red line represents average change at the six CCPs in the default fund assessments of each CCP's five largest clearing memebers. The thinner blue line shows the same average change for the remaining CMs across the CCPs.
Source: Quantitative disclosures accessed via Clarus CCPView. Data through 2021:Q3.

4.4 Default Fund Contributions and Assessments

While VM and IM calls are the most common ways in which CCPs draw in resources from their clearing members, several other contingent relationships can also come into play, particularly in times of financial stress. One important way in which CCPs can require resources from their members is by calling for contributions to the default fund. Figure 11 shows the average quarterly changes in default fund sizes per clearing member split into the five largest versus the rest.[31] Since 2015, average changes in default fund contributions have been small on average, however the average of averages

[31] This figure only includes the derivatives CCPs, as the securities CCPs do not have clearing funds that are distinct from IM.

masks volatility, particularly for larger clearing members, where even the quarterly average across the CCPs has been over $100 million. The largest call at an individual CCP was over $540 million, and even clearing members not among the five largest faced a call of nearly $160 million. Although likely less sensitive than margin, default funds can change due to changing portfolios, market risk, and risk-management practices. The fact that the change in default fund assessments was not more pronounced during the first quarter of 2020 likely reflects a lag in updating stress scenarios; such a lag across multiple CCPs may indicate a correlated weakness in CCP stress testing and therefore a higher level of systemic risk.

There are three distinct circumstances under which CCPs may call on clearing members to add resources to the default fund. The first is intra-month calls for contributions in response to changes in risk that the CCP perceives. Typically, CCPs levy default fund assessments monthly to reflect changes in market conditions and concentration that occur over the course of any given month. In times of market stress, however, CCPs may not wait for the end of the month. Intra-month default fund calls may be issued following sudden changes in market volatility.[32]

The second situation in which unscheduled payments to the default fund may occur is when a clearing member defaults and that member's IM and own default fund contribution is insufficient to cover the liquidation value of its position. An example occurred in 2013 when a topping-up of the default fund was required at the Korean CCP KRX following a trading error that resulted in a clearing member losing $45 million before a margin call could be issued. A more recent example is the previously mentioned 2018 default of a large clearing member at Nasdaq Clearing, which ultimately consumed approximately two-thirds of the mutualized default fund resources. Although both of these defaults were idiosyncratic and occurred during times of relative market calm, the obligation to replenish the default fund is a source of procyclicality, because both the likelihood of defaults and the potential impact of a realized default increase during times of market stress, when the necessary capital may be scarce. Because clearing member defaults that are

[32]Increased default fund needs could alternatively be addressed by issuing margin calls.

Figure 12. Unfunded CM Commitments as a Percent of Total Default Resources Excluding IM

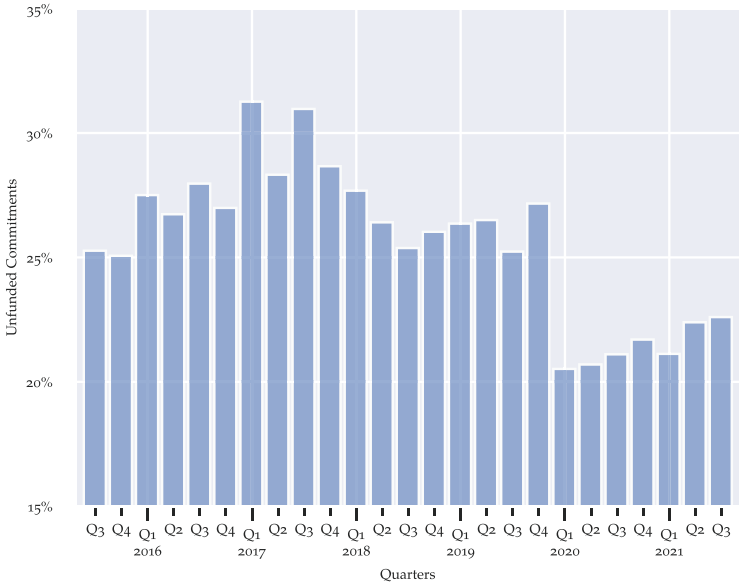

Source: Quantitative. disclosures accessed via Clarus CCPView. Data through 2021:Q3.

large enough to breach a CCP's mutualized default resources are quite uncommon and fortunately did not occur during the pandemic, there is also a risk that market participants may not be attuned to or prepared for the resource demands such defaults could generate. In addition, and from the perspective of a given clearing member, there is a distinction between the risks posed by VM and IM calls relative to calls for additional mutualized default fund resources, regardless of either the frequency or size of the associated calls, since the former two relate entirely to the portfolio that a given clearing member (CM) brings to the clearing house, while the latter, to a degree, depends upon the portfolios brought by other CMs, which are outside of the given CM's immediate control.

Finally, CCPs have powers of assessment in the event that the prefunded portion of the default fund is exhausted. Figure 12 shows the unfunded commitments that clearing members could have to pay CCPs in such circumstances as a percentage of total default

resources excluding IM. The exercise by a CCP of its unfunded commitments is a very rare situation, since prefunded resources are calibrated to cover the losses of a CCP's single or two largest clearing members. Thus, in principle, CCPs should only need to draw on unfunded commitments in cases in which three or more members default nearly simultaneously and have insufficient margin to make up for the loss associated with their positions. The caveat "in principle" reflects that one or two member defaults also could exhaust prefunded resources if the CCP's models are inaccurate or if liquidation of defaulted positions proves unexpectedly challenging. Clearly, an environment in which those defaults occurred would be associated with very high levels of market stress and liquidity demand and could strain clearing members' ability to meet their payment obligations. Notably, the increase in prefunded resources held following March 2020 has reduced the unfunded commitments as a percentage. However, this indicates that assessment powers have not necessarily expanded even as CCPs seem to view risk as being higher, as reflected in their increased holding of prefunded resources.

4.5 Absorbing Defaulting Member Positions

In the event of a clearing member default, the remaining clearing members may also have other obligations. In particular, they may acquire some or all of the defaulting member's positions and may also become responsible for the positions of the defaulting member's clients. As part of its default-management procedures, an OTC derivatives CCP would typically auction part or all of the defaulting member's house portfolio. Depending on the rules of the CCP, surviving clearing members may be obligated to participate in the auction. In addition, the positions of the defaulting member's clients must be transferred to remaining clearing members or be liquidated. Many market participants who clear indirectly have established backup relationships with one or more direct clearing members that they could activate in the event of the default of their primary clearer.

Absorbing this additional business places an added burden on clearing members' resources. In particular, clearing members that are banks or broker-dealers are required to hold capital against both

Figure 13. Largest Total Client Position of CM as a Percentage of Other CMs' Excess Capital

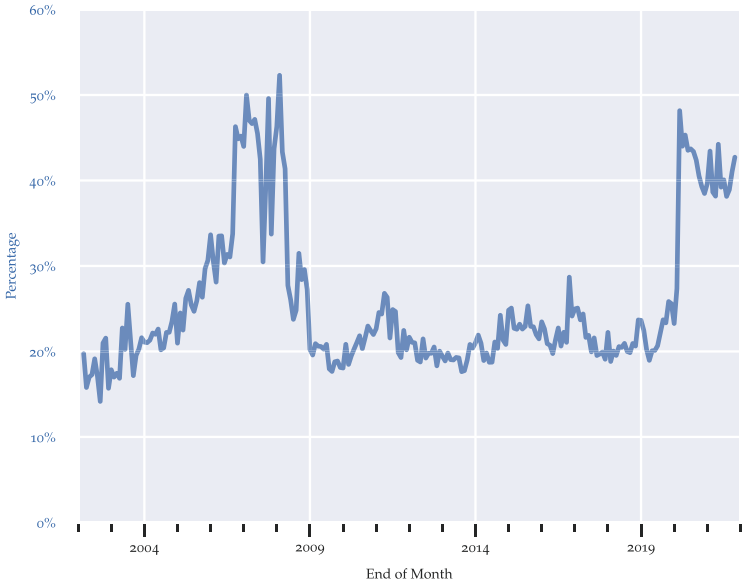

Source: End-of-month data March 2002–November 2021 from CFTC.

their own positions and those of their clients. Although institution-level data on house positions are not available, Figure 13 shows the size of the largest client positions at any clearing member, relative to the excess capital at the remaining clearing members. (The data are from the CFTC and only cover the derivatives CCPs.) "Excess Capital" represents the amount of capital that clearing members have available to support additional positions. Thus, the ratio in the graph is a measure of members' ability to absorb the client business of another member. The ratio was subdued in the years leading up to the COVID-19 crisis, partly reflecting the high levels of capital in the banking sector, but it rose quickly when volatility spiked, which is also when the probability of a default and the possible need to port client positions likely increased.

We note that, unlike many of the other procyclical resource demands noted in this paper, the need for banks to absorb house and client positions of defaulting members is largely an issue of capital. However, during stressful periods, capital and liquidity adequacy

can become intertwined. Moreover, the absorption of new positions also requires bank liquidity, because it requires posting additional contributions to CCP default funds and reserving additional liquidity buffers to meet future margin calls. Again, given that at least one member has defaulted, these commitments are likely to occur in a time when the capital and liquidity of other members are already stretched thin. Furthermore, adding new positions and posting new collateral can affect regulatory liquidity ratio requirements. Even if not binding, a bank's willingness to absorb positions could be influenced by concerns over negatively affecting its liquidity ratio, which could be reflected in its valuation of the defaulted portfolio.[33]

4.6 Liquidity Provision

Finally, CCPs maintain liquidity arrangements with many large banks. These arrangements include committed lines of credit, repo facilities, and foreign-exchange swap agreements. Ideally, the counterparties on the other side of these contracts are liquidity providers that do not face the CCP in other types of transactions. In practice, most of the institutions that are in a position to commit to providing significant liquidity are large banks that also participate directly in CCPs, clearing high volumes of derivatives and securities transactions. Consequently, since large banks tend to be members of multiple CCPs and offer liquidity services to each, the supply of these services to the overall system of CCPs tends to be concentrated. FSB (2018a) reported that 27 percent of clearing members surveyed across 26 CCPs also provide credit lines that provide liquidity to the CCP. Many of these clearing members provide such lines to multiple CCPs.[34]

[33]These concerns motivated in part the recent announcement of a revision to the treatment of client margin in calculating the liquidity coverage ratio (BCBS 2019).

[34]In the event of default, liquidity demands may be particularly large at the two securities CCPs, because such CCPs guarantee settlement of the full purchase price of securities. (The settlement value of derivatives contracts is typically a small fraction of their notional value.) For example, in the few days following the bankruptcy of Lehman Brothers in 2008, NSCC and FICC settled (without loss) over $300 billion of securities transactions that Lehman had executed with its customers and other counterparties. In contrast, the market value of Lehman's obligations under derivative contracts was about $45 billion (Valukas 2010).

Figure 14. Committed Lines of Credit as a Percent of Large Bank Liquid Assets

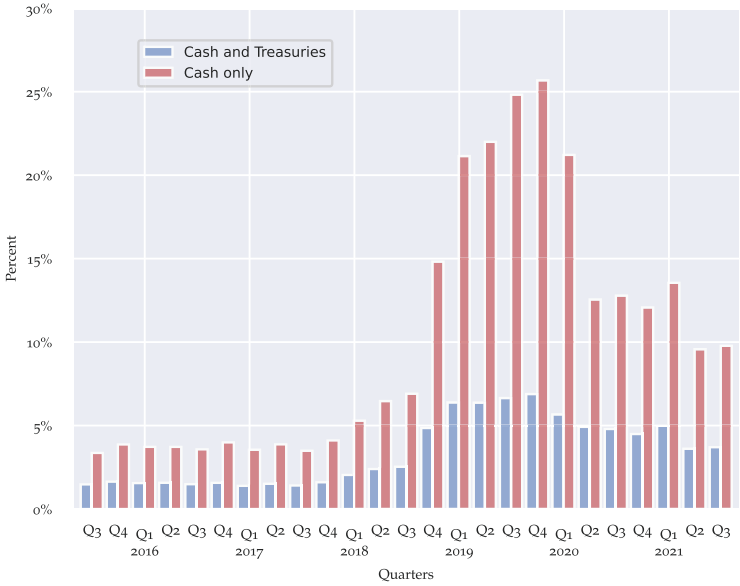

Note: Large banks are domestic banks with more than $250 billion in assets. Liquid assets includes cash and other highly liquid assets.
Source: Quantitative disclosures accessed via Clarus CCPView, and Federal Reserve H.8 release. Data through 2021:Q3.

Figure 14 shows total committed lines of credit and repo arrangements at CCPs as a fraction of the holdings of cash and cash plus liquid securities at large U.S. banks.[35] Although not all of the CCP liquidity providers are large U.S. banks, the comparison nonetheless shows that if these facilities were suddenly and significantly drawn upon, the cash demand could be material. Interestingly, starting in 2018, the significance rose particularly relative to cash holdings, both because CCPs worked to improve their access to liquidity and because banks reduced cash holdings. Fortunately, CCPs did not have to draw on these resources during the COVID-19 crisis, and the size exposures have declined somewhat as banks have maintained

[35]Large banks includes those with over $250 billion in assets.

higher cash holdings. But if needed, even if banks were able to meet this demand, doing so could well put pressure on short-term funding markets, including the interbank market and the repo market for high-quality collateral. Again, the CCPs are only likely to need to access their liquidity arrangements following the default of a very large clearing member, so it is probable that funding markets would already be under some pressure. For example, CCPs would also most likely be withdrawing large sums from their cash deposit accounts at banks in such a situation.

4.7 The Scale of CCP Resources

Tables 2–5 show how the main resources absorbed and potentially demanded by CCPs compare with various measures of the financial sector's capacity to handle those demands. In particular, we scale CCP resource draws by two measures of intermediary capital and four measures of intermediary liquidity. The measures of capital are the total Tier 1 capital at commercial banks and the capital in excess of required at the clearing members of U.S. futures commission merchants. The liquidity measures are the liquid assets (cash and Treasury securities) of U.S. broker-dealers, the cash held by large domestically chartered commercial banks, the total amount of high-quality liquid assets (HQLA) held by the six largest bank holding companies, and the amount of HQLA in excess of the liquidity coverage ratio (LCR) requirement.[36]

Table 2 compares the stock of resources held by CCPs as of 2021:Q1 with the levels of the balance sheet measures at intermediaries at that time. The six CCPs collectively held $700 billion in IM and default funds, an amount equal to several times the excess capital of clearing members' or dealers' holdings of liquid assets.

Tables 3 and 4 consider the situation in terms of flows, rather than stocks. Table 3 reports the largest daily variation and initial

[36]Total bank capital comes from the FDIC Quarterly Banking Profile. Clearing member capital comes from the monthly CFTC report on Financial Data for FCMs. Broker-dealer liquid assets comes from the Financial Accounts of the United States. Large-bank liquid assets comes from the Federal Reserve H.8 Report. Information about HQLA and LCR requirements is taken from the individual regulatory filings of JP Morgan Chase, Bank of America, Citibank, Wells Fargo, Goldman Sachs, and Morgan Stanley.

Table 2. Stock of Resources Absorbed by CCPs

Resources at Six CCPs	Total Value 3/31/21 ($ Bil.)	Relative to Financial Intermediary Balance Sheet Quantities, 3/31/21					
		Tier 1 Capital (All Banks)	Excess Capital (FCMs)	Liquid Assets (Dealers)	Cash (Large Banks)	HQLA (Large Banks)	HQLA Above Min. LCR (Large Banks)
Initial Margin	$572	30%	426%	247%	30%	22%	141%
Default Fund	$128	7%	95%	55%	7%	5%	32%
Total	$700	36%	521%	302%	37%	27%	173%

Source: Quantitative disclosures via Clarus CCPView; FDIC Quarterly Banking Profile; CFTC Financial Data for FCMs; Financial Accounts of the United States; Federal Reserve H.8 Report; regulatory filings of JP Morgan Chase, Bank of America, Citibank, Wells Fargo, Goldman Sachs, and Morgan Stanley. Data as of 2022:Q1.

Table 3. Flow of Resources to CCPs — Maximum Daily

Resources Demanded by Six CCPs	Total Value 3/31/21 ($ Bil.)	Relative to Daily Std. Dev. of Intermediary Balance Sheet Quantities, 2015–21					
		Tier 1 Capital (All Banks)	Excess Capital (FCMs)	Liquid Assets (Dealers)	Cash (Large Banks)	HQLA (Large Banks)	HQLA Above Min. LCR (Large Banks)
Max. Daily VM Call at a Single CCP	$26	1644%	2070%	431%	328%	227%	440%
Max. Daily IM Call at a Single CCP	$32	1984%	2498%	520%	396%	274%	532%

Source: Quantitative disclosures via Clarus CCPView; FDIC Quarterly Banking Profile; CFTC Financial Data for FCMs; Financial Accounts of the United States; Federal Reserve H.8 Report; regulatory filings of JP Morgan Chase, Bank of America, Citibank, Wells Fargo, Goldman Sachs, and Morgan Stanley. Data through 2021:Q3.

Table 4. Flow of Resources to CCPs — Maximum Quarterly

Resources Demanded by Six CCPs	Total Value 3/31/21 ($ Bil.)	Relative to Quarterly Std. Dev. of Intermediary Balance Sheet Quantities, 2015–21					
		Tier 1 Capital (All Banks)	Excess Capital (FCMs)	Liquid Assets (Dealers)	Cash (Large Banks)	HQLA (Large Banks)	HQLA Above Min. LCR (Large Banks)
Max. Qtrly. IM Call	$199	1538%	1938%	403%	307%	213%	412%
Max. Qtrly. Def. Fund Call	$22	170%	214%	45%	34%	24%	46%

Source: Quantitative disclosures via Clarus CCPView; FDIC Quarterly Banking Profile; CFTC Financial Data for FCMs; Financial Accounts of the United States; Federal Reserve H.8 Report; regulatory filings of JP Morgan Chase, Bank of America, Citibank, Wells Fargo, Goldman Sachs, and Morgan Stanley. Data as of 2022:Q1.

margin calls reported at any of the six CCPs over the sample period, relative to estimates of the standard deviation of daily changes in the balance sheet categories.[37] These margin calls represented changes of several daily standard deviations in each of our capital and liquidity measures. Note that these are calls at *individual* CCPs. The data do not allow us to measure the maximum aggregate daily margin calls, but they were necessarily larger and likely significantly so.

Table 4 shows the maximum quarterly changes in initial margin and default funds, compared with the quarterly standard deviation of changes in the balance sheet measures. Again, the maximum aggregate IM call in our sample was equal to a change of several standard deviations in intermediary cash and liquidity positions at the quarterly frequency. The maximum default fund call over this sample was smaller, but it still represented a large flow relative to typical quarterly changes in intermediary capital.

Finally, Table 5 reports the explicit claims that CCPs had on financial institutions as of 2021:Q1. These include lines of credit, deposits (both the CCPs' own house accounts and the portion of initial margin held as bank deposits), and unfunded commitments that the CCPs can call from members in the event of a depletion of the default fund. Because these items all represent potential flows, we again scale them by the quarterly standard deviation of intermediary balance sheet changes. Again, it is clear that these potential resource draws could be large relative to the typical fluctuations banks and dealers see in their balance sheets. To take the extreme case, if all of these CCP commitments were drawn at once, the result could be equal to a 7-standard-deviation move in excess HQLA at large banks and a 32-standard-deviation move in CM excess capital. This would be on top of the VM, IM, and default fund calls that would likely occur at the same time.

The upshot of this discussion is that the demand by CCPs for liquid resources that has been observed over the last several years has been quantitatively large relative to market participants' capital and liquidity cushions, and it has the potential to be even larger in future periods of stress. Because such periods are also associated with other

[37]The estimates are constructed by calculating the variance of *quarterly* balance sheet changes and assuming independence of changes at the business-day frequency within quarters.

Table 5. Claims on Intermediaries, as of 3/31/21

Resources Demanded by Six CCPs	Total Value 3/31/21 ($ Bil.)	Relative to Quarterly Std. Dev. of Intermediary Balance Sheet Quantities, 2015–21					
		Tier 1 Capital (All Banks)	Excess Capital (FCMs)	Liquid Assets (Dealers)	Cash (Large Banks)	HQLA (Large Banks)	HQLA Above Min. LCR (Large Banks)
Lines of Credit	$232	1797%	2263%	471%	359%	248%	482%
Bank Deposits (CCP Acct. + IM Held)	$70	539%	678%	141%	108%	4%	144%
Unfunded Commitments	$30	232%	292%	61%	46%	32%	62%

Source: Quantitative disclosures via Clarus CCPView; FDIC Quarterly Banking Profile; CFTC Financial Data for FCMs; Financial Accounts of the United States; Federal Reserve H.8 Report; regulatory filings of JP Morgan Chase, Bank of America, Citibank, Wells Fargo, Goldman Sachs, and Morgan Stanley. Data as of 2022:Q1.

pressures on financial institutions' balance sheets, CCP procyclicality could be material in inducing systemic liquidity shortfalls.

4.8 The Importance of CCP Incentives

Although their procedures are governed by detailed rule books, CCPs retain some discretion in the extent and timing of their calls for resources. Two examples will illustrate this point.

Bignon and Vuillemey (2020) describe one of the very few failures of a CCP, the Caisse de Liquidation des Affairesen Marchandises in Paris. In 1974, following a collapse of sugar prices, the CCP faced the default of a clearing member that had cleared a large volume of sugar futures. Because the CCP's margining practices were inadequate, the default would have created a loss for the CCP's members, but would have allowed it to continue operating. But the CCP engaged in risk-shifting: it delayed declaring the underwater member in default, apparently in the hopes that its position would right itself. Instead, further losses accrued, eventually resulting in a shortfall so large that the CCP was forced to shut down.

More recently, as Heckinger (2014) describes, prior to the failure of MF Global in 2011, at least two CCPs (ICE Clear US and FICC) refused to release margin, totaling about $100 million, that was due to the teetering broker-dealer. The CCPs may have feared that MF Global might subsequently take losses on positions cleared through them and not be able to make up the shortfall. Although these actions afforded the CCPs an additional layer of protection, they contributed to MF Global's liquidity shortfalls and ultimate demise.

As illustrated by these episodes and others noted above, CCPs can bend or circumvent rules, at least for a time, when it is in their interest to do so. Moreover, there is no guarantee that the interests of CCP owners will align with the advancement of financial stability. Some CCPs are owned by their members or exchanges, while others are owned by publicly traded companies. The ability of these owners to exercise discretion makes it important to consider the incentives that CCPs might face in stress situations. The CCPs discussed herein all have their own capital exposed in their default waterfalls. Being exposed to loss provides CCPs an incentive to manage risks, but could also incentivize defensive moves—like choosing

to withhold MF Global's margin—during a stressful default. For CCPs that are member owned, the potential that clearing members might need to replenish lost capital can be an additional procyclical liquidity risk that could affect incentives even if a loss is not actually realized. The broader point is that, when facing the possibility of a severe liquidity or solvency threat, CCPs may have the ability and incentive to hoard even more resources than their normal practice suggests.[38]

5. Liquidity Stress Tests: Individual and Macroprudential

One of the main ways that the PFMI raised standards for CCPs was the explicit requirements around managing their liquidity risk. A CCP must be able to make all of its payment obligations on time in all relevant currencies with a high degree of confidence. The PFMI established the expectation that CCPs establish a robust framework to identify, measure, monitor, and manage liquidity risk from participants, settlement and custodian banks, liquidity providers, and any other relevant entities. The PFMI specified that CCPs had to be able to make their payments under the default of the clearing member that would generate the largest aggregate liquidity obligation for the CCP in extreme but plausible market conditions. Furthermore, the PFMI defined what resources should qualify for the purposes of meeting such requirements. The ability of the CCP to meet its potential liquidity demands with its liquid resources must be tested daily through rigorous stress testing similar to how the total financial resources are stress-tested daily.[39]

[38] CCP discretion can also go the other way, providing market participants more time to meet payment requirements, or even overriding margin requirements. Such exercises of discretion can result in smaller burdens on clearing members than would otherwise have occurred, but potentially expose the CCP to higher risk during a period of market stress. More generally, CCP discretion means that it may be difficult to predict CCP actions.

[39] In addition to CCPs' own liquidity stress tests, the U.S. Commodity Futures Trading Commission (CFTC) and the European Securities and Markets Authority (ESMA) both have conducted liquidity stress testing using their own scenarios applied jointly to sets of CCPs (see CFTC 2017 and ESMA 2018a, respectively).

The requirement to stress-test liquidity, in addition to being newer, adds new complexities. First, the number of relevant parties goes beyond just the clearing member function and includes those entities who generate a liquidity exposure. Second, multiple roles played by a clearing member—for example, if they were also a liquidity provider—need to be considered. Third, the ability to make payments, often in cash, in different currencies must be tested and met. Fourth, the time dimension matters, because it is not enough to have sufficient cash or other liquid resources in aggregate over a close-out period, but rather the CCP needs to be able to make payments on time when due.

The PFMI significantly enhanced expectations for liquidity risk management at CCPs and advanced the stress testing of each CCP's particular liquidity needs. By their nature, however, these micro-level stress tests cannot measure liquidity demands that may arise across the financial system in a market-stress event. While particular payment obligations are isolated within individual CCPs, the resources available to make such payments extend beyond the CCPs' boundaries. The resulting interdependencies are difficult if not impossible for an individual CCP to disentangle, evaluate, and stress-test.

The simplest relationship illustrating dependencies across CCPs is the committed lines of credit extended to CCPs by liquidity providers. As noted above, a few of the largest clearing members provide such lines of credit, often to multiple CCPs. As shown in FSB (2018a, Figure 11 on p. 21), apart from a few outliers, there is a positive correlation between the amount of prefunded resources a clearing member provides in aggregate and the amount of aggregate resources it provides as a liquidity backstop. The implication is that larger CCP members also provide significant liquid resources to the same CCPs. Individual CCPs do not necessarily have a view to the obligations its clearing members have to other CCPs. The default of a large clearing member likely would entail a simultaneous default across multiple CCPs, and therefore multiple CCPs activating liquidity relationships. Stress testing at individual CCPs cannot capture this interdependency.

A coordinated, macroprudential supervisory stress test across multiple CCPs can complement the micro-oriented stress tests conducted by individual CCPs. The goal would be to evaluate the

collective impact the participating CCPs have on the broader financial system during such a default. The results would inform regulators and participants on potential liquidity demands and perhaps avoid the liquidity hoarding that increased systemic risk during the 2008 crisis (Ashcraft, McAndrews, and Skeie 2011, Acharya and Merrouche 2012, Berrospide 2021). This objective stands in contrast to the CCPs' own stress tests, which look at their individual resiliency. The test would provide a sense of the size of the systemic liquidity risk that a central bank could face in extreme circumstances either in its role as the lender of last resort to the banking system or potentially through direct lending to CCPs.[40]

The testing would be designed in a manner consistent with the international framework for supervisory stress testing published by CPMI-IOSCO in 2017.[41] But we are more focused than the framework in explicitly calling for stressing liquidity needs and focusing on the overall impact on the system. The results would greatly enhance market participants' and regulators' understanding of how liquidity risks could arise in the CCP network and potentially affect the rest of the financial system, going beyond prior analysis, such as Heath et al. (2016). Such tests may require iterative defaults to trace the potential for stress to spread through the network of liquidity need and provision. Such iteration would be different than supervisory stress testing for banks, which is a firmly established regulatory tool after the 2008 crisis, and generally is viewed as effective.[42] Testing the resiliency of the clearing system likely would similarly improve our understanding of the feedback loops discussed in Faruqui, Huang,

[40]Central banks' ability and willingness to lend directly to CCPs varies widely. The Federal Reserve has a constrained ability to lend directly to CCPs that have been designated to be systemically important, but only in unusual or exigent circumstances and if other provisions are met (Baker 2012). In contrast, the Bank of England explicitly includes CCPs in its regular lender-of-last-resort function.

[41]See CPMI-IOSCO (2017). Anderson, Cerezetti, and Manning (2020) discuss such macroprudential CCP stress tests further, including discussing their rationale and objectives; see also Tompaidis (2012).

[42]See Drehman (2008), Petrella and Resti (2013), Borio, Drehmann, and Tsatsaronis (2014), Schuermann (2014), Flannery, Hirtle, and Kovner (2017), and Fernandes, Igan, and Pinheiro (2020). Although the generation of such tests could create incentives for CCPs to adjust their risk and liquidity management, their macroprudential nature and the fact that CCPs do not trade into their own positions likely eliminates the "window-dressing" effect observed by Cornett et al. (2020) in bank stress testing by supervisors.

and Takáts (2018). It would expand on the recent supervisory stress tests conducted by the CFTC on CCPs it regulates (CFTC 2016, 2017), particularly by including in the analysis securities CCPs, which, as noted above, require large amounts of liquidity to effect settlement.

Although macroprudential stress testing of CCPs' liquidity risk would be an important step, it is likely not sufficient to fully identify the scope of risks associated with their interconnectedness. To do so, regulators should work to integrate macroprudential stress testing of CCPs with more established bank stress testing, which has generally ignored banks' large exposures to CCPs (Borio, Drehmann, and Tsatsaronis 2014, Schuermann 2014). Such an integrated approach goes beyond the CPMI-IOSCO framework for CCPs.

6. Conclusion

By taking both sides of derivatives and securities trades, CCPs absorb counterparty credit risk. In so doing, they generally improve financial stability through multilateral netting, risk mutualization, and margining. Following the success of CCPs in managing risk and preventing contagion during the 2008 financial crisis, regulatory reforms have moved even more activity to central clearing, in particular through clearing mandates for the most common OTC derivatives.

Despite their roles in promoting financial stability, and particularly in reducing contagion, large CCPs are concentrated and interconnected and pose risks of their own. While attention has generally focused on the potentially disastrous consequences of a failure or severe disruption of a CCP, we have highlighted a difficulty that may occur in a much less remote state of the world. Namely, to protect themselves, CCPs necessarily require liquidity from large banks and other market participants. From the perspective of clearing members, the counterparty risk that is mitigated with central clearing is, in a sense, replaced with liquidity risk. The demand for resources can take the form of VM and IM calls, default fund assessments, draws on liquidity lines, the transfer of positions associated with defaulting members, and other obligations. This liquidity risk spans banks and funding markets, creating risks due to interconnectedness. Furthermore, the strength of the interconnectedness increases under

market stress. From a systemic perspective, the trade-off between reducing counterparty credit exposure while potentially increasing interconnected liquidity needs is expected to reduce systemic risk. Nevertheless, because resources are most likely to be called for by a CCP at times when bank liquidity positions are already under stress, they are inherently procyclical with respect to market conditions. The procyclicality of liquidity risk must be managed. Although CCPs, banks, and other financial institutions successfully managed the unprecedented margin calls during the market stress driven by COVID-19, not all the potential liquidity risk was realized, as the few defaults that occurred were relatively non-impactful. Consequently, the systemic risk may yet remain under-estimated.

The expansion of central clearing at well-managed CCPs strengthens market function and resiliency. But, as more activity becomes concentrated in CCPs, the possibility that CCP liquidity demands could strain banks and other market participants looms larger. On the policy front, as we have discussed, liquidity-focused macroprudential stress tests could help to assess the impact of shocks to CCPs on systemwide liquidity. Such tests also would be a step towards integrating supervisory stress testing of banks and CCPs. From a research perspective, more work is needed to understand how the liquidity risks posed by CCPs fit into the broader context of demand for safe and liquid assets. The answer to that question has implications for the measurement of system-wide liquidity, for the modeling of funding markets, and for our understanding of the propagation of financial crises.

References

Acharya, V. V., and O. Merrouche. 2012. "Precautionary Hoarding of Liquidity and Interbank Markets: Evidence from the Subprime Crisis." *European Finance Review* 17 (1): 107–60. doi:10.1093/rof/rfs022.

Acharya, V. V., H. S. Shin, and T. Yorulmazer. 2011. "Crisis Resolution and Bank Liquidity." *Review of Financial Studies* 24 (6): 2166–2205. doi:10.1093/rfs/hhq073.

Adrian, T., and H. S. Shin. 2008. "Liquidity, Monetary Policy, and Financial Cycles." *Current Issues in Economics and Finance*

14 (1): 1–7. https://www.newyorkfed.org/research/current_issues/ci14-1.html.

———. 2010. "Liquidity and Leverage." *Journal of Financial Intermediation* 19 (3): 418–37. doi:10.1016/j.jfi.2008.12.002.

Alfranseder, E., P. Fiedor, S. Lapschies, L. Orszaghova, and P. Sobolewski. 2018. "Indicators for the Monitoring of Central Counterparties in the EU." Occasional Paper No. 14, European Systemic Risk Board, European System of Financial Supervision. https://www.esrb.europa.eu/pub/pdf/occasional/esrb.op14.en.pdf.

Allen, F., E. Carletti, and D. Gale. 2009. "Interbank Market Liquidity and Central Bank Intervention." *Journal of Monetary Economics* 56 (5): 639–52. doi:10.1016/j.jmoneco.2009.04.003.

Amini, H., R. Cont, and A. Minca. 2016. "Resilience to Contagion in Financial Networks." *Mathematical Finance* 26 (2): 329–65. doi:10.1111/mafi.12051.

Anderson, E., F. Cerezetti, and M. Manning. 2020. "Supervisory Stress Testing for CCPs: A Macro-prudential, Two-Tier Approach." *Journal of Financial Market Infrastructures* 8 (1): 1–25. doi:10.21314/JFMI.2019.115.

Armakolla, A., and J. P. Laurent. 2017. "CCP Resilience and Clearing Membership." doi:10.2139/ssrn.2625579.

Ashcraft, A., J. McAndrews, and D. Skeie. 2011. "Precautionary Reserves and the Interbank Market." *Journal of Money, Credit and Banking* 43 (s2): 311–48. doi:10.1111/j.1538-4616.2011.00438.x.

Bai, J., A. Krishnamurthy, and C.-H. Weymuller. 2018. "Measuring Liquidity Mismatch in the Banking Sector." *Journal of Finance* 73 (1): 51–93. doi:10.1111/jofi.12591.

Baker, C. 2012. "The Federal Reserve as Last Resort." *University of Michigan Journal of Law Reform* 46 (1): 69–133. https://repository.law.umich.edu/mjlr/vol46/iss1/2.

Barker, R., A. Dickinson, A. Lipton, and R. Virmani. 2017. "Systemic Risks in CCP Networks." Risk.net, January 5. http://www.risk.net/cutting-edge/banking/2479766/systemic-risks-in-ccp-networks.

Basel Committee on Banking Supervision (BCBS). 2019. "Leverage Ratio Treatment of Client Cleared Derivatives." https://www.bis.org/bcbs/publ/d467.htm.

BCBS-CPMI-IOSCO. 2022. "Review of Margining Practices." https://www.bis.org/bcbs/publ/d537.htm.

BCBS-IOSCO. 2015. "Margin Requirements for Non-centrally Cleared Derivatives." https://www.bis.org/bcbs/publ/d317.pdf.

Berger, A. N., and C. H. S. Bouwman. 2009. "Bank Liquidity Creation." *Review of Financial Studies* 22 (9): 3779–3837. doi:10.1093/rfs/hhn104.

———. 2017. "Bank Liquidity Creation, Monetary Policy, and Financial Crises." *Journal of Financial Stability* 30 (June): 139–55. doi:10.1016/j.jfs.2017.05.001.

Bernanke, B. S. 1990. "Clearing and Settlement During the Crash." *Review of Financial Studies* 3 (1): 133–51. http://www.jstor.org/stable/2961962.

Berrospide, J. 2021. "Bank Liquidity Hoarding and the Financial Crisis: An Empirical Evaluation." *Quarterly Journal of Finance* 11 (4): 1–35. doi: 10.1142/S2010139221500208.

Bignon, V., and G. Vuillemey. 2020. "The Failure of a Clearinghouse: Empirical Evidence." *Review of Finance* 24 (1): 99–128. doi:10.1093/rof/rfy039.

Borio, C., M. Drehmann, and K. Tsatsaronis. 2014. "Stress-testing Macro Stress Testing: Does It Live up to Expectations?" *Journal of Financial Stability* 12 (June): 3–15. doi:10.1016/j.jfs.2013.06.001.

Breach, T., and T. King. 2018. "Securities Financing and Asset Markets: New Evidence." Working Paper No. 2018-22, Federal Reserve Bank of Chicago. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3301242.

Brunnermeier, M. K. 2009. "Deciphering the Liquidity and Credit Crunch 2007–2008." *Journal of Economic Perspectives* 23 (1): 77–100. doi:10.1257/jep.23.1.77.

CCP12. 2020. "CCPs Again Demonstrate Strong Resilience in Times of Crisis." Technical Report. https://ccp12.org/papers/.

Chicago Board Options Exchange. "CBOE Volatility Index: VIX [VIXCLS]." Retrieved from FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/VIXCLS.

Clarus CCPView. Clarus Financial Technology Ltd. https://ccpview.clarusft.com.

Commodity Futures Trading Commission (CFTC). 2016. "Supervisory Stress Test of Clearinghouses." https://www.cftc.gov/sites/

default/files/idc/groups/public/@newsroom/documents/file/cftcstresstest111516.pdf.

———. 2017. "Evaluation of Clearinghouse Liquidity." https://www.cftc.gov/sites/default/files/idc/groups/public/@newsroom/documents/file/dcr_ecl1017.pdf.

Cont, R. 2017. "Central Clearing and Risk Transformation." *Financial Stability Review* (Banque de France) 21: 127–38. https://publications.banque-france.fr/en/april-2017.

Cornett, M. M., J. J. McNutt, P. E. Strahan, and H. Tehranian. 2011. "Liquidity Risk Management and Credit Supply in the Financial Crisis." *Journal of Financial Economics* 101 (2): 297–312. doi:10.1016/j.jfineco.2011.03.001.

Cornett, M. M., K. Minnick, P. J. Schorno, and H. Tehranian. 2020. "An Examination of Bank Behavior around Federal Reserve Stress Tests." *Journal of Financial Intermediation* 41 (January): Article 100789. doi:10.1016/j.jfi.2018.05.001.

Cox, R. T., and R. S. Steigerwald. 2018. "A CCP is a CCP is a CCP." *Journal of Financial Market Infrastructures* 6 (4): 1–18. doi:10.21314/JFMI.2018.085.

CPMI-IOSCO. 2012. "Principles for Financial Market Infrastructures." http://www.bis.org/cpmi/publ/d101.htm.

———. 2014. "Recovery of Financial Market Infrastructures." https://www.bis.org/cpmi/publ/d121.htm.

———. 2016. "Resilience and Recovery of Central Counterparties (CCPs): Further Guidance on the PFMI." http://www.bis.org/cpmi/publ/d149.htm.

———. 2017. "Framework for Supervisory Stress Testing of Central Counterparties (CCPs)." http://www.bis.org/cpmi/publ/d161.htm.

Culp, C. L. 2010. "OTC-Cleared Derivatives: Benefits, Costs, and Implications of the 'Dodd-Frank Wall Street Reform and Consumer Protection Act'." *Journal of Applied Finance* 20 (2): 1–27. https://ssrn.com/abstract=2693059.

De Bandt, O., and P. Hartmann. 2019. "Systemic Risk in Banking After the Great Financial Crisis." In *Oxford Handbook of Banking*, 3rd edition, ed. A. N. Berger, Philip Molyneux, and J. O. S. Wilson, 847–89 (chapter 26). Oxford University Press. doi:10.1093/oxfordhb/9780198824633.013.27.

Drehmann, M. 2008. "Stress Tests: Objectives, Challenges and Modelling Choices." *Economic Review* (Sveriges Riksbank) (2): 60–92. http://archive.riksbank.se/Upload/Dokument_riksbank/Kat_publicerat/PoV_sve/eng/2008/er2008_2_ny.pdf.

Duffie, D. 2014. "Financial Market Infrastructure: Too Important to Fail." In *Across the Great Divide: New Perspectives on the Financial Crisis*, ed. J. B. Taylor and M. N. Baily, 251–57 (chapter 11). Hoover Institution Press. https://www.hoover.org/sites/default/files/across-the-great-divide-ch11.pdf.

Duffie, D., and H. Zhu. 2011. "Does a Central Clearing Counterparty Reduce Counterparty Risk?" *Review of Asset Pricing Studies* 1 (1): 74–95. doi:10.1093/rapstu/rar001.

European Securities and Markets Authority (ESMA). 2018a. "EU-wide CCP Stress Test 2017." http://firds.esma.europa.eu/webst/ESMA70-151-1154%20EU-wide%20CCP%20Stress%20Test%202017%20Report.pdf.

———. 2018b. "Guidelines on EMIR Anti-Procyclicality Margin Measures for Central Counterparties." https://www.esma.europa.eu/sites/default/files/library/esma70-151-1293_final_report_on_guidelines_on_ccp_apc_margin_measures.pdf.

Faruqui, U., W. Huang, and E. Takáts. 2018. "Clearing Risks in OTC and Derivatives Markets: The CCP-Bank Nexus." *BIS Quarterly Review* (December 16): 73–90. https://www.bis.org/publ/qtrpdf/r_qt1812h.htm.

Fernandes, M., D. Igan, and M. Pinheiro. 2020. "March Madness in Wall Street: (What) Does the Market Learn from Stress Tests?" *Journal of Banking and Finance* 112 (March): Article 105250. doi:10.1016/j.jbankfin.2017.11.005.

Financial Stability Board. 2017a. "Analysis of Central Clearing Interdependencies." http://www.fsb.org/wp-content/uploads/P050717-2.pdf.

———. 2017b. "Review of OTC Derivatives Market Reforms: Effectiveness and Broader Effects of the Reforms." http://www.fsb.org/wp-content/uploads/P290617-1.pdf.

———. 2018a. "Analysis of Central Clearing Interdependencies." http://www.fsb.org/wp-content/uploads/P090818.pdf.

———. 2018b. "Incentives to Centrally Clear Over-the-Counter (OTC) Derivatives: A Post-Implementation Evaluation of the Effects of the G20 Financial Regulatory Reforms." http://www.

fsb.org/2018/11/incentives-to-centrally-clear-over-the-counter-otc-derivatives-2/.

———. 2021. "Lessons Learnt from the COVID-19 Pandemic from a Financial Stability Perspective: Final Report." https://www.fsb.org/2021/10/lessons-learnt-from-the-covid-19-pandemic-from-a-financial-stability-perspective-final-report/.

Flannery, M., B. Hirtle, and A. Kovner. 2017. "Evaluating the Information in the Federal Reserve Stress Tests." *Journal of Financial Intermediation* 29 (January): 1–18. doi:10.1016/j.jfi.2016.08.001.

Fleming, M. J., and A. Sarkar. 2014. "The Failure Resolution of Lehman Brothers." *Economic Policy Review* (Federal Reserve Bank of New York) 20 (2): 175–206. https://www.newyorkfed.org/research/epr/2014/1412flem.html.

Futures Industry Association (FIA). 2020. "Revisiting Procyclicality: The Impact of the COVID Crisis on CCP Margin Requirements." Technical Report. https://www.fia.org/resources/fia-issues-white-paper-impact-pandemic-volatility-ccp-margin-requirements.

Gibson, R., and C. Murawski. 2013. "Margining in Derivatives Markets and the Stability of the Banking Sector." *Journal of Banking and Finance* 37 (4): 1119–32. doi:10.1016/j.jbankfin.2012.10.005.

Glasserman, P., and Q. Wu. 2018. "Persistence and Procyclicality in Margin Requirements." *Management Science* 64 (12): 5705–24. doi:10.1287/mnsc.2017.2915.

Greenwood, R., A. Landier, and D. Thesmar. 2015. "Vulnerable Banks." *Journal of Financial Economics* 115 (3): 471–85. doi:10.1016/j.jfineco.2014.11.006.

Haver Analytics. http://www.haver.com/our_data.html.

Heath, A., G. Kelly, M. Manning, S. Markose, and A. R. Shaghaghi. 2016. "CCPs and Network Stability in OTC Derivatives Markets." *Journal of Financial Stability* 27 (December): 217–33. doi:10.1016/j.jfs.2015.12.004.

Heber, G., A. Lunde, N. Shephard, and K. K. Sheppard. 2009. "Oxford-Man Institute's Realized Library." Version No. 0.3. University of Oxford. https://realized.oxfordman.ox.ac.uk/.

Heckinger, R. 2014. "MF Global: A Case Study of Liquidity Risks." *Journal of Financial Market Infrastructures* 3 (2): 79–96. doi:10.21314/JFMI.2014.037.

Heckinger, R., R. T. Cox, and D Marshall. 2017. "Cleared Margin Setting at Selected Central Counterparties." *Journal of Financial Market Infrastructures* 5 (4): 1–21. doi:10.21314/JFMI.2017.075.

Heider, F., M. Hoerova, and C. Holthausen. 2015. "Liquidity Hoarding and Interbank Market Rates: The Role of Counterparty Risk." *Journal of Financial Economics* 118 (2): 336–54. doi:10.1016/j.jfineco.2015.07.002.

Heller, D., and N. Vause. 2012. "Collateral Requirements for Mandatory Central Clearing of Over-the-Counter Derivatives." BIS Working Paper No. 373. http://www.bis.org/publ/work373.pdf.

Houllier, M., and D. Murphy. 2017. "Initial Margin Model Sensitivity Analysis and Volatility Estimation." *Journal of Financial Market Infrastructures* 5 (4): 77–103. doi:10.21314/JFMI.2017.078.

Huber, P. J. 1981. *Robust Statistics*. Wiley Series in Probability and Mathematical Statistics. Hoboken, NJ: John Wiley & Sons.

Ippolito, F., J.-L. Peydró, A. Polo, and E. Sette. 2016. "Double Bank Runs and Liquidity Risk Management." *Journal of Financial Economics* 122 (1): 135–54. doi:10.1016/j.jfineco.2015.11.004.

ISDA Clearing Member Committee. 2021. "COVID-19 and CCP Risk Management Frameworks." Technical Report. https://www.isda.org/2021/01/06/covid-19-and-ccp-risk-management-frameworks/.

Ivashina, V., and D. Scharfstein. 2010. "Bank Lending during the Financial Crisis of 2008." *Journal of Financial Economics* 97 (3): 319–38. doi:10.1016/j.jfineco.2009.12.001.

Iyer, R., J.-L. Peydró, S. da Rocha-Lopes, and A. Schoar. 2014. "Interbank Liquidity Crunch and the Firm Credit Crunch: Evidence from the 2007–2009 Crisis." *Review of Financial Studies* 27 (1): 347–72. doi:10.1093/rfs/hht056.

Kashyap, A. K., R. Rajan, and J. C. Stein. 2002. "Banks as Liquidity Providers: An Explanation for the Coexistence of Lending and Deposit-taking." *Journal of Finance* 57 (1): 33–73. doi:10.1111/1540-6261.00415.

Kiff, J., R. Dodd, A. Gullo, E. Kazarian, I. Lustgarten, C. Sampic, and M. Singh. 2010. "Making Over-the-Counter Derivatives Safer: The Role of Central Counterparties." In *Global*

Financial Stability Report: Meeting New Challenges to Stability and Building a Safer System, 1–27 (chapter 3). International Monetary Fund. https://www.imf.org/external/pubs/ft/gfsr/2010/01/pdf/chap3.pdf.

King, T., T. D. Nesmith, A. Paulson, and T. Prono. 2020. "Central Clearing and Systemic Liquidity Risk." Finance and Economics Discussion Series No. 2020-009, Board of Governors of the Federal Reserve System. doi:0.17016/FEDS.2020.009.

Kroszner, R. S. 1999. "Can the Financial Markets Privately Regulate Risk? The Development of Derivatives Clearinghouses and Recent Over-the-Counter Innovations." *Journal of Money, Credit and Banking* 31 (3): 596–618. doi:10.2307/2601077.

———. 2000. "Lessons from Financial Crises: The Role of Clearinghouses." *Journal of Financial Services Research* 18 (2): 157–71. doi:10.1023/A:1026534619637.

LCH.Clearnet. 2008. "$9 Trillion Lehman OTC Interest Rate Swap Default Successfully Resolved." Press Release, October 8. http://secure-area.lchclearnet.com/media_centre/press_releases/2008-10-08.asp.

Lewandowska, O., and F. Glaser. 2017. "The Recent Crises and Central Counterparty Risk Practices in the Light of Procyclicality: Empirical Evidence." *Journal of Financial Market Infrastructures* 5 (3): 1–24. doi:10.21314/JFMI.2017.072.

Madigan, P., D. Wood, and L. Becker. 2016. "LCH under Scrutiny after Outsized Brexit Margin Calls." Risk.net, October 19. https://www.risk.net/risk-management/2474560/lch-under-scrutiny-after-outsized-brexit-margin-calls.

Manning, M. J., and D. Hughes. 2016. "Central Counterparties and Banks: Vive La Difference." *Journal of Financial Market Infrastructures* 4 (3): 1–24. doi:10.21314/JFMI.2016.058.

Maruyama, A., and F. Cerezetti. 2019. "Central Counterparty Antiprocyclicality Tools: A Closer Assessment." *Journal of Financial Market Infrastructures* 7 (4): 1–25. doi:10.21314/JFMI.2018.110.

Mourselas, C. 2021. "Hero or Villain? NSCC Draws Fire for Robinhood Margin Waiver." Risk.net, March 31. https://www.risk.net/risk-management/7817001/hero-or-villain-nscc-draws-fire-for-robinhood-margin-waiver.

Murphy, D. 2013. *OTC Derivatives: Bilateral Trading and Central Clearing*. Palgrave Macmillan. doi:10.1057/9781137293862.

Murphy, D., M. Vasios, and N. Vause. 2014. "An Investigation into the Procyclicality of Risk-Based Initial Margin Models." Financial Stability Paper No. 29, Bank of England. https://www.bankofengland.co.uk/-/media/boe/files/financial-stability-paper/2014/an-investigation-into-the-procyclicality-of-risk-based-initial-margin-models.pdf.

O'Neill, C., and N. Vause. 2018. "Macroprudential Margins: A New Countercyclical Tool?" Staff Working Paper No. 765, Bank of England. https://www.bankofengland.co.uk/working-paper/2018/macroprudential-margins-a-new-countercyclical-tool.

Petrella, G., and A. Resti. 2013. "Supervisors as Information Producers: Do Stress Tests Reduce Bank Opaqueness?" *Journal of Banking and Finance* 37 (12): 5406–20. doi:10.1016/j.jbankfin.2013.01.005.

Presidential Task Force on Market Mechanisms. 1988. "Report of the Presidential Task Force on Market Mechanisms." Nicholas Brady (Chairman).

Raykov, R. 2012. "Reducing Margin Procyclicality at Central Counterparties." *Journal of Financial Market Infrastructures* 7 (2): 43–59. doi:10.21314/JFMI.2018.106.

Schuermann, T. 2014. "Stress Testing Banks." *International Journal of Forecasting* 30 (3): 717–28. doi:10.1016/j.ijforecast.2013.10.003.

Shleifer, A., and R. W. Vishny. 2011. "Fire Sales in Finance and Macroeconomics." *Journal of Economic Perspectives* 25 (1): 29–48. doi:10.1257/jep.25.1.29.

Sidanius, C., and F. Zikes. 2012. "OTC Derivatives Reform and Collateral Demand Impact." Financial Stability Paper No. 18, Bank of England. https://www.bankofengland.co.uk/financial-stability-paper/2012/otc-derivatives-reform-and-collateral-demand-impact.

Tompaidis, S. 2012. "Measuring System-wide Resilience of Central Counterparties." *Journal of Financial Market Infrastructures* 6 (4): 41–54. doi:10.21314/JFMI.2018.098.

Valukas, A. R. 2010. "Report of Examiner, United States Bankruptcy Court, Southern District of New York, in re. Lehman Brothers Holdings, Inc., et al., Chapter 11 case no. 08-13555."

Jenner & Block LLP. https://web.stanford.edu/~jbulow/Lehmandocs/menu.html.

Wiggins, R. Z., and A. Metrick. 2019. "The Lehman Brothers Bankruptcy G: The Special Case of Derivatives." *Journal of Financial Crises* 1 (1): 151–71. https://elischolar.library.yale.edu/journal-of-financial-crises/vol1/iss1/8.

Wong, M., and P. Ge Pei. 2017. "Performance Testing of Margin Models Using Time Series Similarity." *Journal of Financial Market Infrastructures* 5 (4): 51–75. doi:10.21314/JFMI.2017.076.

Government Assistance and Banks' Funding Cost*

Md Jahir Uddin Palas[a,b] and Fernando Moreira[b]
[a]University of Dhaka
[b]University of Edinburgh Business School

This study examines the association of government support with recipient banks' funding cost. The U.S. government's Capital Purchase Program (CPP) is utilized as a case study of government assistance. The study's sample includes quarterly data from 2009 to 2018 on 8,327 U.S. financial institutions in the banking sector. The results suggest that government assistance has a significant relationship with the recipient banks' lower funding cost. Augmented public confidence in recipient banks could be a plausible channel to explain the government assistance–funding cost relationship. The findings are robust to alternative funding cost definitions, samples, estimation methods, and model specifications.

JEL Codes: G21, G28.

1. Introduction

Since the great economic depression of the 1930s and possibly next to the recent economic impact of COVID-19, the financial crisis of 2007–08 can be treated as the worst scenario of global economic meltdown. Investors' confidence in the financial markets had faded and financial institutions were struggling to survive, driving the overall U.S. financial system to the edge of collapse. To safeguard the U.S. financial system from a potential collapse and avoid the further impact of contagion, the U.S. government intervened with

*The authors thank Steven Ongena (the editor) and anonymous referees for their constructive suggestions that have greatly improved the quality of the paper. The authors also thank Jake Ansell, Galina Andreeva, Yizhe Dong, Maria Michou, Belen Martin-Barragan, and seminar participants at the University of Edinburgh for their helpful and constructive comments. Corresponding author: Md Jahir Uddin Palas. E-mail: jahir@du.ac.bd.

assistance programs to revive the financial institutions through capital injections.

The effect of government assistance has long been a controversial issue for market participants and in academic research. The extant literature provides mixed conclusions while evaluating the influence of government assistance in different facets of financial institutions. Government assistance programs can be important to society and the economy, possibly to prevent the spread of financial crises (Anderson and Bluedorn 2017). Current empirical research offers indications that government assistance programs in the United States aided in mitigating the severe impact of the financial crisis through increased lending activities (Li 2010) and creating real economic value (Bayazitova and Shivdasani 2011). Veronesi and Zingales (2010) concluded that the Troubled Asset Relief Program (TARP) offered positive economic gains to the U.S. economy by increasing the value of banks' financial claims where the net value ranges from $86 billion to $109 billion. TARP-induced spillover effects in the corporate sector positively contributed to borrowing firms' stock returns (Norden, Roosenboom, and Wang 2013). U.S. policymakers also considered the government assistance program a success, as the U.S. Treasury claimed that over $204 billion funds were repaid out of the total assistance provided, emphasizing $30 billion in revenue for taxpayers (U.S. Department of the Treasury 2018).

However, negative consequences of government intervention are also possible. Academic literature criticized the U.S. government assistance program in the 2007–08 financial crisis for not learning from failed precedence in Japan's banking crisis of the 1990s (Hoshi and Kashyap 2010), reducing equity shareholders' potential by the U.S. Treasury's stance in preferred shares (Bayazitova and Shivdasani 2011), and banks' increased risk-taking (Duchin and Sosyura 2010; Black and Hazelwood 2013). However, Berger, Roman, and Sedunov (2020) concluded that large and safer banks with better local economic conditions significantly reduced systemic risk contributions after receiving government aid through TARP. European banks that received the government's support at the beginning of the financial crisis displayed the apparent ineffectiveness of the governmental efforts in significantly enhancing their performance (Gerhardt and Vander Vennet 2017). Harris, Huerta, and Ngo (2013)

document a significant reduction in the operational efficiency of TARP recipient banks due to moral hazards issues. Similarly, banks that received government support through capital injection are found to be significantly less efficient (Palas and Moreira 2022). Chavaz and Rose (2019) linked TARP political influences with banks' lending behavior and concluded that TARP funds adversely affected both the quality and quantity of banks' lending.

Although existing government assistance literature investigated the bank lending and risk-taking areas quite largely, little is explored about banks' funding aspects. Understanding a bank's vulnerability to funding cost fluctuations is crucial since funding cost can represent a bank's counterparty credit risk, which regulators take into account to determine the appropriate capital buffer (Schmitz, Sigmund, and Valderrama 2017). Besides current capital levels, funding costs are also connected to a bank's future required capital due to "adverse dynamics" where the effect's magnitude depends on the bank's response to rising funding costs. An increased cost of fund can negatively affect banks' profitability if banks decide to leave the lending rates unchanged (Beau et al. 2014). Alternatively, banks can raise lending rates that may trigger loan default, resulting in credit losses and driving down profitability. Either way, a substantial erosion of profit can have a severe impact on the banks' capital buffer. Therefore, higher funding cost shrinks the capital buffer by reducing profitability in the short term (Schmitz, Sigmund, and Valderrama 2017). In the long term, the capital buffer can deplete further, as investors would require higher compensation to offset increased risk, extending the persistence of high funding costs. Banks' funding cost also has important implications for monetary policy, as it can influence economic growth and inflation outlook. Such implications became noticeable in the 2007–08 financial crisis when banks' funding costs increased remarkably, causing liquidity shortages and higher lending rates (Beau et al. 2014).

In broader literature, studies on banks' funding cost are mostly concentrated around its association with solvency. The common consensus in the literature supports that solvency has a significant negative relationship with funding costs (Arnould et al. 2022). Solvency, if explained through regulatory capital, acts as deposit insurance and, therefore, the market (debtholders and depositors) imposes higher funding costs on banks with lower capital (Carvalho and

Dantas 2020). The extent of such market discipline is greater for banks suffering from solvency issues in conjunction with lower asset quality (Acharya and Mora 2015). Measuring solvency risk through expected capital shortfall, Pierret (2014) found that higher solvency risk restricts access to short-term funding. Similarly, wholesale market participants demand a higher funding rate if their perception of a bank's solvency is negatively affected (Dent, Hoke, and Panagiotopoulas 2021). Compared with the average funding cost, the sensitivity of market reaction is greater for the interbank funding cost (Aymanns et al. 2016). The overall conclusions remain the same when solvency is defined in terms of market-based leverage (Annaert et al. 2013; Hasan, Liu, and Zhang 2016).

Likewise, an increase in bank capital is found to be significantly associated with lower funding costs (Babihuga and Spaltro 2014; Schmitz, Sigmund, and Valderrama 2017; Moreira 2020). Banks having stronger fundamentals enjoy favorable interest rates in both borrowing and lending activities (Barajas and Stein 2000). Depositors prefer well-capitalized and highly liquid banks and, thus, adjust their funding volumes based on banks' capital and liquidity position (Ungan, Caner, and Özyıldırım 2008). As a result, depositors establish market discipline by requiring higher returns from banks with weaker fundamentals. This market discipline also applies to undercapitalized large banks with low liquidity. Schmitz, Sigmund, and Valderrama (2017) concluded that the influence of regulatory capital on funding costs is stronger than the relationship from the opposite direction, i.e., funding cost to capital.

In the case of government assistance through capital infusion, market participants may perceive the assisted banks with extra capital as safer (safety channel) due to current and/or future, if required, government intervention (Berger et al. 2020). Therefore, funding supply may increase for the bailed-out banks. Furthermore, assisted banks can reduce funding demand due to an inflow of funding from the government source. The banks can also reduce funding demand by getting rid of riskier assets in the post-assistance period. Funding costs would reduce under these high-supply and low-demand scenarios.

However, the non-linear nature of such a negative relationship is also pronounced in literature, which makes it possible for capital and funding cost to have a positive relationship. In the event of a

solvency shock, funding costs' responsiveness against lower solvency is greater at the initial levels (Aymanns et al. 2016; Dent, Hoke, and Panagiotopoulos 2021). Additionally, Arnould et al. (2022) found that the non-linearity is convex such that solvency beyond a certain threshold may positively influence funding costs of senior bond and term deposit. Cummings and Wright (2016) confirmed such a convex relationship by concluding that Australian banks may face increased (ranging from 8 to 24 basis points) funding costs in the long run due to higher capital requirements.

Different bailout scenarios are also likely. Banks receiving capital injection may increase their demand for other funding sources to repay the government funds (Berger et al. 2020), especially if the assistance fund is costlier. In an attempt to recover from the crisis, assisted banks may want to increase the size of their asset portfolio, which would require additional funding. Such increased demand can raise funding costs.

Since TARP-CPP funds were aimed at increasing bank capital, a relevant strand of literature could be the impact of capital on competition because capital improves banks' competitive position in the deposit and lending markets (Calomiris and Mason 2003; Calomiris and Wilson 2004). Empirical findings of Berger and Roman (2015) suggest that TARP recipient banks gained competitive advantages in both market power and market share. A higher TARP bailout probability is found to be positively associated with the post-crisis market power of a bank (Koetter and Noth 2016). The conclusions corroborate with the theoretical explanations of Allen, Carletti, and Marquez (2011) and Mehran and Thakor (2011) regarding a positive association between capital and market share. Berger and Bouwman (2013) found that higher capital positively influences the market share for small banks, but, for medium and large banks, such influence holds in the crisis periods only.

Alternative arguments explain how government assistance through capital infusion can lead to decreased competitive advantage (in terms of funding costs) for the recipient banks. First, through the quiet-life and/or charter-value channel (Hicks 1935; Keeley 1990; Cordella and Yeyati 2003), government assistance may augment the charter value of banks and/or motivate less active business practices, i.e., quiet-life that can have adverse funding cost implications, e.g., missing out on cheaper funding opportunities due to

fewer interactions with banking network partners or not actively monitoring the funding market.

Second, through the stigma channel (Berger and Roman 2015), market participants may consider the banks that applied for and/or accepted the government assistance as financially distressed or likely to fail. Such negative market perception can have adverse impacts on both the volume and cost of funding. Third, through the cost-disadvantage channel (Berger and Roman 2015), if capital from government aid is costlier, the recipient banks will experience increased total funding cost. Fourth, through the inefficiency channel,[1] the government's support through capital infusion can increase its managerial decision power on the assisted banks and trigger changes in the banks' internal managerial practices. If the government is inefficient in managing banks, funding costs can rise. In this connection, related literature provides evidence that government support significantly reduces the recipient banks' operational efficiency (Harris, Huerta, and Ngo 2013) and overall efficiency (Palas and Moreira 2022).

Finally, the government may impose certain restrictions on recipient banks' management, which can cause disagreement or dissatisfaction (conflict channel) between the owner (e.g., government) and management (e.g., executives). For example, the U.S. government put a limit of tax-deductible executive pay to $500,000 for all the CPP recipient banks. Therefore, banks with higher CEO pay were less inclined to accept the CPP funds (Cadman, Carter, and Lynch 2012), possibly to avoid potential CEO dissatisfaction. The conflict channel is further strengthened by the social stigma since bank's high CEO compensation was widely discussed by the news, blogs, and politicians concerning the U.S. government's assistance programs in 2007–08 (Wilson and Wu 2012). Management dissatisfaction in banks, especially in terms of pay, can lead to reduced productivity, inactive market monitoring, loss of skilled human resources, etc., which can adversely affect funding costs. The inefficiency and conflict channels may explain why many recipient banks attempted to repay the CPP fund quickly, e.g., Berger and Roman (2015).

[1]We thank an anonymous reviewer for highlighting the relevant possibilities.

Building on the capital-funding cost and capital-competition literature, possible research hypotheses could be formed as below:

Ha: Government assistance through capital infusion results in the recipient banks' lower funding costs.
Hb: Government assistance through capital infusion results in the recipient banks' higher funding costs.

Besides bank-specific factors, uncertainty in the financial market, the central bank's monetary policy directions, and sovereign risk play important roles in driving banks' funding costs (Arnould et al. 2022). The financial crisis of 2007–08 is of particular interest in this discussion, as it provoked disruptions in both the retail and wholesale funding markets for banks. Aymanns et al. (2016) documented evidence that the solvency–funding cost relationship in the interbank funding market is highly sensitive in periods of economic stress. Similarly, funding costs are influenced by financial market shocks, e.g., 2007–08 financial crisis (Babihuga and Spaltro 2014), and conditions in the market for bank lending (Kiser 2003), e.g., high default rate in the 2007–08 financial crisis. Moreover, the economic recession followed by the financial crisis triggered investors' concern about public finance, which leads to increased funding costs (Panetta et al. 2011). One of the primary motivations of the U.S. government's assistance program (TARP-CPP) in the financial crisis of 2007–08 was to stabilize the financial markets with adequate capital such that the rising tensions around banks' funding positions are attenuated.

Despite being an important research agenda, the question of whether and how government assistance as a crisis response affects funding costs has remained unanswered to date. This paper is an attempt to fill the vacuum in this context by investigating the effect of the U.S. government's Capital Purchase Program (CPP) on the recipient banks' funding costs. The results suggest that government assistance has a significant negative relationship with funding costs. The findings can be interpreted as evidence of increased public confidence in the recipient banks' financial health. This research contributes to the literature in several ways. Firstly, this study is possibly the earliest attempt to provide empirical evidence on the relationship between government assistance and banks' funding cost. Understanding the effect of government support through

the lens of funding cost will provide more insight into the debate about the consequences of governmental assistance. Secondly, unlike most studies in the relevant literature, this study incorporates a relatively long-term evaluation of the government intervention. Thirdly, current funding cost literature focuses mostly on large banks and uses a market-based approach, e.g., Babihuga and Spaltro (2014), Schmitz, Sigmund, and Valderrama (2017), and Dent, Hoke, and Panagiotopoulas (2021), which limits small banks since they have limited or no market-related activities. This study uses balance-sheet-based funding cost measures and provides valuable findings on the government assistance–funding cost relationship by incorporating a generous share of small banks. Most importantly, this study not only investigates whether and how government assistance relates to funding cost but also attempts to explain the possible channels of such relationship.

The study will be important to (i) regulators, as they will get a clearer understanding of the link between government assistance and banks' funding cost that will also facilitate improved intervention programs; (ii) depositors and investors, as they will be in a better position to analyze the implications of government assistance and to manage their funds on the recipient or non-recipient banks accordingly; (iii) practitioners, as they will have enhanced knowledge to guide banks' management on government-assistance-related decisions while considering possible funding cost outcomes; and (iv) finally, banks' management-level executives, who will know what to expect when participating in a government assistance program and how to adjust the bank's business strategies to cope with the aftermaths.

The rest of the paper is organized as follows. Section 2 explains the data and empirical methods used. The results are presented and discussed in Section 3. Section 4 concludes.

2. Data and Method

2.1 Institutional Background and Sample

The Troubled Assets Relief Program (TARP), organized by the Emergency Economic Stabilization Act, was the financial support program undertaken by the U.S. government as a response to the

2007–08 financial crisis. TARP was approved in October 2008 to ensure the overall soundness of the U.S. financial system (Ng, Vasvari, and Wittenberg-Moerman 2011). Several sub-programs were hosted under the broader umbrella of the TARP program, and the noteworthy one was the Capital Purchase Program (CPP), announced in October 2008. The main purpose of CPP was to reinstall stability and robustness of the U.S. financial system. To secure these two pillars and regain investors' lost confidence, CPP provided capital support to financial institutions of all sizes that met the stated criteria (U.S. Department of the Treasury 2018). On numbers, CPP constituted more than 33 percent of total TARP funding, about $250 billion of government assistance, which, however, was reduced to $218 billion in March 2009. By the end of the government assistance period, $204.9 billion was invested under the CPP scheme by the U.S. Treasury. Moreover, during the implementation of the CPP program, the U.S. Treasury provided capital to 707 financial institutions that include small and community banks as well as certified community development financial institutions. The U.S. government maintained transparency in the CPP program and regularly updated data on how the U.S. Treasury was using the fund, who the recipients were, and the latest information on the fund's status (Bayazitova and Shivdasani 2011). The detailed information about CPP beneficiaries is accessible on the U.S. Treasury's website (U.S. Department of the Treasury 2018).

This study concentrates on the banking sector, which was the major recipient of the CPP program. The relevant quarterly data for bank holding companies (BHCs) and commercial banks (CBs) are collected from the FDIC (Federal Deposit Insurance Corporation) Call Reports and S&P Capital IQ Pro database. The BHCs' data are not aggregated at the parent level, since this study is aimed at examining the potentially heterogeneous association between CPP assistance and different subsidiaries of a BHC. Especially when banks' names are not the same, many fund providers (especially retail depositors) may not know that some commercial banks belong to a particular BHC. This study accounts for such an important possibility because funding costs are sensitive and highly connected to the depositors' and/or investors' perceptions.

Savings and loan institutions, and other thrift institutions, are excluded since these institutions report data and compete in the

market differently compared with commercial banks. For convenience, the term bank refers to both types of banks (BHC and CB) from this point onwards, unless stated separately. The sample period for the analysis is from 2009 to 2018, considering that CPP assignments started in the last quarter of 2008. Inactive banks are retained to avoid survival bias. Excluding banks that failed before 2009, 8,408 banks remained from the raw sample of 27,687 FDIC banks.

The CPP transaction data and list of CPP recipients are obtained from the U.S. Treasury's website. The CPP recipient banks from the U.S. Treasury list are matched with the initial sample (8,408 banks) by using the banks' name, city, state, and CPP amount received. In case of acquisition, merger, and name changes, banks are investigated further and updated with the appropriate identification details. Observations with incomplete or missing data on total assets or common equity, and/or negative data on income statement items related to funding costs, e.g., interest expense, are excluded. Following Berger and Bouwman (2013) and Berger and Roman (2015), the equity-to-total-asset ratio (if less than 1 percent) is replaced with 1 percent to avoid distortions. The final sample includes 8,327 banks, among which 414 banks are CPP recipients and the remaining 7,913 banks are non-CPP recipients.[2] The sample contains 1,093 BHCs and 7,234 CBs, among which 332 and 82 are CPP banks, respectively.[3]

The 414 CPP banks were granted $184.90 billion, which is around 90.24 percent of the total CPP fund ($204.9 billion). As compared with the CPP banks' capital structure before CPP, the CPP capital injections can be considered significant. On average, CPP infusion amounts to 28 percent of the CPP banks' pre-CPP total capital and 3 percent of pre-CPP risk-weighted assets. The mean and median of the CPP banks' Tier 1 ratio improved noticeably after the CPP injections. Table 1 provides further details on CPP volume and its significance on the CPP banks' capital structure.

[2]From this point onward, the CPP recipients and non-CPP recipients groups will be termed CPP banks and non-CPP banks, respectively.

[3]Recall that the BHCs' data are not aggregated at the parent level.

Table 1. Details of Capital Purchase Program (CPP)

	Mean	Median	25th Percentile	75th Percentile	Min.	Max.
CPP Amount (in USD Millions)	448	22	11	65	0.301	25,000
CPP/Pre-CPP Total Capital	0.28	0.33	0.22	0.32	0.08	1.36
CPP/Pre-CPP RWA	0.03	0.03	0.02	0.03	0.01	0.16
Pre-CPP Tier 1 Ratio	11.68	10.43	9.40	12.55	7.82	23.73
Post-CPP Tier 1 Ratio	12.33	12.01	10.59	13.7	.22	25.19

Note: This table reports CPP-related summary statistics based on the recipient (CPP) banks. Yearly values are considered for Total Capital and RWA (Risk-Weighted Assets). Since capital infusion under CPP was initiated from the last quarter of 2008, 2007 and 2009 are considered as the pre- and post-CPP years, respectively. Total capital includes common equity, preferred equity, and non-redeemable non-controlling interest of a company. Tier 1 ratio represents equity capital plus minority interests less portion of perpetual preferred stock and goodwill as a percent of adjusted risk-weighted assets.

2.2 Models

2.2.1 The Baseline Equation

$$FC_{i,t} = \beta_0 + \beta_1 CPP_i + \sum_{n=2}^{8} \beta_n X_{n,i,t} + \sum_{n=9}^{12} \beta_n C_{n,s,t} + \tau_t + \alpha_s + \varepsilon_{i,t} \tag{1}$$

In Equation (1), $FC_{i,t}$ stands for the funding cost of bank i at time t. This study uses three different balance-sheet-based indicators of funding cost, namely the cost of fund, cost of deposit, and cost of liabilities. Although banks' funding costs can be identified through multiple measures, typical empirical literature focuses largely on market-based indicators, e.g., Babihuga and Spaltro (2014), Schmitz, Sigmund, and Valderrama (2017), and Dent, Hoke, and Panagiotopoulas (2021). However, market-based measures could be imperfect since they may not offer a fair representation of the banks' actual funding cost (Arnould et al. 2022). These measures ignore that different types of bank liabilities can have varied sensitiveness to stress situations. The commonly exploited market-based measure in the funding cost literature is CDS spreads, which are more related to the risk of bank failure, i.e., credit default, but not necessarily to the funding costs. Even when predicting credit defaults, Grossman and Hansen (2010) concluded that the CDS spreads were unreliable during the financial crisis of 2007–08.

Funding cost studies using CDS spreads assume that investors are risk neutral, do not require risk premium, and are insusceptible to changes in risk-aversion sentiment (Schmitz, Sigmund, and Valderrama 2017). These strong assumptions rarely hold in real-life scenarios. Alternatively, studies employing CDS spreads rely on the strong assumption that the market (or investors) can accurately assess banks' risk and demand interest rates accordingly. Again, such a strong assumption can be doubtful due to the opaqueness of banks and their funding activities. Even if investors' assessment is correct, they could discount part of the risk in light of other factors, e.g., too big to fail. In this case, even if investors perceive a bank as very risky (high CDS spreads), they may expect that some banks, mostly the large ones, would be bailed out and the funding costs of those banks would not increase proportionally. In terms of bond yields, the

arguments presented may also hold besides the fact that bond prices and, consequently, yields are affected by many other issues. For example, bond characteristics are unstable and can change significantly over the bond's life.

This study's context also played an important role in favoring the balance-sheet-based funding cost measures. In stress conditions such as the 2007–08 financial crisis, banks' opportunity to access funds from the market may shrink significantly due to market-wide illiquidity. Therefore, market-based measures may not be a good fit for studies based on crisis times. Besides the context, such measures would not suit properly for this study's sample, since it includes a significant share of small banks that have no or limited market-based activity. Aymanns et al. (2016) suggested that the balance-sheet-based approach would allow rich panel data for a large sample and the study findings would also be directly implementable to banks' stress tests, which depend on balance sheet data.

For an effective measure of funding cost, Schmitz, Sigmund, and Valderrama (2017) recommended considering banks' actual funding structure and the cost of alternative funding sources. A growing number of studies are using a balance-sheet-based approach for funding cost analysis—for example, Ungan, Caner, and Özyıldırım (2008), Aymanns et al. (2016), Carvalho and Dantas (2020), Moreira (2020), and Arnould et al. (2022). Considering the significant issues surrounding the market-based approach and this study's specificity in terms of context and sample, a balance-sheet-based approach is adopted in calculating the funding cost measures.

CPP_i is a dummy equal to 1 if bank i received support from the Capital Purchase Program (CPP) and 0 otherwise. X includes seven control variables focused on bank characteristics (size, deposit volume, and CAMEL variables—capital adequacy, assets quality, management quality, earnings, and liquidity). C represents four control variables based on regional economic conditions (competition, personal income, gross domestic product, and unemployment). The selection of bank-related control variables is inspired by existing literature where size (Altunbas, Gambacorta, and Marques-Ibanez 2010; Anbar and Alper 2011) and volume of liabilities (De Nicoló and Loukoianova 2007; Schaeck 2008; Gropp and Heider 2010; Gambacorta et al. 2017) are extensively used to represent bank characteristics. In addition, Barajas and Stein (2000) concluded that depositors select banks based on strong fundamentals, and banks

having stronger fundamentals enjoy funding cost benefits. Therefore, banks' fundamentals (which can be proxied by CAMELS) are important in funding cost discussion. The item "S" in CAMELS, i.e., "sensitivity to market risk," is excluded due to substantial missing values in the sample banks. Details on the variables' definitions are available in Table 2.

Time (τ_t) and state (α_s) fixed effects are used to control for variations relevant to time and states. Bank fixed effects are not applied, as they would exclude the main variable of interest CPP due to its time-invariant nature. ε_{it} is the error term. Equation (1) is estimated using a linear random-effects model, as it is more suitable in this study's context given that it includes a variable (CPP) that does not vary over the sample period (Baltagi 2008; Wooldridge 2010).

2.2.2 Fixed Effects

This study further investigates the government assistance–funding cost relationship by considering the potential influence of bank-level heterogeneity in terms of the CPP funds allocated. Instead of the CPP dummy, using the CPP amount specific to each bank allows incorporating bank fixed effects into the model, which is an additional methodological advantage. Furthermore, outstanding CPP amount is also examined. In this case, CPP amount of each bank is replaced with its CPP outstanding amount, where outstanding refers to the actual CPP amount in each quarter after adjusting for repayments made by the bank. The following Equations (2) and (3) represent the fixed-effect models:

$$FC_{i,t} = \beta_0 + \beta_1 CPPamount_{i,t} + \sum_{n=2}^{8} \beta_n X_{n,i,t} + \sum_{n=9}^{12} \beta_n C_{n,s,t} + \mu_i + \tau_t + \alpha_s + \varepsilon_{i,t} \tag{2}$$

$$FC_{i,t} = \beta_0 + \beta_1 CPPoutstanding_{i,t} + \sum_{n=2}^{8} \beta_n X_{n,i,t} + \sum_{n=9}^{12} \beta_n C_{n,s,t} + \mu_i + \tau_t + \alpha_s + \varepsilon_{i,t}. \tag{3}$$

Table 2. Variables' Notations and Definitions

Variable	Notation	Definition/Calculation
Cost of Fund	CF	Total interest expense as a percent of the sum of average interest-bearing liabilities and average non-interest-bearing deposits
Cost of Deposit	CD	Total interest expense on deposits as a percent of two-point average interest-bearing deposits. The two-point average is based on the current and previous calendar quarters.
Cost of Liabilities	CL	Total interest expense as a percent of average interest-bearing liabilities
Government Assistance	CPP	1 if the bank received CPP funds; 0 otherwise
CPP Amount	CPPamount	CPP amount received by a bank divided by its risk-weighted assets
CPP Amount Outstanding	CPPoutstanding	CPP amount outstanding by a bank divided by its risk-weighted assets. CPP amount outstanding refers to the actual CPP amount in each quarter after adjusting for the repayments made by the bank.
Capital Adequacy	CA	(Common stock + preferred stock) divided by total assets
Asset Quality	AQ	Non-performing assets divided by total assets
Management Quality	MQ	Cost-to-income ratio = (interest and related expense + non-interest expense) divided by (interest income + non-interest income)
Earnings	ER	Net income divided by average total assets
Liquidity	LQ	Cash and due from banks divided by total assets
Deposit	DP	Total deposit divided by total assets
Size	Size	The logarithm of total assets
Competition	CM	The logarithm of the total number of bank branches in the state
Personal Income	PI	The growth rate of per capita personal income of the state
Gross Domestic Product	GDP	The growth rate of the real gross domestic product of the state
Unemployment	UN	The unemployment rate of the state

Here, $CPPamount_{i,t}$ and $CPPoutstanding_{i,t}$ denote the CPP amount and CPP outstanding amount, respectively, as a fraction of risk-weighted assets of bank i at time t. μ_i is the bank fixed effect. Explanations for the other variables and subscripts are presented in the previous section.

2.2.3 Matching CPP and Non-CPP Banks

The propensity score matching (PSM) technique is used to identify comparable non-CPP banks for the CPP banks. Ten years of pre-CPP data (1999–2008) on bank characteristics (capital adequacy, asset quality, management quality, earnings, liquidity, deposit volume, and size) are used to calculate the propensity score in levels and dynamics (changes in the variables). PSM provides two stricter (caliper 0.01) 1:1 matched samples besides the unmatched sample.[4] Figure 1 shows how the difference between CPP and non-CPP banks in terms of funding cost (in levels and dynamics) changes across the matched and unmatched samples over the sample period.

To strengthen the matching exercise, the mean difference test between the CPP and non-CPP banks is conducted using 10 years' quarterly data before the CPP assignment. Table 3 presents the results of the mean difference test, which show that the two groups (CPP and non-CPP banks) in the sample matched in dynamics[5] are not statistically different in terms of capital adequacy, asset quality, management quality, earnings, liquidity, and size.

3. Results

3.1 Descriptive Statistics

Table 4 provides the summary statistics of the study's variables of interest. All the values are in the expected ranges. The negative minimum values of "Personal Income" and "GDP" are explained by the

[4]The respective sample sizes are 1:1 matched sample in levels (total 710 banks, 376 CPP banks, and 334 non-CPP banks), 1:1 matched sample in dynamics (total 736 banks, 377 CPP banks, and 359 non-CPP banks), and unmatched sample (total 8,327 banks, 414 CPP banks, and 7,913 non-CPP banks).

[5]The sample matched in levels provides similar results (unreported) for the mean difference test.

Figure 1. Funding Cost Difference between CPP and Non-CPP Banks (matched and unmatched samples)

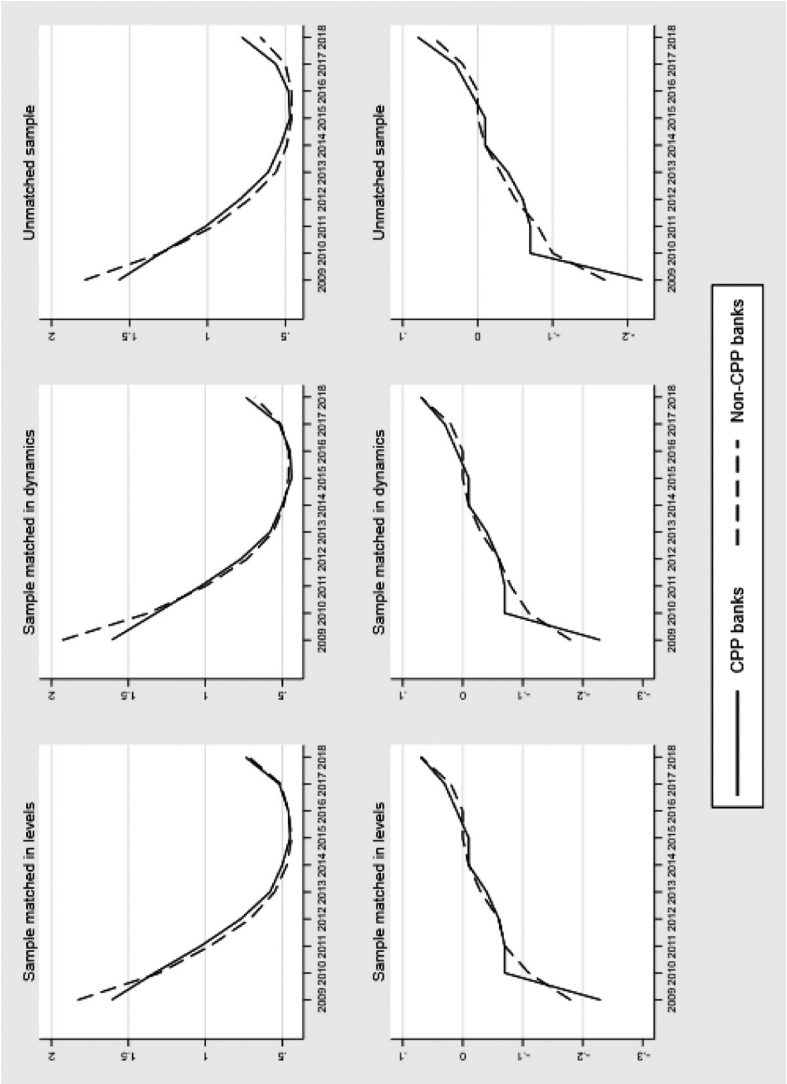

Note: The top row shows the difference between CPP and non-CPP banks in terms of funding cost measured in levels. The bottom row shows the difference between those banks when funding cost is measured in dynamics. To avoid confusion, keep in mind that the in-level and in-dynamics approaches are used both in the funding cost calculation and in the sample matching.

Table 3. Mean Difference Test

	Mean CPP Banks	Mean Non-CPP Banks	Normalized Difference in Means	t-stat	Median CPP Banks	Median Non-CPP Banks
Capital Adequacy	0.02	0.02	0.02	1.13	0.01	0.01
Asset Quality	0.01	0.01	0.02	1.61	0.00	0.00
Management Quality	0.79	0.79	0.01	0.40	0.78	0.76
Earnings	0.78	0.80	0.02	1.23	0.98	1.09
Liquidity	0.03	0.03	0.01	0.46	0.03	0.03
Deposit	0.77	0.79	0.20***	14.96	0.79	0.82
Size	13.58	13.56	0.02	1.22	13.58	13.35
Number of Banks	377	359				

Note: CPP means Capital Purchase Program. Banks that received funds from the Capital Purchase Program are termed CPP banks in this table. *** indicates statistical significance at the 1 percent level. This table reports the mean difference of bank characteristics between CPP and non-CPP banks in the sample matched in dynamics based on 10 years' quarterly data from 1999 to 2008 before the CPP.

Table 4. Summary Statistics

	Mean	Standard Deviation	Minimum	25th Percentile	Median	75th Percentile	Maximum
Cost of Fund	0.84	0.60	0.00	0.39	0.66	1.11	4.43
Cost of Deposit	0.94	0.63	0.01	0.47	0.75	1.23	5.41
Cost of Liabilities	1.01	0.65	0.17	0.52	0.82	1.33	5.55
Capital Adequacy	0.02	0.02	0.01	0.01	0.01	0.01	0.15
Asset Quality	0.02	0.03	0.00	0.00	0.01	0.03	0.13
Management Quality	0.77	0.18	0.45	0.66	0.74	0.84	1.58
Earnings	0.70	1.23	–5.55	0.43	0.85	1.26	3.65
Liquidity	0.03	0.02	0.00	0.01	0.02	0.03	0.12
Deposit	0.83	0.07	0.45	0.81	0.85	0.88	0.93
Size	12.43	1.47	9.41	11.43	12.21	13.25	17.23
Competition	7.56	0.79	4.73	7.13	7.49	8.23	8.87
Personal Income	2.58	3.38	–10.00	1.00	3.00	5.00	15.00
GDP	1.45	2.50	–9.00	1.00	2.00	3.00	22.00
Unemployment	6.38	2.40	2.00	4.00	6.00	8.00	15.00

Note: This table reports summary statistics for variables of interest based on quarterly data from the year 2009 to 2018.

fact that they are measured in percentage changes. The funding cost variables (cost of fund, cost of deposit, and cost of liabilities) show a considerable variation, which is evident from their respective standard deviations. CPP capital infusion could be one of the possible explanations for the substantial variation in the sample banks' funding costs. Especially, CPP banks' funding costs may considerably change after the CPP. Analyses ahead will explore these possibilities.

Independent variables included in the regression models are examined for potential collinearity issues. Table 5 shows that the variables are not affected by high collinearity. Variance inflation factors (VIFs) are calculated to detect possible multicollinearity among the independent variables. The corresponding VIF values (unreported) are below 5 and, therefore, rule out the possibility of multicollinearity according to the usual rule of thumb (Simon 2004).

3.2 Regression Results

Table 6 presents the results of Equation (1) using the matched samples and includes the three types of funding costs with different model specifications. The results show that government assistance is negatively related to banks' funding costs and the coefficient is statistically significant at the 1 percent and 5 percent levels. The finding remains consistent across all six model specifications. The relatively large coefficients' values, ranging from −0.064 to −0.264, are also noteworthy since they highlight the strength of the association between government assistance and banks' funding costs.

The results are also significant in an economic sense. In the sample matched in levels, CPP banks reduced their cost of fund, cost of deposit, and cost of liabilities by 8.33 percent, 14.57 percent, and 6.34 percent, respectively.[6] The funding cost reductions are significantly stronger in the sample matched in dynamics where the cost of fund, cost of deposit, and cost of liabilities reduce by 20 percent, 28.09 percent, and 16.04 percent, respectively. Compared with other funding costs, cost of deposit has the largest negative coefficients

[6]The coefficients of CPP in Table 6 are evaluated at the average cost of fund (0.84), cost of deposit (0.94), and cost of liabilities (1.01) to calculate the percentage reductions.

Table 5. Correlation Matrix

	CA	AQ	MQ	ER	LQ	DP	Size	CM	PI	GDP
AQ	0.13									
MQ	0.22	0.43								
ER	–0.19	–0.48	–0.74							
LQ	–0.02	–0.08	0.11	–0.03						
DP	–0.07	0.07	0.10	–0.10	0.15					
Size	–0.02	0.04	–0.20	0.06	–0.36	–0.29				
CM	0.17	0.06	0.08	–0.09	–0.02	0.05	0.12			
PI	0.00	–0.02	–0.11	0.14	–0.02	0.05	0.04	0.02		
GDP	0.00	–0.10	–0.14	0.18	0.00	0.06	0.03	0.05	0.66	
UN	0.20	0.33	0.26	–0.30	0.03	–0.04	0.04	0.33	–0.27	–0.39

Note: This table reports correlation among variables of interest based on quarterly data from the year 2009 to 2018. The notations are explained in Table 2.

Table 6. Government Assistance and Funding Costs

	Sample Matched in Levels			Sample Matched in Dynamics		
	Cost of Fund (1)	**Cost of Deposit (2)**	**Cost of Liabilities (3)**	**Cost of Fund (4)**	**Cost of Deposit (5)**	**Cost of Liabilities (6)**
Government Assistance	–0.070**	–0.137***	–0.064**	–0.168***	–0.264***	–0.162***
	(0.029)	(0.026)	(0.029)	(0.039)	(0.038)	(0.039)
Capital Adequacy	–0.650*	–0.275	–0.781**	–1.110***	–0.760**	–1.227***
	(0.339)	(0.350)	(0.347)	(0.308)	(0.337)	(0.320)
Asset Quality	0.472*	0.568**	0.481	0.404	0.509*	0.434
	(0.284)	(0.288)	(0.299)	(0.249)	(0.282)	(0.278)
Management Quality	0.212***	0.159***	0.233***	0.184***	0.151***	0.215***
	(0.033)	(0.035)	(0.034)	(0.029)	(0.032)	(0.031)
Earnings	–0.006*	–0.010***	–0.005	0.000	–0.002	0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Liquidity	–1.707***	–1.651***	–1.722***	–0.556*	–0.272	–0.268
	(0.398)	(0.507)	(0.531)	(0.331)	(0.367)	(0.385)
Deposit	–0.003	0.358**	0.001	–0.167	0.346***	–0.052
	(0.152)	(0.166)	(0.159)	(0.128)	(0.131)	(0.146)
Size	0.011	0.008	0.018	0.059***	0.061***	0.065***
	(0.016)	(0.016)	(0.017)	(0.017)	(0.017)	(0.017)
Competition	0.983***	0.861***	0.740***	0.626***	0.564**	0.328
	(0.232)	(0.262)	(0.253)	(0.229)	(0.240)	(0.246)
Personal Income	0.002	0.004**	0.003*	0.002	0.003	0.004*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
GDP	–0.001	–0.001	–0.001	–0.004*	–0.004	–0.004*
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
Unemployment	–0.008	–0.014**	–0.012*	0.000	–0.000	–0.001
	(0.006)	(0.006)	(0.006)	(0.005)	(0.006)	(0.005)
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.826	0.823	0.829	0.836	0.809	0.837
Bank-Quarter Observations	21,521	21,521	21,521	22,327	22,327	22,327
No. of Banks	710	710	710	736	736	736

Note: The variables are defined in Table 2. The numbers in parentheses are standard errors, robust to heteroskedasticity, and clustered by banks. ***, **, and * indicate statistical significance at 1 percent, 5 percent, and 10 percent levels, respectively. The results are obtained using the linear random-effect estimation of Equation (1) with 10 years of quarterly data from 2009 to 2018. 2009 is considered the starting point since banks started to receive government assistance from the last quarter of 2008.

in both matched samples. The finding also holds in terms of economic effect, as CPP banks have the largest reduction in the cost of deposit. Consistent with the findings of Berger et al. (2020), depositors (supply side) may have decreased market discipline towards CPP banks, considering these banks as safer and/or with a higher possibility of future bailouts (if needed). Also, CPP banks (demand side) may have decreased deposit demand due to having additional fund from government assistance programs and/or a reduction in asset portfolios. The combined effect (high supply and low demand) may have reduced the cost of deposit significantly. On average, CPP banks' funding costs reduced remarkably after the CPP assignment.

Regarding the control variables, the findings related to the CAMEL variables suggest that banks with better capital adequacy, asset quality, management quality, earnings, and liquidity would attract funding with lower costs. Such findings are consistent with Barajas and Stein (2000), who conclude that banks with stronger fundamentals enjoy funding cost benefits. Regarding deposit volume, the coefficients are only significant in the cost of deposit models, which can be expected because the other two types of funding costs are not focused primarily on deposits. Estimates suggest that banks with higher deposit volume and larger size are more likely to have higher funding costs. Since the findings contradict the usual expectations in literature, e.g., Berger and Roman (2015), and indicate a possibility of non-linearity in the relationship between bank size and funding cost, the analyses are rerun with large banks with assets over $3 billion only. The results (unreported) provide no statistically significant evidence that size and deposit volume positively associate with funding costs. Therefore, banks may have funding cost advantage only after reaching a certain threshold of assets and deposit volume, consistent with Jacewitz and Pogach (2018), who conclude that funding cost advantages are for the largest banks only. Among the regional control variables, the results of competition deserve attention due to the considerably large and significant coefficients indicating that banks facing higher regional competition are more likely to experience significantly higher funding costs.

As discussed in Section 2.2.2, this study takes into account the bank-level heterogeneity in terms of the CPP funds allocated in a fixed-effect model. Table 7 shows the results of Equation (2) using the two matched samples for the three types of funding costs.

Table 7. Government Assistance and Funding Costs: CPP Amount Received

	Sample Matched in Levels			Sample Matched in Dynamics		
	Cost of Fund (1)	Cost of Deposit (2)	Cost of Liabilities (3)	Cost of Fund (4)	Cost of Deposit (5)	Cost of Liabilities (6)
CPP Amount	–0.002***	–0.003***	–0.000	–0.001	–0.003***	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Capital Adequacy	–0.433	0.039	–0.533	–1.024***	–0.622*	–1.121***
	(0.378)	(0.405)	(0.039)	(0.335)	(0.375)	(0.349)
Asset Quality	0.208	0.259	0.195	0.217	0.277	0.227
	(0.286)	(0.292)	(0.302)	(0.253)	(0.290)	(0.283)
Management Quality	0.221***	0.181***	0.247***	0.198***	0.180***	0.234***
	(0.034)	(0.035)	(0.035)	(0.030)	(0.032)	(0.031)
Earnings	–0.004	–0.007**	–0.003	0.003	0.002	0.007**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Liquidity	–1.541***	–1.479***	–1.565***	–0.488	–0.183	–0.181
	(0.405)	(0.524)	(0.544)	(0.335)	(0.369)	(0.392)
Deposit	0.030	0.347*	0.030	–0.174	0.315**	–0.052
	(0.161)	(0.177)	(0.169)	(0.131)	(0.134)	(0.151)
Size	0.064**	0.088***	0.081***	0.112***	0.143***	0.128***
	(0.027)	(0.029)	(0.029)	(0.025)	(0.026)	(0.026)
Competition	0.962***	0.840***	0.715***	0.620***	0.559**	0.320
	(0.234)	(0.265)	(0.254)	(0.230)	(0.244)	(0.247)
Personal Income	0.002	0.004*	0.003	0.002	0.003	0.003*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
GDP	–0.001	–0.001	–0.001	–0.004*	–0.004	–0.004*
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
Unemployment	–0.007	–0.013**	–0.010*	0.001	0.001	0.000
	(0.006)	(0.006)	(0.006)	(0.005)	(0.006)	(0.005)
Bank Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.827	0.824	0.830	0.837	0.810	0.838
Bank-Quarter Observations	21,521	21,521	21,521	22,327	22,327	22,327
No. of Banks	710	710	710	736	736	736

Note: The variables are defined in Table 2. The numbers in parentheses are standard errors, robust to heteroskedasticity, and clustered by banks. ***, **, and * indicate statistical significance at 1 percent, 5 percent, and 10 percent levels, respectively. The results are obtained using the fixed-effect estimation of Equation (2) with 10 years of quarterly data from 2009 to 2018. 2009 is considered the starting point since banks started to receive government assistance from the last quarter of 2008.

Here, the study's main independent variable of interest, i.e., CPP, is replaced with each bank's CPP amount as a ratio of its risk-weighted assets (RWA). Although results are not significant in all the columns, the key finding remains consistent in the significant coefficients, suggesting that government assistance is negatively related to funding costs. In an economic sense, banks reduce 0.24 percent and 0.32 percent of their cost of fund and cost of deposit, respectively, for each unit increase in their CPP allocated amount/RWA ratio. Simply put, higher CPP fund allocation may have resulted in lower funding costs for banks. In line with the conclusion in Table 6, depositors' response remains higher compared with other types of funding costs.

To understand the importance of CPP amount further, this study utilizes another bank-level granular measure that considers the banks' outstanding CPP amount[7] as opposed to the initially allocated CPP amount. Table 8 presents the results of Equation (3), which considers CPP outstanding amount as a ratio of risk-weighted assets as the main independent variable. The main findings remain consistent and, compared with results in Table 7, the coefficients demonstrate a stronger effect.

Market (depositors' and investors') response to CPP banks may be driven by the timing of the CPP fund's repayment. To examine this possibility, CPP banks are grouped into two categories according to their repayment behavior. CPP banks that repaid the CPP fund in full within 2010 are categorized as "repaid early." The other category contains CPP banks that did not repay the full CPP amount within 2010. Table 9 shows the results for the two categories of banks. Banks that repaid early demonstrate significant negative coefficients in cost of fund and cost of deposit, indicating that funding cost benefits may only be applicable for the early repayers. The finding is consistent with relevant research by Berger and Roman (2015), who conclude that recipients of government assistance that repaid early hold a significant competitive advantage, while the other recipients do not show such results.

The CPP banks' level of capital before CPP injection can be linked with their funding cost afterward. To test this possibility,

[7]The actual CPP amount in each quarter after adjusting for repayments made by banks.

Table 8. Government Assistance and Funding Costs: CPP Amount Outstanding

	Sample Matched in Levels			Sample Matched in Dynamics		
	Cost of Fund (1)	Cost of Deposit (2)	Cost of Liabilities (3)	Cost of Fund (4)	Cost of Deposit (5)	Cost of Liabilities (6)
CPP Outstanding	–0.259	–0.549*	–0.355	–0.854**	–1.291***	–1.057**
	(0.277)	(0.286)	(0.296)	(0.344)	(0.413)	(0.414)
Capital Adequacy	–0.391	0.126	–0.478	–0.878***	–0.400	–0.940***
	(0.379)	(0.409)	(0.391)	(0.331)	(0.377)	(0.348)
Asset Quality	0.239	0.324	0.237	0.301	0.405	0.331
	(0.289)	(0.294)	(0.304)	(0.252)	(0.287)	(0.282)
Management Quality	0.222***	0.182***	0.247***	0.200***	0.182***	0.236***
	(0.034)	(0.035)	(0.035)	(0.029)	(0.031)	(0.031)
Earnings	–0.004	–0.007**	–0.003	0.003	0.002	0.007**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Liquidity	–1.544***	–1.486***	–1.570***	–0.494	–0.191	–0.189
	(0.405)	(0.522)	(0.543)	(0.335)	(0.369)	(0.392)
Deposit	0.032	0.352**	0.033	–0.167	0.325**	–0.044
	(0.161)	(0.177)	(0.168)	(0.131)	(0.133)	(0.150)
Size	0.064**	0.089***	0.082***	0.112***	0.144***	0.128***
	(0.027)	(0.029)	(0.029)	(0.024)	(0.026)	(0.026)
Competition	0.972***	0.860***	0.729***	0.659***	0.617**	0.368
	(0.233)	(0.265)	(0.254)	(0.228)	(0.243)	(0.247)
Personal Income	0.002	0.004*	0.003	0.002	0.003	0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
GDP	–0.001	–0.001	–0.001	–0.003*	–0.004	–0.004*
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
Unemployment	–0.007	–0.013**	–0.010*	0.001	0.001	0.000
	(0.006)	(0.006)	(0.006)	(0.005)	(0.006)	(0.005)
Bank Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.827	0.824	0.830	0.837	0.811	0.838
Bank-Quarter Observations	21,521	21,521	21,521	22,327	22,327	22,327
No. of Banks	710	710	710	736	736	736

Note: The variables are defined in Table 2. The numbers in parentheses are standard errors, robust to heteroskedasticity, and clustered by banks. ***, **, and * indicate statistical significance at 1 percent, 5 percent, and 10 percent levels, respectively. The results are obtained using the fixed-effect estimation of Equation (3) with 10 years of quarterly data from 2009 to 2018. 2009 is considered the starting point since banks started to receive government assistance from the last quarter of 2008.

Table 9. Government Assistance and Funding Costs: CPP Repayment Behavior

	Banks Repaid Early			Banks Did Not Repay Early		
	Cost of Fund (1)	Cost of Deposit (2)	Cost of Liabilities (3)	Cost of Fund (4)	Cost of Deposit (5)	Cost of Liabilities (6)
CPP Amount	–0.001**	–0.002***	0.001	18.455*	12.572	18.317
	(0.001)	(0.001)	(0.001)	(9.433)	(8.165)	(11.168)
Capital Adequacy	–0.765*	–0.294	–0.863**	–2.244	–1.973	–2.465
	(0.409)	(0.447)	(0.425)	(2.205)	(2.681)	(2.662)
Asset Quality	0.538*	0.773**	0.465	2.067	2.242	1.934
	(0.298)	(0.311)	(0.333)	(1.751)	(1.444)	(1.895)
Management Quality	0.113***	0.044	0.135***	0.235***	0.114	0.243**
	(0.029)	(0.030)	(0.032)	(0.077)	(0.070)	(0.089)
Earnings	–0.005	–0.010***	–0.002	0.006	–0.001	0.009
	(0.003)	(0.003)	(0.003)	(0.008)	(0.006)	(0.010)
Liquidity	–1.032***	–0.483	–0.567	–0.057	–0.785	–0.286
	(0.364)	(0.402)	(0.431)	(0.820)	(0.830)	(0.944)
Deposit	–0.224	0.234	–0.247	1.526	0.874	1.670
	(0.148)	(0.146)	(0.164)	(1.000)	(0.860)	(1.125)
Size	0.038	0.066**	0.057*	1.280**	0.905**	1.129*
	(0.027)	(0.029)	(0.031)	(0.506)	(0.427)	(0.595)
Competition	0.818***	0.763**	0.572*	0.184	–0.202	0.351
	(0.277)	(0.349)	(0.333)	(1.509)	(1.012)	(1.493)
Personal Income	0.001	0.002	0.003	0.002	–0.015	0.006
	(0.002)	(0.002)	(0.002)	(0.020)	(0.016)	(0.022)
GDP	0.000	–0.001	–0.001	–0.004	–0.004	–0.005
	(0.002)	(0.003)	(0.003)	(0.009)	(0.006)	(0.010)
Unemployment	–0.004	–0.007	–0.006	–0.019	–0.021	–0.030
	(0.005)	(0.006)	(0.006)	(0.024)	(0.023)	(0.025)
Bank Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.872	0.874	0.868	0.837	0.899	0.832
Bank-Quarter Observations	11,375	11,375	11,375	222	222	222
No. of Banks	349	349	349	27	27	27

Note: The variables are defined in Table 2. The numbers in parentheses are standard errors, robust to heteroskedasticity, and clustered by banks. ***, **, and * indicate statistical significance at 1 percent, 5 percent, and 10 percent levels, respectively. The results are obtained using the fixed-effect estimation of Equation (2) for the sample matched in dynamics with 10 years of quarterly data from 2009 to 2018. 2009 is considered the starting point since banks started to receive government assistance from the last quarter of 2008. CPP banks that repaid the CPP fund in full within 2010 are categorized as repaid early.

the CPP banks are grouped by the risk-based capital ratio[8] at the time of CPP infusion. Since the average risk-based capital ratio of this study's sample is 16.55 percent in 2018:Q3, banks with such a ratio less than or equal to 16.55 percent in 2008:Q3 are grouped as low-equity banks, otherwise high-equity banks. The results in Table 10 show that low-equity banks experience funding cost benefits after the CPP infusion. This could be due to increased public confidence in such banks after receiving government assistance. However, CPP banks having higher capital before CPP demonstrate no significant results. Possibly, the market perception of high-equity banks remained unchanged after CPP. Existing literature provides support for the findings on low- versus high-equity CPP banks from the non-linearity angle. Aymanns et al. (2016) conclude that funding cost and solvency have a non-linear relationship such that the sensitivity of funding cost is higher at the lower level of solvency. The conclusion is further reinforced by Dent, Hoke, and Panagiotopoulos (2021), who study solvency shocks and found a similar non-linear negative relationship where funding costs' response is greater at the lower solvency levels. Besides the market perception perspective, such a non-linear relationship could be another reason why the low-equity CPP banks demonstrate significant changes in funding costs after CPP allocation.

Regional market concentration can have important implications on the banks' funding costs. Although the main analyses control for regional competition in terms of the number of bank branches, analyses on sub-samples based on the market concentration of deposits can offer further insights. Using the Herfindahl-Hirschman Index (HHI) deposits[9] as of June 30, 2009 (time around CPP), the sample banks are grouped by following the U.S. Department of Justice's

[8]Risk-based capital ratio is total regulatory capital as a percent of risk-adjusted assets. For Call Report and FRY-9C filers, depending on institution attributes and time period, it represents risk-based capital ratio reported under either the U.S. Basel III (B3) revised regulatory capital rules, advanced approaches rules or otherwise, or the general risk-based (GRB) regulatory capital rules. Preference between the GRB, B3, and B3-Post Parallel Run ratios is given based on the nature of the filing, the lowest of the B3 and B3-Post Parallel Run ratios, when available. Additionally, a general preference is given to B3 ratios over the GRB ratios, where applicable.

[9]"Deposit Market Share Reports," Federal Deposit Insurance Corporation: available at https://www7.fdic.gov/sod/sodMarketBank.asp.

Table 10. Government Assistance and Funding Costs: Low- vs. High-Equity Banks

	Sample Matched in Levels			Sample Matched in Dynamics		
	Cost of Fund (1)	Cost of Deposit (2)	Cost of Liabilities (3)	Cost of Fund (4)	Cost of Deposit (5)	Cost of Liabilities (6)
A. Low-Equity Banks						
Government Assistance	–0.076**	–0.144***	–0.073**	–0.205***	–0.313***	–0.199***
	(0.033)	(0.031)	(0.034)	(0.045)	(0.046)	(0.045)
Bank-Level Controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.839	0.843	0.840	0.845	0.847	0.844
Bank-Quarter Observations	18,023	18,023	18,023	18,392	18,392	18,392
No. of Banks	600	600	600	615	615	615
B. High-Equity Banks						
Government Assistance	–0.029	–0.032	–0.003	0.061	–0.038	0.038
	(0.071)	(0.063)	(0.067)	(0.064)	(0.064)	(0.069)
Bank-Level Controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.783	0.757	0.790	0.805	0.664	0.811
Bank-Quarter Observations	3,498	3,498	3,498	3,933	3,933	3,933
No. of Banks	110	110	110	121	121	121

Note: The variables are defined in Table 2. The numbers in parentheses are standard errors, robust to heteroskedasticity, and clustered by banks. *** and ** indicate statistical significance at 1 percent and 5 percent levels, respectively. The results are obtained using the linear random-effect estimation of Equation (1) with 10 years of quarterly data from 2009 to 2018. 2009 is considered the starting point since banks started to receive government assistance from the last quarter of 2008. Since the average risk-based capital ratio of this study's sample is 16.55 percent in 2018:Q3, banks with such a ratio less than or equal to 16.55 percent in 2008:Q3 are grouped as low-equity banks, otherwise high-equity banks. Detailed results are available upon request.

guidelines[10] as follows: unconcentrated markets (HHI $\leq$ 1,000) and concentrated markets (HHI $>$ 1,000). The results are reported in Table 11. In most cases, banks in unconcentrated markets show a statistically significant reduction in funding costs after CPP. CPP banks' bargaining power on funding costs could be stronger in regions with low depositors' concentration. From a regulatory perspective, such low market concentration may have useful implications in periods of recovery, e.g., after the 2007–08 crisis. In stress situations, regulators' policy focus aims at maintaining banking system stability and enabling economic recovery with sustained lending exercise. Controlling market concentration through regulations may favorably moderate the effects of government's capital assistance by providing opportunities for banks to raise low-cost fund and fueling the economy with required lending at reasonable interest rates. In addition, low-concentration-led low funding cost can improve banks' profitability (by reducing costs), risk-taking (by reducing high-risk lending to cover high costs), and government assistance fund repayment capacity (by increasing earnings and equity), which are highly important for regulators to maintain banking system stability.

Overall, the results suggest that government assistance through capital injection leads to a decrease in banks' funding costs. The findings can be explained via the safety channel where the market (depositors and investors) may perceive CPP banks as safe and would receive further government support, if necessary. Berger and Roman (2015) advocated such a safety channel while explaining the competitive advantage of banks receiving government assistance. Since CPP banks were government supported and/or government's selection criteria identified these banks as "healthy and viable," depositors and investors might prefer CPP banks and agree to receive lower interest rates, considering these banks as less likely to become financially distressed.

Current literature tends to document relatively more support for the effectiveness of government assistance in boosting the overall market confidence and ensuring stability of the banking system. TARP-CPP funds contributed to reinstating the confidence in financial markets through a considerable increase in bank lending

[10] "Herfindahl-Hirschman Index," U.S. Department of Justice: available at https://www.justice.gov/atr/herfindahl-hirschman-index.

Table 11. Government Assistance and Funding Costs: Unconcentrated vs. Concentrated Markets

	Sample Matched in Levels			Sample Matched in Dynamics		
	Cost of Fund (1)	Cost of Deposit (2)	Cost of Liabilities (3)	Cost of Fund (4)	Cost of Deposit (5)	Cost of Liabilities (6)
A. Banks in Unconcentrated Markets						
Government Assistance	–0.083**	–0.147***	–0.089***	–0.192***	–0.279***	–0.190***
	(0.032)	(0.030)	(0.033)	(0.042)	(0.040)	(0.043)
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Bank-Level Controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional Controls	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.854	0.858	0.856	0.852	0.816	0.849
Bank-Quarter Observations	15,702	15,702	15,702	16,950	16,950	16,950
No. of Banks	518	518	518	557	557	557
B. Banks in Concentrated Markets						
Government Assistance	–0.043	–0.122**	–0.011	–0.085	–0.204**	–0.058
	(0.058)	(0.055)	(0.057)	(0.084)	(0.086)	(0.083)
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Bank-Level Controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional Controls	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.755	0.735	0.762	0.794	0.804	0.808
Bank-Quarter Observations	5,819	5,819	5,819	5,378	5,378	5,378
No. of Banks	192	192	192	179	179	179

Note: The variables are defined in Table 2. The numbers in parentheses are standard errors, robust to heteroskedasticity, and clustered by banks. *** and ** indicate statistical significance at 1 percent and 5 percent levels, respectively. The results are obtained using the linear random-effect estimation of Equation (1) with 10 years of quarterly data from 2009 to 2018. 2009 is considered the starting point since banks started to receive government assistance from the last quarter of 2008. For regional market concentration, banks are grouped using the Herfindahl-Hirschman Index (HHI) deposits as follows: unconcentrated markets (HHI $\leq$ 1,000) and concentrated markets (HHI $>$ 1,000). Detailed results are available upon request.

(Li 2010). The recipient institutions, on average, retained two-thirds of the TARP funds to maintain a stronger capital position and employed the rest in new loan creation. The government assistance is aimed at regaining confidence both at institutional and consumer levels such that the interbank lending rates decrease and greater liquidity is stimulated in the economy (Northehr 2008).

The positive influence of government assistance in restoring market confidence was also demonstrated in the stock market measures. During the 2007–08 crisis, investors' confidence (captured through investor sentiment measures) plummeted due to high anxiety and uncertainty induced by the economy-wide financial distress (Swedberg 2013). General announcements about the government interventions in the 2007–08 crisis are found to have positive impacts on stock returns (Fratianni and Marchionne 2010). Huerta, Perez-Liston, and Jackson (2011) concluded that government intervention through TARP assisted in diminishing investors' anxiety and stock market volatility in the short run. Similarly, Ncube (2016) documented favorable market responses after TARP announcements, suggesting that government assistance improved investors' confidence.

The safety channel may also be augmented by extensive regulatory oversight, stricter policies, and stabilization in both financial markets and banks' financial health that helped regain the lost public confidence. In periods after the CPP assignment, federal regulators established rigorous monitoring procedures and stricter controls for the CPP banks (Shah 2009; Agarwal et al. 2014). The stringent transparency and monitoring practices can be expected to positively influence the CPP banks' financial health and the investors' confidence in these banks.

The findings agree with the related empirical evidence on the funding cost (Aymanns et al. 2016; Schmitz, Sigmund, and Valderrama 2017), lower deposit rates for protected banks (Koetter and Noth 2016), and theoretical discussion of Northehr (2008) about assisted banks regaining lending confidence and decreasing funding costs. The results broadly align with the branch of literature arguing that more capital is negatively related to banks' funding costs from the solvency perspective (Barajas and Stein 2000; Ungan, Caner, and Özyıldırım 2008; Annaert et al. 2013; Babihuga and Spaltro 2014; Pierret 2014; Acharya and Mora 2015; Hasan, Liu, and Zhang

2016; Carvalho and Dantas 2020; Moreira 2020; Dent, Hoke, and Panagiotopoulos 2021; Arnould et al. 2022) and competition perspective (Calomiris and Mason 2003; Calomiris and Wilson 2004; Allen, Carletti, and Marquez 2011; Mehran and Thakor 2011; Berger and Bouwman 2013; Berger and Roman 2015).

3.3 Robustness Tests

Besides multiple types of funding costs and matched samples analyzed in Section 3.2, several robustness tests are conducted to examine if the main findings remain consistent with alternative samples and estimation methods.

Since the baseline equation uses random effects and some control variables that are jointly determined with funding costs, the regression results in Table 6 are analyzed without any control variables and random effects as a robustness test. Another robustness test is conducted with the unmatched sample while keeping the model specifications of Equation (1) intact. The results, presented in panel A of Table 12, agree with the baseline findings.

Although most CPP recipients participated voluntarily, a few banks were forced to join CPP at its inception. However, the involuntary participants[11] left the program early due to stricter regulatory control, CEO compensation issues, and access to cheaper funds (Berger and Roman 2015). To eliminate the possible bias, the main analyses are repeated excluding the nine involuntary participants in this study's sample. The results remain consistent with the main findings.

The Supervisory Capital Assessment Program (SCAP) required mandatory stress tests of banks having assets exceeding $100 billion.[12] These banks cover about two-thirds of the total banking assets and half of the total loans in the United States (Berger

[11]The nine involuntary participants are Goldman Sachs, Morgan Stanley, JP Morgan Chase, Citigroup, Wells Fargo, State Street, Bank of New York Mellon, Bank of America, and Merrill Lynch (acquired by Bank of America Corp.).

[12]The 19 stress-tested banks are Bank of America, Citigroup, Goldman Sachs, JP Morgan Chase, Morgan Stanley, Wells Fargo, Bank of New York Mellon, BB&T (Truist Financial), Fifth Third Bancorp, Keycorp, PNC Financial, Regions Financial, SunTrust Banks, US Bancorp, Ally Financial, American Express Company, Capital One Financial, Metlife, and State Street.

Table 12. Robustness Tests

Panel A									
	Sample Matched in Levels			**Sample Matched in Dynamics**			**Unmatched Sample**		
	Cost of Fund (1)	**Cost of Deposit (2)**	**Cost of Liabilities (3)**	**Cost of Fund (4)**	**Cost of Deposit (5)**	**Cost of Liabilities (6)**	**Cost of Fund (7)**	**Cost of Deposit (8)**	**Cost of Liabilities (9)**
Government Assistance	–0.016***	–0.099***	–0.010*	–0.033***	–0.150***	–0.038***	–0.225***	–0.324***	–0.230***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.023)	(0.024)	(0.023)
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank-Level Controls	No	No	No	No	No	No	Yes	Yes	Yes
Regional Controls	No	No	No	No	No	No	Yes	Yes	Yes
R-squared (within)	0.638	0.653	0.657	0.671	0.668	0.689	0.819	0.837	0.828
Bank-Quarter Observations	21,521	21,521	21,521	22,327	22,327	22,327	263,562	263,562	263,562
No. of Banks	710	710	710	736	736	736	8,327	8,327	8,327

(*continued*)

Table 12. (Continued)

Panel B									
	Excluding Involuntary Banks			**Excluding Stress-Tested Banks**			**Excluding Large Banks**		
	Cost of Fund (1)	**Cost of Deposit (2)**	**Cost of Liabilities (3)**	**Cost of Fund (4)**	**Cost of Deposit (5)**	**Cost of Liabilities (6)**	**Cost of Fund (7)**	**Cost of Deposit (8)**	**Cost of Liabilities (9)**
Government Assistance	–0.171***	–0.268***	–0.166***	–0.172***	–0.267***	–0.168***	–0.145***	–0.243***	–0.144***
	(0.039)	(0.038)	(0.039)	(0.039)	(0.038)	(0.039)	(0.037)	(0.037)	(0.037)
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank-Level Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Regional Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared (within)	0.837	0.809	0.838	0.837	0.810	0.838	0.843	0.812	0.845
Bank-Quarter Observations	22,247	22,247	22,247	21,927	21,927	21,927	19,017	19,017	19,017
No. of Banks	728	728	728	717	717	717	630	630	630

Note: The variables are defined in Table 2. The numbers in parentheses are standard errors, robust to heteroskedasticity, and clustered by banks. *** and * indicate statistical significance at 1 percent and 10 percent levels, respectively. The results are obtained using the linear random-effect estimation (except for columns 1–6 in panel A) of Equation (1) with 10 years of quarterly data from 2009 to 2018. 2009 is considered the starting point since banks started to receive government assistance from the last quarter of 2008. The nine involuntary participants are Goldman Sachs, Morgan Stanley, JP Morgan Chase, Citigroup, Wells Fargo, State Street, Bank of New York Mellon, Bank of America, and Merrill Lynch (acquired by Bank of America Corp.). The 19 stress-tested banks are Bank of America, Citigroup, Goldman Sachs, JP Morgan Chase, Morgan Stanley, Wells Fargo, Bank of New York Mellon, BB&T (Truist Financial), Fifth Third Bancorp, Keycorp, PNC Financial, Regions Financial, SunTrust Banks, US Bancorp, Ally Financial, American Express Company, Capital One Financial, Metlife, and State Street. Following Berger and Bouwman (2013), banks with total assets higher than $3 billion are categorized as large banks. Results in panel B are based on the sample matched in dynamics. Detailed results are available upon request.

and Roman 2015). To ensure public confidence in the financial system, SCAP publicized these banks as too big to fail and confirmed that the U.S. Treasury would assist them with adequate capital in case of an adverse scenario. Such special treatment by SCAP may have an important influence on the stress-tested banks' funding cost advantage, which may cause bias in this study's results. Therefore, the main analyses are re-estimated while excluding the stress-tested banks. The baseline conclusions remain valid after such exclusion.

A bank's size may be considered an economic strength by fund providers. Thus, large banks may have a competitive advantage in the funding market, which may cause bias in this study's result. To eliminate this possibility, robustness of the main results is tested after excluding the large banks from the sample. Following Berger and Bouwman (2013), banks with total assets higher than $3 billion are categorized as large banks. The main results remain valid without the large banks.

In general, the main conclusion that CPP support is associated with lower funding costs for CPP banks is robust to several alternative funding cost definitions, samples, estimation methods, and model specifications. To conserve space, a summary of the six robustness tests is presented in Table 12, where panel B shows the last three robustness tests. Detailed results are available upon request.

4. Conclusion

This study aims at documenting empirical evidence about the relationship of government assistance with recipient banks' funding costs. Understanding the link of government assistance to funding cost is the study's main contribution, as we believe that this is the first work in this research space. The results suggest that government assistance has a significant relationship with the recipient bank's lower funding cost. Depositors' and investors' increased confidence in the recipient banks could be a plausible channel to explain the government assistance–funding cost relationship. Extensive regulatory oversight and stricter policies could also amplify the market confidence in the recipient banks. The results are robust to several alternative funding cost definitions, samples, estimation methods, and model specifications. Due to contextual limitations,

the study's sample considers the U.S. banking sector only. Further studies may investigate a cross-country sample to obtain more generalizable results.

References

Acharya, V. V., and N. Mora. 2015. "A Crisis of Banks as Liquidity Providers." *Journal of Finance* 70 (1): 1–43.

Agarwal, S., D. Lucca, A. Seru, and F. Trebbi. 2014. "Inconsistent Regulators: Evidence from Banking." *Quarterly Journal of Economics* 129 (2): 889–938.

Allen, F., E. Carletti, and R. Marquez. 2011. "Credit Market Competition and Capital Regulation." *Review of Financial Studies* 24 (4): 983–1018.

Altunbas, Y., L. Gambacorta, and D. Marques-Ibanez. 2010. "Bank Risk and Monetary Policy." *Journal of Financial Stability* 6 (3): 121–29.

Anbar, A., and D. Alper. 2011. "Bank Specific and Macroeconomic Determinants of Commercial Bank Profitability: Empirical Evidence from Turkey." *Business and Economics Research Journal* 2 (2): 139–52.

Anderson, H. P., and J. C. Bluedorn. 2017. "Stopping Contagion with Bailouts: Micro-evidence from Pennsylvania Bank Networks during the Panic of 1884." *Journal of Banking and Finance* 76 (3): 139–49.

Annaert, J., M. De Ceuster, P. Van Roy, and C. Vespro. 2013. "What Determines Euro Area Bank CDS Spreads?" *Journal of International Money and Finance* 32 (2): 444–61.

Arnould, G., G. Avignone, C. Pancaro, and D. Żochowski. 2022. "Bank Funding Costs and Solvency." *European Journal of Finance* 28 (10): 931–63.

Aymanns, C., C. Caceres, C. Daniel, and L. Schumacher. 2016. "Bank Solvency and Funding Cost." IMF Working Paper No. 16/64.

Babihuga, R., and M. Spaltro. 2014. "Bank Funding Costs for International Banks." IMF Working Paper No. 14/71.

Baltagi, B. 2008. *Econometric Analysis of Panel Data.* John Wiley & Sons.

Barajas, A., and R. Stein. 2000. "Depositor Behavior and Market Discipline in Colombia." IMF Working Paper No. 00/214.

Bayazitova, D., and A. Shivdasani. 2011. "Assessing Tarp." *Review of Financial Studies* 25 (2): 377–407.

Beau, E., J. Hill, T. Hussain, and D. Nixon. 2014. "Bank Funding Costs: What Are They, What Determines Them and Why Do They Matter?" *Quarterly Bulletin* (Bank of England) 54 (4): 370–84.

Berger, A. N., and C. H. Bouwman. 2013. "How Does Capital Affect Bank Performance during Financial Crises?" *Journal of Financial Economics* 109 (1): 146–76.

Berger, A. N., M. Lamers, R. A. Roman, and K. J. Schoors. 2020. "Unexpected Effects of Bank Bailouts: Depositors Need Not Apply and Need Not Run." Available at SSRN 3712641.

Berger, A. N., and R. A. Roman. 2015. "Did TARP Banks Get Competitive Advantages?" *Journal of Financial and Quantitative Analysis* 50 (6): 1199–1236.

Berger, A. N., R. A. Roman, and J. Sedunov. 2020. "Did TARP Reduce or Increase Systemic Risk? The Effects of Government Aid on Financial System Stability." *Journal of Financial Intermediation* 43 (July): Article 100810.

Black, L. K., and L. N. Hazelwood. 2013. "The Effect of TARP on Bank Risk-taking." *Journal of Financial Stability* 9 (4): 790–803.

Cadman, B., M. E. Carter, and L. J. Lynch. 2012. "Executive Compensation Restrictions: Do They Restrict Firms' Willingness to Participate in TARP?" *Journal of Business Finance and Accounting* 39 (7–8): 997–1027.

Calomiris, C. W., and J. R. Mason. 2003. "Consequences of Bank Distress during the Great Depression." *American Economic Review* 93 (3): 937–47.

Calomiris, C. W., and B. Wilson. 2004. "Bank Capital and Portfolio Management: The 1930s "Capital Crunch" and the Scramble to Shed Risk." *Joumal of Business* 77 (2): 421–56.

Carvalho, J. A. D., and J. A. Dantas. 2020. "Relationship between Market Discipline and Capital Buffers in Brazilian Banks." *Revista Contabilidade & Finanças* 32 (85): 109–25.

Chavaz, M., and A. K. Rose. 2019. "Political Borders and Bank Lending in Post-crisis America." *Review of Finance* 23 (5): 935–59.

Cordella, T., and E. L. Yeyati. 2003. "Bank Bailouts: Moral Hazard vs. Value Effect." *Journal of Financial intermediation* 12 (4): 300–330.

Cummings, J. R., and S. Wright. 2016. "Effect of Higher Capital Requirements on the Funding Costs of Australian Banks." *Australian Economic Review* 49 (1): 44–53.

De Nicoló, M. G., and E. Loukoianova. 2007. "Bank Ownership, Market Structure and Risk." IMF Working Paper No. 07/215.

Dent, K., S. H. Hoke, and A. Panagiotopoulos. 2021. "Solvency and Wholesale Funding Cost Interactions at UK Banks." *Journal of Financial Stability* 52 (February): Article 100799.

Duchin, R., and D. Sosyura. 2010. "TARP Consequences: Lending and Risk Taking." Available at SSRN 1722348.

Fratianni, M. U., and F. Marchionne. 2010. "Banks' Great Bailout of 2008-2009." Available at SSRN 1559413.

Gambacorta, L., G. Ricotti, S. M. Sundaresan, and Z. Wang. 2017. "The Effects of Tax on Bank Liability Structure." BIS Working Paper No. 611. https://www.bis.org/publ/work611.pdf.

Gerhardt, M., and R. Vander Vennet. 2017. "Bank Bailouts in Europe and Bank Performance." *Finance Research Letters* 22 (3): 74–80.

Gropp, R., and F. Heider. 2010. "The Determinants of Bank Capital Structure." *Review of Finance* 14 (4): 587–622.

Grossman, R. J., and M. Hansen. 2010. "CDS Spreads and Default Risk: Interpreting the Signals." Special Report, Fitch Ratings.

Harris, O., D. Huerta, and T. Ngo. 2013. "The Impact of TARP on Bank Efficiency." *Journal of International Financial Markets, Institutions and Money* 24 (2): 85–104.

Hasan, I., L. Liu, and G. Zhang. 2016. "The Determinants of Global Bank Credit-Default-Swap Spreads." *Journal of Financial Services Research* 50 (3): 275–309.

Hicks, J. R. 1935. "Annual Survey of Economic Theory: The Theory of Monopoly." *Econometrica* 3 (1): 1–20.

Hoshi, T., and A. K. Kashyap. 2010. "Will the US Bank Recapitalization Succeed? Eight Lessons from Japan." *Journal of Financial Economics* 97 (3): 398–417.

Huerta, D., D. Perez-Liston, and D. Jackson. 2011. "The Impact of TARP Bailouts on Stock Market Volatility and Investor Fear." *Banking and Finance Review* 3 (1): 45–54.

Jacewitz, S., and J. Pogach. 2018. "Deposit Rate Advantages at the Largest Banks." *Journal of Financial Services Research* 53 (1): 1–35.

Keeley, M. C. 1990. "Deposit Insurance, Risk, and Market Power in Banking." *American Economic Review* 80 (5): 1183–1200.

Kiser, E. K. 2003. "Modeling the Whole Firm: The Effect of Multiple Inputs and Financial Intermediation on Bank Deposit Rates." FEDS Working Paper No. 2004-07. https://ssrn.com/abstract=515863.

Koetter, M., and F. Noth. 2016. "Did TARP Distort Competition among Sound Unsupported Banks?" *Economic Inquiry* 54 (2): 994–1020.

Li, L. 2010. "TARP Funds Distribution and Bank Loan Growth." https://ssrn.com/abstract=1567073.

Mehran, H., and A. Thakor. 2011. "Bank Capital and Value in the Cross-Section." *Review of Financial Studies* 24 (4): 1019–67.

Moreira, F. 2020. "Financial Institutions' Funding Cost: Do Capital and Risk-taking Matter?" *Journal of Risk Management in Financial Institutions* 14 (1): 96–106.

Ncube, M. 2016. "Bank Bailout under TARP in the US." Working Paper No. 2016/010, Blavatnik School of Government, University of Oxford.

Ng, J., F. P. Vasvari, and R. Wittenberg-Moerman. 2011. "The Impact of TARP's Capital Purchase Program on the Stock Market Valuation of Participating Banks." Chicago Booth Research Paper No. 10.

Norden, L., P. Roosenboom, and T. Wang. 2013. "The Impact of Government Intervention in Banks on Corporate Borrowers' Stock Returns." *Journal of Financial and Quantitative Analysis* 48 (5): 1635–62.

Northehr, E. 2008. "Emergency Economic Stabilization Act of 2008." *University of Iowa for International Finance and Development.*

Palas, M. J. U., and F. Moreira. 2022. "The Impact of Government Assistance on Banks' Efficiency." *International Journal of Finance and Economics* 27 (2): 1731–48.

Panetta, F., R. Correa, M. Davies, A. Di Cesare, J.-M. Marques, F. Nadal de Simone, F. Signoretti, C. Vespro, S. Vildo, and M. Wieland. 2011. "The Impact of Sovereign Credit

Risk on Bank Funding Conditions." MPRA Paper No. 32581. https://mpra.ub.uni-muenchen.de/32581/.

Pierret, D. 2014. "Systemic Risk and the Solvency-liquidity Nexus of Banks." Memo, NYU Stern School of Business, Volatility and Risk Institute.

Schaeck, K. 2008. "Bank Liability Structure, FDIC Loss, and Time to Failure: A Quantile Regression Approach." *Journal of Financial Services Research* 33 (3): 163–79.

Schmitz, M. S. W., M. Sigmund, and M. L. Valderrama. 2017. "Bank Solvency and Funding Cost: New Data and New Results." IMF Working Paper No. 17/116.

Shah, A. 2009. "Emergency Economic Stabilization Act of 2008." *Harvard Journal on Legislation* 46 (1): 569–84.

Simon, L. J. 2004. "Detecting Multicollinearity Using Variance Inflation Factors." Penn State Department of Statistics, Pennsylvania State University.

Swedberg, R. 2013. "The Financial Crisis in the US 2008–2009: Losing and Restoring Confidence." *Socio-Economic Review* 11 (3): 501–23.

Ungan, E., S. Caner, and S. Özyıldırım. 2008. "Depositors' Assessment of Bank Riskiness in the Russian Federation." *Journal of Financial Services Research* 33 (2): 77–100.

U.S. Department of the Treasury. 2018. *Capital Purchase Program.* https://home.treasury.gov/data/troubled-assets-relief-program (accessed: April 6, 2018).

Veronesi, P., and L. Zingales. 2010. "Paulson's Gift." *Journal of Financial Economics* 97 (3): 339–68.

Wilson, L., and Y. W. Wu. 2012. "Escaping Tarp." *Journal of Financial Stability* 8 (1): 32–42.

Wooldridge, J. M. 2010. *Econometric Analysis of Cross Section and Panel Data.* Cambridge, MA: MIT Press.

Rise of the Central Bank Digital Currencies*

Raphael Auer,[a] Giulio Cornelli,[a,c] and Jon Frost[a,b]
[a]Bank for International Settlements
[b]Cambridge Centre for Alternative Finance
[c]University of Zurich

Central banks around the world are researching and developing central bank digital currencies (CBDCs). Yet the motivations for issuance, policy approaches, and technical designs differ across countries. We set out a comprehensive database of CBDC projects and technical approaches, and investigate the economic and institutional factors that correlate with CBDC project efforts. Most projects are found in economies with high mobile phone use and a high capacity for innovation. Work on retail CBDCs is more advanced where the informal economy is larger. Many central banks are considering architectures in which a CBDC is a direct cash-like claim on the central bank, but the private sector handles all retail services.

JEL Codes: E42, E44, E51, E58, G21, G28, F31.

*The views expressed here are those of the authors and do not necessarily reflect those of the Bank for International Settlements. For comments and input, we thank David Archer, Douglas Arner, Codruta Boar, Harro Boven, Haiwei Cao, Stijn Claessens, Carl-Andreas Claussen, Sebastian Doerr, Huberto Ennis, Umar Faruqui, Leonardo Gambacorta, Philippe Haene, Henry Holden, Martin Hood, Linda Jeng, Taejin Park, Mateo Piccolo, Martin Summer, Tara Rice, Amber Wadsworth, Peter Wierts, Hye-Rim Yoo, Nouran Youssef, Dirk Zetzsche, an anonymous referee, and participants at a BIS research meeting, the American Economic Association annual meetings, an Arab Monetary Fund Fintech Working Group meeting, the CEMLA/Central Bank of Colombia conference on payments, and a meeting of the OECD Experts Group on Finance and Digitalisation. We thank Mu Changchun, Lyu Yuan, Scott Hendry, Francisco Rivadeneyra, Dinesh Shah, Gabriela Guibourg, Martin Johansson, Hanna Armelius, and Stig Johansson for conversations regarding the CBDC approaches of the People's Bank of China, Bank of Canada, and Sveriges Riksbank. Author contact: Raphael Auer, Mainzer Landstraße 16, 60325 Frankfurt am Main, Germany/115 rue Réaumur 75002 Paris, France; e-mail: Raphael.Auer@bisih.org. Giulio Cornelli, Centralbahnplatz 2, 4002 Basel, Switzerland; e-mail: Giulio.Cornelli@bis.org. Jon Frost, Rubén Dario 281, Polanco, Lomas de Chapultepec, Miguel Hidalgo 11580 Ciudad de México, CDMX, Mexico; e-mail: Jon.Frost@bis.org.

1. Introduction

Digital technologies are disrupting sector after sector of the economy, and money and payments are no exception. As economic activity increasingly moves online, technological innovations are also affecting the way consumers pay and the underlying infrastructure used to serve them. Even new attempts to create money have arisen, in the form of private cryptocurrencies and stablecoins. These have so far not taken off as a means of payment, yet their emergence has opened a debate on what money should look like in the digital age, and who should issue it. In this light, central banks around the world are researching and developing central bank digital currencies (CBDCs).[1]

CBDCs can be meant either for wholesale use—i.e., only for transactions between financial institutions—or retail use, meaning they are open to the general public. Wholesale CBDCs allow for new ways to make central bank money available to regulated financial institutions. Retail CBDCs differ fundamentally from today's electronic money in the hands of households and non-financial firms, which is a liability on a financial institution. Retail CBDCs, like cash, are a direct claim on the central bank.

While the notion of providing electronic central bank money directly to the public is not new,[2] it has been gaining traction recently. Attitudes about whether central banks should issue CBDCs—in particular, retail CBDCs—have changed noticeably since 2019. Only a few years ago, most central banks had considered CBDCs but expressed concern about systemic implications that warranted caution (Barontini and Holden 2019). But over time, the need to respond to the declining use of cash in some countries came to the fore, and a number of central banks have warmed to the idea of issuing a CBDC.[3] A tipping point was the announcement of

[1]CBDCs are defined as digital payment instruments that are denominated in the national unit of account and a direct liability of the central bank (Bank for International Settlements 2020, 2021; Group of Central Banks 2021).

[2]Indeed, Tobin (1987) argued that the central bank should make a safe "deposited currency" accessible to the public.

[3]Neither electronic money nor the discussion on the central bank's role in providing it directly to the people is new. In the context of CBDCs, Broadbent (2016), Liikanen (2016), Menon (2016), Mersch (2016), Nakaso (2016), Skingsley

Figure 1. CBDCs: The Next Hype or the Future of Payments?

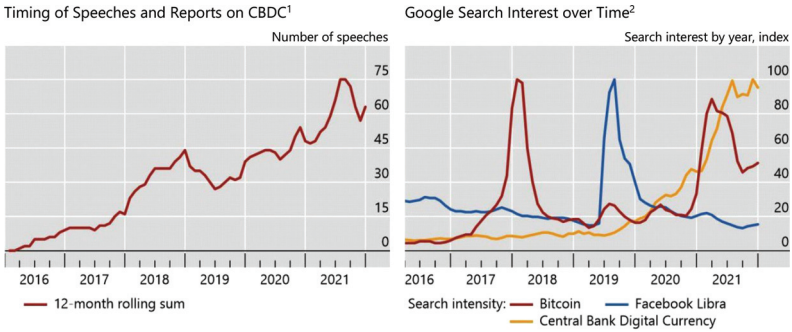

[1]Twelve-month moving sum of the count of central bankers' speeches resulting from a case-insensitive search for any of the following words/phrases: CBDC; central bank digital currency; digital currency and digital money.
[2]Three-month moving average of worldwide search interest. The data have been normalized to the three-month moving average peak of each series. The search was run on search terms "Bitcoin" and "Facebook Libra" and topic "Central Bank Digital Currency." Data accessed on January 16, 2022.
Sources: BIS Central Bankers' Speeches; central banks' websites; Google Trends; authors' calculations.

Facebook's Libra (later renamed to Diem) and the ensuing public sector response. In late 2020, central banks representing a fifth of the world's population reported that they were likely to issue CBDCs very soon (Boar and Wehrli 2021). During the COVID-19 pandemic, social distancing measures, public concerns that cash may transmit the COVID-19 virus, and new government-to-person payment schemes further sped up the shift toward digital payments (see Auer, Cornelli, and Frost 2020a). This gave a further impetus to CBDC projects in many countries. Meanwhile, the need to improve cross-border payments and securities settlement has remained a driver for wholesale CBDC work.[4]

CBDCs have seized global attention and feature broadly in central bank communications and public search interest (Figure 1).

(2016), and Wilkins (2016) were among high-level policymakers who argued early on that the idea should be taken seriously.

[4]The issuance of wholesale CBDCs is much less contentious. See Bech et al. (2020) and Pfister (2020).

As of early 2022, three retail CBDCs have been launched (in the Bahamas, the Eastern Caribbean, and Nigeria), a large-scale pilot is ongoing in China, and major reserve central banks like those of the euro area, Japan, and the United States are stepping up their research and experimentation.[5] Yet many open questions are the focus of a rapidly growing literature (Auer et al. 2022). One aspect is how central banks should create money and whether CBDCs are desirable in that context.[6] Another aspect is the systemic implications of CBDCs, e.g., whether they would disintermediate private banks, deposit taking, and lending to the real economy, and how to cope with these effects.[7] There is a budding literature on the international dimensions of CBDC issuance, including the potential for changes to monetary policy effectiveness, "digital dollarization," and international reserve currency competition.[8] Another strand of literature looks at the case for CBDCs to maintain privacy in payments.[9] Finally, the technology of retail CBDCs and how they relate to private sector proposals is hotly contested (see Brunnermeier, James, and Landau 2019; Clark and Mihailov 2019; Vives 2019; Auer and Böhme 2020, 2021; and Klein, Gross, and Sandner 2020).

To shed light on these issues, this study analyzes the cross-country economic and institutional factors that correlate with CBDC projects. A first step is to understand the status, policy approaches, and technical design of the various projects, and next to look for commonalities and differences across countries. The questions this paper aims to answer are as follows: Which economic and institutional factors are associated with central bank work on CBDCs? What are the technical solutions sought?

[5]See Bank of Japan (BOJ) (2020), Fan (2020), ECB (2020), People's Bank of China (PBC) (2021), and Federal Reserve Board (FRB) (2022).

[6]See Kim and Kang (2019), Armelius, Clausen, and Hendry (2020), Armelius et al. (2020), Jackson and Pennacchi (2021), and Keister and Sanches (2023).

[7]See Brunnermeier and Niepelt (2019), Bindseil (2020), Carletti et al. (2020), Keister and Monnet (2020), Kwon, Lee, and Park (2020), Andolfatto (2021), and Fernández-Villaverde et al. (2021).

[8]See Auer et al. (2021), Chorzempa (2021), and Ferrari, Mehl, and Stracca (2022).

[9]See, e.g., Garratt and Lee (2021), Garratt and van Oordt (2021), and Agur, Ari, and Dell'Ariccia (2022).

To assess the former question, we develop a novel CBDC project index based on central bank research and development (R&D) projects. We empirically investigate common factors in countries that are investigating and piloting CBDCs, of either the wholesale or retail variant. We find that higher mobile phone use and higher innovation capacity are positively associated with the likelihood that a country is currently researching or developing a CBDC. Retail CBDCs are more likely where there is a larger informal economy, and wholesale CBDCs are more advanced in economies that have higher financial development.

To assess technical solutions, we look at four attributes of CBDC technical designs, following the taxonomy of Auer and Böhme (2020), the CBDC pyramid. We show that many central banks are considering "hybrid" or "intermediated" architectures where the CBDC is a cash-like direct claim on the central bank, but the private sector manages customer-facing activity. Only a few central banks have considered designs in which the central bank takes on an important operational role in the customer-facing side of payments, generally as a complement to services by the private sector. None of the central bank reports favor a design with indirect claims on the central bank (referred to as an "indirect" architecture). Central banks are considering both distributed ledger technology (DLT) and conventional technological infrastructures, with many considering multiple technological options simultaneously. Nevertheless, access frameworks tend to be based more often on account identification rather than allowing for token-based fully anonymous access. More and more CBDC projects explicitly envisage international use, e.g., by non-resident visitors and for cross-border payments.

The rest of the paper is organized as follows. Section 2 describes our data and empirical analysis on CBDC projects. Section 3 discusses policy approaches and technical design. Section 4 concludes with policy implications and avenues for future research.

2. Cross-Country Factors Behind CBDC Development

Several global developments—including the digitalization of commerce, the proliferation of private digital currencies, and specific policy concerns around financial inclusion, informality, or data privacy—have recently driven increased interest in CBDCs (Auer et

al. 2020). Yet the economic and institutional motivations for issuance vary across countries. In this section, we first develop a novel CBDC database on central bank CBDC projects, alongside speeches, Internet search interest, and a range of economic and institutional variables. Next, we investigate the factors that correlate with CBDC projects. Specifically, we want to find commonalities in *why* central banks choose to embark on—or step up—CBDC efforts in some countries more than in others, using cross-section regressions. This will also help us to understand *how* they design CBDC projects.

2.1 A Novel CBDC Database

We start by generating a novel global index measuring central banks' progress toward the development of a retail or wholesale CBDC: the *CBDC project index*. This index captures publicly announced work by the central bank on CBDC projects. We construct this index based on publicly available reports by central banks. This extends the initial stocktaking exercise of Auer, Cornelli, and Frost (2020b), which focused only on retail CBDC as of early 2020.[10] The information on CBDC projects was collected through desk research and with the help of contacts at several individual central banks.[11]

Starting from these reports, we assign the CBDC project index score based on the following rules:

- Score of 0 for jurisdictions without any CBDC work publicly announced by the central bank.
- Score of 1 for research reports. These are CBDC projects that are still in early stages. The reports are just a theoretical discussion of CBDC in general or of the potential CBDC model envisaged, without any practical testing or implementation. Additionally, at this stage there is no third-party or external actor involvement.

[10]Compared with that earlier stocktake, we include wholesale CBDCs and newer projects (now through January 2022). Moreover, we complement this with data on speeches and search interest (discussed below), and we conduct empirical analysis.

[11]The list of projects is broadly consistent with other stocktakes, such as Kiff et al. (2020), Atlantic Council (2020), and Mikhalev et al. (2021). Unlike these sources, we only take account of official central bank communications, not press articles.

- Score of 2 for pilot projects. These projects flesh out the details of a pilot or proof-of-concept (PoC). This stage requires applied technological development, experimentation, and the involvement of external actors. We can confidently say that these reports indicate a more advanced stage of CBDC work.[12]
- Score of 3 for live CBDCs. This score is given when a CBDC is rolled out for large-scale use. As of our cut-off date, three central banks had live retail CBDCs, while no central bank had yet rolled out a wholesale CBDC.

Each score is sequential—i.e., to reach a score of 3 a central bank had gone through the pilot or PoC stage and the research stage. For each jurisdiction, the overall index is the maximum of the retail and wholesale sub-indices. In the following, we use data as of January 2022.[13]

Construction of the index requires some judgment. For instance, we consider only jurisdictions that have a central bank or monetary authority that could in theory develop a CBDC.[14] Currency unions require special consideration. Currency unions without national central banks are considered one observation, with all independent variables calculated as weighted averages according to 2018 GDP. For instance, the Eastern Caribbean Currency Union (ECCU), served by the Eastern Caribbean Central Bank (ECCB), is included as a single observation comprising the eight member states.[15] For the euro area, which features national central banks in addition to the ECB, we include an observation for the euro area as a whole (with a project score of 1), and each of the 19 euro-area members (with 0

[12]We include completed pilots, e.g., those by the Central Bank of Uruguay in 2017/8 and the National Bank of Ukraine in 2018. We also include Ecuador's Dinero Electrónico, which was launched in 2014 and discontinued in 2018. See Arauz, Garratt, and Ramos (2021).

[13]Regular updates to the database are made available on the authors' websites (https://www.bis.org/publ/work880.htm).

[14]Countries without a central bank are dropped. This means, for instance, that the Marshall Islands are not considered. The SOV project, which involves a private sector developer, is generally not understood as a CBDC. See International Monetary Fund (IMF) (2018).

[15]Similarly, the countries of the West African Economic and Monetary Union (WAEMU) are consolidated as one observation (project index of 0), as are the members of the Economic Community of Central African States (ECCAS).

or 1, depending on whether the national central bank has published any CBDC research).[16] Full links to public sources are available as part of the background documentation.

In the total sample of 176 countries or currency areas, 68 had a non-zero value for the CBDC project index as of January 2022. This included 66 retail CBDC projects. In the other 108 countries or currency areas without any communication on CBDC, the project index takes the value of 0.

The information on CBDC projects is complemented by a *central bank speech score*, which reflects the stance on CBDCs in speeches by governors and board members of central banks. This score is obtained by classifying the stance of each central banker speech containing at least one keyword from the following list: "CBDC", "Central Bank Digital Currency", "digital currency," or "digital money" (with a manual check to ensure it refers to CBDC and not private digital currencies). Speeches come from the Bank for International Settlements (BIS) central bankers' speeches database (www.bis.org/cbspeeches/), a comprehensive database collecting central bankers' speeches as published on the BIS website for a wide selection of central banks and international organizations.[17] Speeches are added to the database regularly, and many are translated from other languages into English by the domestic central bank. As of our cut-off date, the database counted 17,448 speeches, covered a period of more than 25 years (1997–2022), and had a wide geographical coverage (104 countries and 116 institutions). A query yielded a set of 242 speeches that contained at least one of the keywords of interest. The resulting sample covers the period December 2013–January 2022 and speeches from 45 central banks, including those of the euro area and several of its member countries.[18]

After compiling relevant speeches, we go through each and classify them by interpreting the stance of the speech towards adoption

[16]Empirical results are robust to treating the euro area as one observation.

[17]In addition to speeches by central bankers, many governments and parliaments have published reports or made statements on CBDCs. See, for instance, the U.K. House of Lords (2022) or the White House (2022). These reports and statements are not included in our database.

[18]When the speaker was an ECB official, we labeled the speech as euro area. Conversely, if the speaker was an official of a national central bank member of the Eurosystem, we labeled the speech as the corresponding country.

of CBDC or CBDC more in general. Each speech score can take a value of either −1, 0, or +1 according to the specific speech stance. A value of −1 is given if the speech stance was clearly negative or if it was explicitly said that there was no specific plan at present to issue a digital currency. The score takes a value of 0 in the case of a neutral stance. Finally, it takes a value of +1 if the speech stance was clearly positive or a project was launched or was in the pipeline. The country-level speech score is calculated as a simple average of the individual scores to date.

Finally, in order to gauge public interest in CBDCs, our database also includes an *Internet search interest* score. The score reflects both interest by residents in the idea of a CBDC, and how widely the public knows about any central bank plans to introduce a CBDC. The score is estimated as a simple average of the interest score from Google Trends for the keywords "CBDC" (search word) and "Central Bank Digital Currency" (topic) over the period January 2013–January 2022. The resulting two values for each country range between 0 (no searches) and 100 (maximum level) and are averaged to arrive at the score. For China, we used the Baidu index for keywords "Central Bank Digital Currency" and "DC/EP" (the Digital Currency/Electronic Payment program) or "e-CNY" (electronic Chinese yuan). We have rescaled the values to make them comparable to Google Trends figures (i.e., values range between 0 and 100) and applied the same procedure described above.

For each of the indicators described above, we replace country-level missing observations with zeros. This choice is consistent with the absence of a project (research or pilot), a neutral stance towards the development of a CBDC (speech score), or a lack of public interest (as captured by the search intensity score).

These indicators—made available with the paper—can help to gauge the project work on CBDCs in specific countries and to compare it with communications by central banks and public interest. The CBDC project index, and the speech and search interest scores, display substantial variance in the cross-section. Naturally, the three variables are correlated with one another, as central bank board members often use speeches to broadcast project work, and public search interest is higher where central banks have communicated that they are working on a project (see pairwise correlations in Table 1).

Table 1. Pairwise Correlations between CBDC Indicators

	CBDC Project Index[1] (Overall)	Central Bank Speech Score[2]	Search Interest Index (Google/Baidu)[3]
CBDC Project Index[1] (Overall)	1		
Central Bank Speech Score[2]	0.47***	1	
Search Interest Index (Google/Baidu)[3]	0.39***	0.26***	1

[1]The project stance is equal to 0 when there is no known work on retail or wholesale CBDC, 1 in the case of research output, and 2 in the case of an active or completed retail or wholesale CBDC pilot. [2]Search on keywords "CBDC," "digital currency," and "digital money." The classification is based on the authors' judgment. The score takes a value of −1 if the speech stance was clearly negative or in case it was explicitly said that there was no specific plan at present to issue digital currencies. It takes a value of +1 if the speech stance was clearly positive or a project/pilot was launched or was in the pipeline. Other speeches have been classified as neutral. Normalized and winsorized at the 5 percent level. [3]Data have been normalized.
Note: *** denotes significance at the 1 percent level.
Sources: Baidu; central banks' websites; BIS Central Bankers' Speeches; Google Trends; authors' calculations.

There are necessarily caveats to these measures. For instance, many central banks have not publicly released reports on their ongoing CBDC projects. Some central banks (e.g., the People's Bank of China) have quite advanced projects, but have given relatively few speeches on their plans. In some jurisdictions, Google (or Baidu) is not widely used for Internet searches. Still, the index can provide a comparable yardstick to assess changes across countries and over time. Moreover, this can provide a useful complement to the anonymized responses of central banks through official surveys. In the next section, we try to explain the cross-country heterogeneity in the CBDC project index.

2.2 Examining the Cross-Country Factors Behind CBDC Projects

In this section we investigate the factors that correlate with the CBDC project index. To complement central bank surveys and official motivations, we look at "revealed policy preferences," i.e., the economic and institutional factors that are associated with central banks' actual work on overall, retail, or wholesale CBDCs. Our cross-section estimations use an ordered probit approach (McKelvey and Zavoina 1975) and take the form of

$$Prob(CBDCPI_i = 0, 1, 2, 3|x_i) = F\left(\alpha + \beta x_i + \varepsilon_i\right),$$

where $Prob(CBDCPI_i = 0, 1, 2, 3|x_i)$ is the probability that the CBDC project index (overall, or for retail or wholesale projects) in jurisdiction i equals 0 (no project), 1 (research), 2 (pilot), or 3 (live CBDC); $F()$ is the functional form of ordered probit; X_i is one or more variables from a vector of potential factors; α and β are estimated coefficients; and ε_i is an error term.

Some potential factors behind CBDC development can be related to a country's technological capability to develop and deploy a CBDC. Focusing on indicators from reliable sources that are available for a wide cross-section of countries, we include in our analysis the following indicators:

- *Digital Infrastructure:* Jurisdictions with greater *mobile phone use* (mobile cellular subscriptions per 100 people) or *Internet use* (fixed-line broadband subscriptions per 100 people) may

have a more developed infrastructure for the central bank to develop CBDCs. Data on both come from the World Bank.

- *Innovation Capacity:* Jurisdictions with a higher *innovation score* overall, and hence the ingenuity and R&D potential to support central banks in designing a new CBDC ecosystem, may be more likely to see CBDCs. Data come from the World Intellectual Property Organization (WIPO) Global Innovation Index, which aggregates measures in the political environment, education, infrastructure, and business sophistication (WIPO 2018). To look at the innovation capacity of the central bank itself, we have a dummy for countries that have in place or plan to institute a retail fast payment system (FPS). Data for this come from Bech and Boar (2019).
- *Institutional Quality:* Jurisdictions with higher *government effectiveness* may be more likely to launch CBDC projects. Data come from the World Bank. Conversely, central banks in jurisdictions with a large *informal ("shadow") economy* may have greater interest in creating a data trail for transactions and thus promoting use of a digital currency. Estimates of the size of the informal economy come from Medina and Schneider (2019).

On the other side, countries may differ in their perceived demand for a CBDC. To proxy these factors, we include the following indicators:

- *Development and Financial Inclusion:* Countries that are more developed, as measured by *GDP per capita*, may see a higher demand for new digital payment methods. Yet all else equal, jurisdictions with lower *access to transaction accounts* may see a greater need for retail CBDCs as a financial inclusion policy. Data come from the World Bank Findex.[19]

[19]We have also looked at various measures of cash use, such as small-denomination bank notes to GDP. These are available from the CPMI Red Book statistics, but unfortunately for only 18 jurisdictions. A broader measure of cash in circulation, available for 149 jurisdictions, is positively associated with the CBDC project index in the univariate setting, implying that CBDC projects are more advanced where cash in circulation is increasing. Yet this measure may include large bills held as a store of value. Moreover, it is insignificant when

Meanwhile, jurisdictions with higher *financial development* may have greater demands on innovative solutions for wholesale settlement; data for this are available from Svirydzenka (2016). Finally, countries with greater market power in the banking sector, as measured by the *Lerner index* of banking sector markups, may see greater need for policy interventions to increase competition in payments. Data on the Lerner index come from Igan et al. (2021).

- *Cross-Border Transactions:* while most CBDCs serve a domestic purpose, one could expect that some types of CBDCs (e.g., wholesale projects for cross-border interbank settlement or migrant remittances) may be more likely in more internationally integrated economies. *Trade openness* (the sum of imports and exports over GDP) can proxy for cross-border demand for new payment options for goods and services. *Remittance flows* (inflows and outflows divided by GDP) gauge the economic importance of migrants' remittances. Again, both series come from the World Bank.

Table 2 gives descriptive statistics for our sample.

The CBDC project index has 176 observations and it takes values ranging from 0 to 3. Mobile cellular subscriptions range from 13 per 100 people (North Korea) to 321 (Macao), and GDP per capita ranges from USD 281 (Burundi) to USD 110,343 (Luxembourg). For some key variables (e.g., the innovation output score, estimates of the informal economy, account ownership, and remittances) coverage is lower but generally still well above 100 jurisdictions.

Table 3 displays our univariate regression results. We can confirm that the CBDC project index is strongly associated with higher mobile phone and Internet use, a higher innovation capacity, an existing or planned FPS, and greater government effectiveness. Somewhat against our expectations, there is a negative association with the informal economy in these univariate estimations; as we will see below using a multivariate approach, this relates to the

controlling for other factors. Overall, declining cash use is not a consistent indicator of CBDC work across our global sample—even if it may be important in individual jurisdictions (e.g., Sweden).

Table 2. Descriptive Statistics[1]

	Observations	Mean	Standard Deviation	Min.	Max.
Dependent Variables:					
Overall CBDC Project Index	176	0.58	0.82	0	3
Retail CBDC Project Index	176	0.51	0.74	0	3
Wholesale CBDC Project Index	176	0.23	0.59	0	2
Independent Variables:					
Mobile Cellular Subscriptions (per 100 People)	170	109.66	39.17	12.60	320.55
Broadband Subscriptions (Fixed Line, per 100 People)	168	14.30	13.67	0	47.16
Innovation Output Score (WIPO)[2]	118	29.67	12.69	7.90	67.13
Fast Payment System Dummy	176	0.36	0.48	0	1.00
Government Effectiveness	175	0.08	0.99	–2.25	2.22
Informal Economy (% of GDP)	122	26.08	11.62	5.43	55.78
GDP per Capita (USD)	169	16,857.15	22,027.54	391.56	114,009.80
Account Ownership (1% Age 15+)	135	60.39	27.96	6.45	99.96
Financial Development Index[3]	158	0.36	0.22	0.06	0.93
Lerner Index[4]	76	0.31	0.15	0.06	1
Remittances[5] to GDP	110	5.75	7.50	0.18	42.03
Trade Openness[6]	134	80.06	48.87	0.00	345.69
Search Interest Index (Google/Baidu)[7]	176	0.13	1.12	–0.40	9.78
Central Bankers' Speech Stance Index[8]	176	0.09	0.90	–0.23	3.08

[1]For all the independent variables, average over the period 2013–20, subject to data availability. [2]Data for 2018. [3]Svirydzenka (2016). [4]The Lernex index of banking sector markups in economy i reflects market power by incumbent banks. For more details, see Cornelli et al. (2020). [5]Sum of inflows and outflows. [6]Sum of imports and exports divided by the country GDP. Data for 2018. [7]Data have been normalized. [8]Normalized and winsorized at the 5 percent level.

Sources: Bech et al. (2020); G. Cornelli, J. Frost, L. Gambacorta, R. Rau, R. Wardrop, and T. Ziegler, "Fintech and Big Tech Credit: A New Database," BIS Working Paper No. 887, September 2020. Media and Schneider (2019); Svirydzenka (2016); WIPO (2018); IMF, *World Economic Outlook*; World Bank, *Remittance Prices Worldwide*, remittanceprices.worldbank.org; World Bank; Baidu; central banks' websites; BIS Central Bankers' Speeches; Datastream; Google Trends; authors' calculations.

Table 3. Univariate Ordered Probit Regressions on Overall CBDC Project Index[1]

Digital Infrastructure:												
Mobile Cellular Subscriptions (per 100 People)	0.008*** (0.003)											
Broadband Subscriptions (Fixed Line, per 100 People)		0.027*** (0.006)										
Innovation Capacity:												
Innovation Output Score (WIPO)[2]			0.031*** (0.009)									
Fast Payment System (FPS) Dummy				1.009*** (0.196)								
Institutional Characteristics:												
Government Effectiveness					0.465*** (0.096)							
Informal Economy (% of GDP)						–0.016 (0.011)						
Development and Financial Inclusion:												
Ln(GDP per Capita)							0.293*** (0.067)					
Account Ownership (% Age 15+)								0.017*** (0.004)				
Financial Development Index[3]									2.606*** (0.408)			
Lerner Index[4]										0.072 (0.930)		
Cross-Border Transactions:												
Remittances[5] to GDP											–0.06*** (0.022)	
Trade Openness[6]												–0.001 (0.003)
Number of Observations	170	168	118	176	175	122	169	135	158	76	110	134
Pseudo R^2	0.035	0.047	0.054	0.084	0.07	0.013	0.054	0.071	0.111	0.000	0.042	0.001

[1]For all the independent variables, average over the period 2013–20, subject to data availability. [2]Data for 2018. [3]Svirydzenka (2016). [4]The Lerner index of banking sector markups in economy *i* reflects market power by incumbent banks. For more details, see Cornelli et al. (2020). [5]Sum of inflows and outflows. [6]Sum of imports and exports divided by the country GDP. Data for 2018.

Note: Robust standard errors in parentheses; ***/**/* denotes results that are significant at the 1/5/10 percent level.

Sources: Bech et al. (2020a); G. Cornelli, J. Frost, L. Gambacorta, R. Rau, R. Wardrop, and T. Ziegler, "Fintech and Big Tech Credit: A New Database," BIS Working Paper No. 887, September 2020; WIPO (2018); Medina and Schneider (2019); Svirydzenka (2016); IMF, *World Economic Outlook*; World Bank, *Remittance Prices Worldwide*, remittanceprices.worldbank.org; World Bank; Baidu; central banks' websites; Datastream; Google Trends; authors' calculations.

correlation of this variable with mobile phone use and other positively associated covariates. Further, when it comes to those factors potentially affecting the demand for CBDC, we find CBDC projects to be more advanced where there is higher GDP per capita, higher account ownership, and greater financial development. Banking sector markups show no significant association with CBDC project work, and remittances are negatively correlated. Univariate results are also very similar for the retail and wholesale indices separately (unreported for brevity but available upon request).

Of course, these simple regression coefficients need to be interpreted with great care, as many of the regressors are collinear. More advanced economies tend to have stronger digital infrastructures, to be more innovative, and to feature more effective goverments and smaller informal economies. Moreover, isolating individual factors is complicated by the fact that sample size for some indicators is more limited, thus not allowing us to include all possible regressors at the same time.

To better control for multiple country characteristics, Table 4 displays multivariate ordered probit regression results for the overall CBDC project index, and for retail and wholesale CBDCs. The results confirm that overall projects are more likely where there is greater use of mobile phones and greater innovation capacity (column 1). We further find that, controlling for mobile phone use, there is a significant positive association with the size of the informal economy and financial development for the overall project index (column 2). We do not find a significant link with trade openess.

Retail CBDCs also appear to be more advanced in jurisdictions with high innovation capacity and where the informal economy is larger, all else equal (columns 3 and 4). There is actually a negative association with trade openness, ceteris paribus; this could reflect that retail CBDC projects are easier to implement in jurisdictions with fewer international trade ties and a more domestically focused economy.

Wholesale CBDCs are positively correlated with financial development, which could reflect the focus of such projects on increasing the efficiency of wholesale settlement (column 5).[20] This is also

[20]The innovation output score is not included, given its high correlation with financial development (81 percent).

Table 4. Multivariate Ordered Probit Regressions on CBDC Project Indices[1]

	Overall CBDC Project Index		Retail CBDC Project Index		Wholesale CBDC Project Index	
Mobile Cellular Subscriptions (per 100 People)	0.008**	0.009**		0.012**		0.006
	(0.004)	(0.005)		(0.005)		(0.007)
Innovation Output Score (WIPO)[2]	0.027***		0.050***	0.060***		
	(0.009)		(0.013)	(0.013)		
Informal Economy (% of GDP)		0.038***	0.033**	0.042***		0.019
		(0.014)	(0.015)	(0.015)		(0.018)
Financial Development Index[3]		3.604***			2.959***	3.788***
		(0.706)			(0.620)	(0.966)
Trade Openness[4]		–0.004		–0.009***	0.002	0.001
		(0.003)		(0.003)	(0.002)	(0.003)
Number of Observations	118	105	110	100	132	105
Pseudo R^2	0.07	0.163	0.068	0.131	0.184	0.223

[1]For all the independent variables, average over the period 2013–20, subject to data availability. [2]Data for 2018. [3]Svirydzenka (2016). [4]Sum of imports and exports divided by the country GDP. Data for 2018.

Note: Robust standard errors in parentheses; ***/**/* denotes results that are significant at the 1/5/10 percent level. Constants are not reported.

Sources: Medina and Schneider (2019); WIPO (2018); IMF, *World Economic Outlook*; World Bank; central banks' websites; Datastream; authors' calculations.

apparent in the more parsimonious specification (column 6).[21] The link with trade openness is not significant. While many wholesale projects focus on the cross-border dimension, it thus seems that domestic financial development is more consistently correlated with the stage of wholesale CBDC project work.

To quantify the economic significance of these results, we report in Table 5 the predicted probabilities at mean values and after a one-standard-deviation increase in each variable at a time. For example, as shown in column 1, at mean values of all variables, a country has a 49 percent probability of not doing any work on CBDCs, a 31 percent probability of conducting research, a 20 percent probability of a testing pilot, and a 0.6 percent probability of rolling out a live CBDC. A country with a one-standard-deviation increase in mobile phone subscriptions (ceteris paribus) has a 33 percent probability of research, a 26 percent probability of a pilot, and a 1.1 percent probability of a live CBDC.[22] A one-standard-deviation increase in the innovation output score is associated with probabilities of 33 percent, 30 percent, and 1.5 percent, respectively.

With respect to retail CBDCs, a one-standard-deviation increase in the size of the informal economy is associated with a 46 percent probability for research (6 percentage points higher than mean values), 19 percent for a pilot (8 percentage points higher than mean values), and 1.4 percent for a live CBDC (1 percentage point higher than mean values) (column 3). This result, obtained only when controlling for other factors, could relate to a desire by authorities to have a data trail for transactions, as discussed above. An increase by one standard deviation in financial development is linked to a 13–16 percent probability of wholesale research and 21–28 percent probability of a wholesale pilot, depending on the specification (columns 5 and 6).

[21]The Lerner index of banking sector markups is also positively associated with greater wholesale CBDC work, at the 95 percent confidence level—but its inclusion implies a much lower number of observations (unreported but available upon request). This implies that wholesale CBDC work is more advanced where the potential efficiency gains from greater competition in the banking sector are larger, other factors equal.

[22]There are necessarily caveats to this simple calculation given the non-linear nature of the ordered probit and correlation between the independent variables. For a discussion of interpretation issues in logits, probits, and other non-linear probability models, see Breen, Karlson, and Holm (2018).

Table 5. Effect of a One-Standard-Deviation Increase on Predicted Probabilities (oprobit)[1]

		Overall CBDC Project Index		Retail CBDC Project Index		Wholesale CBDC Project Index	
	Probability of	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Predicted Probabilities at Mean Values	No Activity	49.1%	45.2%	48.8%	45.8%	86.3%	85.1%
	Research	30.7%	36.8%	39.8%	44.6%	6.9%	7.8%
	Pilot	19.6%	17.8%	10.9%	9.3%	6.9%	7.1%
	Live CBDC	0.6%	0.2%	0.5%	0.3%	—	—
Mobile Cellular Subscriptions (per 100 People)	No Activity	40.1%	33.5%		31.9%		80.4%
	Research	32.7%	39.4%		50.7%		9.7%
	Pilot	26.1%	26.5%		16.5%		10%
	Live CBDC	1.1%	0.5%		0.9%		—
Innovation Output Score (WIPO)[2]	No Activity	35.6%		25.2%	19.1%		
	Research	33.1%		46.3%	51.3%		
	Pilot	29.8%		25.9%	27.2%		
	Live CBDC	1.5%		2.6%	2.4%		
Informal Economy (% of GDP)	No Activity		28.4%	33.8%	27.5%		79.3%
	Research		39.5%	45.5%	51.7%		10.1%
	Pilot		31.3%	19.2%	19.6%		10.6%
	Live CBDC		0.8%	1.4%	1.2%		—
Financial Development Index[3]	No Activity		16.5%			65.8%	55.9%
	Research		36%			13.1%	15.9%
	Pilot		45.2%			21.1%	28.2%
	Live CBDC		2.2%			—	—
Trade Openness[4]	No Activity		52.6%		63.7%	83.7%	83.5%
	Research		33.9%		32.4%	7.9%	8.5%
	Pilot		13.4%		3.8%	8.4%	8%
	Live CBDC		0.1%		0.1%	—	—

[1]For all the independent variables, average over the period 2013–20, subject to data availability. [2]Data for 2018. [3]Svirydzenka (2016). [4]Sum of imports and exports divided by the country GDP. Data for 2018.

Note: The table reports the predicted probabilities for a one-standard-deviation shock in each of the individual variables keeping all the other variables at their individual mean values.

Sources: Medina and Schneider (2019); WIPO (2018); IMF, *World Economic Outlook*; World Bank; central banks' websites; Datastream; authors' calculations.

3. Policy Approaches and Technical Design of Retail CBDCs

We have thus far established that CBDCs are more likely to be under research and development in jurisdictions with high mobile phone use, innovation capacity, and financial development, with some differences across retail and wholesale CBDCs, notably with respect to the informal economy. We have also noted that CBDC projects differ across countries, with regard to their economic and institutional motivations, policy approach, and their technical design.

In what follows, we focus on the 66 retail CBDC projects in our sample. We explore four key technological attributes of retail CBDC projects.

3.1 Attributes of Retail CBDC Projects

Approaches to CBDC design are heterogeneous across countries, requiring us to distill the main design choices and the dimensions along which national approaches differ. One way to classify design approaches is the "CBDC pyramid" (see Auer and Böhme 2020). This approach starts from the consumer needs that a retail CBDC could address, identifies associated technical design trade-offs, and then derives the design choices. The scheme of design choices forms a hierarchy in which the lower, initial layers represent design decisions that feed into subsequent, higher-level decisions. To reflect this hierarchy, the choices can be thought of as a pyramid.

The first and foundational design choice is the *architecture*, i.e., which operational role the central bank and private intermediaries take on in a CBDC. Intermediaries can run into technical difficulties or solvency issues. A CBDC should be safe from such outages. Yet payment intermediaries offer valuable services to consumers, which are needed to ensure the same level of convenience, innovation, and efficiency as in today's payments. The architecture needs to balance these two considerations.

We draw on the classification in Auer and Böhme (2021) by classifying various proposals for CBDC design into *three CBDC architectures and a fully backed alternative*. These differ in the structure of legal claims and the record kept by the central bank. They are as follows:

- *Direct CBDC:* a payment system operated by the central bank, which offers retail services. A CBDC is a direct claim on the central bank. The central bank maintains the ledger of all transactions and offers retail payment services.
- *Hybrid CBDC:* an intermediate solution that runs on two engines. Intermediaries handle retail payments, but the CBDC is a direct claim on the central bank, which also keeps a central ledger of all transactions and operates a backup technical infrastructure, allowing it to restart the payment system if intermediaries fail.
- *Intermediated CBDC:* an architecture that is a variant of the hybrid CBDC but in which the central bank maintains only a wholesale ledger, rather than a central ledger of all retail transactions. Again, the CBDC is a claim on the central bank and private intermediaries execute payments. For the purposes of this paper, this will be considered alongside the hybrid model in our stocktake.

In addition to these three generally recognized retail CBDC architectures, another approach is the indirect provision of CBDC to financial intermediaries. Thus:

- *Indirect Architecture:* a payment system operated by intermediaries that resemble narrow payment banks. Consumers have claims on these intermediaries, which operate all retail payments. These intermediaries need to fully back all liabilities to retail clients with claims on the central bank.[23]

We note that, as this does not allow the consumer to directly access central bank money, many central banks do not recognize this architecture as a retail CBDC.[24]

[23]See Adrian and Mancini-Griffoli (2019), who refer to this as "synthetic CBDC," and Kumhof and Noone (2019).

[24]For example, Bank of England (2020), Sveriges Riksbank (2020), and for the case of Canada, Shah et al. (2020) only consider architectures featuring direct claims on the central bank. The Group of Central Banks (2020), which includes central banks from the G-7 countries and others, states explicitly that "synthetic CBDC is not a CBDC."

The second technical design choice regards the *infrastructure.* A CBDC must be secure from outages at the central bank. The infrastructure can be based on a conventional centralized database or instead on DLT. These technologies differ in their efficiency and degree of protection from single points of failure. DLT often aims to replace trust in intermediaries with trust in an underlying technology. Calle and Eidan (2020) describe some of these proofs-of-concept in detail. Also noteworthy is that all central banks experimenting with DLT use permissioned variants, where operators can decide who is admitted to the network. No central bank report examined in this study has ventured to rely on permissionless DLT, as used for Bitcoin and many other private cryptocurrencies.[25]

The third choice concerns how consumers can *access* the CBDC. Account-based CBDCs are tied to identification, which can serve as the basis for well-functioning payments with sound law enforcement. Yet access is likely to be difficult for one core target group: the unbanked and individuals who rely on cash. There may be challenges to match the qualities of cash as an inclusive, crisis-proof, and anonymous means of payment (Pichler, Summer, and Weber 2019). An alternative is to base access on so-called digital tokens.[26] This allows for value-based payment options—for example, pre-paid CBDC bank notes that can be exchanged both physically and digitally. Yet this also brings new risks of illicit activity and counterfeiting.

Closely tied to the domestic access framework is the fourth design choice, the use of CBDC for *cross-border payments*, which relates to international interlinkages in a CBDC's design and its accessibility for residents versus non-residents. Token-based domestic access would naturally be open to anyone, including non-residents. But

[25]See Auer (2019) for a discussion of the inefficient economics of permissionless models, Ali and Narula (2020) for a specific analysis of permissioned and permissionless DLT in the context of CBDCs, and Auer, Monnet, and Shin (2021) for a general discussion of the economics of distributed ledgers.

[26]Importantly, this definition of token versus accounts must not be confused with the one used in the field of computer science. Rather, it follows Kahn and Roberds (2009). As put by Kahn (2016), the distinction between accounts and tokens are the identification requirements: "In a token-based system, the thing that must be identified for the payee to be satisfied with the validity of the payment is the 'thing' being transferred—'is this thing counterfeit or legitimate?' In an account-based system, however, the identification is of the customer—'Is this person who she says she is? Does she really have an account with us?' "

Figure 2. Number of Retail CBDC Projects Investigating Each Design Option

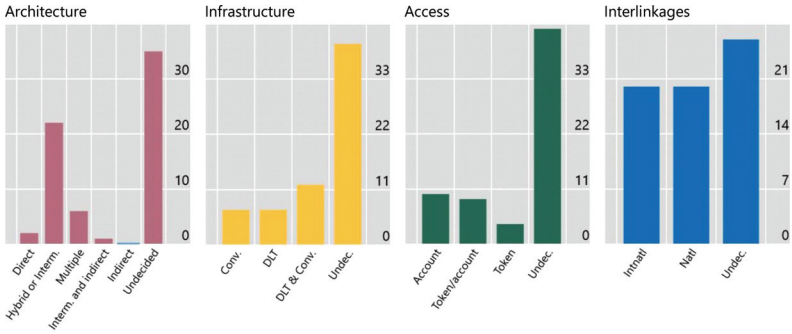

Note: Interm. = intermediated; Multiple = two or more options among direct, hybrid, and intermediated under consideration; Undecided = undecided/unspecified or multiple options under consideration; DLT = distributed ledger technology; Conv. = Conventional; Token/account = tiering of token- and account-based access; Natl = national use; Intnatl = international use.
Source: authors' calculations, as of January 2022.

central banks may allow for use by non-residents (while visiting the jurisdiction, or from abroad) in systems based on identification, as well. Some central banks may even see a role for CBDCs to support lower-cost and faster cross-border payments.

3.2 Diversity of Retail CBDC Designs

Figure 2 classifies the attributes of ongoing retail CBDC projects. Among the retail CBDC projects in our sample, we find a wide variety of approaches to architecture, infrastructure, access, and international interlinkages. On architecture, we find that only two central banks have considered a direct model, both as part of early-stage research projects. A much larger group—22 central banks—are considering the hybrid or intermediated options, with a strong role for the private sector in retail services. Six central banks are considering multiple options (i.e., hybrid or intermediated, direct); in some of these cases, the direct model is being considered as a potential "public option" that could be offered alongside services by private sector intermediaries, rather than a purely public CBDC system.

Meanwhile, one central bank (HKMA) is considering both the intermediated and the indirect model. A group of 35 central banks has not yet specified the architecture. Overall, most central banks that have made a decision are considering options in which the CBDC is a direct cash-like claim on the central bank, but where the private sector handles all customer-facing activity.

Regarding infrastructure, we find 7 central banks running their prototypes on DLT, 7 with conventional technology, and 12 considering both (Shah et al. 2020). Yet these infrastructure choices are often for first proofs-of-concept or pilots. Only time will tell if the same choices are made for large-scale designs. In practice, designs may feature elements inspired by DLT but rely on trusted intermediaries. Among access methods, account-based access appears to be the most common to date, with 10 central banks clearly leaning toward account based, 4 looking at token based, and a further 9 looking at both account- and token-based access.

Finally, while many of the retail CBDC projects in our sample were initially focused on domestic use, a growing number explicitly target use by non-residents or the potential of CBDCs to support cross-border payments. Additionally, several projects—by the ECB; the central banks of France, Spain, and the Netherlands; and the ECCB—are by construction focused on cross-border use among the members of a multi-country currency area.

4. Conclusion

This paper has examined the rise of central bank digital currencies, a new payment technology studied by central banks around the world. We have presented a novel CBDC project index. We have shown that this index is higher in jurisdictions with higher mobile phone usage and higher innovation capacity. Especially retail CBDCs are more likely where there is a larger informal economy, and wholesale CBDCs are more advanced in economies that have higher financial development. We have also noted that CBDC projects differ across countries, both in their motivations and in their economic and technical design. Many central banks are pursuing models where a CBDC is a direct claim on the central bank, but with retail payment services performed by private intermediaries.

Given the novelty of CBDC, and the scope for "clean-slate" thinking on the nature and provision of money, it is natural that the approaches will differ across countries, in line with economic circumstances and users' priorities. In countries where digital payments are already very advanced, central banks may respond in particular to ensure greater competition and privacy and the ongoing availability of a public-sector-provided means of payment. In countries with a lower penetration of digital payments, financial inclusion may be an important driver. The choice of architectures, infrastructures, access, and interlinkages will be tailored to fit local circumstances.

Yet our overview has also shown some key common features. In particular, none of the designs we survey is intended to replace cash; all are intended to complement it. Most still involve a strong role for intermediaries—although potentially in parallel to direct provision of some services by central banks. And a growing number of central banks are considering use of CBDCs for cross-border payments. We believe that by sharing information on the drivers, approaches, and technologies, central banks can learn from one another, thus complementing international policy work in this area.

Going forward, events such as the COVID-19 pandemic highlight the value of access to diverse means of digital payments, and the need for any payment method to be both inclusive and resilient against a broad range of threats, just as cash is. While it is difficult to anticipate the range of challenges ahead, central banks will continue to take a long-term view and carefully consider the role of CBDCs in a range of potential future scenarios.

References

Adrian, T., and T. Mancini-Griffoli. 2019. "The Rise of Digital Money." IMF Note No. 19/001 (July).

Agur, I., A. Ari, and G. Dell'Ariccia. 2022. "Designing Central Bank Digital Currencies." *Journal of Monetary Economics* 125 (January): 62–79.

Ali, R., and N. Narula. 2020. "Redesigning Digital Money: What Can We Learn from a Decade of Cryptocurrencies?" MIT DCI Working Paper (January).

Andolfatto, D. 2021. "Assessing the Impact of Central Bank Digital Currency on Private Banks." *Economic Journal* 131 (634): 525–40.

Arauz, A., R. Garratt, and D. F. Ramos. 2021. "Dinero Electrónico: The Rise and Fall of Ecuador's Central Bank Digital Currency." *Latin American Journal of Central Banking* 2 (2): 1–10.

Armelius, H., A. Clausen, and S. Hendry. 2020. "Is Central Bank Currency Fundamental to the Monetary System?" *Economic Review* (Sveriges Riksbank) 2 (June): 19–32.

Armelius, H., G. Guibourg, S. Johansson, and J. Schmalholz. 2020. "E-krona Design Models: Pros, Cons and Trade-offs." *Economic Review* (Sveriges Riksbank) 2 (June): 80–96.

Atlantic Council. 2020. "Central Bank Digital Currency Tracker." https://www.atlanticcouncil.org/cbdctracker/. Accessed December 1.

Auer, R. 2019. "Beyond the Doomsday Economics of 'Proof-of-Work' in Cryptocurrencies." BIS Working Paper No. 765 (January).

Auer, R., and R. Böhme. 2020. "The Technology of Retail Central Bank Digital Currency." *BIS Quarterly Review* (March): 85–100.

———. 2021. "Central Bank Digital Currency: The Quest for Minimally Invasive Technology." BIS Working Paper No. 948 (June).

Auer, R., C. Boar, G. Cornelli, J. Frost, H. Holden, and A. Wehrli. 2021. "CBDCs Beyond Borders: Results from a Survey of Central Banks." *BIS Papers* 116 (June).

Auer, R., G. Cornelli, and J. Frost. 2020a. "Covid-19, Cash, and the Future of Payments." BIS Bulletin No. 3 (April).

———. 2020b. "Taking Stock: Ongoing Retail CBDC Projects." *BIS Quarterly Review* (March): 97–98.

Auer, R., J. Frost, L. Gambacorta, C. Monnet, T. Rice, and H. S. Shin. 2022. "Central Bank Digital Currencies: Motives, Economic Implications and the Research Frontier." *Annual Review of Economics* 14: 697–721.

Auer, R., J. Frost, T. Lammer, T. Rice, and A. Wadsworth. 2020. "Inclusive Payments for the Post-Pandemic World." SUERF Policy Note (August).

Auer, R., C. Monnet, and H. S. Shin. 2021. "Permissioned Distributed Ledgers and the Governance of Money." BIS Working Paper No. 924.

Bank of England. 2020. "Central Bank Digital Currency: Opportunities, Challenges and Design." (March).

Bank of Japan. 2020. "The Bank of Japan's Approach to Central Bank Digital Currency." (October 9).

Bank for International Settlements (BIS). 2020. "Central Banks and Payments in the Digital Era." In *Annual Economic Report*, Chapter III (June).

———. 2021. "CBDCs: An Opportunity for the Monetary System." In *Annual Economic Report*, Chapter III (June).

Barontini, C., and H. Holden. 2019. "Proceeding with Caution — A Survey on Central Bank Digital Currency." *BIS Papers* 101 (January).

Bech, M., and C. Boar. 2019. "Shaping the Future of Payments." In *Analysis of the 2018 Red Book Statistics* (November).

Bech, M., J. Hancock, T. Rice, A. Wadsworth. 2020. "On the Future of Securities Settlement." *BIS Quarterly Review* (March): 67–83.

Bech, M., J. Hancock, and W. Zhang. 2020. "Fast Retail Payment Systems." *BIS Quarterly Review* (March): 28–29.

Bindseil, U. 2020. "Tiered CBDC and the Financial System." ECB Working Paper No. 2351 (January).

Boar, C., and A. Wehrli. 2021. "Ready, Steady, Go? — Results of the Third BIS Survey on Central Bank Digital Currency." *BIS Papers* 114 (January).

Breen, R., K. B. Karlson, and A. Holm. 2018. "Interpreting and Understanding Logits, Probits, and other Nonlinear Probability Models." *Annual Review of Sociology* 44: 39–54.

Broadbent, B. 2016. "Central Banks and Digital Currencies." Speech at the London School of Economics, London, March 2.

Brunnermeier, M., H. James, and J. P. Landau. 2019. "The Digitalization of Money." NBER Working Paper No. 26300.

Brunnermeier, M., and D. Niepelt. 2019. "On the Equivalence of Private and Public Money." *Journal of Monetary Economics* 106 (October): 27–41.

Calle, G., and D. Eidan. 2020. "Central Bank Digital Currency: An Innovation in Payment." R3 White Paper (April).

Carletti, E., S. Claessens, A. Fatás, and X. Vives. 2020. "The Bank Business Model in the Post-Covid-19 World." VoxEU e-Book (June 19).

Chorzempa, M. 2021. "China, the United States, and Central Bank Digital Currencies: How Important is It to be First?" *China Economic Journal* 14 (1): 102–15.

Clark, A., and A. Mihailov. 2019. "Why Private Cryptocurrencies Cannot Serve as International Reserves but Central Bank Digital Currencies Can." Discussion Paper No. 2019-09, University of Reading–Department of Economics.

European Central Bank (ECB). 2020. "Report on a Digital Euro." (October).

Fan, Y. 2020. "Some Thoughts on CBDC Operations in China." Speech, April 1.

Federal Reserve Board (FRB). 2022. *Money and Payments: The U.S. Dollar in the Age of Digital Transformation.* (January).

Fernández-Villaverde, J., D. Sanches, L. Schilling, and H. Uhlig. 2021. "Central Bank Digital Currency: Central Banking for All?" *Review of Economic Dynamics* 41 (July): 225–42.

Ferrari, M., M. Mehl, and L. Stracca. 2022. "Central Bank Digital Currency in an Open Economy." *Journal of Monetary Economics* 127 (April): 54–68.

Garratt, R., and M. Lee. 2021. "Monetizing Privacy." Staff Report No. 958, Federal Reserve Bank of New York.

Garratt, R., and M. van Oordt. 2021. "Privacy as a Public Good: A Case for Electronic Cash." *Journal of Political Economy* 129 (7): 2157–80.

Group of Central Banks. 2020. "Central Bank Digital Currencies: Foundational Principles and Core Features." Joint Report No. 1 (October).

Igan, D., M. S. Martinez Peria, N. Pierri, and A. Presbitero. 2021. "When They Go Low, We Go High? Bank Market Power and Interest Rates." IMF Working Paper No. 21/149.

International Monetary Fund (IMF). 2018. "Republic of the Marshall Islands: Selected Issues." IMF Country Report No. 18/271 (September).

Jackson, T., and G. Pennacchi. 2021. "How Should Governments Create Liquidity?" *Journal of Monetary Economics* 118 (March): 281–95.

Kahn, C. 2016. "How Are Payment Accounts Special?" Paper prepared for the Payments Innovation Symposium, Federal Reserve Bank of Chicago, October 12–13.

Kahn, C., and W. Roberds. 2009. "Why Pay? An Introduction to Payments Economics." *Journal of Financial Intermediation* 18 (1): 1–23.

Keister, T., and C. Monnet. 2020. "Central Bank Digital Currency: Stability and Information." Mimeo.

Keister, T., and D. Sanches. 2023. "Should Central Banks Issue Digital Currency?" *Review of Economic Studies* 90 (1): 404–31.

Kiff, J., J. Alwazir, S. Davidovic, A. Farias, A. Khan, T. Khiaonarong, M. Malaika, H. Monroe, N. Sugimoto, H. Tourpe, and P. Zhou. 2020. "A Survey of Research on Retail Central Bank Digital Currency." IMF Working Paper No. 20/104, June.

Kim, J., and J. Kang. 2019. "Money, to be Publicly Issued, or not to be, That is the Question." (인터넷전자상거래연구) *Journal of Internet Electronic Commerce Research* 19 (5): 77–91.

Klein, M., J. Gross, and P. Sandner. 2020. "The Digital Euro and the Role of DLT for Central Bank Digital Currencies." FSBC Working Paper.

Kumhof, M., and C. Noone. 2018. "Central Bank Digital Currencies — Design Principles and Balance Sheet Implications." Working Paper No. 725, Bank of England.

Kwon, O., S. Lee, and J. Park. 2020. "Central Bank Digital Currency, Inflation Tax, and Central Bank Independence." Mimeo.

Liikanen, E. 2016. "Cash and the Central Bank." Opening Remarks at the Bank of Finland Conference on Going Cashless, Helsinki, June 14.

McKelvey, R., and W. Zavoina. 1975. "A Statistical Model for the Analysis of Ordinal Level Dependent Variables." *Journal of Mathematical Sociology* 4 (1): 103–20.

Medina, L., and F. Schneider. 2019. "Shedding Light on the Shadow Economy: A Global Database and the Interaction with the Official One." CESifo Working Paper No. 7981.

Menon, R. 2016. "Singapore's FinTech Journey — Where We Are, What Is Next?" Speech at the Singapore FinTech Festival, Singapore, November 16.

Mersch, Y. 2016. "Distributed Ledger Technology — Panacea or Flash in the Pan?" Speech, Frankfurt am Main, April 25.

Mikhalev, I., K. Burchardi, I. Struchkov, B. H. Song, and J. Gross. 2021. "CBDC Tracker." Mimeo.

Nakaso, H. 2016. "FinTech — Its Impacts on Finance, Economies and Central Banking." Speech at the conference on "FinTech and the Future of Money," Tokyo, November 18.

National Bank of Ukraine (NBU). 2019. *Analytical Report on the E-hryvnia Pilot Project.* (February).

People's Bank of China (PBC). 2021. "Progress of Research & Development of E-CNY in China." (July).

Pfister, C. 2020. "Central Bank Digital Currency: A Primer." SUERF Policy Note No. 143 (March).

Pichler, P., M. Summer, and B. Weber. 2019. "Does Digitalization Require Central Bank Digital Currencies for the General Public?" *Monetary Policy and the Economy* (Oesterreichische Nationalbank) Q4/19: 40–56.

Shah, D., R. Arora, H. Du, S. Darbha, J. Miedema, and C. Minwalla. 2020. "Technology Approach for a CBDC." Staff Analytical Note No. 2020-6, Bank of Canada.

Skingsley, C. 2016. "Should the Riksbank Issue e-krona?" Speech at FinTech Stockholm, Stockholm, November 16.

Sveriges Riksbank. 2020. *The Riksbank's e-krona Pilot.* (February).

Svirydzenka, K. 2016. "Introducing a New Broad-Based Index of Financial Development." IMF Working Paper No. 16/15.

Tobin, J. 1987. "The Case for Preserving Regulatory Distinctions." In *Restructuring the Financial System.* Proceedings of the Economic Policy Symposium, sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyoming, August 20–22.

United Kingdom (UK) House of Lords. 2022. *Central Bank Digital Currencies: A Solution in Search of a Problem?* (January).

Vives, X. 2019. "Digital Disruption in Banking." *Annual Review of Financial Economics* 11: 243–72.

White House. 2022. *Executive Order on Ensuring Responsible Development of Digital Assets.* (March).

Wilkins, C. 2016. "Fintech and the Financial Ecosystem — Evolution or Revolution?" Remarks at Payments Canada, Calgary, Alberta, June 17.

World Intellectual Property Organization (WIPO). 2018. *Global Innovation Index 2018: Energizing the World with Innovation.* (July).

Why Are Inflation Forecasts Sticky? Theory and Application to France and Germany*

Frédérique Bec,[a] Raouf Boucekkine,[b] and Caroline Jardet[c]
[a]Thema, CY Cergy Paris University, and CREST, Paris
[b]Centre for Unframed Thinking (CUT),
Rennes School of Business, France
[c]Banque de France, DGSEI-DECI, Paris

This paper proposes to adapt the model of pricing decisions developed by Alvarez, Lippi, and Paciello (2011) to the decision process of forecasters. The model features both a fixed cost of announcing a revised forecast and a fixed cost of updating the information set and adapting the forecast accordingly. Basically, the former fixed communication costs determine state dependence, which implies that the forecaster changes its forecast only when it is far enough from the optimal forecast, i.e., beyond a fixed threshold; the latter fixed information costs determine time dependence, which implies that the forecaster updates its information set only every other T periods, where T is optimally chosen. We show that survey data of inflation forecast updates as well as the last known monthly inflation rates can be used to estimate the threshold implied by the theoretical model. This threshold estimate is then crucial to uncover the existence of both types of costs as well as an upper bound of

*We are grateful to two anonymous referees and the Editor (Òscar Jordà) for their helpful comments and suggestions. Previous versions of this work have also benefited from many fruitful remarks by Anurag Banerjee, Mélika Ben Salem, Guillaume Chevillon, Annabelle De Gaye, Neil Ericsson, Fabian Gouret, David Hendry, Hervé Le Bihan, Francesco Lippi, Sarah Mouabbi, Muriel Roger, Pierre Sicsic, and Mark Watson, as well as by participants at the CIRET 2016 conference in Copenhagen, the 8th French Econometrics Conference, the 25th Symposium of SNDE in Cergy-Pontoise, the Durham University Economics Department's seminar, the 19th OxMetrics Conference in Paris, and the Banque De France seminar. F. Bec acknowledges support by Labex MME-DII (ANR-11-LBX-0023-01). The authors thank Isabelle Capelle and Céline Garnier for helpful research assistance. Of course, all remaining errors are ours. This paper reflects our opinion and does not necessarily express the views of the Banque de France. Author e-mails: bec@ensae.fr; raouf.boucekkine@rennes-sb.com; caroline.jardet@banque-france.fr.

the optimal time between two information observations. French and German data suggest that the maximum optimal time to next observation is six months, while the observation cost is at most twice as large as the communication cost.

JEL Codes: C23, D8, E31.

1. Introduction

Recent empirical evidence from forecast surveys panel data reveals forecasts stickiness: Coibion (2010), Coibion and Gorodnichenko (2012, 2015), Andrade and Le Bihan (2013), Dovern (2013), Loungani, Stekler, and Tamirisa (2013), and Dovern et al. (2015), among others, all point to the evidence that forecasters fail to systematically update their forecasts and/or their information set. This is also the case in our monthly forecasts survey data set, where only a proportion of around 43 percent of professional forecasters update their forecast between two consecutive months. Similar orders of magnitude are found from monthly survey data by Dovern (2013) and Dovern et al. (2015) for forecasts of real GDP growth. Using quarterly survey data from the European Survey of Professional Forecasters, Andrade and Le Bihan (2013) find that only 75 percent of these forecasters update their one-year or two-year forecasts each quarter. This figure is also very close to the proportion of forecasters who update every three months in our sample. As stressed by these authors, even if the proportion of updaters is much larger from micro-level survey data than from aggregate forecasts data, it still remains remarkable that professionals do not systematically update their forecasts.

Models with information rigidities have been proposed to solve this forecast stickiness puzzle. Basically, they can be divided into two strands of research: one which considers information stickiness and another one which considers information noisiness. In the first branch, made popular by Mankiw and Reis (2002), agents face fixed costs to acquiring and processing information and therefore update their information sets infrequently. In the second branch, initiated by Woodford (2002) and Sims (2003), agents are continuously updating their information set, but they receive noisy information about

underlying macroeconomic conditions. However, both types of information rigidity models fail to reproduce the degree of stickiness observed in forecasts data (see, e.g., Andrade and Le Bihan 2013). Furthermore, they imply a decision to update the forecast which is mainly time dependent, in contradiction with some empirical evidence of state dependence found in the data by, e.g., Dovern (2013), Loungani, Stekler, and Tamirisa (2013), or Coibion and Gorodnichenko (2015).[1] This state dependence is confirmed by the results of the special questionnaire for participants in the European Central Bank (ECB) Survey of Professional Forecasters (SPF) conducted and published every five years by ECB. This special survey aims at exploring the forecast processes and methodologies underlying the contributions made to the regular quarterly SPF. In its third edition,[2] published in 2019, an average of 73 percent of respondents mentioned "internal timetable" as the determinant of the timing of their forecast updates. Many forecasters, however, reported that data releases are also an important trigger of a full update of their (short-term) forecasts. Some of them also provided qualitative comments explaining that if new data materially affected their view on the economy, they would react by updating their forecasts sooner than their regular timetable might imply.

Our first contribution to this literature is to introduce and evaluate what we consider as the missing ingredient in these models: a communication cost. Indeed, a forecaster can pay the observation cost (and refresh his/her forecast) but choose not to announce or communicate officially this new forecast due to the cost it involves in terms of meetings, reports, interviews, discussions with the hierarchy as well as with the customers, etc. We claim that taking into account this additional cost, as small as it might be, reinforces the forecast stickiness and generates state dependence in the decision rules.

We first elaborate on the model developed by Alvarez, Lippi, and Paciello (2011) to design a theoretical framework allowing for

[1] Evidence of both state and time dependence is also found by Magnani, Gorry, and Oprea (2016) from laboratory decision experiment data.

[2] See the document "Results of the Third Special Questionnaire for Participants in the ECB Survey of Professional Forecasters," downloadable from the ECB website: https://www.ecb.europa.eu/stats/ecb_surveys/survey_of_professional_forecasters.

optimal inflation forecasting in the presence of both observation and communication costs. Even though our framework is similar to Alvarez, Lippi, and Paciello (2011), there are two main differences. The first one comes from the fact that the timing of the next observation is bounded by the forecast horizon. There is no upper bound of this sort in Alvarez, Lippi, and Paciello (2011), because their problem is about price setters who do not bear any time horizon constraint. However, this is a pure control constraint, the simplest one from the mathematical point of view, and we show in our empirical section that it is never binding in practice.[3] The second one has to do with the underlying stochastic structure. In Alvarez, Lippi, and Paciello (2011), the latter concerns nominal price law of motion, while in our setting, we focus on the inflation forecast law of motion. Therefore, the underlying stochastic structure has to be adapted.

Our original setup can generate optimal forecast stickiness under rational expectations hypothesis. In this model, to make a long story short, fixed communication costs generate state dependence, which implies that the forecaster announces his/her forecast change only when it is large enough, i.e., larger than a given threshold; fixed information costs cause time dependence, which implies that the forecaster updates its information set only every other T periods, where T is optimally chosen by the forecaster. Finally, we also show how the threshold can be used to evaluate the relative importance of observation and communication costs.

A second contribution of this paper is to provide an estimate of the latter. To this end we use survey data of inflation forecast updates as well as last-known monthly inflation rates to estimate the threshold implied by the theoretical model. As will be shown below, this threshold estimate is in turn crucial to uncover the existence of both types of costs as well as an upper bound of the optimal time between two information observations. The main difficulty is that we only observe forecast changes which are announced—in other words, the communication activity. Indeed, no direct data on the forecasters' observation activity is available, which rules out the direct evaluation of its cost. To circumvent this issue, a logit model is fitted to

[3]Without loss of generality, we shall omit it hereafter to unburden the presentation.

the forecast's revision probability as a function of the time elapsed since the last revision as well as a proxy for the forecast gap to capture the state-dependent dimension. In this model, the threshold is endogenously determined as the one maximizing the likelihood function. Our French and German data suggest that the optimal time to next observation is about six months, while the observation cost is at most twice as large as the update communication cost.

The theoretical model is sketched in Section 2, along with its main implications and our empirical testing strategy. Section 3 presents the data. Section 4 describes the threshold estimation method and reports the results of the estimated threshold logit models. Their implications regarding the relevance of our theoretical model are then discussed, while Section 5 concludes.

2. A Theoretical Model of Forecast Formation

This section proposes a non-technical description of our model, while more details can be found in the appendix. After a brief presentation of the main assumptions, we proceed with the description of the decision maker's behavior. We then turn to the model's properties and finally present our testing strategy.

2.1 Preamble

The following notation will be used throughout the paper. Time is continuous and $t \in (0, 1)$. The object to forecast is the annual inflation rate, $\pi(1)$. $\pi_f(t)$ is its forecast value at time t, with $0 < t \leq 1$, of the announced forecast of $\pi(1)$. $\pi_f^*(t)$ denotes the forecast target, namely the optimal forecast which would prevail if the information set was up to date in a frictionless setup. The "forecast gap," denoted $\tilde{\pi}_f(t)$, is defined as the difference $\pi_f(t) - \pi_f^*(t)$.

Assumption 1. *The law of motion of the target is subject to i.i.d. innovations, i.e., a Brownian motion (BM) without drift, so in continuous time that is $d\pi^*(t) = \sigma dW(t)$ where $W(t)$ is a standard BM and σ^2 the variance per unit of time.*

This assumption is in line with, e.g., Stock and Watson (2007)'s finding that the univariate inflation process is well characterized by a

unit root. Note also that we depart from Alvarez, Lippi, and Paciello (2011), who assume that the law of motion of the targeted price has a time drift ($\mu \geq 0$).[4] While non-zero drifts make sense in their framework (as they are concerned with price setting and the time drift corresponds indeed to inflation), they do not in ours.

Assumption 2. *The instantaneous loss faced by the professional forecaster is a quadratic function of the forecast gap* $\tilde{\pi}_f(t)^2$.

Thus, the more the forecast deviates from the optimal forecast, the greater the loss incurred by the forecaster or, more generally, for the institution from which the forecast originates. This loss can be interpreted in several ways. Firstly, it may represent a loss in terms of reputation that the institution producing the forecast incurs when providing sub-optimal forecast. Of course, judging whether a forecast is optimal or not is not obvious as it is not observable. While agents do not expect the forecast to be exactly equal to the true ex post realized value, they however may expect that it should be close to it. So agents should be able to form an opinion on the quality/optimality of an institution's forecasts by looking at its ex post forecast errors. Therefore, an institution that would provide forecasts very different to ex post realization, notably in comparison with other institutions, could suffer from a loss in credibility and reputation. Secondly, one may think that these forecasts feed partly economic and financial decisions made by these institutions. Therefore, a sub-optimal forecast could lead to sub-optimal decision-making and induce an additional loss (financial and/or reputational) for the institution.

Assumption 3. *Each time the forecaster observes the state* $\pi_f^*(t)$, *it incurs a fixed observation cost* $\theta \in R^+$. *This cost includes, for instance, the time it takes in terms of hours worked to collect and process all relevant information, and the time it takes to implement forecasting models.*

Assumption 4. *It is costly to communicate about the revised forecast, to announce it "officially." Each revision announcement incurs*

[4]See Equation (1) in Alvarez, Lippi, and Paciello (2011, p. 1917).

a fixed communication cost $\psi \in R^+$. This cost includes, e.g., the extra writing and publication of reports involved by the external communication process with the public, customers, and media. Indeed it is necessary to justify thoroughly the revision of the forecast, even more so if the forecaster is the only one who publishes a significant forecast revision. Actually, this revision is likely to attract a lot of attention. In addition, we assume that this cost also includes the loss of forecaster's credibility implied by too frequent forecast revisions.

2.2 Sketch of the Model

The problem considered here is the one of a decisionmaker (a forecasting unit) subject to the observation and communication costs, as defined in Assumptions 3 and 4 above. This is an adaptation of Alvarez, Lippi, and Paciello (2011)'s model of firms price setting to the forecast formation matter.[5] More precisely, the forecaster tracks a target, the optimal forecast $\pi_f^*(t)$, which by Assumption 1 is a BM. By Assumption 2, the forecaster's losses are quadratic in the forecast gap $\tilde{\pi}_f(t)$. Figure 1 illustrates the forecaster's problem, where $\bar{\pi} \in \mathbb{R}^+$ is a threshold value and $T(\cdot)$ is the optimal time until next observation as a function of the forecast gap $\tilde{\pi}_f(t)$—to be described later in Figure 2.

In this setting, an observation is defined as the forecasters' act to pay a fixed cost θ and retrieve the observation on the current $\pi_f^*(t)$ and hence on $\tilde{\pi}_f(t)$ too. This is illustrated in the middle of Figure 1. The forecaster's decision then depends on the size of the forecast gap, measured by $|\tilde{\pi}_f(t)|$, compared with the threshold $\bar{\pi}$. If the size of the gap is greater than $\bar{\pi}$, they choose to pay another fixed cost ψ in order to communicate on a forecast revision, namely to adjust $\pi_f(t)$, as indicated in the left part of Figure 1. Since the target has no drift (i.e., $\pi_f^*(t)$ is a unit root until the next forecast revision announcement), it is straightforward that the optimal target revision occurs at the same time as the actual observation of the target. In other words, the forecast gap is set back to zero as noted in the top box of Figure 1.

[5]It turns out to be a special case with no trend in the forecast target and, by contrast with Alvarez, Lippi, and Paciello (2011)'s model, our problem has a finite horizon.

Figure 1. Forecaster's Problem

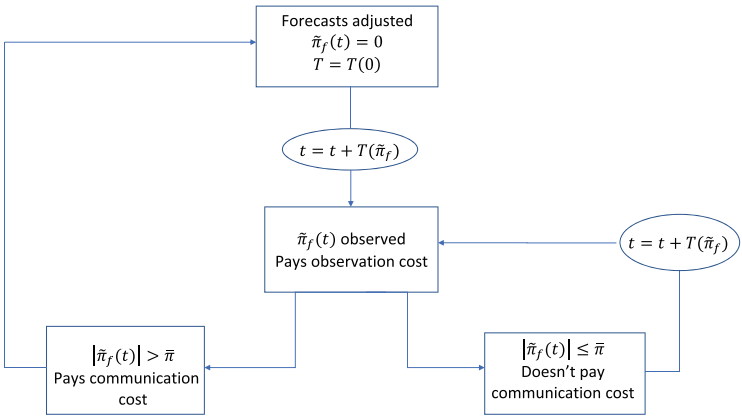

Otherwise, if the size of the gap is smaller than $\bar{\pi}$, they do not communicate, as it is not worth paying the fixed cost ψ. Instead, they wait until they choose to observe again. Therefore, the presence of a communication cost implies that not all observations are followed by a forecast revision announcement. Technically, conditionally on observing the gap, there is an Ss interval for gaps where inaction is optimal, here defined by $(-\bar{\pi}, \bar{\pi})$. This is illustrated in the bottom right box of Figure 1. In this case, the decisionmaker has to plan optimally the time to next observation by choosing $T(\tilde{\pi}_f)$.

As illustrated in Figure 2, which represents this optimal time until next observation as a function of the forecast gap, when the forecaster pays the communication cost and closes the gap (so that $\tilde{\pi}_f(t) = 0$), the time to next observation reaches its maximum, $T(0)$. Then, as innovations accumulate in $\pi_f^*(t)$, the forecast gap approaches one of the boundaries of the inaction set, either $\bar{\pi}$ or $-\bar{\pi}$, on the x-axis of Figure 2.[6] As the forecast gap widens, a new observation by the forecaster becomes more and more likely. As a consequence, the relationship between the optimal time to next observation and the forecast gap within the inaction set is hump-shaped. As soon as one of these boundaries is hit, the communication

[6]The symmetry around zero stems from the quadratic loss function of the forecaster, Assumption 2.

Figure 2. Time Elapsed Until Next Observation

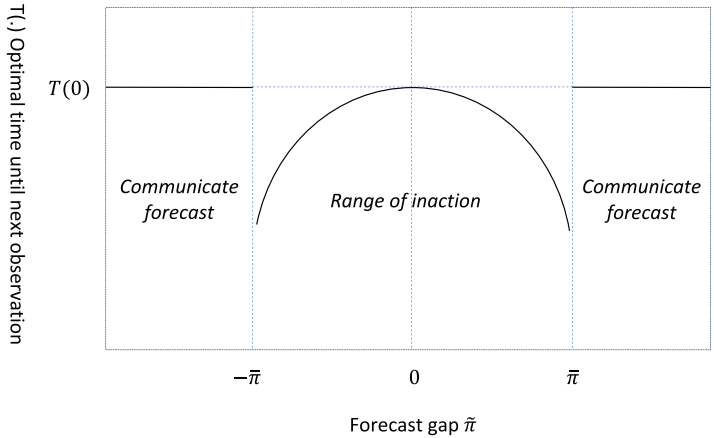

cost is paid and the gap is closed, which drives the optimal time to next observation back to its maximum $T(0)$. So the optimal decision involves both a time-dependent component (the time elapsed since the last observation) and a state-dependent decision (the forecast gap at the last observation).

2.3 The Model's Main Properties

As exposed to a larger extent in Appendix A, the optimal control problem of the forecaster is to choose the time until next observation, $0 < T \leq 1$, as well as the number and size of forecast revisions between two observations. This model's main properties are summarized below. More properties are derived in the appendix.

PROPERTY 1. *Optimal number and timing of forecast updates announcements:*

(i) There is at most one forecast revision communication between two observations, and it occurs if and only if $|\tilde{\pi}_f| > \bar{\pi}$, i.e., when the forecast gap if large enough to compensate for the communication cost (bottom left box of Figure 1).

(ii) If there is a forecast revision communication, it is immediate after observation and it closes the gap: $\pi_f = \pi_f^*$, so that $\tilde{\pi}_f = 0$ (top box of Figure 1).

See Appendix B for more details.

From Property 1, it follows straightforwardly that a forecast revision communication will take place if and only if the observed forecast gap is greater than the threshold $\overline{\pi}$, in absolute value. Let $\lambda(\tilde{\pi}_f)$ denote the indicator variable, which takes on value one if a forecast revision communication occurs and zero otherwise. Accordingly, it is defined as a function of the threshold $\bar{\pi}$ in Equation (1) below.

Property 2. *Optimal forecast revision communication occurrences:*

$$\lambda(\tilde{\pi}_f) = \begin{cases} 0 & \text{if } \tilde{\pi}_f \in (-\overline{\pi}, \overline{\pi}) \\ 1 & \text{otherwise.} \end{cases} \tag{1}$$

See Appendix A for more details. So, $\lambda(\tilde{\pi}_f)$ is state dependent. Yet, since $\tilde{\pi}_f$ is a function of $T(0)$, the maximum time elapsed until next observation as represented in Figure 2—see Proposition D.1 in Appendix D—$\lambda(\tilde{\pi}_f)$ is time dependent too. Indeed, the forecast gap is more likely to be bridged as the time elapsed until next observation approaches its maximum because, as long as it is not corrected, it behaves like a Brownian motion and, as such, it accumulates innovations—see Equation (A.2) in Appendix A. This in turn makes the forecast gap more likely to cross one of the inaction zone boundaries.

Then, the optimal time between two observations, $T(\tilde{\pi}_f)$, can be shown to be also a threshold regime-switching process with the forecast gap as the switching variable. The optimal rule for the time of the next observation of the forecast gap is given below.

Property 3. *Optimal time between two observations:*

$$T(\tilde{\pi}_f) = \begin{cases} T(0) - \left(\frac{\tilde{\pi}_f}{\sigma}\right)^2 + o(|\tilde{\pi}_f^3|) & \text{if } \tilde{\pi}_f \in (-\overline{\pi}, \overline{\pi}) \\ T(0), & \text{otherwise.} \end{cases} \tag{2}$$

See Appendix C for more details.

Moreover, as can be seen above, within the inaction set where $\tilde{\pi}_f \in (-\overline{\pi}, \overline{\pi})$, the closer to the threshold $|\overline{\pi}|$, the shorter the time to next observation: This is the theoretical ground of Figure 2.

Finally, as noticed by Alvarez, Lippi, and Paciello (2011), there is a rather simple approximate relation between the ratio of the observation to communication frequencies and the ratio of ψ to θ.[7]

PROPERTY 4. *The relative costs approximate evaluation:*

$$\mu = \frac{\psi}{\theta} \simeq \left(\frac{n_o}{n_c} - 1\right)^2, \tag{3}$$

where n_o and n_c denote the observation and communication frequencies, respectively.

Equation (3) is instrumental to ground the relevance of our approach, as the coexistence of both observation and communication costs implies that the costs ratio is strictly positive and finite. Indeed, if there is no communication cost, then this ratio is zero, while if there is no observation cost, it goes to infinity. The relation between the costs ratio μ and the ratio of frequencies n_o/n_c is qualitatively intuitive. When the ratio μ is zero, which corresponds to the zero communication cost case, the forecaster will communicate at each observation, leading to $n_o/n_c = 1$. As μ rises, the ratio n_o/n_c also goes up: When the communication cost increases more than the observation's, the forecaster will observe more frequently than she will communicate, implying a rising frequencies ratio n_o/n_c. On a more quantitative ground, the functional relation between μ and the ratio of frequencies is approximately quadratic within the analytical framework of Alvarez, Lippi, and Paciello (2011). In other words, $n_o/n_c - 1$ is well approximated by the square root of μ (so the elasticity with respect to the ratio of costs is near 1/2). As shown in Alvarez, Lippi, and Paciello (2011), the lower the communication cost, the better the approximation.

As stressed above, the costs ratio is zero if the absence of communication cost, and it tends to infinity if there is no observation cost.

[7]See the discussion in Section VI.A, pp. 1936–38, in Alvarez, Lippi, and Paciello (2011).

In practice, the latter case should not be observed, as it requires that either $n_c = 0$ or $n_o \to \infty$. Nevertheless, an arbitrarily large value of the ratio could be found for arbitrarily low (respectively, large) value of n_c (respectively, n_o) and would question the presence of an observation cost. By contrast, a ratio close to zero cannot be ruled out a priori from an empirical point of view, as n_c could be found equal to n_0, hence questioning the presence of a communication cost. Therefore, we believe that the empirical evaluation of this costs ratio is important, as it is likely to provide information about the relevance of our theoretical setting.

We will show here below how these properties can be used to estimate the inflation gap threshold which triggers action, $\bar{\pi}$, along with the maximum optimal time to next observation, $T(0)$. In turn, as explained in the next subsection, this allows to infer the observation to communication costs ratio.

2.4 Our Estimation Strategy

Our estimation strategy runs in two main steps.

Firstly, we fit a threshold logit regression based on Equation (1), in order to estimate the forecast update communication probability as a function of the forecast gap and time dependence. However, as the forecast gap is not observable, we need to use an "observable" proxy instead in order to achieve estimation of the model. The choice of this proxy will be discussed in Section 3. Finally, this first step allow us to obtain the maximum-likelihood estimate of the threshold, $\bar{\pi}$.

Secondly, with the threshold estimate from the first step, it is then possible to check the existence of both types of costs grounding time and state dependence of the probability to announce a forecast revision. The computation of the communication frequency, n_c in Equation (3), is straightforward: It is simply the share of forecasters who update in the sample. The computation of the observation frequency, n_o, is more tricky. This is the share of observers, which includes the ones who communicate, i.e., n_c,[8] and the ones who observe and do not communicate a revised forecast.

[8]One can safely assume that when a forecast update is announced, it has first been preceded by an observation.

It is necessary to circumvent the main caveat of our data: The observation activity of the forecasters is not observable.[9] Hence, it is impossible to disentangle among the ones who do not communicate the ones who have observed and the ones who haven't. However we propose to compute n_o by considering the share of the ones who do not update, conditionally to $\tilde{\pi}_f < \overline{\pi}$, i.e., the forecast gap is lesser than the threshold. This is where an estimate of $\overline{\pi}$ is clearly required. In fact, this measure of n_0 should overestimate the true value of the observation frequency because in addition to those who observe but do not adjust for good reasons (the ones we want to count), this population also includes those who just do not observe. Unfortunately, these two categories are observationally equivalent with the data at hand: They both correspond to $\lambda(\tilde{\pi}_f) = 0$ *and* $\tilde{\pi}_f < \overline{\pi}$. Consequently, our computation of n_o will overestimate its true value.

3. Data

3.1 Forecasts Updates Data

The empirical analysis below is based on forecasts of the annual inflation rate from the monthly survey data set compiled by Consensus Economics Inc. and made by private and public economic institutions such as banks and research institutes. Since Consensus Economics Inc. reports forecasts of the current and next calendar years, the data set is a three-dimensional panel: forecasters indexed by i, target years indexed by t, and horizons indexed by h. For each target year, the data set contains a sequence of 24 forecasts from each forecaster made from January of the year before the target year to December of the target year. The latter will be labeled $h = 0$ since this corresponds to nowcasting. Consequently, the former will correspond to $h = 23$.

Let $\pi_{i,t|h}$ denote agent i forecast of year t inflation rate, made h months before December of year t. The unconditional probability of updating a forecast for horizon h, denoted $\lambda(i, h)$ after Dovern (2013), is given by the probability that a forecast for the same

[9] Indeed, from forecasts data, the only activity which is tracked is the forecast revision communication.

Table 1. Descriptive Statistics of $\hat{\lambda}(h)$

	Average	**Min. [h]**	**Max. [h]**	**$h = 22$**	**$h = 0$**
France	45.0%	27.4% [16]	63.3% [9]	28.8%	32.4%
Germany	41.7%	29.9% [19]	61.5% [12]	30.0%	34.2%

object, e.g., year t inflation rate, is revised by institution i between two consecutive forecast horizons: $Pr(\pi_{i,t|h} \neq \pi_{i,t|h+1})$. As noted by Dovern (2013), assuming that the probability is the same for each institution, it can be estimated for each horizon as

$$\hat{\lambda}(h) = \frac{1}{T} \sum_{t=1}^{T} \frac{1}{N_{t|h}} \sum_{i=1}^{N_{t|h}} 1[\pi_{i,t|h} \neq \pi_{i,t|h+1}], \tag{4}$$

where T is the total number of target years in the sample, $N_{t|h}$ is the number of observed forecasts for target year t with forecast horizon h, and $1[.]$ is the indicator function which takes on value one if the condition into brackets is verified and zero otherwise.

Our data set includes yearly inflation rate forecasts for the current and next years, made monthly by individual professional forecasters since January 1998 for target years from 2000 to 2015. In the subsequent analysis, we consider country-specific forecast data. This guarantees indeed that the series to forecast, namely the national inflation rate, is homogenous across forecasting units. Moreover, this allows to reveal similarities and differences in forecasting behaviors across countries, if any. For France and Germany, the sample includes, respectively, 36 and 51 forecasters. Table 1 reports the main descriptive statistics of the average probabilities of updating between two consecutive months, the expression for which is given in Equation (4).[10] Figure 3 plots these update probabilities as a function of the forecast horizon.

These basic statistics look quite homogenous amongst countries. On average over all horizons, 41.7 percent (Germany) to 45 percent

[10]The panel is heavily unbalanced since a large part of the "individuals" have given their forecast in an irregular manner. Consequently, a significant part of forecasts observations could not be used to compute revisions because they are adjacent to missing values. As a result, we are left with a total of 4,990 usable points for France and 8,202 for Germany.

Figure 3. Average Revision Probability as a Function of the Forecast Horizon

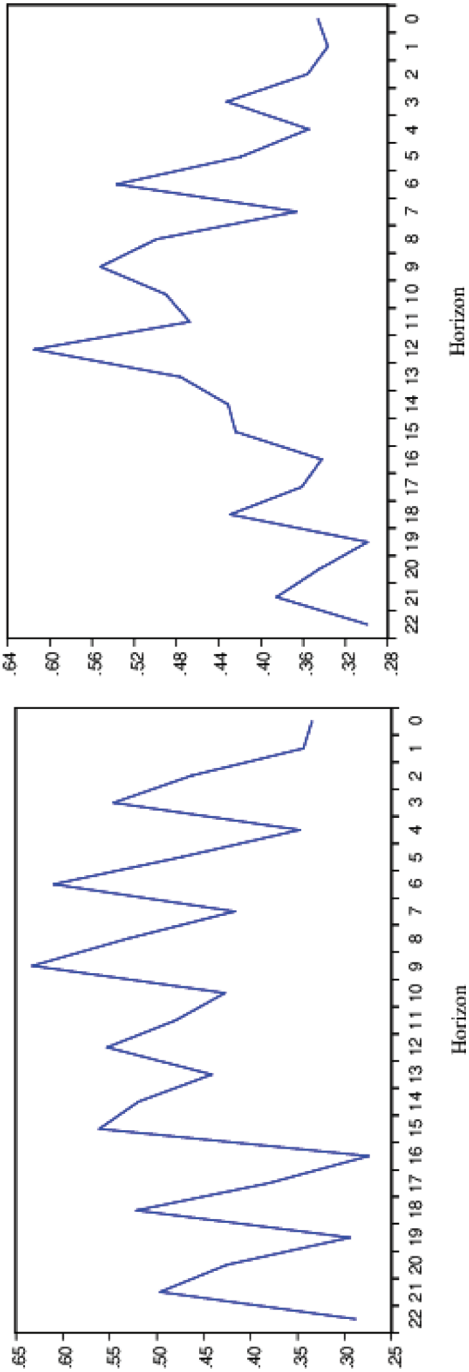

(France) of forecasters revise their forecasts between two consecutive months. This result is slightly less than Dovern's (2013) average estimation of nearly 50 percent over 14 advanced economies: The European professional forecasters considered in our study are not the most attentive ones among advanced countries. Nevertheless, this implies a degree of information rigidity which is much lower than the one obtained from the aggregate (or average/median) forecast. For instance, Coibion and Gorodnichenko (2012) find that the average inflation forecast across 40 U.S. agents surveyed by SPF is updated every six to seven quarters.[11] Similarly, high degrees of information rigidity is found by Coibion (2010). Our data show that 43.3 percent of the professional forecasters update their forecasts as frequently as every month.[12]

As can be seen from Figure 3, there is no obvious linear trend in the revision probability as horizon increases or decreases, which differs from Dovern's [2013] finding based on GDP growth forecast survey data. It can be seen that attentiveness increases until $h = 12$ (Germany) and $h = 9$ (France) and decreases afterwards. The pattern of French update rate shows pronounced seasonality with clear quarterly peaks in March, June, September, and December. This could stem from two reasons. First, it is likely that the forecasting exercise is done every quarter instead of every month. Second, over the period considered in this analysis, the French National Institute of Statistics and Economics Studies releases its first estimate of quarterly real GDP growth 45 days after the end of the quarter, which is around the middle of the second month of each quarter. Since the monthly Consensus Economics survey is completed before the twelfth day of each month, this piece of information is not available when producing the forecast of the quarter's second month. It is incorporated in the quarter's third month instead, enhancing the revision probability. This quarterly pattern is also present in

[11]By contrast, in our French sample for instance, the update share reaches 78 percent when the probability of updating at least once over the last three months is considered.

[12]The relative frequencies of individual-specific unconditional probability of updating can be obtained by computing the average of $1[\pi_{i,t|h} \neq \pi_{i,t|h+1}]$ across target years and forecast horizons for each forecaster. To save space, descriptive statistics of these $\hat{\lambda}(i)$'s are not reported here but are available upon request. Note that their mean is again found to be close to 45 percent.

Germany, but to a lesser extent and mainly during the target year itself, i.e., for $h = 12$, 9, 6, and 3. These seasonal patterns will be captured by the inclusion of a horizon fixed effect in our regression equation.

3.2 Threshold and Explanatory Variables

According to the model presented in Section 2, some variables must be allowed to switch from an inaction regime—defined by $\tilde{\pi}_f \in (-\bar{\pi}, \bar{\pi})$—to an action regime elsewhere. Let us gather in a vector denoted $X^S_{i,t}$ the variables allowed to switch. As discussed in Section 2—see Equation (2) above and Equation (D.9) in Appendix D—the dynamics of the time elapsed between two observations ($T(\tilde{\pi}_f)$) and of the forecast gap itself ($\tilde{\pi}_f$) are expected to switch according to the size of the forecast gap. In our empirical counterpart, the time dependence ($T(\tilde{\pi}_f)$) will be captured by a dummy variable denoted D(d=m), for $m = 2, 3 \ldots$, which is equal to one if the unit's last revision occurred m months ago and zero otherwise. Since very few observations belong to each D(d=m) for $m \geq 8$, the latter are gathered into D(d≥8).[13]

Then, let us turn to the threshold variable, namely the forecast gap. Unfortunately, this is not observable, since the optimal forecast, from which it is evaluated, is unobservable. As mentioned earlier, to circumvent that issue, we use an "observable" proxy instead. This proxy should be available at the time of the forecast exercise. Hence, ex post realized annual inflation cannot be used for that purpose, as it is known only at the end of the target year. By contrast, the last release of monthly inflation rate is available, and even if the optimal forecast should depend on a lot of factors, we can reasonably assume that updates of optimal forecast should at least incorporate this new information. Under the strong assumption that forecasts of monthly inflation rate are constant, large values of realized monthly inflation rate should trigger large values of the forecast gap. Therefore, in the following, we consider the last known monthly inflation rate, seasonally adjusted and weighted by its mechanical contribution to the annual inflation rate forecast, as a proxy for the unobservable forecast gap. As it is released at the

[13] We have checked that this simplification does not affect our conclusions.

end of last month/very beginning of the current month (Eurostat release), it is considered with one month lag and denoted $\pi_m(-1)$ hereafter. From our theoretical model, we expect updates to be triggered only by large changes in the state variable: $|\pi_m(-1)|$. Accordingly, a switching function denoted s_t is defined as the indicator function $s_t = 1(|\pi_m(-1)| > \hat{\bar{\pi}})$, taking on value one for absolute values of $\pi_m(-1)$ greater than the estimated threshold $\hat{\bar{\pi}}$ and zero otherwise. Finally, this allows to distinguish regressors in the outer regime, $s_t X^S_{i,t}$, and their analogues in the inaction band, $(1-s_t)X^S_{i,t}$, with $X^S_{i,t} = (D(d=1), D(d=2), \cdots, D(d=8), \pi_m(-1))$. As will be seen in the next section, the threshold will be estimated from a grid search set defined by the 25 percent and 75 percent quantiles of $|\pi_m(-1)|$ as its lower and upper bounds.

The variables with coefficients that are assumed to be the same across regimes, gathered in $X^0_{i,t}$, aim at controlling for various effects which have been shown to influence the forecast updating behavior in previous empirical studies.[14] Firstly, we control for the seasonal pattern of the forecast updates—noticed in Figure 3—by introducing a forecast horizon fixed effect,[15] denoted c_h. Then, in order to take into account a business cycle effect, we introduce a target year fixed effect, denoted c_y. Thirdly, we introduce the percentage of forecasting units which have revised their forecast last month so as to capture a "cascade" effect ($Sha(-1)$), as described, e.g., in Banerjee (1992) or Graham (1999). Fourthly, a variable called $|Dev(-1)|$, which measures the previous period deviation of the unit's forecast from the average in absolute value, is considered to capture a possible "herding" effect.[16] Finally, an individual fixed effect, c_i, is also introduced. As a total, on top of the horizon, target year, and individual fixed effects, the vector $X^0_{i,t} = (Sha(-1), |Dev(-1)|)$ is introduced in the non-switching part of the regressors.

[14]See, for instance, Dovern (2013, 2015) and references therein.

[15]This horizon fixed effect completely captures a potential institutional pattern of updating within a given quarter, so that dummies indicating the second and third month of a quarter would be completely redundant. This is also the case for a dummy indicating August, when most people are having their summer vacation in the countries considered.

[16]See, e.g., Lamont (2002) or Pons-Novell (2003).

4. Empirical Testing of Our Model's Properties

4.1 The Threshold Logit Model

The empirical testing of our model's properties in terms of time and state dependence of the inflation forecast updates will rely on the estimation of threshold binary choice models for panel data. More precisely, the traditional random- or fixed-effects logit models are considered to analyze the conditional probability of forecasts revisions. Using a latent variable framework, where $\lambda^*_{i,t,h}$ is a continuous variable that is not observed, it may be written as

$$\lambda^*_{i,t,h} = s_t X^S_{i,t}\beta_{out} + (1-s_t)X^S_{i,t}\beta_{in} + X^0_{i,t}\beta + c_i + c_y + c_h + u_{it} \quad (5)$$

$$s_t = 1(|\tilde{\pi}_f(t)| > \bar{\pi}) \quad (6)$$

$$\lambda_{i,t,h} = 1(\lambda^*_{i,t,h} > 0), \quad (7)$$

where s_t is the regime-switching indicator function presented in the previous section. It is equal to one when the forecast gap is greater than a threshold $\bar{\pi}$ and zero otherwise. $X^S_{i,t}$ is the vector including the regressors which coefficients are allowed to switch across regimes, and β_{out} and β_{in} are the corresponding switching coefficients in the outer and inner regimes, respectively—the latter being the so-called inaction set defined by small $|\tilde{\pi}_f(t)|$'s. β is the vector of non-switching coefficients corresponding to the explanatory variables gathered in $X^0_{i,t}$. c_j, for $j = i, y, h$, are unobserved individual, target year, and forecast horizon fixed effects, respectively. Then, $\lambda_{i,t,h}$ is an indicator for forecast updates and

$$Pr(\lambda_{i,t,h} = 1|X_{i,t}, c_i) = G(s_t X^S_{i,t}\beta^S_{out} + (1-s_t)X^S_{i,t}\beta^S_{in} + X^0_{i,t}\beta + c_i + c_y + c_h),$$

where $G(.)$ is the cumulative distribution function (CDF) for the logit model and the standard normal CDF for the probit model. In the random-effect (RE) version, it is assumed that $(X^S_{i,t}, X^0_{i,t})$ and the c_j's, for $j = i, y, h$, are independent and that the latter have a Gaussian distribution with zero mean and constant variance so that it is possible to estimate the model by maximum-likelihood techniques, integrating out the unobserved c_j's. The assumption of

independence between $(X^S_{i,t}, X^0_{i,t})$ and the c_j's is relaxed in the fixed-effect (FE) version, and, in this case, the conditional maximum-likelihood estimator is used.[17] The threshold estimate $\hat{\tilde{\pi}}$ is obtained by grid search of all possible values of $|\tilde{\pi}_f(t)|$ over $\Gamma = [\gamma_{25\%}, \gamma_{75\%}]$, as the one which maximizes the log-likelihood of the threshold logit model described above. Here, $\gamma_{25\%}$ and $\gamma_{75\%}$ denote the 25 percent and 75 percent quantiles of $|\tilde{\pi}_f(t)|$.[18]

4.2 Estimation Results

The estimation results are reported in Table 2. The models' estimates presented below include both horizon, target year, and individual fixed effects: The corresponding Hausman statistic p-value strongly rejects the null that both RE and FE models yields the same estimates.

The average probability of update predicted by the model is rather close to the observed one in both countries, even though slightly overestimated: It is found to be 46.6 percent in France and 44 percent in Germany, whereas their observed counterparts are 45 percent and 41.7 percent, respectively. So, this model correctly accounts for a slightly larger proportion of updaters in France compared with Germany.

Let us begin with the non-switching variables, at the bottom panel of the regressors list. It appears that the update probability is increased by the cascade effect (variable $Sha(-1)$), and particularly so in the German case where the estimated coefficient is 0.219, and it is significantly different from zero at the 1 percent level. In France, this estimated coefficient is 0.08, and it is significantly different from zero at the 10 percent level only. Contrarily, the herding

[17] All estimations have been performed using version 14.1 of Stata software. In order to estimate a robust variance–covariance matrix, a clustered sandwich estimator is used to allow for intragroup correlation, relaxing the usual requirement that the observations be independent. In other words, the observations are independent across groups but not necessarily within groups.

[18] These quantiles, which make the search set Γ smaller than what is typically considered in time-series threshold models applications, have been chosen so as to save computation time since we do not expect the inaction zone to be very large. Indeed, recall that almost 45 percent of our forecasters change their forecast between two consecutive months. Hence, the threshold has to be crossed quite often.

Table 2. Threshold Logit Models with Fixed Effects

	FR		GE	
	Coefficient	**p-value**	**Coefficient**	**p-value**
$\lvert\pi_m(-1)\rvert > \hat{\bar{\pi}}$:				
s_tD(d=2)	0.025	(0.353)	0.013	(0.362)
s_tD(d=3)	0.009	(0.740)	0.067***	(0.004)
s_tD(d=4)	0.047	(0.350)	0.036	(0.415)
s_tD(d=5)	–0.036	(0.564)	–0.011	(0.771)
s_tD(d=6)	0.192**	(0.030)	0.076	(0.186)
s_tD(d=7)	0.084	(0.562)	–0.108	(0.281)
s_tD(d≥8)	0.194	(0.175)	–0.024	(0.779)
$s_t\lvert\pi_m(-1)\rvert$	0.219**	(0.045)	0.350***	(0.000)
$\lvert\pi_m(-1)\rvert \leq \hat{\bar{\pi}}$:				
$(1-s_t)$D(d=2)	–0.004	(0.847)	0.012	(0.505)
$(1-s_t)$D(d=3)	0.119***	(0.000)	0.133***	(0.000)
$(1-s_t)$D(d=4)	–0.065	(0.215)	–0.040	(0.248)
$(1-s_t)$D(d=5)	–0.063	(0.364)	–0.072*	(0.076)
$(1-s_t)$D(d=6)	0.332***	(0.000)	0.172***	(0.000)
$(1-s_t)$D(d=7)	0.064	(0.701)	–0.24	(0.738)
$(1-s_t)$D(d≥8)	–0.155	(0.494)	–0.080	(0.220)
$(1-s_t)\lvert\pi_m(-1)\rvert$	0.580	(0.292)	–0.551**	(0.050)
$Sha(-1)$	0.080*	(0.083)	0.219***	(0.000)
$\lvert Dev(-1)\rvert$	0.348***	(0.000)	0.030	(0.594)
Horizon FE	Yes		Yes	
Target Year FE	Yes		Yes	
Individual FE	Yes		Yes	
$\hat{\bar{\pi}}$ (Quantile)	0.064 (54%)		0.085 (57%)	
$\hat{\bar{\lambda}}$, Predicted Prob.	46.6%		44%	
Log-Likelihood	–3,187.13		–5,176.56	
No. Obs.	4,990		8,202	
No. i	36		51	

Note: Numbers are marginal effects computed at sample means for continuous covariates. For dummies, they show the effect of moving from one discrete state to the other one. z-statistics p-values, given in parentheses, are computed using a robust variance-covariance matrix (from a clustered sandwich estimator). *, **, and *** denote 10, 5, and 1 percent significance levels, respectively.

effect (variable $|Dev(-1)|$) is found to be much stronger in France, with an estimated coefficient of 0.35, than in Germany, where it is not even significantly different from zero. So, in France, the probability to update is significantly increased when a forecaster's previous forecast is far from last-period average forecast.[19]

Let us now turn to the threshold effect. First, it is worth noticing that the threshold models' likelihoods are improved compared with their non-switching analogues, reported in Appendix E. As can be seen there, their log-likelihoods are −3,195.6 and −5,188.1 for France and Germany, respectively. So as to test for the presence of a threshold effect, we have performed a SupWald test of the non-threshold model under the null versus our threshold model under the alternative. Since the threshold is a nuisance parameter under the non-switching null, the test statistics is not distributed as a chi-squared. Consequently, the SupWald p-value is instead calculated by bootstrap following the fixed-regressor bootstrap method of Hansen (1996, 2000b) or Hansen and Seo (2002). We ran 1,000 simulations, each with only 25 equally spaced thresholds values between the 25 percent and 75 percent quantiles of $|\tilde{\pi}_f(t)|$ under the alternative.[20] It turns out that the null is rejected at the 6 percent level for Germany and at the 31 percent level for France. This result for France might be partly due to the smaller sample size available in this country than in Germany (about 60 percent of the German sample size). Accordingly, French results should be interpreted with caution. Then, the threshold estimates correspond to the 54 percent and 57 percent quantiles of $|\Pi_m(-1)|$ distribution in France and Germany, respectively. This means that the forecast gap does not need to be very large to trigger more updates among the forecast units of our sample: The gap may just be slightly above the median.

[19]Even though the comparison with Dovern (2013)'s results is not straightforward, as this author considers GDP forecasts from a cross-country survey data, it is worth noting that the estimates of our non-threshold regressions reported in Appendix E support his findings regarding the influence of the forecast horizon, the last revisions occurring three and six months ago (D(d=3) and D(d=6)), as well as the "herding" ($|Dev(-1)|$) and "cascade" ($Sha(-1)$) effects.

[20]A more precise grid search of the threshold values would probably result in slightly lower p-values, but our actual choice seems to be a good compromise between accuracy and computational time burden concerns.

Regarding the time dependence of updates, empirical evidence supporting this view comes from the significant increase in the probability of updating if the last revision has been made three or six months ago. With our theoretical model's property that forecast update is triggered immediately after observation if the forecast gap is large enough, this gives an upper bound estimate for $T(0)$ of six months in both regimes, as predicted by the theoretical model. The six-month pattern is even stronger than the three-month one. This is particularly true in France, where the probability of update is increased by 19.2 percent if the last revision was six months ago, compared with 1 percent if it occurred three months ago in the outer regime. This result holds despite the inclusion of a forecast horizon fixed effect, hence revealing a pure time dependence effect which is not exclusively imputable to, e.g., forecasters' institutional framework. Also, as expected given the model (see Figure 2), the optimal time to next observation can be shorter in the inaction band, as the estimated coefficients associated with $(1-s_t)\mathrm{D(d{=}3)}$ are larger than the ones related to $s_t\mathrm{D(d{=}3)}$.[21]

Finally, our empirical results also support the state dependence of the forecast update decision rule. Actually, the update probability is significantly increased when the forecast gap crosses the threshold, by 22 percent in France and 35 percent in Germany. Contrarily, this probability is strongly decreased in the aptly named inaction band for Germany (-55 percent). French updates are not significantly influenced by $(1-s_t)|\pi_m(-1)|$, as the z-test p-value is close to 30 percent. Again, these conclusions are robust to the random-effect version of the logit model.

4.3 Empirical Evidence of Coexistence of Observation and Adjustment Costs

We turn now to Property 4 of the theoretical model so as to evaluate the relative costs. Recall that the ratio of adjustment to observation

[21] The conclusions of these models remain unchanged from random-effect logit model estimation and also if a probit model is substituted for the logit model. The exclusion of institutions that provide fewer than 20 usable observations does not affect the conclusions either. These results are not reported to save space but are available upon request.

Table 3. An Evaluation of Relative Costs

	$\hat{\bar{\pi}}$	n_o	n_c	ψ/θ
FR	0.064	76.6%	45.0%	49.1%
	[0.046;0.074]	[71.8%;79.7%]	—	[35.4%;59.3%]
GE	0.085	75.9%	41.7%	67.1%
	[0.081;0.106]	[76.6%;82.1%]	—	[70.1%;93.7%]

costs, ψ/θ, is a non-linear function of observations to adjustment frequencies ratio, the former being calculated thanks to the threshold estimates obtained in Table 2, along the way described in Section 2. Using these approximations, we obtain $\psi/\theta = 49.1$ percent in France and 67.1 percent in Germany, as reported in Table 3. In order to calculate the 95 percent confidence interval of these estimates, the method developed in Hansen (1999, 2000a) is used. More precisely, for all threshold values included in the grid search interval, i.e., $\forall|\tilde{\pi}_f(t)| \in \Gamma = [\gamma_{25\%}, \gamma_{75\%}]$, his proposed likelihood ratio test $LR_1(\gamma)$ is calculated: it compares the likelihood obtained with the estimated threshold $\hat{\bar{\pi}}$ reported in Table 2 with the ones obtained with all other values in Γ. The asymptotic distribution of $LR_1(\gamma)$ is shown to be highly non-standard but free of nuisance parameters and its critical value at the 5 percent level is 7.35 (see Hansen 1999, 2000a). Hence, the "no-rejection region" of confidence level 95 percent is the set of values γ such that $LR_1(\gamma) \leq 7.35$. With this confidence interval for the threshold estimate at hand, it is straightforward to compute the n_o and ψ/θ boundaries accordingly.[22] These confidence intervals are given in square brackets in Table 3.

It is worth noting that this ratio is neither close to zero nor arbitrarily large, hence confirming coexistence of both costs: As can be seen from Table 3, the smallest lower boundary of the confidence interval for the costs ratio is 35.4 percent in France, while the highest upper boundary is 93.7 percent in Germany. In addition, the costs ratio is smaller than one in both countries, meaning that observation costs dominate communication costs. Even though applied to price-setting rules, conclusions obtained by Zbaracki et al.

[22]Remember that n_c is observed, hence assumed known and fixed.

(2004) have the same flavor as ours in the sense that the managerial costs of a large U.S. industrial manufacturer (including information gathering, decisionmaking, and internal communication costs) are found to be larger than its menu costs (comparable to the external communication costs of our model).

5. Conclusion

This paper first develops a theoretical model of forecasts formation which incorporates separate observation and communication costs. As a result, the forecast update decision rule is found to be both time and state dependent. The model's main properties for the forecasts update process are the following. First, the forecast update is communicated immediately after observation if the forecast gap upon observation is large enough in absolute value. Second, the forecast revision communication is a non-linear function of the forecast gap: For small gaps $\in (-\overline{\pi}, \overline{\pi})$, there is no communication, whereas there is one as soon as the absolute value of the gap exceeds $\overline{\pi}$. If so, the gap is closed by the update. Third, the optimal time between two observations is also a non-linear function of the forecast gap, and the closer to the boundaries $-\overline{\pi}$ or $\overline{\pi}$, the sooner the next observation is. Fourth, the time between two observations reaches a maximum when the gap is closed.[23]

The time and state dependence of the observation and forecast revision communication implied by this model are then tested using inflation forecast updates of professional forecasters from recent Consensus Economics panel data for France and Germany. For this purpose, conditional forecast revision communication probabilities are estimated from binary choice models incorporating regime-switching features. This allows to estimate the threshold value $\overline{\pi}$ along all other parameters.

Our findings clearly support time-dependent updates, a result which is compatible with the observation cost assumption as introduced by Mankiw and Reis (2002). Indeed, they point to a strong positive effect when the last update has occurred three and/or six

[23]This maximum time increases in both observation and communication costs; see Proposition C.2 in Appendix C.

months ago, even after controlling for the institution's periodic forecast framework (quarterly or biannual). This gives a maximum of six months for the optimal time between two observations. This confirms Coibion and Gorodnichenko (2015)'s estimate of the average duration between information updates.

Evidence of update state dependence is also provided. In fact, a strong positive and significant effect is found on updates when the forecast gap is larger than the estimated threshold, as proxied by the last known monthly inflation rate, weighted by its mechanical contribution to the yearly inflation rate forecast.

Finally, our results confirm the co-existence of both types of costs with a forecast communication cost smaller than the observation cost. This result is very much in line with Zbaracki et al. (2004)'s findings for firms' price-setting decision rules.

All in all, this work is a first attempt to explain why not communicating forecast updates every month can be optimal for professional forecasters. The theoretical model grounding it is of course highly stylized. We believe that it is desirable to improve this model so as to make it closer to the true forecaster's environment and hence to test it directly instead of estimating a reduced-form regression. This is on our research agenda.

Appendix A. The Optimal Control Problem

Suppose we are in a base period, say $t = 0$, where the forecast is equal to $\pi_f(0)$. From there on, we seek to forecast the inflation rate at date $t = 1$. As stated in Assumption 2, the instantaneous quadratic loss function faced by a representative professional forecaster is $\tilde{\pi}_f(t)^2$. As long as a new forecast has not been publicly adjusted and communicated, $\pi_f(t)$ is constant, equal to $\pi_f(0)$. It is assumed that the forecaster's objective is to produce the best possible forecast, which amounts to minimize the distance between his forecast and the optimal forecast. The latter is a BM, according to Assumption 1.

We assume that at time $t = 0$, the forecaster pays the cost θ and observes $\pi_f^*(0)$. Then, until the next observation of the target forecast, $E(\pi_f^*(h)|\mathcal{I}_0) = \pi_f^*(0)$, where $\mathcal{I}_0$ is the information set observed at time $t = 0$. Indeed, with $\pi_f^*(0) = E(\pi(1)|\mathcal{I}_0)$, where $\pi(1)$ is the inflation rate at $t = 1$, we have

$$E\left(\pi_f^*(h)|\mathcal{I}_0\right) = E\left(E\left(\pi(1)|\mathcal{I}_h\right)|\mathcal{I}_0\right) = E\left(\pi(1)|\mathcal{I}_0\right)$$
$$= \pi_f^*(0). \tag{A.1}$$

The target forecast dynamics is given by Assumption 1.

We define the "uncontrolled" forecast gap as the forecast gap $\tilde{\pi}_f(t)$ between two observations of $\pi_f^*(t)$ and before a new adjustment of $\pi_f(t)$. It follows that starting from $\tilde{\pi}_f(0)$ at time $t = 0$, the uncontrolled forecast gap evolves as

$$d\tilde{\pi}(t) = -\sigma dW(t). \tag{A.2}$$

As will be shown below, at optimum there will be adjustments at finite time in accordance with the loss-minimization objective, which will make the forecast gap process globally stationary.

In this setup, the problem of the forecaster is to choose the time elapsed until the next observation of the target forecast, $0 < T \leq 1$, which will cost θ, as well as the number of forecast update communications between two observations (at $t = 0$ and T, respectively), $J \in N$, occurring at successive dates $0 \leq t_1, t_2, \ldots, t_J < T$, each one incurring an adjustment cost ψ. He also chooses the size of the forecast update so that the expected value of the forecast gap on adjustment is $\hat{\pi}_{f,j}$, with $j = 1, \ldots, J$.

As it is stated, the problem is formally similar to the one treated by Alvarez, Lippi, and Paciello (2011). There are two differences:

(i) The first difference comes from the fact that the timing of the next observation, that is the choice of T, is bounded by the forecast horizon, 1. There is no upper bound of this sort in Alvarez, Lippi, and Paciello (2011). However, this is a pure control constraint, the simplest one from the mathematical point of view, and we show in our empirical section that it is never binding in practice. Without loss of generality, we shall omit it hereafter to unburden the presentation.

(ii) The second concerns the law of motion assumed by Alvarez, Lippi, and Paciello (2011) for the target: It has a time drift ($\mu \geq 0$), which also shows up in the law of motion of the

analogous stochastic gap.[24] While non-zero drifts make sense in their framework (they are studying *price* setting), they do not in ours (focused on *inflation forecast* behavior). As a result, from this point of view, our problem corresponds to a special case in Alvarez, Lippi, and Paciello (2011), explicitly treated in their Section V, pp. 1928–34.

Henceforth, we shall use the same formalism and methodology as in Alvarez, Lippi, and Paciello (2011). Let $V(\tilde{\pi}_f)$ denote the value function of the forecaster at the time of an observation of the forecast gap $\tilde{\pi}_f$, and $V_J(\tilde{\pi}_f)$ the best value that the forecaster can reach by making J forecast updates between observations. Then, note that

$$V(\tilde{\pi}_f) = \min_{J \geq 0} V_J(\tilde{\pi}_f), \quad \text{so that} \quad J^\star(\tilde{\pi}_f) = \arg\min_{J \geq 0} V_J(\tilde{\pi}_f), \tag{A.3}$$

where $J^\star(\tilde{\pi}_f)$ is the optimal number of forecast updates. The corresponding Bellman equations can be written accordingly. To ease the exposition, suppose that $J \in \{0, 1\}$, that is, at most one forecast update is optimal. This will turn out to be true for our forecast formation model. Let's denote $\tilde{\pi}_f = \tilde{\pi}_f(0)$ the value of the forecast gap at time $t = 0$.

For $J = 0$, the conditional value function for the forecaster problem writes

$$\begin{aligned} V_0(\tilde{\pi}_f) = \theta + \min_T \int_0^T & e^{-\rho t}(\tilde{\pi}_f^2 + \sigma^2 t)dt \\ & + e^{-\rho T} \int_{-\infty}^{\infty} V(\tilde{\pi}_f - s\sigma\sqrt{T})dN(s), \end{aligned} \tag{A.4}$$

where s is a standard normal. The first component of the right member of $V_0(\tilde{\pi}_f)$ is of course the information observation cost, while the second one is the time $t = 0$ expected loss between $t = 0$ and T. In the third (continuation) component, $N(.)$ is the probability density function of a Gaussian distribution.

[24] See Equation (1) in Alvarez, Lippi, and Paciello (2011, p. 1917).

For $J = 1$ it becomes (with the simplified notation of the adjusted forecast at time t_1 as $\hat{\pi}_f = \hat{\pi}_{f,1}$)

$$\begin{aligned} V_1(\tilde{\pi}_f) = \theta + \min_{T,\hat{\pi}_f,t_1} \int_0^{t_1} e^{-\rho t}\tilde{\pi}_f^2 dt + \int_0^T e^{-\rho t}\sigma^2 t dt \\ + e^{-\rho t_1}\left[\psi + \int_0^{T-t_1} e^{-\rho t}\hat{\pi}_f^2\right] dt \\ + e^{-\rho T}\int_{-\infty}^{\infty} V(\hat{\pi}_f - s\sigma\sqrt{T})dN(s). \end{aligned} \tag{A.5}$$

The unconditional value function is hence given by

$$V(\tilde{\pi}_f) = min\{V_0(\tilde{\pi}_f), V_1(\tilde{\pi}_f)\}. \tag{A.6}$$

Because of the quadratic functions involved, it is possible to provide with a partial analytical characterization of the solutions to these Bellman equations.[25] Let us first focus on the optimal forecast revisions before turning to the optimal time between observations.

Appendix B. The Optimal Number and Timing of Forecast Updates Announcements

Using Proposition 1, p. 1921, in Alvarez, Lippi, and Paciello (2011), Proposition B.1 below can be established straightforwardly.[26]

Proposition B.1. *Let $\theta > 0$, $\psi > 0$ and $\sigma > 0$.*

(i) $J^\star(\tilde{\pi}_f) \leq 1, \forall \tilde{\pi}_f \in R$.

(ii) If $\tilde{\pi}_f$ is such that $J^\star(\tilde{\pi}_f) = 1$, then the forecast update is full and occurs instantaneously ($\hat{\pi}_f = 0$ and $t_1 = 0$).

[25] Alvarez, Lippi, and Paciello (2011) do solve the more general case with positive time drift in Equation (2). As mentioned above, only the zero drift case is relevant here.

[26] Actually, Alvarez, Lippi, and Paciello (2011) get the same characterization even if the law of motion of their forecast target has a time drift, provided its absolute value is small enough. Our case—zero time drift—is indeed a very elementary special case.

This proposition means that for the quadratic loss function and the forecast gap process defined in Equation (A.2), there will be at most one forecast update communication between two forecast target observations and, if there is one, it will be fully and instantaneously adjusted to the updated forecast target. The intuition of this result is as follows. If ψ and σ are strictly positive, it is immediate to see that in case of inflation forecast adjustment, $\hat{\pi}_f = 0$. Then the optimal choice of the forecaster is to reset the forecast gap to zero. Indeed, it follows from Equation (A.2) that the expected value of the forecast gap remains at zero between observations. Consequently, there are no gains to expect from a forecast update without the new pieces of information acquired thanks to an observation. By contrast, due to the fixed communication cost $\psi > 0$, the losses involved by a forecast update release are strictly positive. So, it is not optimal to adjust the forecast between information set updates. Using the same type of arguments, one can also readily get why the update communication, if any, should not only be full but also instantaneous. Indeed, if $t_1 > 0$, then the forecaster will incur losses in the time interval $[0; t_1[$. Because the presence of "menu costs" typically induces that adjustments only occur if the gaps are sufficiently large, delaying cannot be optimal, as it involves starting the adjustment at such values of the gap: The corresponding trajectory of losses is clearly dominated by a trajectory where adjustment is made instantaneously. Hence $\hat{\pi}_f = 0$ and $t_1 = 0$. A first result of our theoretical model of forecasts formation is that forecasts communication plans, i.e., progressive release of forecasts' revisions between two updates of the information set, are not optimal. The new forecast release, if any, is full and instantaneous at the observation time. Notice that the resulting Bellman equation (A.5) can be rewritten as

$$V_1 = \psi + \theta + \min_{T, \hat{\pi}_f} \int_0^T e^{-\rho t} (\hat{\pi}_f^2 + \sigma^2 t) dt + e^{-\rho T} \int_{-\infty}^{\infty} V(\hat{\pi}_f - s\sigma\sqrt{T}) dN(s). \tag{B.1}$$

Note that V_1 does not depend on $\tilde{\pi}_f$ since $t_1 = 0$. Hence the value function is given by

$$V(\tilde{\pi}_f) = min\{V_0(\tilde{\pi}_f), V_1\}. \tag{B.2}$$

One can see immediately that the value function V is symmetric around $\tilde{\pi}_f = 0$ and increasing for $|\tilde{\pi}_f| < \bar{\pi}$, where $\bar{\pi}$ is a threshold value such that $V_1 > V_0(\tilde{\pi})$ for $\tilde{\pi}_f \in (-\bar{\pi}, \bar{\pi})$. Hence $(-\bar{\pi}, \bar{\pi})$ defines the range of inaction in which no forecast adjustment is communicated. It means that when $\tilde{\pi}_f$ is smaller than $\bar{\pi}$ in absolute value, then the inflation forecast is not adjusted. Since V is not differentiable at $\tilde{\pi}_f = \overline{\pi}$, the value function is discontinuous, non-smooth at the boundaries of the inaction band. These properties are summarized in the next proposition:[27]

Proposition B.2. *The value function V is symmetric around $\tilde{\pi}_f = 0$, and V is strictly increasing in $\tilde{\pi}_f$ for $0 < \tilde{\pi}_f < \bar{\pi}$. $V'(\tilde{\pi}) = 0$ for $\tilde{\pi}_f > \bar{\pi}$ and V is not differentiable at $\tilde{\pi}_f = \bar{\pi}$.*

Appendix C. The Optimal Time between Observations

Since $\hat{\pi}_f = 0$ at the forecaster optimum, then the optimal time between observations, conditional on adjustment, is $T(0)$. It is possible to characterize much more closely function $T(.)$: Indeed, Alvarez, Lippi, and Paciello (2011) show that it has a maximum at $\tilde{\pi}_f = 0$ and is symmetric around 0 with an inverted U-shape.[28] The function $T(\cdot)$ is discontinuous and not differentiable at $\tilde{\pi}_f = \overline{\pi}$, as stated in Proposition C.1.[29]

Proposition C.1. *As $\rho \downarrow 0$, the optimal rule for the time of the next observation of the forecast gap is given by Equation (2).*

Indeed, Equation (2) is obtained from a third-order expansion of $T(.)$ around zero, which requires the condition $\rho \downarrow 0$ appearing in Proposition C.1. Note that the optimal rule for the time of the

[27]The proof of Proposition B.2 follows the same steps as the proof of Proposition 3, p. 1928, and Lemma 1, p. 1945, in Alvarez, Lippi, and Paciello (2011).

[28]See Alvarez, Lippi, and Paciello (2011), Proposition 4, p. 1929.

[29]Because the proof of Proposition C.1 follows the same steps as the proof of Proposition 4, p. 1929, in Alvarez, Lippi, and Paciello (2011), we refer the reader to the latter in order to lighten the presentation.

next observation is a discontinuous, threshold function of the forecast gap: it switches from $T(0)$ for all $|\tilde{\pi}_f| > \bar{\pi}$ to a function of $\tilde{\pi}_f$ otherwise. It is also worth noticing that if there is no forecast gap after an observation, i.e., $\tilde{\pi}_f = 0$, then T(0) is optimal. Finally, $T(\tilde{\pi}_f)$ decreases with $|\tilde{\pi}_f|$. When close to the boundaries of the range of inaction, the forecaster plans an early observation since the target is likely to cross the threshold.

Finally, the next proposition can easily be shown to hold[30]:

Proposition C.2. *Let $\theta > 0$, $\psi > 0$ and $\sigma > 0$, $\frac{\psi}{\theta} < 5.5$. As $\rho \downarrow 0$, there exists a unique solution for $T(0)$ and $\overline{\pi}$ and it is such that*

(i) $T(0)$ is increasing in ψ and θ

(ii) $\overline{\pi}$ is increasing in ψ and decreasing in θ.

As expected, the time to the next observation after a forecast update, $T(0)$, is increasing in the observation cost, and the width of the inaction band defined by $(-\overline{\pi}, \overline{\pi})$ is increasing in the adjustment cost.

Appendix D. The Optimal Forecast Gap Dynamics

The decision rules described by the threshold $\overline{\pi}$ and the function $T(\cdot)$ imply a stationary Markov process for the dynamics of the forecast gap on observation and before communication. It follows from Proposition B.1 in Appendix B that the forecast gap adjustment rule $\Delta(\tilde{\pi}_f)$ is zero in the inaction zone and $-\tilde{\pi}_f$ otherwise. Assume that the next observation occurs in $T' = T(\tilde{\pi}_f)$ periods and the corresponding forecast gap is $\tilde{\pi}'_f$. Then, the proposition below follows immediately from Proposition C.1.[31]

Proposition D.1. *Optimal forecast gap dynamics:*

$$\tilde{\pi}'_f = \begin{cases} \tilde{\pi}_f - s\sigma\sqrt{T(0) - \left(\frac{\tilde{\pi}_f}{\sigma}\right)^2}, & \text{if } \tilde{\pi}_f \in (-\overline{\pi}, \overline{\pi}) \\ -s\sigma\sqrt{T(0)}, & \text{otherwise.} \end{cases} \tag{D.1}$$

[30]Here, we use Proposition 5, p. 1931, in Alvarez, Lippi, and Paciello (2011) as well as on results reported p. 1932 therein.

[31]See also Equation (16), p. 1934, in Alvarez, Lippi, and Paciello (2011).

Again, this is a threshold regime-switching process: If the forecast gap is large enough, i.e., larger than $\overline{\pi}$ in absolute value, then it is corrected by setting the forecast to its target and the remaining gap is entirely imputable to unexpected shocks. $\tilde{\pi}'_f$ increases with $|\tilde{\pi}_f|$. When close to the boundaries of the range of inaction but still inside it, the forecast gap reaches its largest values until it hits the boundary and then goes back to its minimum. This is the Ss rule.

Appendix E. Logit Models Estimates without Threshold Effect

Table E.1. Logit Models: Estimates without Threshold Effect

	FR		GE	
	Coefficient	**p-value**	**Coefficient**	**p-value**
D(d=2)	0.010	(0.541)	0.013	(0.218)
D(d=3)	0.062***	(0.009)	0.101***	(0.000)
D(d=4)	–0.012	(0.713)	–0.002	(0.937)
D(d=5)	–0.048	(0.369)	–0.033	(0.314)
D(d=6)	0.245***	(0.001)	0.126***	(0.002)
D(d=7)	0.077	(0.537)	–0.701	(0.297)
D(d≥8)	0.069	(0.473)	–0.043	(0.529)
$\lvert\pi_m(-1)\rvert$	0.183**	(0.026)	0.428***	(0.000)
$Sha(-1)$	0.089**	(0.050)	0.197***	(0.000)
$\lvert Dev(-1)\rvert$	0.354***	(0.000)	0.029	(0.611)
Horizon FE	Yes		Yes	
Target Year FE	Yes		Yes	
Individual FE	Yes		Yes	
No. Obs.	4,990		8,202	
Log-Likelihood	–3,195.57		–5,188.12	
No. i	36		51	

Note: Numbers are marginal effects computed at sample means for continuous covariates. For dummies, they show the effect of moving from one discrete state to the other one. z-statistics p-values, given in parentheses, are computed using a robust variance-covariance matrix (from a clustered sandwich estimator). *, **, and *** denote 10, 5, and 1 percent significance levels, respectively.

References

Alvarez, F., F. Lippi, and L. Paciello. 2011. "Optimal Price Setting With Observation and Menu Costs." *Quarterly Journal of Economics* 126 (4): 1909–60.

Andrade, P., and H. Le Bihan. 2013. "Inattentive Professional Forecasters." *Journal of Monetary Economics* 60 (8): 967–82.

Banerjee, A. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107 (3): 797–817.

Coibion, O. 2010. "Testing the Sticky Information Phillips Curve." *Review of Economics and Statistics* 92 (1): 87–101.

Coibion, O., and Y. Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1): 116–59.

———. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8): 2644–78.

Dovern, J. 2013. "When Are GDP Forecasts Updated? Evidence from a Large International Panel." *Economics Letters* 120 (3): 521–24.

Dovern, J., U. Fritsche, P. Loungani, and N. Tamirisa. 2015. "Information Rigidities: Comparing Average and Individual Forecasts for a Large International Panel." *International Journal of Forecasting* 31 (1): 144–54.

Graham, J. 1999. "Herding among Investment Newsletters: Theory and Evidence." *Journal of Finance* 54 (1): 237–68.

Hansen, B. 1996. "Inference when a Nuisance Parameter Is Not Identified Under the Null Hypothesis." *Econometrica* 64 (2): 413–30.

———. 1999. "Threshold Effects in Non-dynamic Panels: Estimation, Testing, and Inference." *Journal of Econometrics* 93 (2): 345–68.

———. 2000a. "Sample Splitting and Threshold Estimation." *Econometrica* 68 (3): 575–604.

———. 2000b. "Testing for Structural Change in Conditional Model." *Journal of Econometrics* 97 (1): 93–115.

Hansen, B., and B. Seo. 2002. "Testing for Two-Regime Threshold Cointegration in Vector Error-Correction Models." *Journal of Econometrics* 110 (2): 293–318.

Lamont, O. 2002. "Macroeconomic Forecasts and Microeconomic Forecasters." *Journal of Economic Behavior and Organization* 48 (3): 265–80.

Loungani, P., H. Stekler, and N. Tamirisa. 2013. "Information Rigidity in Growth Forecasts: Some Cross-Country Evidence." *International Journal of Forecasting* 29 (4): 605–21.

Magnani, J., A. Gorry, and R. Oprea. 2016. "Time and State Dependence in an Ss Decision Experiment." *American Economic Journal: Macroeconomics* 8 (1): 285–310.

Mankiw, N., and R. Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.

Pons-Novell, J. 2003. "Strategic Bias, Herding Behaviour and Economic Forecasts." *Journal of Forecasting* 22 (1): 67–77.

Sims, C. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.

Stock, J., and M. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (s1): 3–33.

Woodford, M. 2002. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information and Expectations in Modern Macroeconomics,* ed. P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, 25–58. Princeton University Press.

Zbaracki, M., M. Ritson, D. Levy, S. Dutta, and M. Bergen. 2004. "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets." *Review of Economics and Statistics* 86 (2): 514–33.

Countering Appreciation Pressure with Unconventional Monetary Policy: The Role of Financial Frictions*

Nicole Aregger[a] and Jessica Leutert[b]
[a]UBS
[b]Swiss National Bank

In a simple two-country framework with imperfect financial intermediation we analyze and compare the effectiveness of two unconventional monetary policy measures: foreign exchange interventions and credit easing. Central bank interventions only have real effects when banks are financially constrained. For a country facing excess demand for its bonds, we study three external sources of appreciation pressure: increased financial frictions in the international credit market, an increase in capital inflows, and increased financial frictions in the foreign investment market. Only in the first two cases, foreign exchange interventions can reverse the appreciation and the resulting misallocation of capital. Under certain conditions, credit easing is a substitute for foreign exchange interventions.

JEL Codes: E52, F31, F32, F41, G15, G20.

*We are grateful to the editor, Pierpaolo Benigno, and two anonymous referees for their constructive comments. We would also like to thank Philippe Bacchetta, Ambrogio Cesa-Bianchi, Harris Dellas, Matthias Lutz, Matteo Maggiori, Jean-Paul Renne, Cédric Tille, Christoph Winter, and seminar participants at the University of Bern, the University of Lausanne, the SSES Annual Congress, the Annual Inflation Targeting Seminar of the Central Bank of Brazil, the Young Swiss Economists Meeting, the ECOBATE, the Swiss Macro-Finance Workshop, and the Spring Meeting of Young Economists for valuable comments and suggestions. The views expressed in this paper are those of the authors and do not represent those of the Swiss National Bank or UBS AG and its affiliates. Author contact: Nicole Aregger, UBS, Europaallee 21, 8004 Zurich, Switzerland. E-mail: nicole_aregger@bluewin.ch. Jessica Leutert, Swiss National Bank, Boersenstrasse 15, Postfach, 8022 Zurich, Switzerland. E-mail: jessica.leutert@snb.ch.

1. Introduction

In response to the global financial crisis of 2007–09, the world's leading central banks lowered their policy rates to very low levels. However, the liquidity they provided did not reach the private sector because financial intermediation was severely disrupted. At the same time, economies like Switzerland, Denmark, and Israel were facing intense appreciation pressure on their currencies, which meant a further tightening of monetary conditions. With low short-term interest rates proving to be of limited effectiveness and policy rates close to the effective lower bound, their central banks started engaging in unconventional monetary policy.

For a central bank wishing to dampen appreciation pressure on its currency by using unconventional policy tools, a number of questions arise before it can take action: What is the source of the appreciation? What tools would be effective at reducing the appreciation pressure in the first place? And if there are several tools, which one would be most suitable? In this paper, we propose a simple framework to address these questions. We extend a simplified, real version of the closed-economy financial frictions model of Gertler and Karadi (2013) by adding an open-economy dimension and by incorporating the idea of Gabaix and Maggiori (2015) that financial frictions lead to deviations from interest parity. We focus on two of the most prominent unconventional measures: foreign exchange (FX) interventions and credit easing (CE), i.e., a special form of quantitative easing (QE).[1] We examine under what conditions a central bank can use these two tools to reduce financial frictions and, in particular, respond to appreciation pressure.

There already exists substantial research on unconventional monetary policy tools. On the theoretical side, both in the literature on asset purchase programs and in the literature on FX interventions, there are models that introduce financial frictions as limited commitment of financial intermediaries. In the past few years, financial

[1]Credit easing (central bank purchases of bonds issued by private sector borrowers) is a special form of QE if it involves increasing the monetary base. However, the intentions of credit-easing programs differ from those of QE programs. The goal of credit easing is to change the asset composition on the central bank's balance sheet in order to reduce specific interest rates or restore the functioning of specific markets (see Bernanke 2009).

frictions and imperfect financial intermediation have become an important modeling tool, since they are considered to be crucial for the strong spread of the 2007–09 financial crisis. They provide a plausible explanation for credit spreads and deviations from interest parity and generate a portfolio balance channel through which QE/CE and FX interventions are effective. The key feature for the portfolio balance channel to work is imperfect substitutability between the assets the central bank purchases and those that it uses to finance these purchases. In QE/CE models, different types of domestic assets are modeled as imperfect substitutes, while in FX interventions models, domestic and foreign assets are imperfectly substitutable. Despite the similarities in the literature on QE/CE and FX interventions, what is missing to our knowledge is a framework in which the two tools are directly compared. Our paper presents a first contribution to fill this gap.

From an empirical point of view, there are many arguments in favor of studying the effectiveness of foreign exchange interventions and credit easing in a unified framework. Crisis times tend to be reflected in both the foreign exchange market and credit markets. During periods of high uncertainty, credit spreads tend to widen (see, for example, Kwon 2020). At the same time, an increase in uncertainty tends to trigger exchange rate movements. Safe-haven currencies like the U.S. dollar (USD) and the Swiss franc, for example, typically appreciate and see an increase in their safety premia, indicating an increase in the deviation from uncovered interest parity (see Maggiori 2013 for the USD and Leutert 2018 for the Swiss franc). As we show in Appendix A, at least in periods of high uncertainty, changes in credit spreads in the United States and in Switzerland are significantly correlated with changes in the exchange rate, suggesting that there is a co-movement between credit spreads and deviations from uncovered interest parity. This potentially tight link between the domestic credit market and the foreign exchange market has important implications for the transmission channels of unconventional monetary policies. On the one hand, there is a large literature documenting that QE/CE leads to a depreciation of the domestic currency.[2] For the case of the European Central

[2]For an overview of the literature, see for example Dedola et al. (2021). While most contributions assess the exchange rate effects of QE in general, Saadi Sedik,

Bank (ECB) and Federal Reserve QE and credit-easing programs, Dedola et al. (2021) find that adjustments in deviations from covered and uncovered interest parity (UIP) add to explaining this depreciation. On the other hand, empirical evidence suggests that foreign exchange interventions, in turn, have an impact on domestic credit market conditions. Fuhrer, Nitschka, and Wunderli (2021) look at the case of Switzerland where, since the global financial crisis, large-scale foreign exchange interventions have led to a strong increase in central bank reserves in the financial system. They provide evidence that this increase in reserves has lowered banks' lending spreads. In this paper, we provide a theoretical framework that can explain both the potential co-movement between credit spreads and deviations from interest parity and the potential spillovers of unconventional monetary policies.

Banks are at the core of our model. An agency problem between borrowers and lenders generates the portfolio balance channel: Limited commitment of banks leads to an endogenous credit constraint and results in limits to arbitrage in both the domestic investment markets and the international credit market, which is reflected in excess returns on the corresponding assets. Compared with a frictionless equilibrium, capital costs in the investment markets are higher and there is a deviation from interest parity in the international credit market.

For a country facing an excess demand for its bonds, we identify three external sources of appreciation pressure related to financial frictions. The first is financial frictions in the international credit market: Banks are less able to bear exchange rate risk and to absorb the excess supply of foreign bonds resulting from trade and financial imbalances. Less intermediation in the international credit market leads to a deviation from interest parity and causes a home real appreciation. The second source of external appreciation pressure is capital inflows. If banks are credit constrained in the international credit market, they lead to an increase in the deviation from interest parity and therefore an appreciation. This is because such inflows absorb a large part of the banks' limited intermediation capacity. The third source of external appreciation pressure is financial

Jacome H., and Ziegenbein (2018) focus on credit easing and provide evidence that it leads to a depreciation of the domestic currency.

frictions in the foreign investment market. When foreign banks are less able to intermediate funds, investment in the foreign country decreases. The higher relative level of future home output induces a permanent home appreciation but does not lead to a deviation from interest parity.

We show that within our model the home central bank can use unconventional monetary policy to reduce the appreciation pressure in the first two cases only. Both financial frictions in the international credit market and capital inflows lead to a temporary appreciation because of an increase in the deviation from interest parity, i.e., the excess return on foreign bonds. By acquiring foreign bonds and issuing domestic bonds in return, the central bank can increase overall financial intermediation and thereby help to reduce excess returns and thus bring the home country's economy closer to the frictionless state. In contrast, since financial frictions in the foreign investment market lead to a permanent home appreciation, unconventional policy, if effective at all, can lower the appreciation today only at the price of a future appreciation.

Finally, we show that credit easing can achieve the same goal as foreign exchange interventions if banks are not only credit constrained in the international credit market but also in the domestic investment market. Interventions in one market make banks shift their assets to the other market, reducing the excess returns in both markets. In this case, central bank intervention should target the market that exhibits the highest excess returns. This ensures that the balance sheet extension of the central bank needed to reach its goal is minimized.

Most of the literature on asset purchase programs emerged after the financial crisis of 2007–09. Overall, there is broad empirical evidence that QE programs were successful at flattening the yield curve, while the precise mechanism through which they work remains unclear. Theory suggests two main channels: the signaling channel (the intervention is a signal about the future stance of monetary policy; see, e.g., Eggertsson and Woodford 2003) and the portfolio balance channel.[3] The portfolio balance channel was first described

[3]Although the central bank extends the quantity of reserves under QE, only a few studies discuss the liquidity channel through which QE may work (see, e.g., Christensen and Krogstrup 2019).

by Tobin (1958, 1969). Since then, various theoretical foundations for imperfect asset substitutability have emerged.[4] Gertler and Karadi (2013) follow Gertler and Karadi (2011) and Gertler and Kiyotaki (2011) by looking at QE as a form of financial intermediation, performed by the central bank. The central bank acquires assets by issuing interest-bearing short-term debt. The assets that the central bank purchases and those that it issues are imperfectly substitutable because of limits to arbitrage in private financial intermediation caused by financial frictions. Such limits to arbitrage lead to extraordinary returns on assets, and thus generate a role for central bank intermediation to drive these returns down. As the central bank can intermediate both long-term government bonds and short-term private securities, the model presents a unified approach to analyze a variety of programs used in practice. While most of the theoretical literature models asset purchase programs in closed economies, Dedola, Karadi, and Lombardo (2013) provide a two-country model to analyze the international dimension of unconventional policies in economies with financial frictions, but do not study exchange rates.

The literature on FX interventions grew rapidly after the end of the Bretton Woods system. It finds that non-sterilized interventions do have an effect on the exchange rate because they change the monetary base. The effectiveness of sterilized interventions is less clear on both theoretical and empirical grounds.[5] Similar to the theory on QE, there are two main channels through which FX interventions may affect the exchange rate: the signaling channel and the portfolio balance channel. Early theoretical foundations for the portfolio balance channel were provided by Kouri (1976), Henderson and Rogoff (1982), and Branson and Henderson (1985). More recent advances are Kumhof (2010) and Gabaix and Maggiori (2015). The latter's model of exchange rate determination is a modern version of the traditional portfolio balance models. It illustrates how gross capital

[4]Andrés, López-Salido, and Nelson (2004) introduce an adjustment to household preferences leading to imperfect asset substitutability between holdings of long-term bonds and money in a New Keynesian model. Similarly, using a partial equilibrium approach, Vayanos and Vila (2009) propose a model where investors have heterogeneous preferences for assets of different maturities ("preferred-habitat" motive).

[5]For a literature review see, e.g., Sarno and Taylor (2001), Neely (2005), or Menkhoff (2010).

flows matter for the determination of exchange rates and what the effects of foreign exchange interventions are. Similar to Gertler and Karadi (2013), they use limited commitment of financial intermediaries to introduce financial frictions and endogenize a deviation from interest parity, reflecting a currency risk premium. Cavallino (2019) derives the optimal foreign exchange intervention policy in a New Keynesian small open economy version of the Gabaix and Maggiori (2015) model. He finds that in response to a capital inflow shock, the optimal foreign exchange intervention leans against the wind to stabilize the path of the exchange rate. In addition, using Swiss data, Cavallino (2019) provides some empirical evidence suggesting that capital inflows lead to an appreciation of the Swiss franc. While our model contains the key mechanisms of Gabaix and Maggiori (2015) and Cavallino (2019) and hence their main results regarding the exchange rate determination, capital flows, and foreign exchange interventions also hold in our setup, our framework incorporates capital and the domestic credit markets as additional elements. This allows us to study the spillovers between the foreign exchange market and domestic credit markets and include credit easing as an additional policy option.

The key mechanisms that are at work in our model are also contained in the ECB's New Area-Wide Model II (NAWM II; see Coenen et al. 2018). In particular, large-scale asset purchases in this open-economy dynamic stochastic general equilibrium (DSGE) framework exert their influence through both a credit and an exchange rate channel. The key contributions of our paper relative to the ECB model are to provide a detailed discussion of the role of financial frictions in the exchange rate determination and the preconditions for asset purchases to be effective, and to contrast credit easing to FX interventions. In addition, we shed light on the potential importance of gross capital flows and gross foreign asset positions in the determination of the exchange rate.

The main building blocks of our model are also similar to Nuguer (2018). Proposing a two-country model with globally acting banks, she studies the international transmission of a financial crisis through the international interbank market and looks at the welfare effects of unconventional credit policies that help to mitigate the effects of a financial disruption. In contrast, in our paper, we mainly focus on the role of financial frictions in the exchange rate determination and

compare the effectiveness of foreign exchange interventions to the effectiveness of credit easing at reducing spreads in financial markets, abstracting from a DSGE framework and a welfare analysis. Yet, relative to Nuguer (2018), our model has some insightful additional dimensions. First, in her model, domestic and foreign banks trade a global bond and the global interbank market is assumed to be frictionless. In our model, we abstract from such a bond but assume that there is not only an agency problem between households and banks, but also between domestic banks and foreign banks.[6] As a result, in our model, there is a straightforward link between deviations from interest parity and any excess supply of foreign bonds in the spirit of Gabaix and Maggiori (2015). Second, as opposed to Nuguer, we allow the friction parameters to differ across asset classes (as in Gertler and Karadi 2013), thereby capturing the fact that not all markets need simultaneously be subject to distortions. As a result, in our model, the effectiveness of different unconventional monetary policies depends strongly on the types of frictions present. Finally, in contrast to Nuguer, we allow households to hold not only domestic bonds but also foreign bonds. Thereby, our model highlights the potentially important role of gross foreign asset positions in the distortion of financial markets and, in particular, the exchange rate.

Finally, even though we use quite a different modeling approach, our model has some similar channels as Adrian et al. (2020). They have developed a New Keynesian model to assess the effects of FX interventions and capital controls. They find these two tools to be useful in improving the trade-offs of conventional monetary policy or when policy rates hit the effective lower bound. Our model can replicate the main features and predictions of their model.[7] Yet,

[6]Nuguer (2018) assumes that foreign banks can only divert assets funded by households, but not assets funded by domestic banks. This seems unrealistic. As pointed out by Kollmann (2016), the interbank market was severely impaired during the global financial crisis.

[7]For example, Adrian et al. (2020) assume the UIP risk premium to rise when net foreign liabilities increase—this is a key result in our model—or corporate spreads in emerging economies to rise as their currencies depreciate. In our model, for a country facing an excess supply of its bonds, increases in corporate spreads go hand in hand with a currency depreciation when there is an increase in risk perception regarding the investment market.

our setup has some additional insightful features. For instance, our model suggests that changes in the UIP risk premium may affect corporate spreads also in advanced economies (not only in emerging economies), namely through the portfolio decision made by the banking sector.[8] As a result, central banks in these economies have another potential tool at their disposal when wishing to counter appreciation pressure by reducing interest rate spreads, namely credit easing. Furthermore, as opposed to Adrian et al. (2020), our model emphasizes the potential importance of gross capital flows and gross foreign asset positions in deviations from UIP and, therefore, the level of the exchange rate.

2. Model

Our general equilibrium model combines the main elements of Gertler and Karadi (2013) and Gabaix and Maggiori (2015). As a starting point, we take a simplified version of the closed-economy model by Gertler and Karadi (2013), breaking it down to two periods ($t = 0, 1$) and assuming a deterministic and real model environment. We extend this setup to a two-country model with one homogeneous traded good, and one non-traded good in each country. The prices of the non-traded goods act as numéraires. There are households, firms, and banks. Later, we also introduce a central bank.

By choosing a real setup, we abstract from exchange rate movements stemming from monetary phenomena and nominal frictions. This allows us to study the credit channels in isolation from any other influences. Even though highly reduced in some respects, our model captures the fundamental structure of the domestic and international bond markets and thus contains the main economic intuitions.

2.1 Households

In each of the two countries, there is a continuum of identical households having unit mass. The representative household in the home

[8] In Adrian et al. (2020), corporate spreads in emerging economies are assumed to increase as their currency depreciates. The authors motivate this feature with a balance sheet channel resulting from large foreign-currency debt.

country works, consumes, and saves. It is endowed with traded goods in the first period and non-traded goods in both periods. It provides labor in two ways: it runs a bank and it works for the non-financial firm (in the second period only). The supply of labor to the non-financial firm L is inelastic. The household saves in the first period by transferring some exogenous amount N_0 of its tradable goods endowment as seed capital to its bank and by buying domestic bonds issued by a bank other than the one it owns. The household consumes the consumption basket C_t, which is a composite of non-tradable goods consumption $C_{NT,t}$ and tradable goods consumption $C_{T,t}$. The consumption index is of Cobb-Douglas form:

$$C_t = \left(C_{NT,t}^{\chi} C_{T,t}\right)^{\frac{1}{1+\chi}}, \tag{1}$$

where χ is a preference parameter. We consider a log utility function. The household's optimization problem is

$$\max_{B_0,(C_{NT,t},C_{T,t})_{t=0,1}} \ln C_0 + \beta \ln C_1 \quad \text{subject to (1) and}$$

$$P_0 C_0 + B_0 = p_0 Y_{T,0} - p_0 N_0 + Y_{NT,0}, \tag{2}$$

$$P_1 C_1 = R_1 B_0 + w_1 L + p_1 N_1 + Y_{NT,1}, \tag{3}$$

where P_t is the price index, $Y_{NT,t}$ and $Y_{T,0}$ are the endowments of the non-traded and the traded good, and B_0 are bond holdings. Note that while these bonds represent a claim on traded goods, they are expressed in terms of the numéraire, i.e., the domestic non-traded good. $p_t N_t$ are transfers to and from the household's bank and taken as given by the household. p_t is the price of the traded good, w_1 is the wage, and R_1 is the gross return on bond holdings, all measured in terms of the numéraire. The price index is defined as the minimum cost of obtaining one unit of the consumption basket. Thus, given the optimal choice of $C_{NT,t}$ and $C_{T,t}$, total consumption expenditure is

$$P_t C_t = C_{NT,t} + p_t C_{T,t}. \tag{4}$$

Solving the household's intertemporal problem yields the standard Euler condition

$$1 = R_1 \Lambda_{0,1}, \tag{5}$$

where $\Lambda_{0,1} = \beta \frac{P_0 C_0}{P_1 C_1}$ is the household's intertemporal marginal rate of substitution.

Solving the intratemporal problem yields the demand functions for tradables and non-tradables:

$$C_{NT} = \frac{\chi}{1+\chi} \left(\frac{1}{P}\right)^{-1} C, \tag{6}$$

$$C_T = \frac{1}{1+\chi} \left(\frac{p}{P}\right)^{-1} C. \tag{7}$$

The foreign household is modeled in an equivalent way. It is the owner of a foreign bank and faces an identical maximization problem as the home household. Foreign country-variables will be denoted with an asterisk (*). The foreign household receives the same amount of endowment and has the same intratemporal preferences as the home household. Yet, the households in the two countries differ in one respect. To induce a trade imbalance, and hence an excess supply of the foreign country's bonds, we assume that $\beta^* < \beta$, i.e., that the foreign household has a relatively higher discount rate. As a consequence, in an otherwise symmetric setup, in equilibrium the home country runs a trade surplus in the first period and, accordingly, there is an excess supply of the foreign country's bonds.

2.2 Non-financial Firms

The traded good is produced by perfectly competitive firms in the second period. The representative firm in the home country operates according to the following constant-returns-to-scale technology:

$$Y_{T,1} = K_1^{\alpha} L_1^{1-\alpha}. \tag{8}$$

Labor L_1 and capital K_1 are internationally immobile. In the first period, the firm can transform traded goods into capital and then use it for production in the second period. One unit of output invested raises capital by one unit. This process is reversible, so that a unit of capital, after having been used to produce output, can be retransformed into the tradable consumption good. The firm obtains the necessary funds for this investment by issuing investment securities $S_{p,0}$ at price q_0. One security finances one unit of capital, so we have $S_{p,0} = I_0 = K_1$ and $q_0 = p_0$. Given our assumption about the

capital transformation process, the price of capital is always equal to the price of output. From now on we will use p_t as the price per investment security and refer to K_1 as the total supply of investment securities.

The firm's first-order conditions are

$$Z_1 = \alpha \left(\frac{L_1}{K_1}\right)^{1-\alpha} p_1, \tag{9}$$

$$w_1 = (1-\alpha) \left(\frac{K_1}{L_1}\right)^{\alpha} p_1, \tag{10}$$

where Z_1 is the cost of capital to the firm, or the profit flow from a security financing one unit of capital to the holder of this security, measured in terms of the domestic numéraire. In period 1, after production, the firm is left with $(1-\delta)K_1$ units of capital, which represent the outstanding claims of the security holders. Therefore, the rate of return on one home investment security is

$$R_{k,1} = \frac{Z_1 + (1-\delta)p_1}{p_0}. \tag{11}$$

Foreign firms are modeled in an equivalent way, i.e., they face the same production technology as the home firms. By assumption, the law of one price holds for the traded good, so $p_0 = e_0 p_0^*$ and $p_1 = e_1 p_1^*$. e_t is the real exchange rate, defined as the price of the foreign numéraire (i.e., the foreign non-traded good) in terms of home numéraire (i.e., the home non-traded good).

2.3 Banks

Banks are at the core of our model. We assume financial markets are segmented, implying that non-financial agents cannot lend funds directly to each other. Banks act as the financial intermediaries in two types of financial markets: the investment market in each country and the international credit market. In the former, banks intermediate funds between households and firms and in the latter between the agents of the two countries by financing trade imbalances (later we will also consider imbalances from financial flows). Due to an agency problem between creditors and banks, however, this financial intermediation is imperfect.

The imperfect nature of financial intermediation will become visible in the form of spreads or excess returns in the two financial markets. In our—for the sake of simplicity—deterministic setup, these spreads reflect arbitrage opportunities. In practice, though, such excess returns might rather be related to risk. For instance, excess returns in the international credit market, i.e., deviations from uncovered interest parity, are the rule rather than the exception and there is broad agreement in the literature that these deviations do not necessarily imply arbitrage opportunities, but rather reflect a fair compensation for holding currency risk (for a review of the literature, see Engel 2014).

In our model, "banks" are meant to capture the large players in the global financial market like investment banks, currency hedge funds, active investment managers, and pension funds. As stressed in Gabaix and Maggiori (2015), these institutions have in common that they often bear the ultimate risk (which arises because household medium-term currency demand is unbalanced). In this sense, the spreads in our model can be interpreted as a compensation that such institutions demand for holding currency risk. In particular, we follow Gabaix and Maggiori (2015) and refer to the deviation from interest parity as a risk premium, even though in our model it does not stem from uncertainty and is therefore not a risk premium in the traditional sense.

Yet, "true" limits to arbitrage that can cause deviations from covered interest parity are not unrealistic neither, even in the huge foreign exchange market where vast amounts of capital are around.[9] Even the biggest players can face financial constraints, depending on their risk-bearing capacities and their existing balance sheet risks. Moreover, as noted by Shleifer and Vishny (1997), textbook-style arbitrage, i.e., involving neither capital nor risk, hardly exists in practice. Thus, in the context of the international credit market, the spreads in our model can be interpreted in either way: deviations from the covered or the uncovered interest parity.

[9]The failure of covered interest parity is, for example, documented and discussed by Borio et al. (2016) and Du, Tepper, and Verdelhan (2018); deviations from covered interest parity were particularly large during the global financial crisis.

Figure 1. Banking Sector: Period-0 Flows

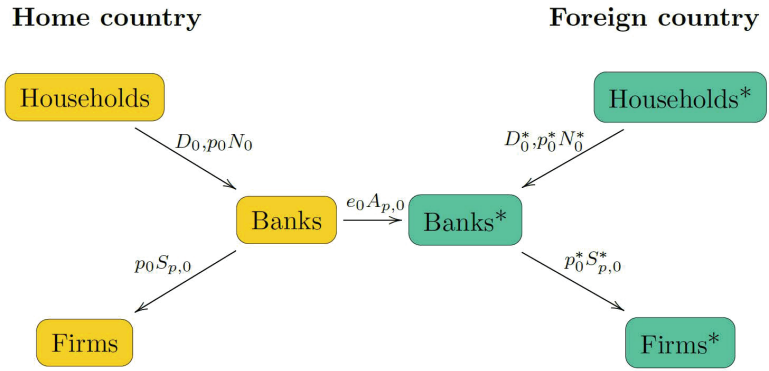

Figure 1 provides an overview of the first-period (period-0) flows in the banking sector. In each country, there is a continuum of identical banks that have unit mass. The representative home and foreign banks are modeled in a similar way. Differences between the two result from their interaction in the international credit market. This interaction in turn is a result of the trade flows (and later also portfolio flows) between the two countries. Since the domestic country runs a trade surplus in the first period, there is an excess supply of foreign bonds, i.e., claims on traded goods issued in the foreign country and denominated in the foreign numéraire. More precisely, a trade surplus in the home country and, hence, a trade deficit in the foreign country imply that foreign households consume relatively more, bringing less of their endowment of traded goods to the foreign bank to save ($D_0 > D_0^*$). As a consequence, the foreign bank has a relatively smaller amount of traded goods available to invest in local firms. To increase these investments it issues foreign bonds to the home bank and receives traded goods in exchange.

Let us now take a more detailed look at the home bank. Its balance sheet is

$$p_0 S_{p,0} + e_0 A_{p,0} = D_0 + p_0 N_0. \tag{12}$$

In addition to obtaining funds from households by issuing domestic bonds D_0, the bank receives seed capital from its owner, $N_0 > 0$. This seed capital represents the bank's net worth in the first period.

The bank uses these funds, first, to invest in domestic firms by buying home investment securities to the amount of $p_0 S_{p,0}$ in the domestic investment market. Second, it purchases foreign bonds, issued by the foreign bank, to the amount of $e_0 A_{p,0}$ in the international credit market. All variables on the home bank's balance sheet are measured in terms of the domestic numéraire, except for $A_{p,0}$, which is measured in terms of the foreign numéraire.

The interaction between the bank and the firm, and hence the transfer of funds from banks to firms, is frictionless. However, this is not the case for the interaction between the bank and its creditors, where financial frictions lead to imperfect financial intermediation. These frictions are the key feature of our model. To generate them endogenously, we introduce a limited commitment problem. Following Gertler and Karadi (2013), we assume that in period 0, after taking positions, the bank can choose to divert a certain fraction of its assets and transfer the proceeds to its owner. In particular, the bank can divert a fraction θ of its holdings of investment securities $p_0 S_{p,0}$ and a fraction Δ of its foreign bond holdings $e_0 A_{p,0}$.[10] The creditors can react by forcing the bank into bankruptcy and claiming the remaining fraction of assets. Like Gertler and Karadi (2013), we allow the friction parameters in the two markets to differ $\left(\Delta \lesseqgtr \theta\right)$. For instance, when $\Delta > \theta$ the bank can more easily divert foreign bonds than investment securities, and vice versa. The intuition behind this is that the performance of some assets in the bank's portfolio may be less transparent for creditors and therefore an easier target for diversion.

Creditors anticipate the possibility of diversion and limit the amount of funds that they lend to the bank. Hence, the bank can only issue domestic bonds as long as it has no incentive to misbehave, i.e., as long as its discounted profit V_0, which is given up under diversion, is greater than or equal to the gain from diversion. Thus, the bank's incentive constraint (IC) is

$$\text{IC:} \qquad V_0 \geq \theta p_0 S_{p,0} + \Delta e_0 A_{p,0}, \tag{13}$$

$$\text{where} \qquad V_0 = \Lambda_{0,1} \left(R_{k,1} p_0 S_{p,0} + R_1^* e_1 A_{p,0} - R_1 D_0 \right). \tag{14}$$

[10] Note that this setup of the banking sector is valid only for a non-negative excess supply of foreign bonds.

$R_{k,1}$, R_1^*, and R_1 are the gross returns on home investment securities, foreign bonds, and domestic bonds, respectively. Using balance sheet equation (12), we can rewrite the discounted profit in terms of excess returns $(R_{k,1} - R_1)$ and $\left(R_1^* \frac{e_1}{e_0} - R_1\right)$:

$$V_0 = \Lambda_{0,1} \left((R_{k,1} - R_1)\, p_0 S_{p,0} + \left(R_1^* \frac{e_1}{e_0} - R_1 \right) e_0 A_{p,0} + R_1 p_0 N_0 \right). \tag{15}$$

The optimization problem of the bank is

$$\max_{S_{p,0}, A_{p,0}} V_0 \qquad \text{subject to (13) and (15)}$$

and the first-order conditions are (λ is the Lagrange multiplier attached to the incentive constraint):

$$\Lambda_{0,1} (R_{k,1} - R_1) = \frac{\lambda}{1+\lambda} \theta, \tag{16}$$

$$\Lambda_{0,1} \left(R_1^* \frac{e_1}{e_0} - R_1 \right) = \frac{\lambda}{1+\lambda} \Delta. \tag{17}$$

If the financial friction parameters (θ, Δ) are small, we are in a frictionless environment. The incentive constraint is not binding and $\lambda = 0$ because the divertable part of the bank's assets is lower than the equity capital it would lose in case of misbehavior. Banks acquire assets up to the point where no arbitrage possibilities are left and excess returns are zero. Firms can borrow at the home interest rate, $R_{k,1} = R_1$, and the interest parity holds, $R_1 = R_1^* \frac{e_1}{e_0}$.

If θ and/or Δ are above a certain threshold, however, the incentive constraint is binding and $\lambda > 0$. We are in a model environment of financial frictions that become visible in the form of positive excess returns in the respective market(s). Compared with the frictionless equilibrium, there is less financial intermediation. The bank would like to borrow more funds and invest them in the respective market(s) to earn the excess returns, but creditors are unwilling to provide them. These limits to arbitrage lead to higher returns on securities in the home investment market, $R_{k,1} > R_1$, and to a deviation from interest parity in the international credit market,

$R_1^* \frac{e_1}{e_0} > R_1$. As explained above, the latter can be interpreted as a premium that the home bank requires in order to be willing to absorb the imbalance in the international credit market and not divert any of its assets. Moreover, the size of this spread reflects the risk of shifting funds from the home to the foreign country. In this sense, the exchange rate change between periods 0 and 1 incorporates a risk premium on foreign bonds (or, equivalently, a safety premium on domestic bonds).

Note that Equations (16) and (17) imply the following no-arbitrage relation:

$$(R_{k,1} - R_1) = \frac{\theta}{\Delta}\left(R_1^* \frac{e_1}{e_0} - R_1\right). \tag{18}$$

The households' willingness to lend, and hence the size of the bank's portfolio, depends not only on the fractions that the bank can divert but also on the size of the bank's equity capital. The limited commitment of the bank generates an endogenous capital constraint (CC) (for the derivation, see Appendix B):

$$\text{CC} = \begin{cases} \frac{\Delta \Lambda_{0,1} R_1}{\Delta - \Lambda_{0,1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \theta p_0 S_{p,0} + \Delta e_0 A_{p,0} & \text{if} \quad \theta \geq 0, \Delta > 0 \\ \frac{\theta \Lambda_{0,1} R_1}{\theta - \Lambda_{0,1}(R_{k,1} - R_1)} p_0 N_0 \geq \theta p_0 S_{p,0} + \Delta e_0 A_{p,0} & \text{if} \quad \theta > 0, \Delta \geq 0 \\ \text{no CC} & \text{if} \quad \theta = 0, \Delta = 0. \end{cases} \tag{19}$$

The left-hand side gives the bank's net worth multiplied by some weight. The right-hand side depicts a measure of the bank's portfolio. The weights θ and Δ represent the weaker (stronger) limits to arbitrage in the investment market if $\theta < \Delta$ ($\theta > \Delta$), which allows the bank to acquire a relatively higher (lower) share of investment securities compared with foreign bonds in its portfolio. The higher N_0, the more assets the bank can buy. Capital constraint (19) reveals that a high Δ stands for a low ability of the bank to intermediate international funds, which, viewed from the aggregate perspective, implies a disruption in the international credit market. Likewise, a high θ means that the bank is limited in its capacity to intermediate investment funds, reflecting a disruption in the home investment market. In this respect, we can also interpret the divertable fractions θ and Δ as a measure of the risk-bearing capacity of the banks. The

higher the fractions are, the lower is the banks' risk-bearing capacity and the lower the funds that they can intermediate in the respective markets.

The setup of the foreign bank is similar to the domestic bank with some differences in the balance sheet, reflecting the interaction of the two in the international credit market. Thus, the balance sheet of the foreign bank is

$$p_0^* S_{p,0}^* = D_0^* + A_{p,0} + p_0^* N_0^*. \tag{20}$$

It can divert a fraction θ^* of its holdings of investment securities $p_0^* S_{p,0}^*$. Accordingly, the foreign bank's optimization problem is

$$\max_{S_{p,0}^*} V_0^* \qquad \text{subject to}$$

$$V_0^* \geq \theta^* p_0^* S_{p,0}^* \quad \text{and} \tag{21}$$

$$V_0^* = \Lambda_{0,1}^* \left((R_{k,1}^* - R_1^*) p_0^* S_{p,0}^* + R_1^* p_0^* N_0^* \right), \tag{22}$$

which yields the following first-order condition:

$$\Lambda_{0,1}^* \left(R_{k,1}^* - R_1^* \right) = \frac{\lambda^*}{1+\lambda^*} \theta^*. \tag{23}$$

The restriction on the foreign bank's portfolio, i.e., the foreign endogenous capital constraint, is

$$\text{CC}^* = \begin{cases} \frac{1}{\theta^* - \frac{1}{R_1^*}(R_{k,1}^* - R_1^*)} p_0^* N_0^* \geq p_0^* S_{p,0}^* & \text{if} \quad \theta^* > 0 \\ \text{no CC}^* & \text{if} \quad \theta^* = 0. \end{cases} \tag{24}$$

2.4 Market Clearing

To close the model, we require the markets for assets, labor, and goods to clear. Therefore, the home and foreign capital markets as well as the markets for home and foreign bonds are characterized by

$$S_{p,0} = K_1; \qquad S_{p,0}^* = K_1^* \quad \text{and} \tag{25}$$

$$B_0 = D_0; \qquad B_0^* = D_0^*. \tag{26}$$

Remember that K_1 and K_1^* are the total supplies of domestic and foreign investment securities. In the labor market, labor demand in

each country needs to equal the inelastic labor supply: $L_1 = L$ and $L_1^* = L^*$. In the goods markets, market clearing for traded goods requires that

$$Y_{T,0} + Y_{T,0}^* = C_{T,0} + C_{T,0}^* + K_1 + K_1^* \quad \text{and} \tag{27}$$

$$Y_{T,1} + Y_{T,1}^* + (1-\delta)K_1 + (1-\delta)K_1^* = C_{T,1} + C_{T,1}^*. \tag{28}$$

Finally, for simplicity we assume that the endowment of non-traded goods is constant across countries and time: $Y_{NT,t} = Y_{NT,t}^* = \chi$. Hence, it must hold that $C_{NT,t} = C_{NT,t}^* = \chi$.

Note that combining the budget constraints of the home households and the home banks in period 0 and using the market clearing condition for domestic bonds, domestic investment securities, and non-traded goods yields the market clearing equation of the international credit market:

$$e_0 A_{p,0} = p_0 \underbrace{(Y_{T,0} - K_1 - C_{T,0})}_{\equiv NX_0}. \tag{29}$$

The right-hand side is equal to p_0 times net exports of the home country in the first period, NX_0. Equilibrium in the international credit market requires that the excess supply of foreign bonds in the amount of $p_0 NX_0$ is fully absorbed by home banks and therefore equal to their demand for foreign bonds $e_0 A_{p,0}$.

A summary of the equilibrium conditions is provided in Appendix B.

3. The Role of Financial Frictions

Even though this is a parsimonious model, a closed-form solution only exists for the frictionless case. With non-zero financial frictions, it is not possible to solve the model analytically. In this section, we use graphical and numerical illustrations to visualize the implications of the different frictions and to obtain an intuition about the mechanisms at work.

Before starting, note that putting the market clearing condition for the non-tradable good into the demand equation (6) yields the general result that total consumption expenditure is constant over time ($P_0 C_0 = P_1 C_1$). From this, it follows that the rate of return on

domestic bonds must always satisfy $R_1 = 1/\beta$ (see Equation (5)). Likewise, the rate of return on foreign bonds is $R_1^* = 1/\beta^*$. It follows that any movements in excess returns must come from a change in $R_{k,1}$, $R_{k,1}^*$, or $\frac{e_1}{e_0}$, respectively.

3.1 Effect of Financial Frictions in the International Credit Market

Starting from the frictionless case, consider first the effect of financial frictions in the international credit market, captured by an increase in Δ, with θ and θ^* set to zero. When Δ is sufficiently large for the domestic capital constraint to become binding ($\lambda > 0$), home banks invest less funds in foreign bonds, which, as explained earlier, results in a deviation from interest parity: $R_1^* \frac{e_1}{e_0} - R_1 > 0$ (see Equation (17)). Furthermore, the level of the home country's net exports is limited as follows (plug (29) into (19)):[11]

$$\frac{1}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq p_0 N X_0. \tag{30}$$

Graphically, the global equilibrium can be illustrated by the Metzler diagram in Figure 2. It depicts how first-period savings ($S_0 = Y_{T,0} - C_{T,0}$ and $S_0^* = Y_{T,0}^* - C_{T,0}^*$) and investment ($I_0 = K_1$ and $I_0^* = K_1^*$) schedules[12] change with real interest rates[13] ($\frac{p_0}{p_1} R_1$ and $\frac{p_0^*}{p_1^*} R_1^* = \frac{e_1}{e_0} \frac{p_0}{p_1} R_1^*$). Starting from the frictionless equilibrium, the consequences of financial frictions in the international credit market are twofold. First, due to the limits to arbitrage, the equilibrium real rate of return is higher in the foreign country than in the home country ($\frac{e_1}{e_0} \frac{p_0}{p_1} R_1^* > \frac{p_0}{p_1} R_1$). Second, there is a slight leftward shift of the home country's savings curve as home households increase their first-period consumption given that they can expect a positive return on the home banks' equity capital they hold: the excess return on foreign bonds is positive and the home banks will make positive profits. Overall, compared with the frictionless case, a larger fraction

[11] As $\lambda > 0$, inequality (30) will hold with equality.

[12] For a formal definition of the savings (SS, SS^*) and investment (KK, KK^*) schedules, see Appendix E.

[13] We use the term *real* to mean in terms of the traded good.

Figure 2. Financial Frictions in the International Credit Market: $\Delta > 0$

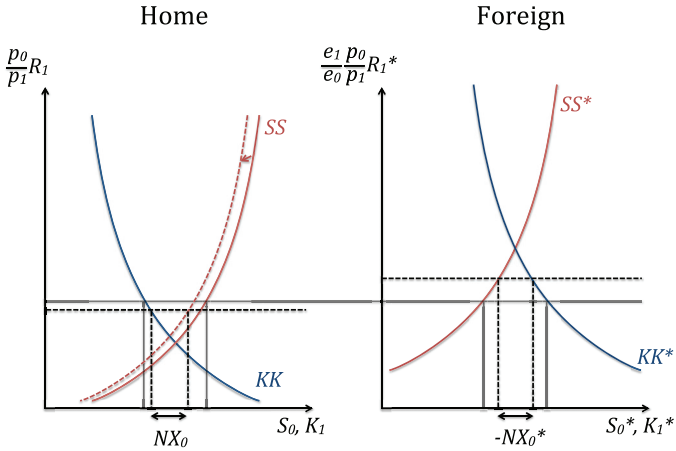

Note: The solid lines represent the frictionless equilibrium, the dashed lines the equilibrium with frictions.

of the home country's first-period endowment is either consumed in period 0 or invested domestically, making its net exports shrink. Frictions in the international credit market thus make intertemporal trade more costly: they act like a tax on capital flows and make the current account shrink. The international mobility of funds is limited, leading to a misallocation of capital: the majority of capital is invested in the home country, whereas in the frictionless case, identical production technologies and equal labor forces imply that both countries invest the same amount.

In this scenario, the home country, i.e., the country facing excess demand for its bonds, experiences an appreciation in period 0, while at the same time the foreign country, i.e., the country facing an excess supply of bonds, experiences a depreciation. There are two mechanisms driving these exchange rate movements. First and foremost, the frictions in the international credit market cause a deviation from interest parity (i.e., a safety premium on domestic bonds), which incorporates a home appreciation in the first and a depreciation in the second period. Moreover, there is an increase in the home country's relative lifetime resources coming from the change in

the allocation of capital. This implies that, relative to the frictionless case, the home country's second-period output increases while the foreign country's second-period output decreases. This change in fundamentals induces further appreciation pressure in the first period (and offsets part of the depreciation in the second period).

Figure 3 provides a numerical illustration of these results. With θ and θ^* set to zero and using the parameterization in Table 1 to calibrate the remaining parameters, it shows the evolution of the model's equilibrium as Δ increases. For small values, the home banks' incentive constraint is not binding. However, as soon as Δ is above a certain threshold, the equilibrium adjusts as described in the graphical analysis above. Any further increase in Δ amplifies these effects, i.e., it leads to further increases in the deviation from interest parity and an additional home appreciation (depreciation) in the first (second) period.

3.2 Effect of Financial Frictions in the Foreign Investment Market

Next, consider the effect of financial frictions in the foreign investment market, captured by an increase in θ^*, with θ and Δ set to zero. When θ^* is sufficiently large for the foreign banks' capital constraint to become binding ($\lambda^* > 0$), foreign banks are unable to exploit all arbitrage opportunities. Excess returns in the foreign investment market thus become positive: $R^*_{k,1} - R^*_1 > 0$ (see Equation (23)). The level of foreign capital is limited by (plug Equation (25) into CC* (24)):[14]

$$\frac{1}{\theta^* - \frac{1}{R^*_1}\left(R^*_{k,1} - R^*_1\right)} p^*_0 N^*_0 \geq p^*_0 K^*_1. \tag{31}$$

In contrast, excess returns in both the home investment market and the international credit market remain zero (see Equations (16) and (17)).

[14] As $\lambda^* > 0$, inquality (31) will hold with equality.

Figure 3. Effect of an Increase in International Credit Market Frictions Δ (Δ on x-axis)

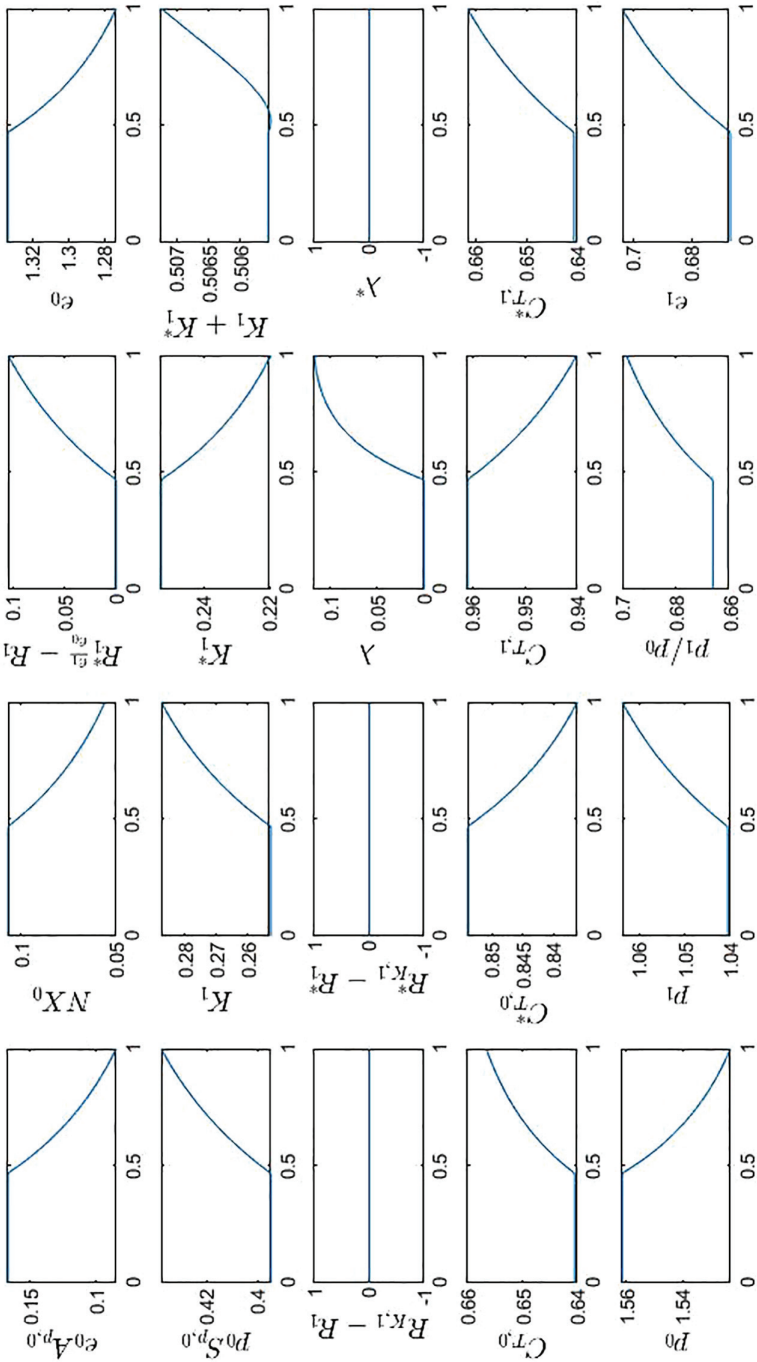

Note: Evolution of the model's equilibrium as the friction parameter in the international credit market increases, starting from a frictionless point. $\theta = 0$, $\theta^* = 0$, $\Delta \geq 0$.

Table 1. Parameterization

β	1	δ	0.33
β^*	0.5	$Y_{T,0}, Y^*_{T,0}$	1
$\frac{1}{1+\chi}$	0.5	α	0.33
L, L^*	1	N_0, N^*_0	0.05

Figure 4. Financial Frictions in the Foreign Investment Market: $\theta^* > 0$

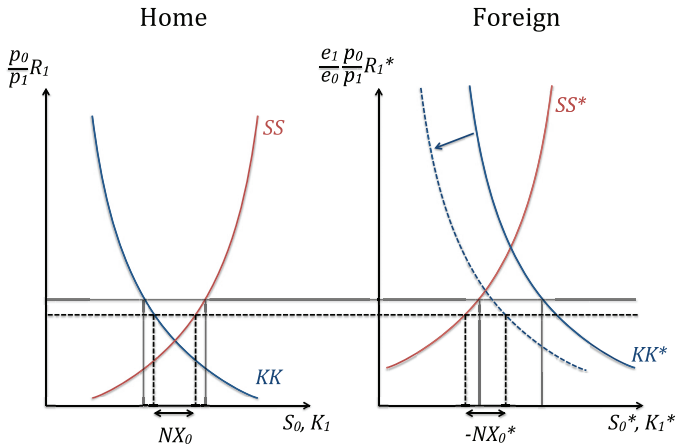

Note: The solid lines represent the frictionless equilibrium, the dashed lines the equilibrium with frictions.

In a Metzler diagram, financial frictions in the foreign investment market shift the foreign investment curve to the left (see Figure 4).[15] Intuitively, for a given real rate of return $\frac{e_1}{e_0}\frac{p_0}{p_1}R_1^*$, investment in the foreign country declines compared with the frictionless case, because foreign banks are now less able to intermediate funds in this market. Costs of capital in the foreign market thus increase. In order to

[15]As opposed to the case of an increase in Δ, an increase in θ or θ^* does not lead to a shift in the savings curve of the respective country, as the higher return on the households' holdings of equity capital (in this case due to the increase in the spreads $R_{k,1} - R_1$ and $R^*_{k,1} - R^*_1$, respectively) is just nullified by a decrease in their second-period labor income due to the lower level of capital. For the formal proofs, see Appendix E.

maintain the global equilibrium, the equilibrium real rate of return has to decrease.

The frictions in the foreign investment market affect savings and investment in both countries. The home country lowers its savings and invests a larger share of its remaining savings domestically, which causes its net exports to decrease. The foreign country also lowers its savings, but since it lowers investment even more, its net imports are reduced. Overall, the foreign investment market frictions lead to a decrease in global savings and, consequently, global investment, implying a lower level of global output in the second period. Furthermore, there is a misallocation of capital in the first period. Investment is not equal in the two countries as it would be in the frictionless case; instead, a majority is invested in the home country. From this change in the investment allocation it follows that, compared with the frictionless case, the second-period output is now higher in the home country but lower in the foreign country, implying a change in the two countries' fundamentals. The relative increase in the home country's lifetime resources induces a home appreciation (foreign depreciation) in both periods.[16]

The charts in Figure 5 provide a numerical illustration of the effects that have just been described. Setting θ and Δ equal to zero, it shows how the model's equilibrium evolves as θ^* increases. For small values, the foreign banks' incentive constraint is not binding yet. However, as soon as θ^* is above a certain threshold, the equilibrium adjusts as described in the graphical analysis above. Any further increase in θ^* has the same effects, i.e., it leads to another

[16]The home country's lifetime resources are apparent in the home households' intertemporal budget constraint:

$$C_{NT,0} + p_0 C_{T,0} + \frac{1}{R_1}\left(C_{NT,1} + p_1 C_{T,1}\right)$$
$$= Y_{NT,0} + p_0 Y_{T,0} + \frac{1}{R_1}(Y_{NT,1} + w_1 L + p_1 N_1) - p_0 N_0.$$

As the endowment of non-traded goods and period-0 traded goods is given and the law of one price holds, any relative changes in lifetime resources between the two countries must either come from relative changes in labor income $w_1 L = (1-\alpha) p_1 Y_{T,1}$ or changes in the relative payoffs to equity capital $p_1 N_1 = (R_{k,1} - R_1)\, p_0 S_{p,0} + \left(R_1^* \frac{e_1}{e_0} - R_1\right) e_0 A_{p,0} + R_1 p_0 N_0$.

Figure 5. Effect of an Increase in Foreign Financial Frictions θ^* (θ^*on x-axis)

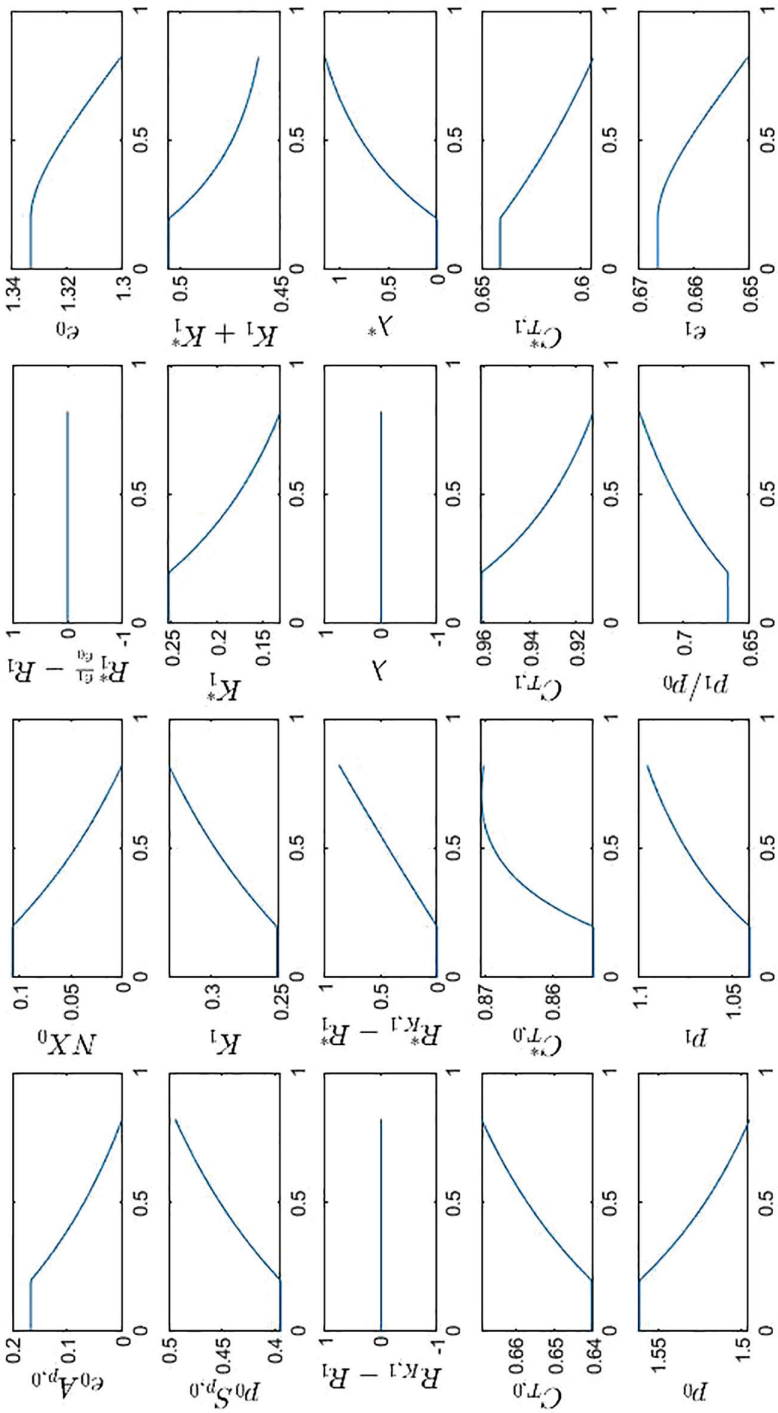

Note: Evolution of the model's equilibrium as the friction parameter in the foreign investment market increases, starting from a frictionless point. $\theta = 0$, $\theta^* \geq 0$, $\Delta = 0$. Given that the setup of the banking sector is only valid with $e_0 A_{p,0} \geq 0$, the plots only cover a limited range of possible values for θ^*.

increase in excess returns in the foreign investment market and an additional home appreciation.

The case of financial frictions in the home investment market, captured by an increase in θ, is symmetric to the one just described. When θ is sufficiently large for the home incentive constraint to become binding, the home investment curve shifts to the left, resulting in a home depreciation (foreign appreciation) in both periods. A detailed analysis is provided in Appendix B.

3.3 General Case

After studying the effects of financial frictions for different markets one at a time, we now turn to a description of the general model in which all friction parameters are positive and banks in both countries face binding incentive constraints. In this case, excess returns are positive in all three markets, i.e., there is a deviation from interest parity and home as well as foreign firms face capital costs above the frictionless level. The banks' capital constraints (see Equations (19) and (24)), combined with market clearing in both the investment markets and in the international credit market (Equations (25) and (29)) describe the restrictions on capital and net exports:[17]

$$\frac{1}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \frac{\theta}{\Delta} p_0 K_1 + p_0 N X_0, \tag{32}$$

$$\frac{1}{\theta^* - \frac{1}{R_1^*}\left(R_{k,1}^* - R_1^*\right)} p_0^* N_0^* \geq p_0^* K_1^*. \tag{33}$$

In line with our earlier reasoning, Equation (32) shows that, when limits to arbitrage are higher in the home investment market than in the international credit market (i.e., when $\theta > \Delta$), intermediating capital in the home country is more constraining than intermediating net exports.

In general, further increases in any of the friction parameters make the banks shift funds away from the respective market. Overall, the mechanisms are the same as those described in Sections 3.1 and 3.2, even though graphically (regarding the shifts in the savings

[17] As by assumption $\lambda > 0$ and $\lambda^* > 0$, both inequalities will hold with equality.

and investment curves of the Metzler diagrams) there can be small differences due to the non-zero Lagrange multipliers λ and λ^* and excess returns. Accordingly, a further increase in Δ leads to a shift of funds away from foreign assets and herewith home net exports and foreign capital towards home assets and home capital. Likewise, a further increase in θ^* leads to a decrease in the total amount of global investment $(K_1 + K_1^*)$ and a shift of funds away from foreign investment securities and thus foreign capital towards home investment securities and thus home capital. A numerical illustration of the corresponding effects is provided in Figures B.3 to B.5 in Appendix B.

Obviously, higher frictions in one of the markets lead to a direct increase in excess returns in this specific market. In contrast, the spreads in the other markets are only affected marginally. Higher frictions in the foreign investment market, for instance, make the home banks' incentive constraint slightly less binding because the foreign banks issue fewer foreign bonds and the supply of foreign assets thus decreases. Home banks thus now need to absorb fewer of these and therefore can invest more of their constrained funds in home investment securities. So, overall, the constraint on home banks is slightly eased and, accordingly, excess returns in both the home investment market and the international credit market decrease.

4. Impact of International Portfolio Flows

So far, we have assumed that households are only able to trade bonds denominated in the numéraire of their own country, and hence that the imbalance in the demand for domestic and foreign bonds that is absorbed by the home banks results from trade flows only. Now, we introduce financial flows other than those resulting from trade imbalances and refer to them as *international portfolio flows*. In particular, we allow households in both countries to hold bonds issued by the other country's banks. Home households are assumed to have an inelastic demand for foreign bonds f that is funded by an offsetting position $e_0 f$ in domestic bonds. Equivalently, foreign households have an exogenous inelastic demand for domestic bonds f^* that is funded by an offsetting position f^*/e_0 in foreign bonds. For simplicity, we set all these portfolio flows to be exogenous, which helps to

avoid mixing up different transmission channels. Hence, one could think of f and f^* as the result of simple noise or liquidity trading or as "deliberate" holdings of foreign bonds motivated, for example, by practical reasons in daily business or foreign direct investment.[18]

With portfolio flows, market clearing conditions (26) for home and foreign bonds change to

$$B_0 + f^* = D_0 + e_0 f; \qquad B_0^* + f = D_0^* + f^*/e_0. \tag{34}$$

Accordingly, the richer set of financial flows also alters equilibrium condition (29) in the international credit market. Home banks now need to absorb the imbalance in the demand for foreign bonds stemming from both trade and portfolio flows:

$$e_0 A_{p,0} = p_0 NX_0 + f^* - e_0 f, \tag{35}$$

where $p_0 NX_0 + f^*$ reflects the total supply of foreign bonds and $e_0 f$ is the domestic households' demand for foreign bonds. All else equal, the larger the inflows f^*, the larger the international funds that the home banks need to intermediate. Hence, as will become evident in the rest of the section, it is not necessarily net exports ($p_0 NX_0$) and thus net foreign assets that matter in the determination of the exchange rate, but rather the excess supply of foreign bonds that need to be absorbed by the private financial sector ($p_0 NX_0 + f^* - e_0 f$). In practice, the two are likely to have the same sign for most countries, given that the excess supply of foreign bonds depends positively on net foreign assets, but they can also have opposite signs. An extreme example is certainly the United States. While the U.S. net exports and current account are persistently negative, the demand for U.S. dollar assets by the rest of the world (in the model captured by f^*) is huge, among other things due to the U.S. dollar's role as a vehicle and reserve currency. Hence, even though

[18]In practice, a large part of the demand for foreign bonds is likely to be endogenous and depend on present and expected future fundamentals, such as interest rate differentials. For this endogenous part, it would be more realistic to have $f = f(R_1, R_1^*, e_0, e_1, \ldots)$ and $f^* = f^*(R_1, R_1^*, e_0, e_1, \ldots)$. For instance, as suggested in Gabaix and Maggiori (2015), a straightforward way to model a popular trading strategy, carry trade, would be to set $f = a + b(R_1 - R_1^*)$ and $f^* = c + d(R_1 - R_1^*)$ for some constants a, b, c, and d. However, given that our main goal is to improve understanding of the effects of international financial flows, but not the origin of these, we abstract from such dependencies.

U.S. net foreign assets are negative, the excess demand for U.S. dollars (and hence the excess supply of foreign currency that need to be absorbed by the private financial sector) is likely to be positive.[19]

When the international credit market frictions parameter Δ is equal to zero, home banks are able to absorb any imbalance in the international credit market. That is, interest parity holds and gross capital flows have no effect on any variable other than $A_{p,0}$. Consider the example of foreign households suddenly wanting to hold a certain amount of domestic bonds f^*. When banks are not constrained in the international credit market and the returns on domestic and foreign bonds are equalized, home banks are willing to issue any additional amount of domestic bonds and, in return, increase their holdings of foreign bonds correspondingly. Hence, home banks increase their holdings of foreign bonds $e_0 A_{p,0}$ one for one with the inflow of capital f^*. It is a trade that concerns only foreign households and home banks and does not affect the rest of the economy. This irrelevance of gross capital flows is a common feature of the traditional international economics literature inspired by Dornbusch (1976) and Obstfeld and Rogoff (1995), where interest parity (or more specifically, uncovered interest parity) is often either directly assumed to hold or imposed in the process of first-order linearization.

Gross capital flows start to matter once the international credit market friction parameter Δ is positive and the home banks' incentive constraint is, or starts to be, binding. When banks are credit constrained in the international credit market, gross capital flows have an impact on the tightness of the capital constraint they are facing:[20]

$$\mathrm{CC} = \begin{cases} \frac{\Delta}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \theta p_0 K_1 & \\ \quad + \Delta(p_0 N X_0 + f^* - e_0 f) & \text{if} \quad \theta \geq 0, \Delta > 0 \\ \frac{\theta}{\theta - \frac{1}{R_1}(R_{k,1} - R_1)} p_0 N_0 \geq \theta p_0 K_1 & \\ \quad + \Delta(p_0 N X_0 + f^* - e_0 f) & \text{if} \quad \theta > 0, \Delta \geq 0 \\ \text{no CC} & \text{if} \quad \theta = 0, \Delta = 0. \end{cases} \tag{36}$$

[19]The U.S. dollar's safe-haven property of appreciating during periods of high uncertainty is thus in line with the predictions of our model.

[20]The inequalities in (36) will hold with equality in this case given that banks are balance sheet constrained.

Note that the critical value of Δ at which the constraint becomes binding is endogenous and depends negatively (positively) on capital inflows f^* (outflows $e_0 f$). For instance, if f^* is very high, i.e., the excess supply of foreign bonds that the home banks need to absorb is large, a relatively low Δ suffices to make the banks' incentive constraint binding.

An increase in capital inflows to the home country or a decline in capital outflows make the home banks' capital constraint more binding, as they need a larger part of their risk-bearing capacity to intermediate these flows. Consider again the example of foreign households suddenly wanting to hold a certain amount f^* of home country bonds. If home banks are constrained in the international credit market, we are at a point at which creditors are not willing to provide them with more funds, since the banks would invest these (at least partially) in foreign bonds, which in turn would result in higher proceeds under diversion and hence a higher incentive to misbehave. Capital inflows from foreign households, however, represent an exogenous increase in the funds available to home banks, but as a consequence there is also a higher amount of foreign bonds that they need to absorb in order to maintain equilibrium in the international credit market (see Equation (35)). Due to the binding balance sheet constraint, home banks are only able to intermediate these exogenous capital inflows if they can simultaneously relax their capital constraint in some other way. This can happen through two channels. The first is an adjustment on the creditor side. Home households will find it optimal to increase their period-0 consumption and decrease their savings, i.e., provide home banks with less funds. Together with the concurrent decrease in net exports (and hence in the demand for domestic bonds), this leads to a relaxation of the capital constraint. The second potential channel is an adjustment in the home banks' portfolio towards home investment securities. The relative importance of these two channels can vary. The first channel always plays an important role, and if home banks are not constrained in the home investment market (i.e., $\theta = 0$), the second channel also becomes significant. If, however, there are frictions in both the international credit market and the home investment market, substituting foreign bonds (and thereby the intermediation of net exports) for additional home investment securities will not necessarily relax the balance sheet constraint. The higher the investment

market frictions, the less such a substitution yields the necessary loosening of the constraint, and hence the stronger the first channel, i.e., the more households will increase their period-0 consumption and reduce the amount of funds they provide the banks with.

In the case where only the international credit market friction parameter is positive, an increase in capital inflows has similar effects to an increase in Δ. While an increase in Δ represents a direct reduction in the banks' ability to intermediate international funds, higher capital inflows imply that a larger part of the banks' risk-bearing capacity is absorbed by these exogenous flows, representing an indirect reduction in their ability to intermediate international funds. Graphically, an increase in capital inflows implies a widening of the spread between the two dashed vertical lines in Figure 2, corresponding to a larger deviation from interest parity and a home appreciation (foreign depreciation) in period 0. Accordingly, the results in the numerical example correspond for the most part to those of an increase in Δ, as can be seen in Figure D.1 in Appendix D. Higher portfolio inflows increase the misallocation of capital that is potentially already present due to frictions in the international credit market. The higher the portfolio inflows that the home banks need to intermediate, the larger the amount of funds that they find optimal to invest in the home country. The resulting increase in the home country's second-period output relative to the foreign country raises the home country's relative lifetime resources, leading to further appreciation pressure in period 0.

Finally, Figure 6 provides a numerical illustration of the case in which both home and foreign investment markets exhibit limits to arbitrage. Note that since we set $\Delta = \theta$, excess returns in the international credit and the home investment market are equally large. Again, exogenous capital inflows trigger an increase in the excess return in the international credit market (i.e., an increase in the safety premium on domestic bonds) and thus an appreciation in period 0. Note that the levels of capital K_1 and K_1^* change only marginally. As the home banks are equally constrained in the home investment market and the international credit market, there is no portfolio adjustment. While net exports decline strongly, variations in the level of investment are of second order only. Hence, relative second-period output in the two countries changes but marginally, if at all. However, compared with foreign households, home households

Figure 6. Effect of an Increase in Capital Inflows f^* (f^* on x-axis)

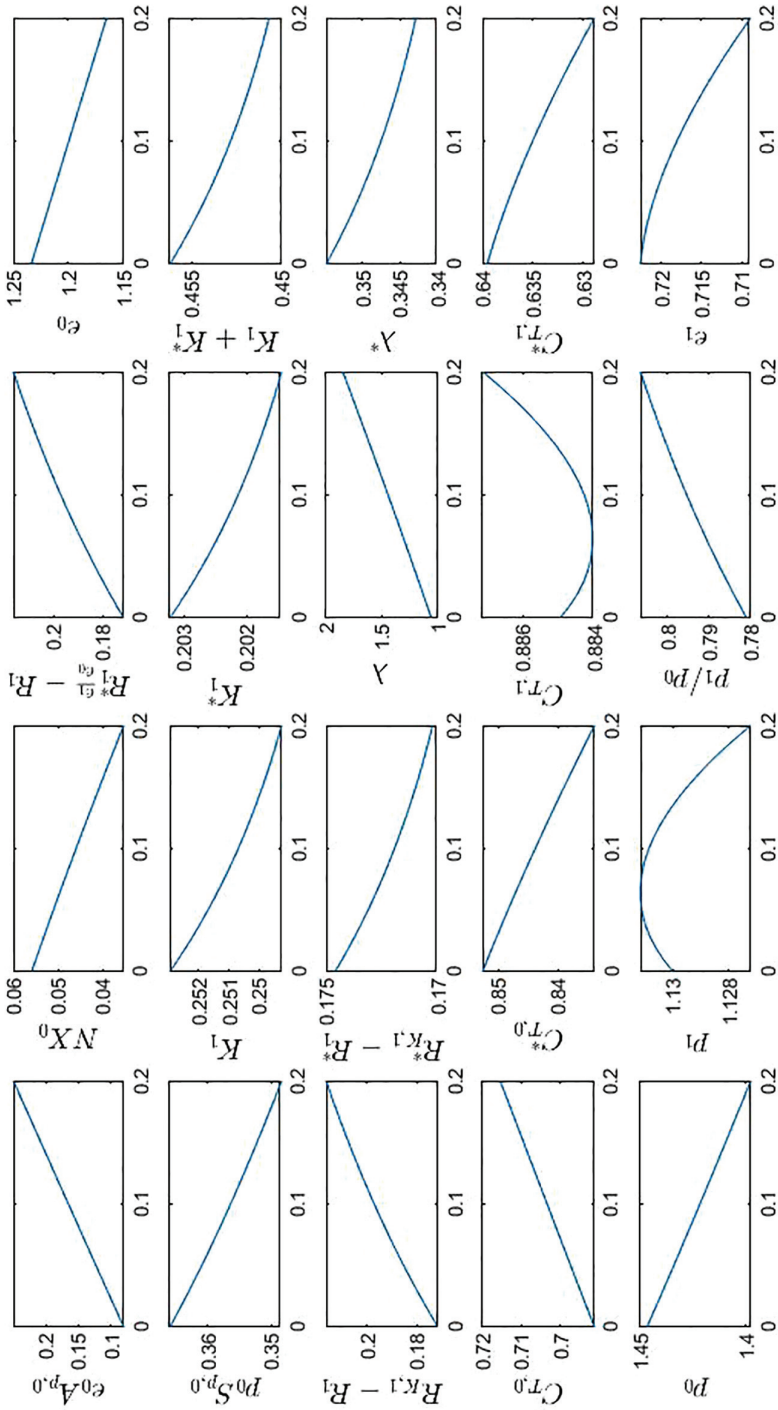

Note: Evolution of the model's equilibrium as capital inflows increase, starting from a point where there are limits to arbitrage in all financial markets. $\theta = 1/3$, $\theta^* = 1/3$, $\Delta = 1/3$.

will have a relatively higher payoff from their equity capital. The implicit decrease in the banks' ability to intermediate funds in the international credit market and in the investment market drives up excess returns in both of these markets (see Equation (18)). Hence, in this case too, there is an increase in the home country's relative lifetime resources, which induces a further appreciation in the first period.

Regarding the foreign country, i.e., the country facing an excess supply of its bonds, note that there an exogenous increase in portfolio inflows (captured by an increase in $e_0 f$) would also generate an appreciation, at least as long as the home banks' capital constraint is binding. Yet, the underlying reason differs across the two countries. In the case of the home country, the appreciation pressure resulting from higher capital inflows is explained by a tightening of the home banks' capital constraint and, therefore, a widening of the interest parity spread. The foreign appreciation in the case of higher capital inflows to the foreign country, on the other hand, stems from a relaxation of the home banks' capital constraint. The excess supply of foreign bonds declines, making the interest parity spread narrow. In other words, higher capital inflows to the foreign country bring the economy closer to its frictionless state and the exchange rate closer to its frictionless value, i.e., the value explained by economic fundamentals.

5. Central Bank Asset Purchases

The previous sections have shed light on how frictions in financial markets—as they are experienced during a financial crisis—and global imbalances in gross flows can lead to excess returns on domestic and foreign bonds and distortions in real activity, namely in the allocation of capital and in international trade. We now analyze a central bank's policy options to counteract such distortions. In this section, we show how large-scale asset purchases by the home central bank can be used to lower excess returns in general. In the next section, we then look at three specific cases of external appreciation pressure to the home country and identify the possible policy responses of the home central bank.

Our model allows us to study and compare two policies that were repeatedly put into practice in the course of the 2007–09

financial crisis: credit easing and foreign exchange (FX) interventions. By applying these policies, the home central bank itself plays the role of an intermediary, reducing the (excess) supply of investment securities and foreign bonds that need to be absorbed by the private intermediaries, and thus relaxes the banks' capital constraint. In our model, the two policy options are implemented as follows. The central bank intervenes in the domestic investment market by purchasing domestic investment securities $p_0 S_{CB,0}$ and in the international credit market by purchasing foreign bonds $e_0 A_{CB,0}$, issuing domestic bonds $D_{CB,0}$ to finance these transactions.[21] Following the baseline scenario of Gertler and Karadi (2013), we assume that the central bank issues these domestic bonds directly to households.[22] The central bank's profits in period 1 are transferred to the home households.

In contrast to private intermediaries, the central bank has the crucial advantage that it is not balance sheet constrained because it is not facing a limited commitment problem. We furthermore make the simplifying assumption that both types of interventions, i.e., FX interventions and credit easing, are costless to the central bank. This implies that it is as efficient as private intermediaries at intermediating funds. This assumption simplifies the analysis but is not critical for our results.[23]

As a result of the interventions, the central bank's balance sheet,

$$p_0 S_{CB,0} + e_0 A_{CB,0} = D_{CB,0}, \tag{37}$$

[21] Theoretically, we could also look at a policy where the home central bank directly acquires foreign investment securities instead of foreign bonds. However, this seems little realistic, as the foreign investment market tends to be the responsibility of the foreign authorities.

[22] As Gertler and Karadi (2013) point out, $D_{CB,0}$ can also be interpreted as reserves held by banks on account at the central bank.

[23] It would be straightforward to introduce relative efficiency costs, as in Gertler and Karadi (2013), to capture the fact that a central bank is less efficient at intermediating funds than ordinary banks. For welfare considerations (from which we abstract in this paper), central bank interventions would then only be desirable when private intermediation is significantly constrained and even then only if efficiency costs are not too large. The latter, however, is a reasonable assumption, so that in the end these costs would not change the qualitative implications of the model.

expands. Accordingly, the market clearing condition (26) for home bonds extends to $B_0 + f^* = D_{p,0} + D_{CB,0} + e_0 f$.

In the home investment market, the market clearing condition (25) changes to

$$S_{p,0} + S_{CB,0} = K_1, \tag{38}$$

which reflects the fact that capital is now partly intermediated by the central bank.

Finally, the consolidation of the home households' first-period budget constraint and the balance sheet equations of the home banks and the central bank yields the new market clearing condition of the international credit market:

$$e_0 A_{p,0} + e_0 A_{CB,0} = p_0 NX_0 + f^* - e_0 f. \tag{39}$$

The excess supply of foreign bonds, determined by the home country's trade surplus as well as the portfolio flows of home and foreign households, is now jointly absorbed by the home banks and central bank.

As the home central bank only intervenes in the domestic investment market and the international credit market, only the home banks' incentive constraint is relevant for determining whether the interventions are effective. When the constraint is not binding, returns in the home investment market and the international credit market are determined by frictionless arbitrage and interventions by the central bank are neutral. Its purchases of either home investment securities or foreign bonds simply replace part of the private intermediation, but have no effect on returns and the exchange rate. However, central bank interventions are non-neutral in markets where the financial friction parameters are high enough to generate limits to arbitrage. When banks are balance sheet constrained in one or both markets, central bank purchases of the respective assets do not just replace private intermediation one for one, but rather expand the total demand for the respective asset type, which in turn drives down the excess return(s).

The precise effect on different excess returns depends on whether the balance sheet constraint is binding in just one or in both markets. In the former case, central bank interventions obviously only

have an effect on excess returns in the market that exhibits limits to arbitrage while the other market is unaffected. In the latter case, purchases of either asset affect excess returns in both markets. This spillover effect results from the no-arbitrage relation (18). For a better intuition, consider the example where the central bank intervenes by purchasing foreign bonds. According to the reasoning above, this reduces excess returns in the international credit market. Home banks then shift part of their funds towards the home investment market where excess returns are still high and therefore more attractive. This in turn also reduces the excess return on investment securities. The banks' portfolio adjustment ends once excess returns in the two markets, adjusted by the weight $\frac{\theta}{\Delta}$, are equalized. Note that, no matter which of the two assets the central bank purchases, its interventions change excess returns most (in absolute terms) in the market where banks are most constrained.

The additional intermediation by the central bank allows for higher levels of home capital and net exports to be intermediated, as can be observed in the new capital constraint:[24]

$$\mathrm{CC} = \begin{cases} \frac{\Delta}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \theta p_0 (K_1 - S_{CB,0}) & \\ \quad + \Delta(p_0 N X_0 + f^* - e_0 f - e_0 A_{CB,0}) & \text{if} \quad \theta \geq 0, \Delta > 0 \\ \frac{\theta}{\theta - \frac{1}{R_1}(R_{k,1} - R_1)} p_0 N_0 \geq \theta p_0 (K_1 - S_{CB,0}) & \\ \quad + \Delta(p_0 N X_0 + f^* - e_0 f - e_0 A_{CB,0}) & \text{if} \quad \theta > 0, \Delta \geq 0 \\ \text{no CC} & \text{if} \quad \theta = 0, \Delta = 0. \end{cases} \tag{40}$$

Central bank asset purchases $p_0 S_{CB,0}$ and $e_0 A_{CB,0}$ reduce the amount of funds that need to be intermediated by home banks, which relaxes their capital constraint and brings the intermediated quantities closer to their frictionless level. In the limit, intermediation by the home central bank can make the excess returns in the home investment market and the international credit market disappear completely. As long as θ^* is small enough for the foreign banks'

[24]The inequalities in (40) hold with equality given that we are considering the case where banks are constrained.

capital constraints not to be binding, the economy would then be back in the frictionless state.[25]

Furthermore, Equation (40) reveals that in terms of the amount of intervention needed to reach a given reduction of excess returns, it matters which asset the central bank buys. If both markets are affected by limits to arbitrage, buying a certain amount of investment securities $p_0 S_{CB,0}$ relaxes the constraint to the same extent, and therefore has exactly the same effect, as buying foreign assets $e_0 A_{CB,0}$ to the amount of $\frac{\theta}{\Delta} p_0 S_{CB,0}$. Intuitively, a central bank intervention involving the issuance of a given amount of domestic bonds frees up a higher amount of bank capital if purchases are made in the market that faces higher limits to arbitrage. This implies that when $\Delta > \theta$, i.e., when the international credit market exhibits higher excess returns than the home investment market, FX interventions have a stronger effect than credit easing and are preferable to the latter in order to avoid an unnecessary expansion of the central bank's balance sheet.[26] Likewise, credit easing is the preferred instrument when $\theta > \Delta$.

An overview of the model's formal equilibrium conditions with portfolio flows and central bank interventions is provided in Appendix D. Here, Figures 7 and 8 provide numerical illustrations of the effects of credit easing and FX interventions, respectively, when banks are equally constrained in both markets ($\theta = \Delta$) and $\theta^* = 0$.[27] As just described, in this case both credit easing and FX interventions reduce excess returns and raise the level of intermediated funds in both markets. However, to make excess returns in both markets disappear completely and return the economy to the frictionless equilibrium, it may not be enough to intervene in just one market. The reason is that the amount of assets outstanding in

[25]Remember that, realistically, we would have to introduce efficiency costs to central bank intermediation, which in turn could make such extreme interventions less desirable.

[26]The larger a central bank's balance sheet, the larger the size of potential losses due to valuation changes on its assets (e.g., because of exchange rate fluctuations). Substantial reductions in a central bank's net worth may be considered problematic in the public and—in the worst case—lead to a loss in confidence. However, as pointed out by Jordan (2011), low or even temporarily negative levels of equity do not constrain a central bank's capacity to act.

[27]Numerical illustrations for when $\Delta = \theta = \theta^* > 0$ are provided in Figures D.2 and D.3 in Appendix D.

Figure 7. Effect of an Increase in $S_{CB,0}$ ($p_0 S_{CB,0}$ on x-axis)

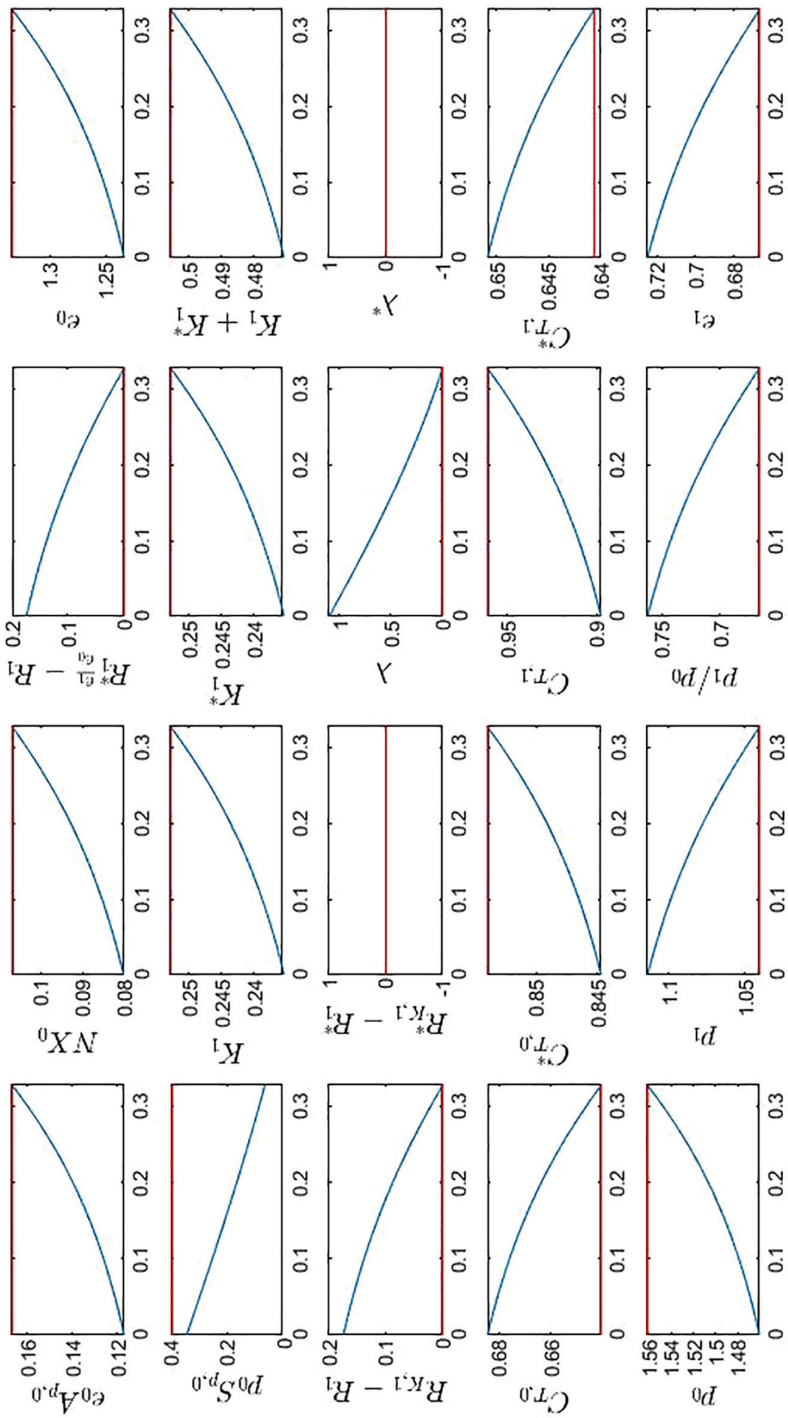

Note: Evolution of the model's equilibrium as central bank intermediation in the domestic investment market increases, starting from a point where there are limits to arbitrage in the domestic investment market and the international credit market. $\theta = 1/3$, $\theta^* = 0$, $\Delta = 1/3$. The red line represents the value in the frictionless equilibrium.

Figure 8. Effect of an Increase in $A_{CB,0}$ ($e_0 A_{CB,0}$ on x-axis)

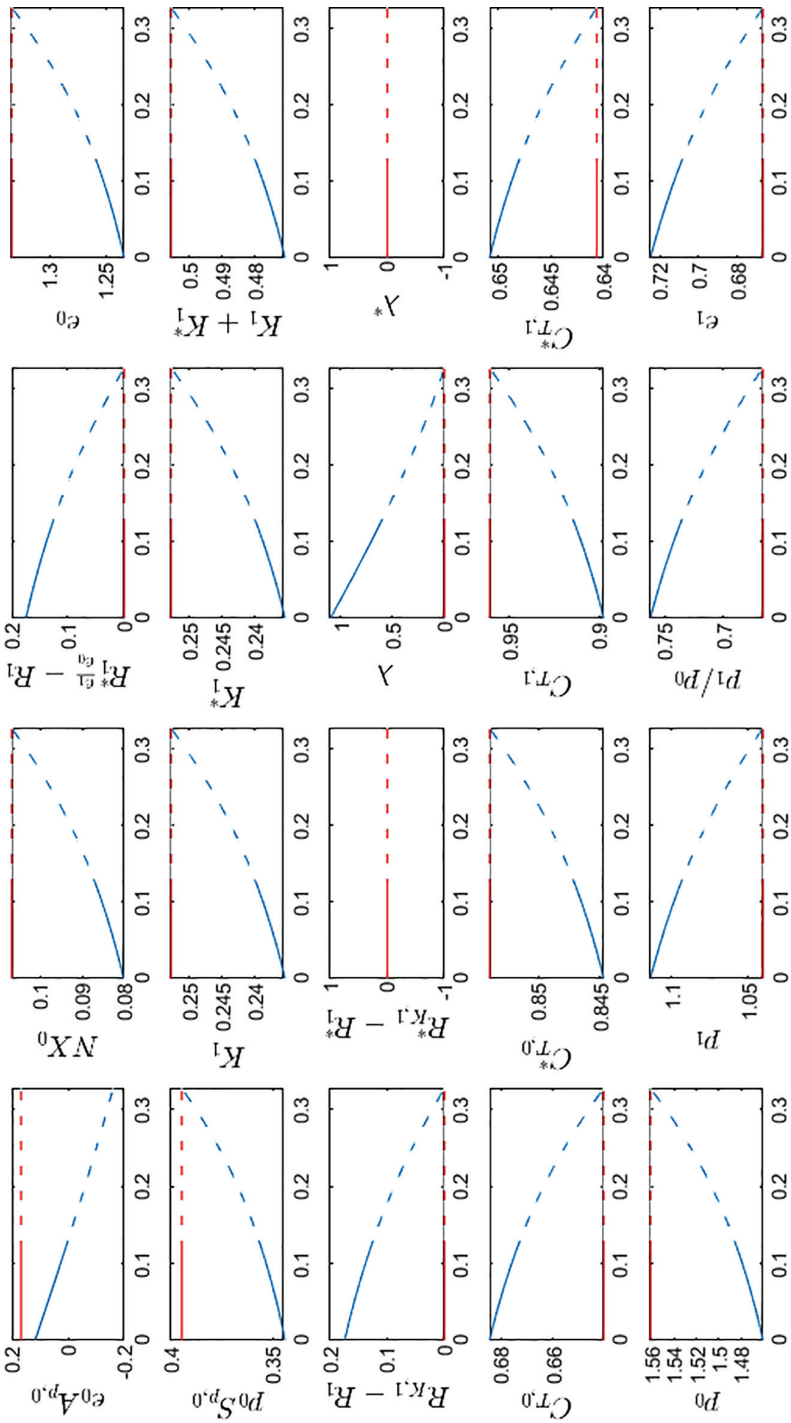

Note: Evolution of the model's equilibrium as central bank intermediation in the international credit market increases, starting from a point where there are limits to arbitrage in the domestic investment market and the international credit market. $\theta = 1/3$, $\theta^* = 0$, $\Delta = 1/3$. The red line represents the value in the frictionless equilibrium. The dashed part of the lines captures the range of FX interventions where the latter would require the home banks to go short in foreign bonds in order to fulfill the central bank's demand for these assets and hence covers a part where our model technically is not valid.

one market alone may be too small so that the economy remains constrained even when the central bank has bought all of them. As an example, consider Figure 8. Even once the central bank has absorbed the entire excess supply of foreign bonds from the home banks ($e_0 A_{CB,0} = p_0 NX_0$ and $e_0 A_{p,0} = 0$), they are still credit constrained ($\lambda > 0$).

Regarding the foreign central bank, i.e., the central bank in the country facing an excess supply of its bonds, note that the set of unconventional monetary policy tools at its disposal to reverse the effects of financial frictions and portfolio flows is more restricted. In the case of frictions in the foreign investment market, which would trigger a depreciation and an outflow of funds, it can, like the home central bank, engage in credit easing, i.e., purchase (foreign) investment securities and issue (foreign) bonds in return. In the case of frictions in the international credit market and exogenous capital outflows, both triggering a foreign depreciation, the foreign central bank would need to purchase foreign bonds for home bonds to reverse the distorting effects. But this obviously requires that it is in the possession of such home bonds, i.e., has sufficient holdings of foreign exchange reserves. Furthermore, unlike in the home economy, intervening in one market does not automatically reduce the spreads in both the international credit and the foreign investment market. In contrast, a relaxation of financial conditions in the international credit market leads to a shift of funds towards the foreign country, increasing the spread in the foreign investment market if foreign banks are balance sheet constrained.[28] Hence, if there are frictions present in both markets, the foreign central bank should engage in both credit easing and foreign exchange interventions, as interventions in one market cannot substitute interventions in the other market.

6. Policy Response to Appreciation Pressure

Even though simple, our model aims to capture the fundamental structure of domestic and international bond markets. This allows

[28] For illustration, see Figure D.3 in Appendix D. While the figure shows purchases of foreign bonds by the home central bank, the effects would be largely identical in the case of purchases of foreign bonds for home bonds by the foreign central bank.

us to draw a number of useful conclusions on the effectiveness of two important unconventional monetary policy tools, credit easing and FX interventions. In this last section, we shed light on three sources of external home appreciation pressure related to credit market frictions and discuss the policy options of the home central bank to reverse the effects on the exchange rate if judged to be unwanted.

Importantly, note that a home appreciation does not necessarily cause damage to the home country in our stylized setup.[29] This is because the typical channels through which a home appreciation would lead to a depression of home output, employment, and income and downward pressure on home inflation are missing.[30] However, we would like to emphasize that the main mechanism of how financial frictions affect the exchange rate in our model, namely the widening of the interest rate spread, would also be at work in a much richer setup and, hence, our conclusions regarding the effectiveness of unconventional monetary policy tools would remain valid.[31]

In practice, a central bank may judge appreciation pressure to be unwanted for a number of reasons. Especially for economies with a high degree of openness, appreciation pressure can pose a major challenge, in particular during periods of economic slack and low inflationary pressure. A stronger currency not only reduces the relative competitiveness of domestic products in the international market, but also weighs on domestic consumer prices through imported

[29]For example, an increase in capital inflows can make consumption by home households increase in both periods (see Figure 6). This is because the higher holdings of foreign bonds and the higher excess return on these bonds boost the home banks' profit and thereby the home households' lifetime income.

[30]The absence of systematically negative welfare effects is mainly because in our model, for simplicity, output is not produced but given exogenously in the first period and because there is only one traded good and no price or wage stickiness. Without these elements, there is, for example, no expenditure-switching effect after an exchange rate movement and no role for the international competitiveness.

[31]Note that, despite the inefficiencies caused by exchange rate fluctuations, Cavallino (2019) finds that it is not optimal for the central bank to fully stabilize the exchange rate. The reason is that the higher excess return that banks earn on foreign assets after capital inflow shocks has positive wealth effects for the home country. By reducing the excess return, foreign exchange interventions reduce this wealth effect. From a welfare point of view, as there are very similar mechanisms at work in our model, in a New Keynesian version of it, a full stabilization of the exchange rate would probably not be optimal either.

Figure 9. Overview of Possible Constellations of Parameters Δ and θ^* for a Given θ

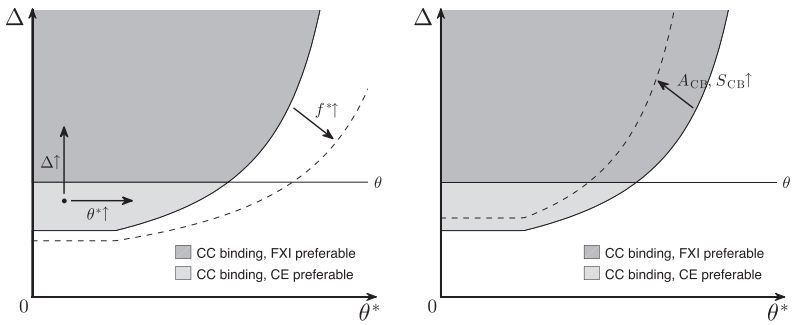

Note: The home banks' capital constraint (CC) is binding in the shaded area only. International portfolio flows affect the general tightness of the CC (see left-side graph) and have an impact on excess returns whenever the CC is binding. Likewise, central bank interventions can be used to relax the CC (see right-side graph). Whenever the CC is binding, such interventions are effective in lowering excess returns. FX interventions (FXI) should be given priority over credit easing (CE) whenever $\Delta > \theta$.

goods—therefore posing an additional risk to price stability, which many central banks are committed to.

The three cases of external appreciation pressure to the home country that we consider are financial frictions in the international credit market ($\Delta \uparrow$), portfolio inflows ($f^* \uparrow$), and financial frictions in the foreign investment market ($\theta^* \uparrow$). In Figure 9, we provide an overview of how these three factors and the two above-mentioned policies affect the capital constraint of the home banks, which lies at the core of our model, for different values of the friction parameters of the international credit and the foreign investment market. For a given level of θ (where by assumption $\theta > 0$), these figures show where in the $\theta^*\Delta$-space the home capital constraint is binding (shaded area), and how a change in the friction parameters θ^* and Δ or a change in portfolio inflows and central bank intermediation affect its general tightness. The line at the border of the shaded area corresponds to the critical values of Δ as a function of θ^* at which the capital constraint starts to bind and hence portfolio flows and central bank intervention become effective. For low levels of θ^*, the capital constraint of the foreign banks is not binding and, hence,

the critical value of Δ does not depend on the level of θ^*. Once foreign banks are constrained, home banks need to absorb a lower excess supply of foreign bonds, which makes them more risk resistant so that they can face higher levels of Δ before being restricted. In the area above this curve, the home capital constraint binds more tightly and this tightness increases further as the economy moves to the upper left.

The first source of external appreciation pressure that we focus on is *financial frictions in the international credit market* that can, for instance, arise when markets lose confidence in a country's economic institutions or future economic development and as a result lose confidence in its currency. An example is the European sovereign debt crisis, which resulted in an appreciation of major currencies against the euro. Within the framework of our model, this scenario can be captured by a (further) increase in Δ. As described in Sections 3.1 and 3.3, in this case, the home appreciation mainly results from an increase in the deviation from interest parity. A country facing appreciation pressure from this type of market imperfection has at least one policy tool it can rely on. As described in Section 5, FX interventions ($A_{CB,0}$) reduce the excess supply of these bonds that needs to be absorbed by private intermediaries. This relaxes the capital constraint of the home banks, which translates into an upward shift of the "critical-Δ" curve in Figure 9 (right-side graph). This, in turn, results in a reduction of the deviation from interest parity and a home depreciation in the first period. Suppose now that, in addition to the international credit market, the home investment market exhibits limits to arbitrage as well. Such a situation may arise when financial frictions reach a global level. An example is the 2007–09 financial crisis, which raised fears of the potential collapse of large financial institutions in many countries. In this case, the home central bank has an additional policy tool at its disposal. As argued in Section 5, credit easing ($S_{CB,0}$) and FX interventions are close substitutes whenever financial frictions are present in both the home investment market and the international credit market. Both tools relax the capital constraint of banks, and the ensuing portfolio rebalancing of private intermediaries makes sure that the relative excess returns in the two markets remain constant. Hence, like FX interventions, credit easing leads to an upward shift of the "critical-Δ" curve in Figure 9. In order to avoid an unnecessarily

large expansion of its balance sheet, the central bank should intervene in the market with higher frictions. Thus, should the limits to arbitrage in the home investment market be larger than in the international credit market, then credit easing is more desirable to reduce appreciation pressure than direct interventions in the international credit market. Note, however, that whichever type of intervention the central bank chooses, it will not be able to bring the economy back to its initial state, i.e., the state before the increase in frictions in the international credit market, unless this initial state corresponds to the frictionless state. The reason is that the increase in Δ has permanently altered the no-arbitrage relation (18) and hence the optimal allocation of funds, from the banks' perspective, across investment market and international credit market.

A second external source of appreciation pressure is an increase in *portfolio inflows* f^*, as experienced by countries like Switzerland, Denmark, and Israel during the global financial crisis and the European sovereign debt crisis.[32] As discussed in Section 4, this involves an increase in the excess supply of foreign bonds that the home banks need to absorb. This results in a larger deviation from interest parity and hence a home appreciation in period 0 if banks are constrained in the international credit market. In Figure 9 (left-side graph), portfolio inflows lead to a downward shift in the "critical-Δ" curve, reflecting that banks will be subject to a higher general restrictiveness of the credit constraint for given levels of the parameters θ, θ^* and Δ. When facing this type of appreciation pressure, the policy options of the home central bank are the same as in the case of an increase in Δ. Whenever capital flows have an impact on the exchange rate, FX interventions, which in the end are just another but special type of capital flow, will as well. By choosing $e_0 A_{CB,0} = f^*$ (where e_0 is equal to the value before the increase in f^*), central bank purchases of foreign bonds can fully reverse the increase in the excess supply of foreign bonds and thus the appreciation. By absorbing all portfolio inflows, the central bank can shift the "critical-Δ" curve in Figure 9 back to its original position. A prominent example of a central bank addressing capital inflows by

[32]Obviously, a decrease in portfolio outflows $e_0 f$ would have the same effects as an increase in portfolio inflows f^*, but would not be classified as an external source of appreciation pressure.

FX interventions is the Swiss National Bank. Starting in 2009, it purchased considerable amounts of foreign assets to counter the upward pressure on the Swiss franc and prevent a tightening of monetary conditions. Ten years later, the SNB's foreign reserves amounted to 112 percent of GDP (average between 2017 and 2019) and thereby exceeded the size of Switzerland's net foreign assets (94 percent).

Again, an interesting result is that credit easing can achieve exactly the same goal as FX interventions if banks are constrained in the home investment market as well ($\theta > 0$). The central bank's acquisitions of home investment securities can likewise free up risk-bearing capacity of home banks, which these in turn can use to absorb the increased excess supply of foreign bonds, reducing the deviation from interest parity. As we can see from Equation (40), purchases of home investment securities to the amount of $p_0 S_{CB,0} = \frac{\Delta}{\theta} f^*$ (where p_0 is equal to the value before the increase in f^*) can even bring the economy back to the state prior to the increase in portfolio inflows. Because both policies act in the same manner, they are both able to shift the "critical-Δ" curve in Figure 9 back to its original position and thereby offset the effect on interest rate spreads. Once again, it is more costly for the central bank to intervene in the market with lower limits to arbitrage, i.e., credit easing is the preferred policy whenever $\theta > \Delta$. However, there is a major limit to credit easing: It cannot exceed the level of home capital $p_0 K_1$—once all domestic capital is owned by the central bank, the policy is no longer feasible. As an example, consider again the case of Switzerland. In 2009, in addition to the interventions in the foreign exchange market, the Swiss National Bank also embarked on a bond purchasing program. The goal was to relax the conditions in capital markets and thereby improve the transmission of monetary policy. Compared with the programs of other countries, the amount of bonds purchased was rather small relative to GDP, though, and the program was stopped in 2010.

The third and last external source of appreciation pressure related to credit market frictions in our model is a financial crisis in the foreign country, taking the form of *financial frictions in the foreign investment market* and captured by an increase in θ^*. During the 2007–09 global crisis, countries like Australia, Canada, and Norway did not experience financial crises themselves, but were negatively affected by the global consequences of the turmoil in U.S. and

European markets, i.e., by financial frictions abroad. As described in Sections 3.2 and 3.3, foreign financial frictions induce a reallocation of capital away from the foreign country towards the home country, leading to a relative increase in the home country's lifetime resources and hence to an appreciation in both periods, i.e., a permanent home appreciation. If home banks are not constrained in the international credit market, there is no further effect on the exchange rate. If home banks are also constrained in the international credit market, an increase in θ^* and the resulting drop in net outflows may take the economy to a state where the home banks' capital constraint is relaxed (see Figure 9, left-side graph) and therefore result in a decrease in excess returns in the international credit market.

Compared with the previous two cases, this type of appreciation pressure is solely driven by a change in economic fundamentals. In this respect, it is "justified," as opposed to any (temporary) appreciation pressure resulting only from an increase in the deviation from interest parity, as is experienced when there are limits to arbitrage in the international credit market or an increase in capital inflows. Even if it wanted to, with the tools discussed here, the home central bank could not counter this last type of appreciation pressure. When home banks are not constrained in the international credit market ($\Delta = 0$ or Δ and θ small enough not to be binding), the central bank has no possibility to affect the exchange rate. When there are limits to arbitrage in the international credit market, it would have this option, but purchases of foreign bonds (or home investment securities, if $\theta > 0$) merely lead to a decrease in the safety premium on domestic bonds: they do reduce appreciation pressure in the first period, but cause additional upward pressure in the second period. Such purchases only address the capital misallocation and exchange rate distortion caused by the international credit market frictions, but not the appreciation pressure caused by the frictions in the foreign investment market.

While the focus of this section has been on appreciation pressure to the home country, our model also allows some interesting statements regarding the exchange rate and unconventional monetary policy tools from the point of view of the foreign country, i.e., the country facing an excess supply of its currency. As can be concluded from the previous sections, in the case of the foreign country, an appreciation is always the result of a relaxation of the (home or

foreign) banks' balance sheet constraints, either due to a decrease in the international credit market friction, a decrease in the foreign investment market friction, or exogenous capital inflows to the foreign country. Hence, for the foreign country, an appreciation always means shifting closer towards the frictionless state and the exchange rate explained by economic fundamentals. Obviously, if the appreciation is the result of a narrowing of the spread in the interest parity, the foreign central bank could always buy home bonds if it wanted to counter such an appreciation. However, the probably more relevant question to answer from the point of view of the foreign central bank is whether it can use interventions to reverse the *depreciation* pressure to the foreign country caused in the three cases discussed above. This issue was discussed in Section 5 where we argued that, as opposed to the home central bank, the foreign central bank's set of unconventional monetary policy tools to respond to distortions caused by financial frictions is more restricted. Above all, it needs to have sufficient holdings of home bonds if it wishes to counter depreciation pressure resulting from frictions in the international credit market and exogenous capital outflows.

7. Conclusions

We provide a simple two-country framework with imperfect financial intermediation to analyze and compare the effectiveness of two unconventional monetary policy measures, foreign exchange interventions and credit easing. International portfolio flows and central bank interventions only have real effects when banks are financially constrained. Our focus is on cases where the financial frictions lead to appreciation pressure. Increased frictions in the international credit market and higher capital inflows both result in an increase in the safety premium on domestic bonds and hence a temporary home appreciation. In these two cases, foreign exchange interventions can reverse the appreciation. An increase in the limits to arbitrage in the foreign investment market also triggers an appreciation, but the appreciation is permanent in this case. It cannot be reversed by central bank purchases of foreign bonds.

Another interesting result concerns the relative effects of the two unconventional policy responses. If, in addition to frictions in the international credit market, there are also frictions in the home

investment market, credit easing is a substitute for foreign exchange interventions. However, the effectiveness of the two policies can differ. Interventions will come at a lower cost if they target the market that faces the highest excess returns.

Appendix A. Stylized Facts

Figure A.1 and Table A.1 suggest that for the United States and for Switzerland, changes in the nominal effective exchange rate (NEER) and changes in credit spreads are significantly correlated, at least in periods of high uncertainty (VIX > 20). The sample is monthly from 1994:1 to 2021:5 for the United States and 2008:2 to 2021:5 for Switzerland. Data on the NEER are from the Bank for International Settlements. Data on bond yields and the VIX are from Datastream. Credit spreads are defined as the yield spread between corporate bonds (AAA-rating, maturity of five to seven years) and government bonds (maturity of five to seven years).

Figure A.1. Changes in NEER versus Changes in Credit Spreads

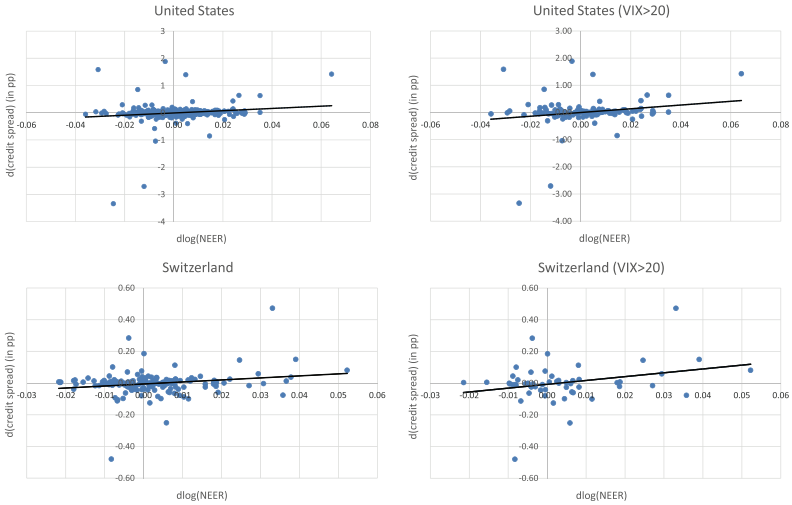

Table A.1. Correlation between Changes in NEER and Changes in Credit Spreads

	United States	Switzerland
Full Sample	0.16***	0.21***
	(N = 328)	(N = 155)
VIX > 20	0.20**	0.29**
	(N = 131)	(N = 55)
VIX ≤ 20	0.10	0.06
	(N = 197)	(N = 100)

Note: ***, **, and * denote significance levels of 1, 5, and 10 percent, respectively, based on a t-test.

Appendix B. Baseline Model: Derivations, Equilibrium Equations, and Numerical Solution

B.1 Derivation of Equilibrium Equations

We can reduce the system of equations to one of 13 equations and 13 unknowns.

- **Euler conditions:**
  Since we set $Y_{NT,0} = Y_{NT,1} = \chi$, it follows from the combination of Equations (4), (6), (7), and the market clearing conditions for the non-traded goods that the domestic households' Euler condition reduces to $R_1 = \frac{1}{\beta}$. Equivalently, we have $R_1^* = \frac{1}{\beta^*}$.
- **Home household's intertemporal budget constraint:**
  The combination of Equations (2) and (3) yields the intertemporal budget constraint of the domestic household,

$$C_{NT,1} + p_1 C_{T,1} = R_1(p_0 Y_{T,0} - p_0 N_0 + Y_{NT,0} - C_{NT,0} - p_0 C_{T,0}) + w_1 L + p_1 N_1 + Y_{NT,1}.$$

  Using the market clearing condition for non-traded goods, the market clearing condition for labor, and Equation (10), the intertemporal budget constraint simplifies to

$$\underbrace{1}_{p_1 C_{T,1}} = R_1 \underbrace{\left(p_0 Y_{T,0} - p_0 N_0 - \underbrace{1}_{p_0 C_{T,0}}\right)}_{B_0} + \underbrace{(1-\alpha)\left(\frac{K_1}{L}\right)^{\alpha} p_1}_{w_1} L + p_1 N_1. \tag{B.1}$$

Consumption expenditure in period 1 depends on the savings in domestic bonds in period 0, on the wage in period 1, and on the profit of the bank the household owns. Note that consumption expenditure on the traded good is constant and equal to 1.

- **Returns on investment securities:**

  Using Equation (9) and the market clearing condition for labor, we can rewrite the return on home securities as follows:

$$R_{k,1} = \left(\alpha \left(\frac{L}{K_1}\right)^{1-\alpha} + (1-\delta)\right) \frac{p_1}{p_0}. \tag{B.2}$$

Equivalently, we can simplify the return on foreign securities:

$$R^*_{k,1} = \left(\alpha \left(\frac{L^*}{K^*_1}\right)^{1-\alpha} + (1-\delta)\right) \frac{p_1 e_0}{p_0 e_1}. \tag{B.3}$$

- **Value of the home bank's equity capital:**

  The value of the home bank's equity capital in period 1 is (from Equation (15))

$$p_1 N_1 = (R_{k,1} - R_1)\, p_0 S_{p,0} + \left(R^*_1 \frac{e_1}{e_0} - R_1\right) e_0 A_{p,0} + R_1 p_0 N_0. \tag{B.4}$$

Using the market clearing condition for domestic investment securities and Equation (29), the value of the home bank's equity capital in period can be rewritten as

$$p_1 N_1 = (R_{k,1} - R_1)\, p_0 K_1 + \left(R_1^* \frac{e_1}{e_0} - R_1\right) \underbrace{(p_0 Y_{T,0} - p_0 K_1 - 1)}_{p_0 N X_0} + R_1 p_0 N_0. \tag{B.5}$$

- **Banks' first-order conditions:**

  The first-order conditions of the domestic bank, Equations (16) and (17), can be slightly simplified to

$$\frac{1}{R_1}(R_{k,1} - R_1) = \frac{\lambda}{1+\lambda}\theta \tag{B.6}$$

$$\frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right) = \frac{\lambda}{1+\lambda}\Delta. \tag{B.7}$$

  And the first-order condition of the foreign bank, Equation (23), is now

$$\frac{1}{R_1^*}\left(R_{k,1}^* - R_1^*\right) = \frac{\lambda^*}{1+\lambda^*}\theta^*. \tag{B.8}$$

- **Home capital constraint:**

  Under the assumption that $\theta > 0$ and $\Delta > 0$, we can rewrite the domestic incentive constraint (13) as a capital constraint (CC) (use Equations (15)–(18)):

$$\begin{aligned}
& V_0 \geq \theta p_0 S_{p,0} + \Delta e_0 A_{p,0} \\
\Leftrightarrow\ & \Lambda_{0,1}\left((R_{k,1} - R_1)\, p_0 S_{p,0} + \left(R_1^* \frac{e_1}{e_0} - R_1\right) e_0 A_{p,0} + R_1 p_0 N_0\right) \\
& \geq \theta p_0 S_{p,0} + \Delta e_0 A_{p,0} \\
\Leftrightarrow\ & \Lambda_{0,1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)\left(\frac{\theta}{\Delta} p_0 S_{p,0} + e_0 A_{p,0}\right) \\
& + \Lambda_{0,1} R_1 p_0 N_0 \geq \Delta\left(\frac{\theta}{\Delta} p_0 S_{p,0} + e_0 A_{p,0}\right) \\
\Leftrightarrow\ & \Lambda_{0,1} R_1 p_0 N_0 \geq \left(\frac{\theta}{\Delta} p_0 S_{p,0} + e_0 A_{p,0}\right)\left(\Delta - \Lambda_{0,1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)\right)
\end{aligned}$$

$$\Leftrightarrow \frac{\Lambda_{0,1} R_1 p_0 N_0}{\Delta - \Lambda_{0,1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} \geq \frac{\theta}{\Delta} p_0 S_{p,0} + e_0 A_{p,0}$$

$$\Leftrightarrow \phi p_0 N_0 \geq \theta p_0 S_{p,0} + \Delta e_0 A_{p,0}, \tag{B.9}$$

where $\phi = \frac{\Delta \Lambda_{0,1} R_1}{\Delta - \Lambda_{0,1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} = \frac{\theta \Lambda_{0,1} R_1}{\theta - \Lambda_{0,1}(R_{k,1} - R_1)}$. If in one of the two markets the friction parameter is set to zero, then this inequality simplifies to

$$\frac{\Lambda_{0,1} R_1}{\Delta - \Lambda_{0,1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq e_0 A_{p,0} \qquad \text{if} \quad \theta = 0 \tag{B.10}$$

$$\frac{\Lambda_{0,1} R_1}{\theta - \Lambda_{0,1}\left(R_{k,1} - R_1\right)} p_0 N_0 \geq p_0 S_{p,0} \qquad \text{if} \quad \Delta = 0. \tag{B.11}$$

Note that if $\theta > 0$ and/or $\Delta > 0$, this does not necessarily imply that the capital constraint is binding ($\lambda > 0$), i.e., that there are limits to arbitrage in at least one market. The capital constraint is only binding if $\theta > \bar{\theta}$ and/or $\Delta > \bar{\Delta}$. We can summarize the capital constraint as follows:

$$\text{CC} = \begin{cases} \frac{\Delta \Lambda_{0,1} R_1}{\Delta - \Lambda_{0,1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \theta p_0 S_{p,0} & \\ \quad + \Delta e_0 A_{p,0} & \text{if} \quad \theta \geq 0, \Delta > 0 \\ \frac{\theta \Lambda_{0,1} R_1}{\theta - \Lambda_{0,1}(R_{k,1} - R_1)} p_0 N_0 \geq \theta p_0 S_{p,0} & \\ \quad + \Delta e_0 A_{p,0} & \text{if} \quad \theta > 0, \Delta \geq 0 \\ \text{no CC} & \text{if} \quad \theta = 0, \Delta = 0. \end{cases} \tag{B.12}$$

If $\theta > 0$ and $\Delta > 0$, it is irrelevant whether the first or the second Equation of (B.12) is considered.

Using the Euler condition, the market clearing condition for domestic investment securities, and Equation (29), we can simplify the capital constraint of the domestic bank as follows:

$$\mathrm{CC} = \begin{cases} \frac{\Delta}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \theta p_0 K_1 & \\ \quad + \Delta(p_0 Y_{T,0} - p_0 K_1 - 1) & \text{if} \quad \theta \geq 0, \Delta > 0 \\ \frac{\theta}{\theta - \frac{1}{R_1}(R_{k,1} - R_1)} p_0 N_0 \geq \theta p_0 K_1 & \\ \quad + \Delta(p_0 Y_{T,0} - p_0 K_1 - 1) & \text{if} \quad \theta > 0, \Delta \geq 0 \\ \text{no CC} & \text{if} \quad \theta = 0, \Delta = 0. \end{cases}$$

Keep in mind that $(p_0 Y_{T,0} - p_0 K_1 - 1) = p_0 N X_0$. For any parameter specification but $\theta = \Delta = 0$, the Karush-Kuhn-Tucker (KKT) conditions need to hold. Define

$$g \equiv \begin{cases} \frac{\Delta}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 - \theta p_0 K_1 & \\ \quad - \Delta(p_0 Y_{T,0} - p_0 K_1 - 1) & \text{if} \quad \theta \geq 0, \Delta > 0 \\ \frac{\theta}{\theta - \frac{1}{R_1}(R_{k,1} - R_1)} p_0 N_0 - \theta p_0 K_1 & \\ \quad - \Delta(p_0 Y_{T,0} - p_0 K_1 - 1) & \text{if} \quad \theta > 0, \Delta \geq 0, \end{cases} \tag{B.13}$$

where g is a function of domestic endogenous variables. Then, the KKT conditions for the inequality constraint of the home bank are

$$g \geq 0 \tag{B.14}$$

$$\lambda \geq 0 \tag{B.15}$$

$$\lambda g = 0. \tag{B.16}$$

- **Foreign capital constraint:**

  Equivalently, for the foreign bank we have

$$\mathrm{CC}^* = \begin{cases} \frac{1}{\theta^* - \frac{1}{R_1^*}(R_{k,1}^* - R_1^*)} \frac{p_0}{e_0} N_0^* \geq \frac{p_0}{e_0} K_1^* & \text{if} \quad \theta^* > 0 \\ \text{no CC}^* & \text{if} \quad \theta^* = 0, \end{cases} \tag{B.17}$$

where we have used the law of one price. For any parameter specification but $\theta^* = 0$, the KKT conditions need to hold. Define

$$g^* \equiv \left\{ \frac{1}{\theta^* - \frac{1}{R_1^*}(R_{k,1}^* - R_1^*)} \frac{p_0}{e_0} N_0^* - \frac{p_0}{e_0} K_1^* \quad \text{if} \quad \theta^* > 0, \right.$$

where g^* is a function of foreign endogenous variables. Then, the KKT conditions for the inequality constraint of the foreign bank are

$$g^* \geq 0 \tag{B.18}$$

$$\lambda^* \geq 0 \tag{B.19}$$

$$\lambda^* g^* = 0. \tag{B.20}$$

- **Market clearing conditions for traded goods:**
  The market clearing condition for traded goods in period 0 simplifies to

$$Y_{T,0} + Y^*_{T,0} = \frac{1}{p_0} + \frac{e_0}{p_0} + K_1 + K^*_1. \tag{B.21}$$

The market clearing condition for traded goods in period 1 is

$$K_1^{\alpha} L^{1-\alpha} + K_1^{*\alpha} L^{*1-\alpha} + (1-\delta)K_1 + (1-\delta)K^*_1 = \frac{1}{p_1} + \frac{e_1}{p_1}, \tag{B.22}$$

where we have used the production function of the domestic and foreign firm.

In sum, we reduce our system of equations to the following *13 equilibrium equations:*

$$R_1 = \frac{1}{\beta} \tag{B.23}$$

$$R^*_1 = \frac{1}{\beta^*} \tag{B.24}$$

$$1 = R_1(p_0 Y_{T,0} - p_0 N_0 - 1) + (1-\alpha)\left(\frac{K_1}{L}\right)^{\alpha} p_1 L + p_1 N_1 \tag{B.25}$$

$$R_{k,1} = \left(\alpha\left(\frac{L}{K_1}\right)^{1-\alpha} + (1-\delta)\right)\frac{p_1}{p_0} \tag{B.26}$$

$$R^*_{k,1} = \left(\alpha\left(\frac{L^*}{K^*_1}\right)^{1-\alpha} + (1-\delta)\right)\frac{p_1 e_0}{p_0 e_1} \tag{B.27}$$

$$p_1 N_1 = (R_{k,1} - R_1)\, p_0 K_1$$
$$+ \left(R_1^* \frac{e_1}{e_0} - R_1 \right) (p_0 Y_{T,0} - p_0 K_1 - 1) + R_1 p_0 N_0 \tag{B.28}$$

$$\frac{1}{R_1} (R_{k,1} - R_1) = \frac{\lambda}{1+\lambda} \theta \tag{B.29}$$

$$\frac{1}{R_1} \left(R_1^* \frac{e_1}{e_0} - R_1 \right) = \frac{\lambda}{1+\lambda} \Delta \tag{B.30}$$

$$\frac{1}{R_1^*} \left(R_{k,1}^* - R_1^* \right) = \frac{\lambda^*}{1+\lambda^*} \theta^* \tag{B.31}$$

$$\mathrm{CC} = \begin{cases} \frac{\Delta}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \theta p_0 K_1 & \\ \quad + \Delta (p_0 Y_{T,0} - p_0 K_1 - 1) & \text{if} \quad \theta \geq 0, \Delta > 0 \\ \frac{\theta}{\theta - \frac{1}{R_1}(R_{k,1} - R_1)} p_0 N_0 \geq \theta p_0 K_1 & \\ \quad + \Delta (p_0 Y_{T,0} - p_0 K_1 - 1) & \text{if} \quad \theta > 0, \Delta \geq 0 \\ \text{no CC}, \lambda = 0 & \text{if} \quad \theta = 0, \Delta = 0 \end{cases} \tag{B.32}$$

$$\mathrm{CC}^* = \begin{cases} \frac{1}{\theta^* - \frac{1}{R_1^*}(R_{k,1}^* - R_1^*)} \frac{p_0}{e_0} N_0^* \geq \frac{p_0}{e_0} K_1^* & \text{if} \quad \theta^* > 0 \\ \text{no CC}^*, \lambda^* = 0 & \text{if} \quad \theta^* = 0 \end{cases} \tag{B.33}$$

$$Y_{T,0} + Y_{T,0}^* = \frac{1}{p_0} + \frac{e_0}{p_0} + K_1 + K_1^* \tag{B.34}$$

$$K_1^\alpha L^{1-\alpha} + K_1^{*\alpha} L^{*1-\alpha} + (1-\delta) K_1 + (1-\delta) K_1^* = \frac{1}{p_1} + \frac{e_1}{p_1}. \tag{B.35}$$

Furthermore, we have to take into account the remaining KKT conditions for the domestic country if the domestic friction parameters are non-zero ($\theta \neq 0$ and $\Delta \neq 0$), Equations (B.15) and (B.16), and the remaining KKT conditions for the foreign country if the foreign friction parameter is non-zero ($\theta^* \neq 0$), Equations (B.19) and (B.20). The 13 unknowns are $e_0, e_1, p_0, p_1, K_1, K_1^*, R_1, R_1^*, R_{k,1}, R_{k,1}^*, N_1, \lambda, \lambda^*$.

B.2 Proof: Properties of the Model with $\theta = \theta^* > 0$ and $\Delta = 0$ (for ICs Binding)

First note that taking the ratio of the two expressions for the returns on the investment securities, $R_{k,1}$ and $R^*_{k,1}$ (see Equations (B.26) and (B.27)), yields the following (generally valid) relationship between K_1 and K_1^* (for $L = L^*$):

$$K_1 = \left(\frac{\frac{p_0}{p_1}\frac{e_1}{e_0} R^*_{k,1} - (1-\delta)}{\frac{p_0}{p_1} R_{k,1} - (1-\delta)} \right)^{\frac{1}{1-\alpha}} K_1^*. \tag{B.36}$$

For finding the relative level of capital for when $\theta = \theta^* > 0$ and $\Delta = 0$ conditional on the ICs being binding, combine each country's (binding) capital constraints (see Equations (B.32) and (B.33)) with the respective expression for the return on the investment securities:

$$\frac{1}{\theta - \frac{1}{R_1}\left(\left(\alpha\left(\frac{L}{K_1}\right)^{1-\alpha} + (1-\delta)\right)\frac{p_1}{p_0} - R_1\right)} N_0 = K_1 \tag{B.37}$$

$$\frac{1}{\theta^* - \frac{1}{R_1^*}\left(\left(\alpha\left(\frac{L^*}{K_1^*}\right)^{1-\alpha} + (1-\delta)\right)\frac{p_1}{p_0}\frac{e_0}{e_1} - R_1^*\right)} N_0^* = K_1^*. \tag{B.38}$$

By assumption, $L^* = L$ and $N_0^* = N_0$. Furthermore, $\Delta = 0$ implies that $R_1 = R_1^* \frac{e_1}{e_0}$, and we have $\theta = \theta^* > 0$. Thus, Equation (B.38) can be written as

$$\frac{1}{\theta - \frac{1}{R_1}\left(\left(\alpha\left(\frac{L}{K_1^*}\right)^{1-\alpha} + (1-\delta)\right)\frac{p_1}{p_0} - R_1\right)} N_0 = K_1^*. \tag{B.39}$$

Looking at Equations (B.37) and (B.39), it becomes clear that it must be the case that $K_1 = K_1^*$. From $K_1 = K_1^*$, in turn, it follows that $R_{k,1} = R^*_{k,1}\frac{e_1}{e_0}$ (see Equation (B.36)), i.e., returns on the investment securities in terms of the home numéraire are equalized.

For evaluating the relative tightness of the two countries' banks' incentive constraints, note that the first-order conditions (16) and

(23) can be written as follows (remember that $\Lambda_{0,1} = \frac{1}{R_1}$ and $\Lambda^*_{0,1} = \frac{1}{R^*_1}$):

$$\frac{R_{k,1}}{R_1} = 1 + \frac{\lambda}{1+\lambda}\theta \tag{B.40}$$

$$\frac{R^*_{k,1}}{R^*_1} = 1 + \frac{\lambda^*}{1+\lambda^*}\theta^*. \tag{B.41}$$

From $R_1 = R^*_1 \frac{e_1}{e_0}$ and $R_{k,1} = R^*_{k,1} \frac{e_1}{e_0}$, it follows that $\frac{R_{k,1}}{R_1} = \frac{R^*_{k,1}}{R^*_1}$. Hence, and given that $\theta = \theta^*$, Equations (B.40) and (B.41) imply that $\lambda = \lambda^*$, i.e., the incentive constraints are equally binding in the two countries.

B.3 Proof: Properties of the Model with $\theta = \theta^* = \Delta > 0$ (for ICs Binding)

An interesting case to have a closer look at is the one where banks are equally constrained in all markets: $\theta = \theta^* = \Delta > 0$: From no-arbitrage relation (18), it follows that home banks will choose their portfolio such that excess returns in the domestic investment market and the international credit market are just equalized. Measured in terms of the home numéraire, we have $R_1 < R_{k,1} = R^*_1 \frac{e_1}{e_0} < R^*_{k,1} \frac{e_1}{e_0}$, which implies that the level of investment is lower in the foreign country as compared with the home country. This results from the fact that home banks are constrained to hold less foreign bonds and hence intermediate fewer net exports relative to the frictionless case. Also note that the home banks' incentive constraint is more binding than the one of the foreign banks ($\lambda > \lambda^*$). Intuitively, when banks are only constrained in the investment markets ($\theta = \theta^* > 0$ and $\Delta = 0$), the real interest rates in the two countries must be equalized and home and foreign banks hold the same amount in investment securities. Consequently, both incentive constraints are equally binding. Once Δ is larger than zero, the restriction of the home banks increases, while tension on the foreign banks is released as less funds flow into the country and demand for intermediation falls.

Formal proof: $\theta = \theta^* = \Delta > 0$ implies that $R_1 < R_{k,1} = R_1^* \frac{e_1}{e_0} < R_{k,1}^* \frac{e_1}{e_0}$ (conditional on the ICs being binding). From $R_{k,1} < R_{k,1}^* \frac{e_1}{e_0}$ and Equation (B.36), it follows that $K_1 > K_1^*$.

For evaluating the relative tightness of the two countries' banks' incentive constraints, combine each country's (binding) capital constraints (see Equations (B.32) and (B.33)) with first-order conditions (16) and (23) and the respective market clearing conditions:

$$\frac{1}{\theta - \frac{\lambda}{1+\lambda}\theta} N_0 = K_1 + NX_0 \tag{B.42}$$

$$\frac{1}{\theta^* - \frac{\lambda^*}{1+\lambda^*}\theta^*} N_0^* = K_1^*. \tag{B.43}$$

Again, by assumption, $N_0^* = N_0$. Furthermore, we have $\theta = \theta^*$. Thus, taking the ratio of Equations (B.42) and (B.43) yields

$$\frac{1+\lambda}{1+\lambda^*} = \frac{K_1 + NX_0}{K_1^*}. \tag{B.44}$$

We know that $K_1 > K_1^*$ and $NX_0 \geq 0$, from which it finally follows that $\lambda > \lambda^*$, i.e., the home banks' incentive constraint binds tighter than the foreign banks' incentive constraint.

B.4 Effect of Financial Frictions in the Home Investment Market

This section discusses the effects of financial frictions in the home investment market, which are captured by an increase in the home investment market friction parameter θ. θ^* and Δ are set to zero. When the home investment market friction parameter is sufficiently large for the home incentive constraint and hence also the endogenous capital constraint to become binding ($\lambda > 0$), home banks are hindered to exploit all arbitrage opportunities and excess returns in the home investment market become positive: $R_{k,1} - R_1 > 0$ (see Equation (16)). Excess returns in the international credit market and the foreign investment market, however, remain zero (see Equations (17) and (23)). Combining the home banks' capital constraint (19) and the investment market clearing condition (25) reveals that

Figure B.1. Financial Frictions in the Home Investment Market: $\theta > 0$

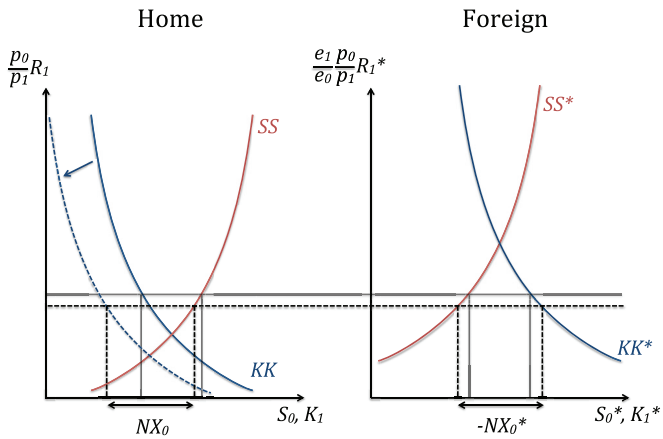

Note: The solid lines represent the frictionless equilibrium, the dashed lines the equilibrium with the fiction.

with the constraint starting to be binding, the level of capital in the home country will obviously be limited:

$$\frac{1}{\theta - \frac{1}{R_1}\left(R_{k,1} - R_1\right)} p_0 N_0 \geq p_0 K_1. \tag{B.45}$$

(As $\lambda > 0$, this equation will hold with equality.)

Graphically, financial frictions in the home investment market shift the home investment curve to the left (see Figure B.1).[33] For a given real rate of return $\frac{p_0}{p_1}R_1$, investment in the home country decreases, as the home banks' ability to intermediate funds in this market has decreased and they face limits to arbitrage. Costs of capital in the home market increase. In order to maintain the world equilibrium, the equilibrium real rate of return has to decrease. Due to the frictions in the home investment market, the home country will in equilibrium slightly decrease its savings, and invest a much larger part abroad: The credit constraint with respect to investments

[33]For a formal proof of how an increase in θ affects the two countries' saving and investment schedules, see Appendix E.

in home capital makes the home banks reallocate their portfolio and invest a larger part in foreign bonds. Altogether, this causes an increase in the home country's net exports. The foreign country, on the other hand, also decreases its savings, but at the same time can increase its investments as the foreign banks obtain a larger amount of funds, which leads to an increase in its net imports. Overall, there is a decrease in world savings and, consequently, world investments, implying a lower level of world output in the second period. Furthermore, the frictions in the home investment market also lead to a misallocation of capital: Now, a majority of capital is invested in the foreign country. This change in the allocation of capital implies that relative to the frictionless level, the home country's output in the second period will decrease while the foreign country's output in the second period will increase, implying that there is a change in the two countries' fundamentals. The relative decrease in the home country's lifetime resources induces a home depreciation in both periods.

Figure B.2 provides a numerical illustration of these results. Setting θ^* and Δ equal to zero and using the calibration of Table 1 for the remaining parameters, it shows how the model's equilibrium evolves as the home investment market friction parameter θ increases. Note that $\theta = 1$ (i.e., banks can divert *all* home investment securities) does not imply that home banks do not hold domestic assets anymore: It just means that in case of misbehavior, the banks could divert and keep the proceeds of all these assets. If, however, the excess returns they can earn on the investment securities when not diverting them are large enough, they still have no incentive to misbehave and the financial markets will work even with $\theta = 1$.

B.5 Numerical Illustration

As mentioned before, there exists no closed-form solution of the model; however, we can solve it numerically in MATLAB. The following figures provide numerical results—in particular, they show the evolution of the model's equilibrium under different specifications of the friction parameters using the calibration of Table 1 for the remaining parameters.

Figure B.2. Effect of an Increase in θ (θ on x-axis)

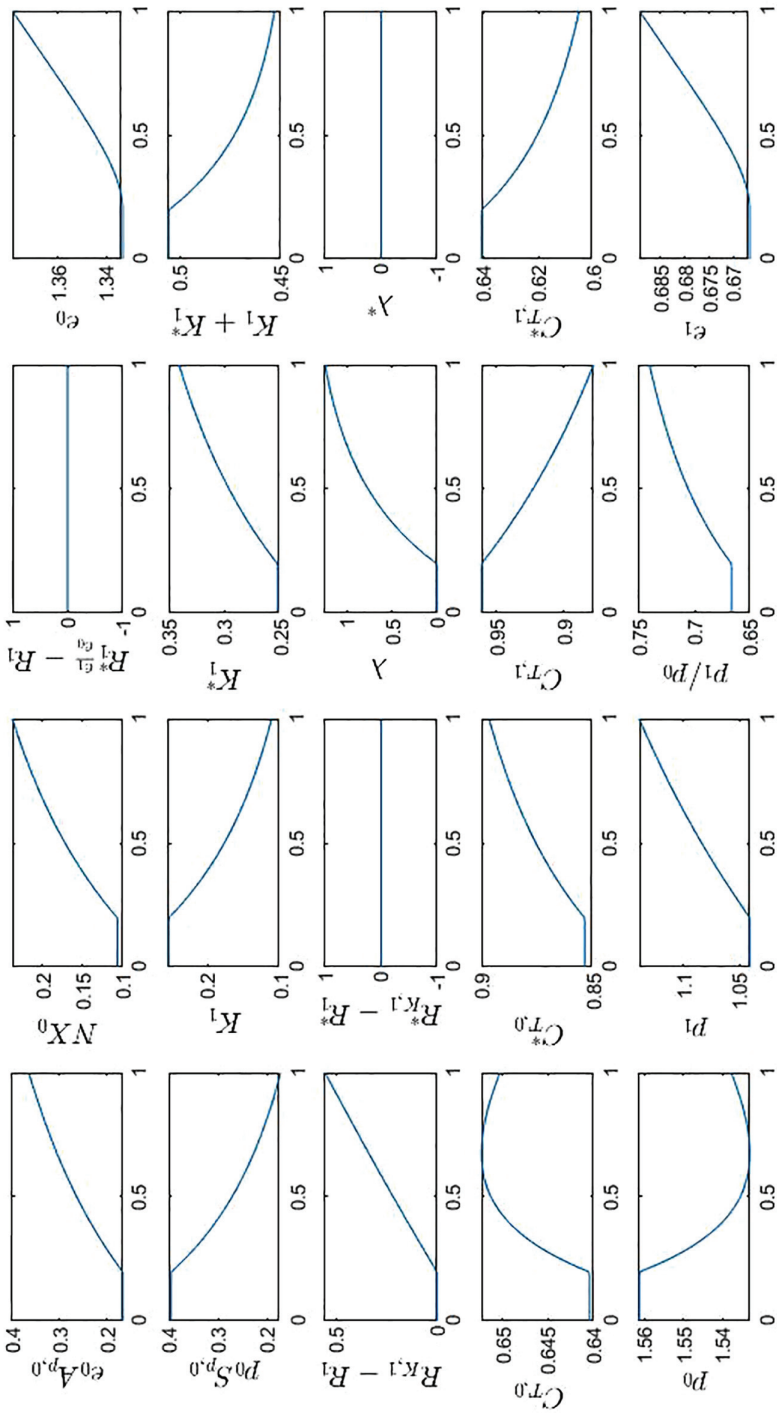

Note: Evolution of the model's equilibrium as the friction parameter in the home investment market increases, starting from a frictionless point. $\theta \geq 0$, $\theta^* = 0$, $\Delta = 0$. The remaining parameter values are summarized in Table 1. The red line represents the value in the frictionless equilibrium.

Figure B.3. Effect of an Increase in Δ (Δ on x-axis)

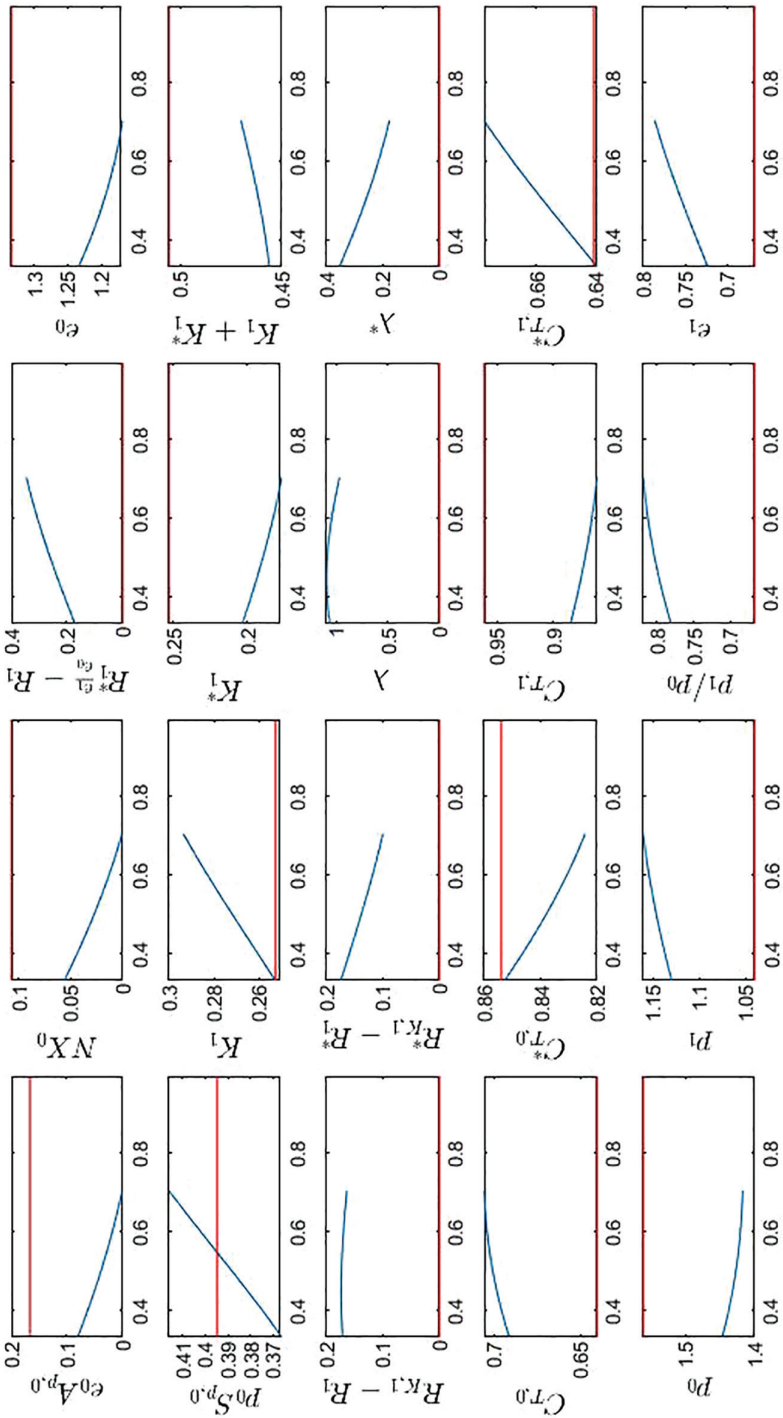

Note: Evolution of the model's equilibrium as the friction parameter in the international credit market increases, starting from a point where there are limits to arbitrage in all financial markets. $\theta = 1/3$, $\theta^* = 1/3$, $\Delta \geq 1/3$. The remaining parameter values are summarized in Table 1. The red line represents the value in the frictionless equilibrium.

Figure B.4. Effect of an Increase in θ^* (θ^* on x-axis)

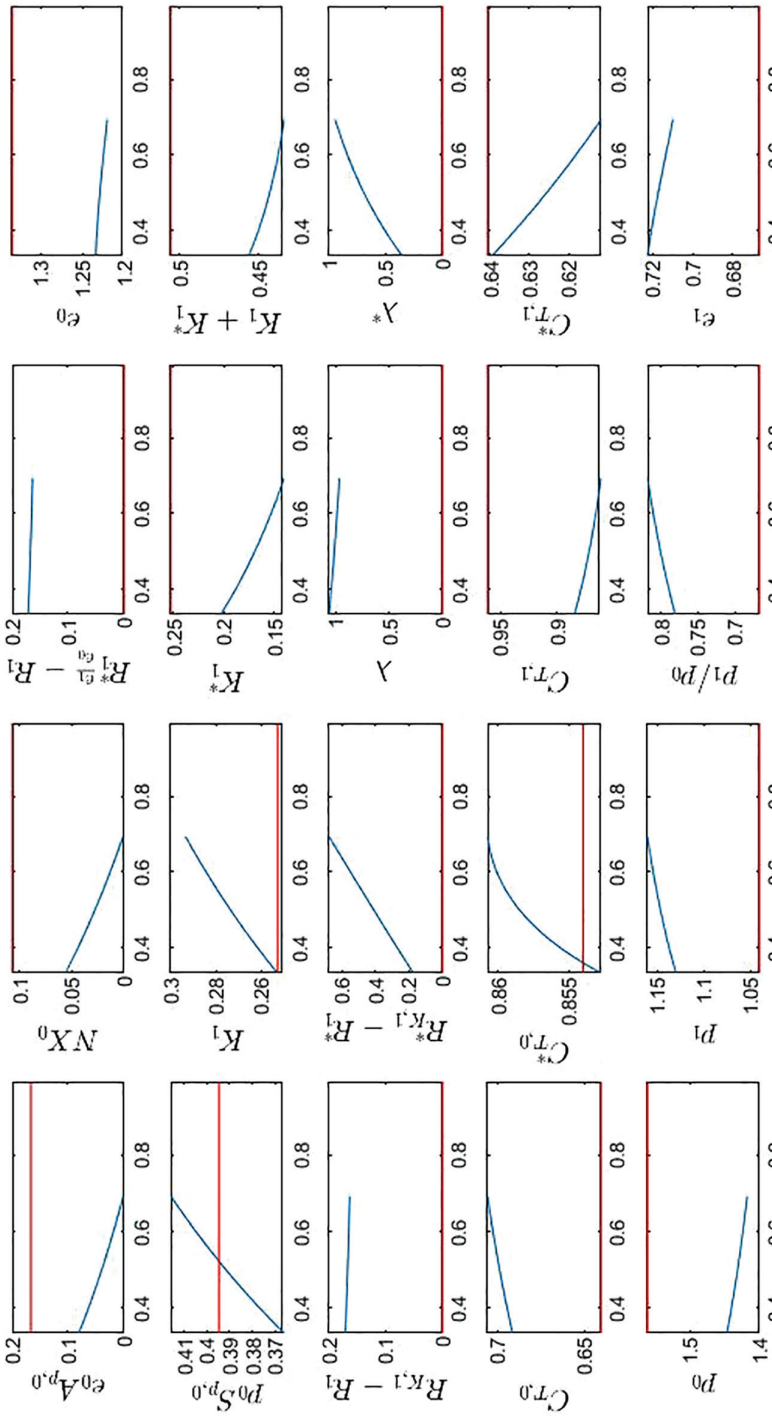

Note: Evolution of the model's equilibrium as the friction parameter in the foreign investment market increases, starting from a point where there are limits to arbitrage in all financial markets. $\theta = 1/3$, $\theta^* \geq 1/3$, $\Delta = 1/3$. The remaining parameter values are summarized in Table 1. The red line represents the value in the frictionless equilibrium. Given that the setup of the banking sector is only valid when $e_0 A_{p,0} \geq 0$, the plots only cover a limited range of possible values for θ^*.

Figure B.5. Effect of an Increase in θ (θ on x-axis)

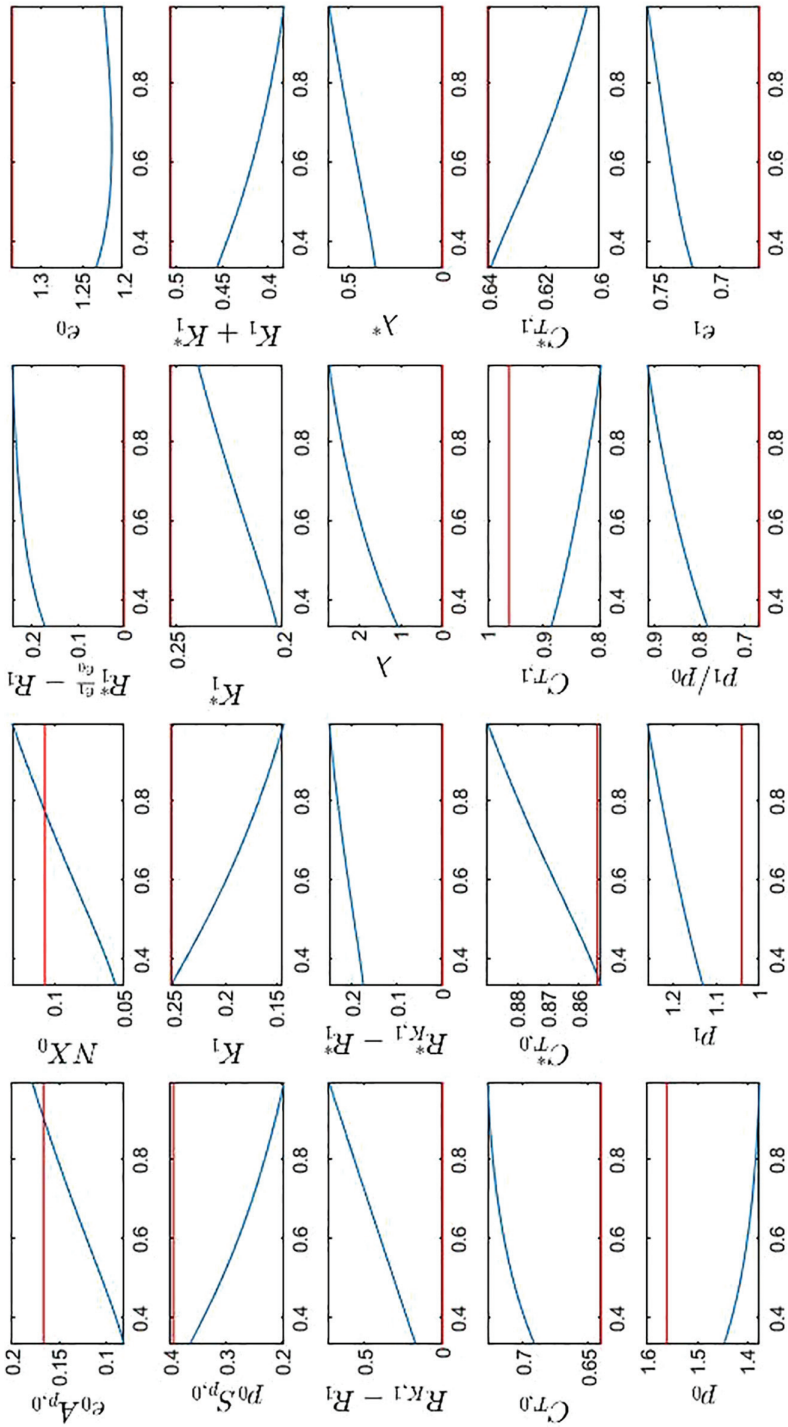

Note: Evolution of the model's equilibrium as the friction parameter in the home investment market increases, starting from a point where there are limits to arbitrage in all financial markets. $\theta \geq 1/3$, $\theta^* = 1/3$, $\Delta = 1/3$. The remaining parameter values are summarized in Table 1. The red line represents the value in the frictionless equilibrium.

Appendix C. Analytical Solution of the Frictionless Model

In the frictionless model the parameters θ, Δ, and θ^* are either equal to zero or sufficiently low for the banks' incentive constraint not to be binding. In the system of 13 equations derived in Appendix B, Equation (B.28) reduces to $p_1 N_1 = R_1 p_0 N_0$ since all excess returns are zero and Equations (B.29)–(B.31) can be replaced by $R_{k,1} = R_1$, $R_1^* \frac{e_1}{e_0} = R_1$ and $R_{k,1}^* = R_1^*$, respectively. In a frictionless environment domestic and foreign banks face no capital constraint, therefore we can omit Equations (B.32) and (B.33), and the variables λ and λ^* and consider the resulting system of 11 equations and 11 unknowns. Under the assumption that $\delta = 1$, it is possible to find an analytical solution of the frictionless model. Given that $Y_{T,0} = Y_{T,0}^*$ and $L = L^*$, solving the system of equations yields

$$e_0 = \frac{1+\beta}{1+\beta^*}, e_1 = \frac{\beta^*}{\beta} e_0, K_1 = K_1^* = \gamma_1 Y_{T,0},$$
$$p_0 = \gamma_2 \frac{1}{Y_{T,0}}, p_1 = \gamma_3 \frac{1}{Y_{T,1}} = \gamma_3 \frac{1}{K_1^\alpha L^{1-\alpha}}, \tag{C.1}$$

where $\gamma_1 \equiv \frac{\alpha\beta(1+\beta^*)+\alpha\beta^*(1+\beta)}{(1+\alpha\beta^*)(1+\beta)+(1+\alpha\beta)(1+\beta^*)}$, $\gamma_2 \equiv \frac{(1+\alpha\beta^*)(1+\beta)+(1+\alpha\beta)(1+\beta^*)}{2(1+\beta^*)}$, and $\gamma_3 \equiv \frac{\beta(1+\beta^*)+\beta^*(1+\beta)}{2\beta(1+\beta^*)}$. The exchange rate only depends on the discount factor of home and foreign agents. Investment, and hence production, is equally high in both countries and is increasing in endowment of traded goods in the first period $Y_{T,0}$. The price of traded goods depends negatively on its supply. The remaining variables of the model can be derived from these six variables. Net exports, e.g., are

$$e_0 A_{p,0} = p_0 NX_0 = \frac{\beta - \beta^*}{2(1+\beta^*)}. \tag{C.2}$$

Appendix D. Equilibrium Equations under International Portfolio Flows and Central Bank Intermediation

Introducing international portfolio flows and central bank intermediation to the baseline model in Section 2 leads to the following changes in the system of 13 equations derived in Appendix B: Equation (B.25) is augmented by the returns on the portfolio flows and

$\Pi_{CB,1}$, the profit of the central bank that is transferred to the domestic household:

$$1 = R_1(p_0 Y_{T,0} - p_0 N_0 - 1) + (1-\alpha)\left(\frac{K_1}{L}\right)^{\alpha} p_1 L + p_1 N_1 + \left(R_1^* \frac{e_1}{e_0} - R_1\right) e_0 f + \Pi_{CB,1} \tag{D.1}$$

$$\text{where } \Pi_{CB,1} = (R_{k,1} - R_1) p_0 S_{CB,0} + \left(R_1^* \frac{e_1}{e_0} - R_1\right) e_0 A_{CB,0}. \tag{D.2}$$

The consolidation of the domestic household's, bank's and central bank's budget constraint (Equation (37)) yields

$$e_0 A_{p,0} + e_0 A_{CB,0} + e_0 f - f^* = p_0 N X_0. \tag{D.3}$$

Equation (B.28) changes to

$$p_1 N_1 = (R_{k,1} - R_1)\, p_0 (K_1 - S_{CB,0}) + \left(R_1^* \frac{e_1}{e_0} - R_1\right)(p_0 N X_0 - e_0 f + f^* - e_0 A_{CB,0}) + R_1 p_0 N_0 \tag{D.4}$$

and Equation (B.32) now looks as follows:

$$\text{CC} = \begin{cases} \frac{\Delta}{\Delta - \frac{1}{R_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right)} p_0 N_0 \geq \theta p_0 (K_1 - S_{CB,0}) & \\ \quad + \Delta(p_0 N X_0 - e_0 f + f^* - e_0 A_{CB,0}) & \text{if } \theta \geq 0, \Delta > 0 \\ \frac{\theta}{\theta - \frac{1}{R_1}(R_{k,1} - R_1)} p_0 N_0 \geq \theta p_0 (K_1 - S_{CB,0}) & \\ \quad + \Delta(p_0 N X_0 - e_0 f + f^* - e_0 A_{CB,0}) & \text{if } \theta > 0, \Delta \geq 0 \\ \text{no CC}, \lambda = 0 & \text{if } \theta = 0, \Delta = 0, \end{cases} \tag{D.5}$$

where $p_0 N X_0$ is substituted by $(p_0 Y_{T,0} - p_0 K_1 - 1)$. The KKT conditions change accordingly. Along with these changes, we have to include one additional equation which is the profit of the central bank in period 1 (Equation (D.2)).

The following figures provide a numerical illustration of the model solution under capital inflows, credit easing, or foreign exchange interventions. They show the evolution of the model's equilibrium under increasing values for one of these variables using different specifications of the friction parameters and the calibration of Table 1 for the remaining parameters.

Figure D.1. Effect of an Increase in f^* (f^* on x-axis)

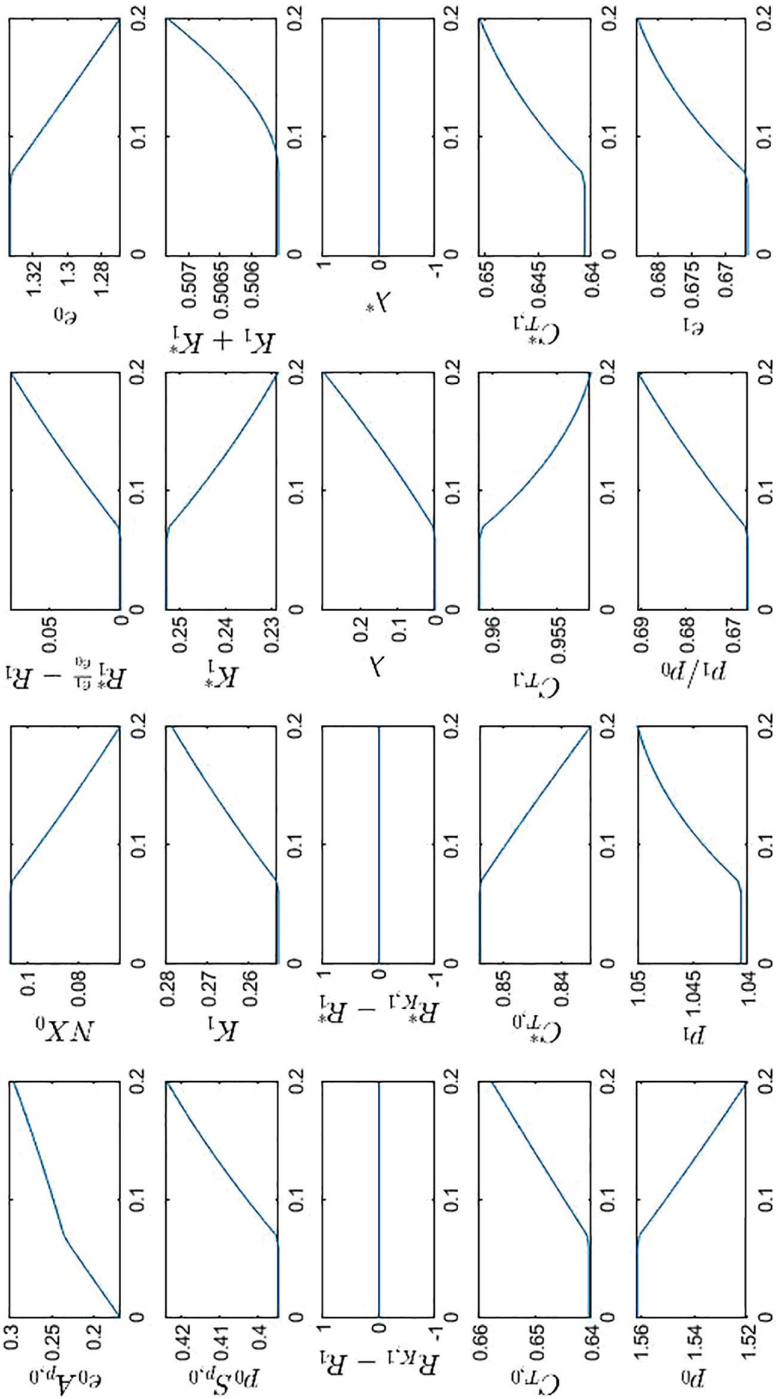

Note: Evolution of the model's equilibrium as capital inflows increase, starting from a frictionless point. $\theta = 0$, $\theta^* = 0$, $\Delta = 1/4$. The remaining parameter values are summarized in Table 1. Note that the home banks only get credit constrained once f^* exceeds a certain value.

Figure D.2. Effect of an Increase in $S_{CB,0}$ ($p_0 S_{CB,0}$ on x-axis)

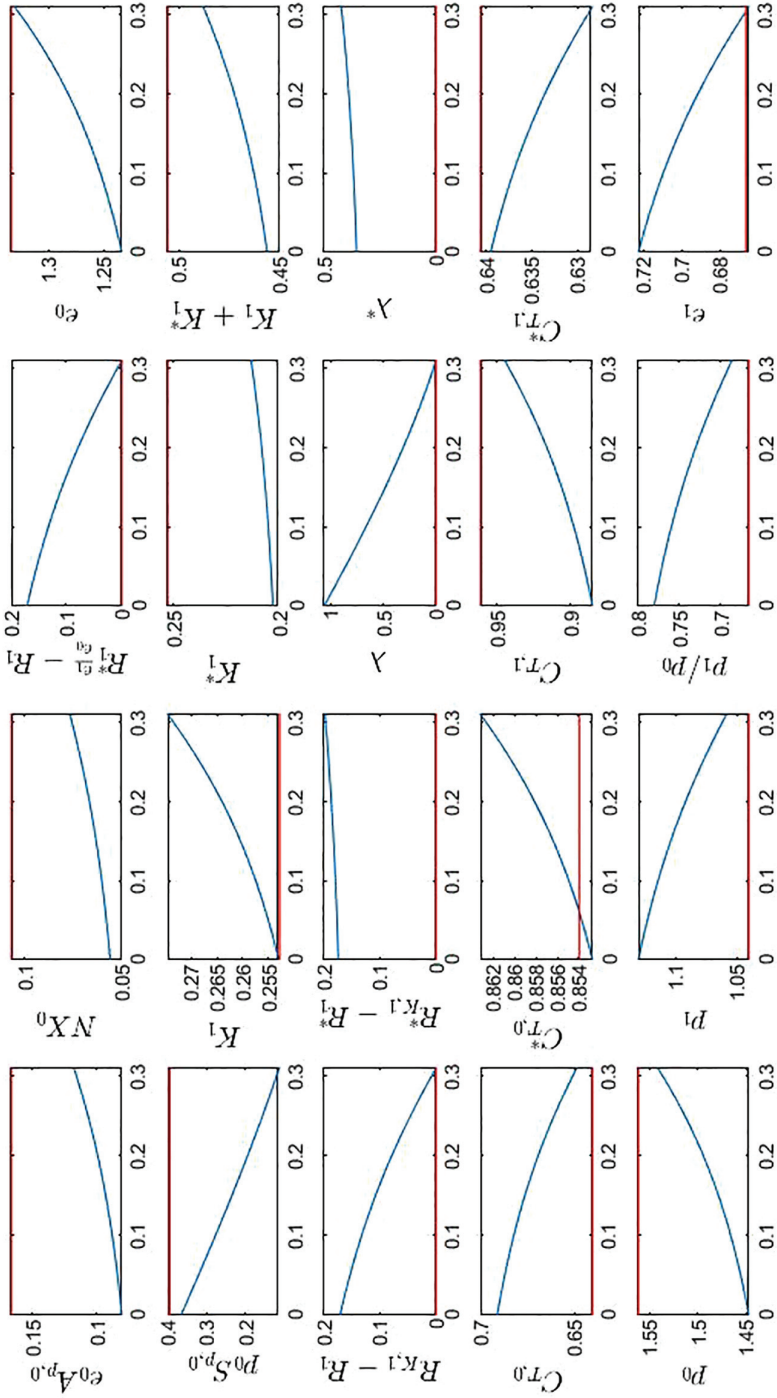

Note: Evolution of the model's equilibrium as central bank intermediation in the domestic investment market increases, starting from a point where there are limits to arbitrage in all financial markets. $\theta = 1/3$, $\theta^* = 1/3$, $\Delta = 1/3$. The remaining parameter values are summarized in Table 1. The red line represents the value in the frictionless equilibrium.

Figure D.3. Effect of an Increase in $A_{CB,0}$ ($e_0 A_{CB,0}$ on x-axis)

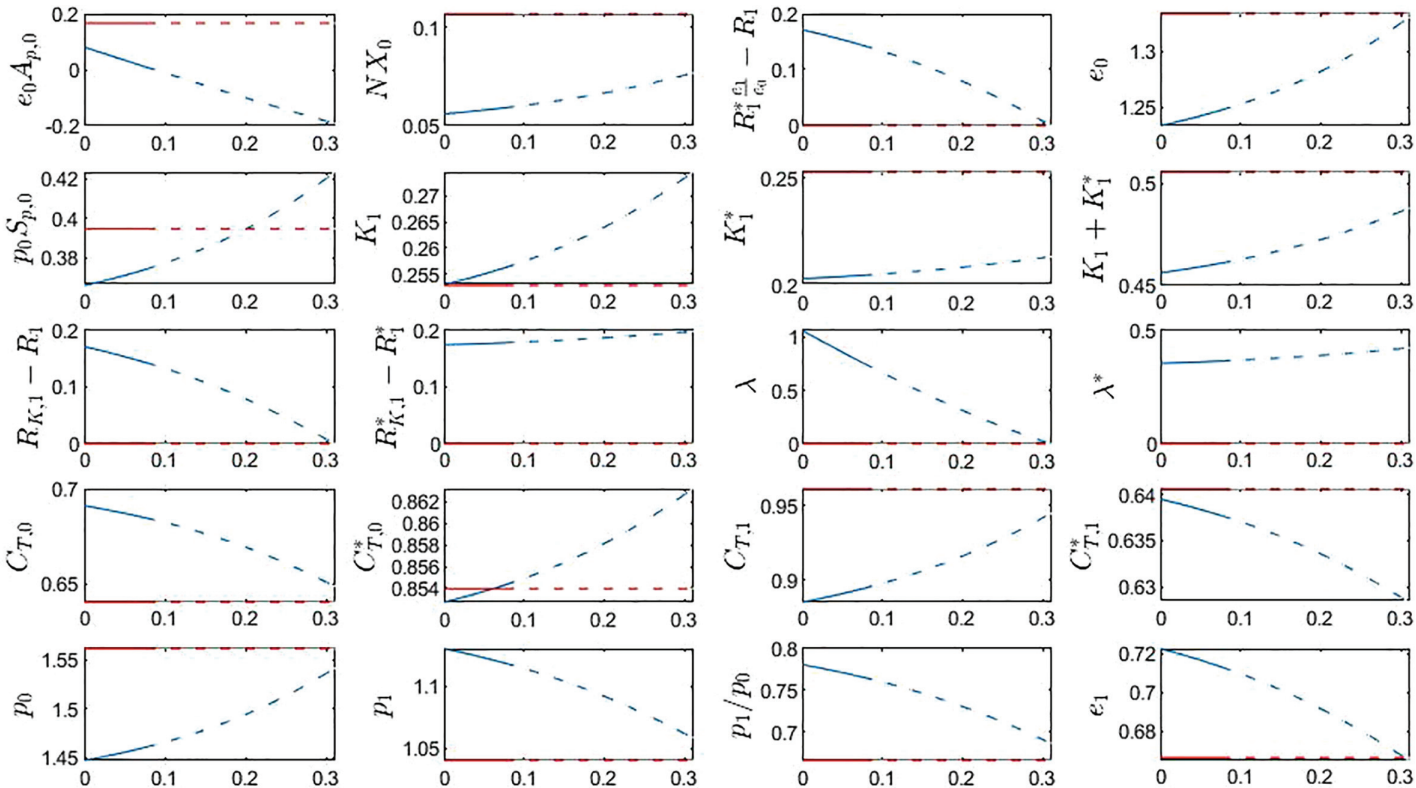

Note: Evolution of the model's equilibrium as central bank intermediation in the international credit market increases, starting from a point where there are limits to arbitrage in all financial markets. $\theta = 1/3$, $\theta^* = 1/3$, $\Delta = 1/3$. The remaining parameter values are summarized in Table 1. The red line represents the value in the frictionless equilibrium. The dashed part of the lines captures the range of FX interventions where the latter would require the home banks to go short in foreign bonds in order to fulfill the central bank's demand for these assets and hence covers a part where our model technically is not valid.

Appendix E. Proof: Metzler Diagram

For convenience, Figures 4, 2, and B.1 depict the reaction of the economy for the case where the respective friction parameter passes from being non-binding ($\lambda = 0$) to being binding ($\lambda > 0$). Due to the banks' positive equity capital, this always happens at some strictly positive value of θ, θ^*, or Δ, respectively, denoted by $\overline{\theta}$, $\overline{\theta^*}$, or $\overline{\Delta}$, which represent the highest possible values where the friction parameters are still non-binding.[34]

E.1 Investment Schedules

The two countries' investment schedules are given by market clearing in the investment markets ($S_{p,0} = K_1$ and $S^*_{p,0} = K^*_1$), by Equations (B.26) and (B.27), which relate the levels of capital and the (real) returns on the investment securities $\frac{p_0}{p_1}R_{k,1}$ and $\frac{p^*_0}{p^*_1}R^*_{K,1}$, and by Equations (5), (16), and (23), which define the relationships between the returns on the investment securities and the return on the bonds (R_1 and R^*_1, respectively) and result from the banks' optimization problem. Thus, the investment schedules (KK) and (KK^*) are

$$(KK) \qquad K_1 = \left(\frac{1}{\alpha}\left(\frac{p_0}{p_1}R_{k,1} - (1-\delta)\right)\right)^{\frac{1}{\alpha-1}} L,$$

$$\text{where} \quad \frac{p_0}{p_1}R_{k,1} = \frac{p_0}{p_1}R_1\left(1 + \frac{\lambda}{1+\lambda}\theta\right)$$

$$(KK^*) \qquad K^*_1 = \left(\frac{1}{\alpha}\left(\frac{p^*_0}{p^*_1}R^*_{K,1} - (1-\delta)\right)\right)^{\frac{1}{\alpha-1}} L^*,$$

$$\text{where} \quad \frac{p^*_0}{p^*_1}R^*_{K,1} = \frac{p^*_0}{p^*_1}R^*_1\left(1 + \frac{\lambda^*}{1+\lambda^*}\theta^*\right).$$

[34] Intuitively, for very low values of θ, θ^*, or Δ, respectively, the divertable part of a bank's assets will inevitably be lower than the bank's equity capital, which it would lose in case of misbehavior. Thus, the banks' incentive constraint will not be binding.

- **Effect of an increase in θ:**

  Using the concept of the total differential, one finds that for a given real interest rate $\frac{p_0}{p_1}R_1$, an increase in θ has the following effect on home capital:

$$\begin{aligned}
&\left.\frac{\mathrm{d}K_1}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&= \frac{1}{\alpha(\alpha-1)}\left(\frac{1}{\alpha}\left(\frac{p_0}{p_1}R_{k,1}-(1-\delta)\right)\right)^{\frac{2-\alpha}{\alpha-1}} \\
&\quad L\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&= -\frac{1}{\alpha(1-\alpha)}\left(\frac{L}{K_1}\right)^{\alpha-2} L\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} .
\end{aligned} \tag{E.1}$$

  The term in front of the final derivative is negative as $0 < \alpha < 1$, while the derivative itself is equal to the following expression:

$$\begin{aligned}
&\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&= \left.\frac{\mathrm{d}}{\mathrm{d}\theta}\left(\frac{p_0}{p_1}R_1\left(1+\frac{\lambda}{1+\lambda}\theta\right)\right)\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&= \frac{p_0}{p_1}R_1\left(\frac{1}{(1+\lambda)^2}\theta\left.\frac{\mathrm{d}\lambda}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} + \frac{\lambda}{1+\lambda}\right) \\
&= \frac{p_0}{p_1}R_1\overline{\theta}\left.\frac{\mathrm{d}\lambda}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} > 0,
\end{aligned} \tag{E.2}$$

  where the last steps follow from the fact that we look at the effect where the friction parameter passes from being non-binding to being binding ($\Rightarrow \lambda = 0$ and $\left.\frac{\mathrm{d}\lambda}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} > 0$),

which happens at the strictly positive value $\theta = \overline{\theta}$. Altogether, this implies that for a given real interest rate $\frac{p_0}{p_1} R_1$, an increase in θ leads to a decrease in the level of home capital, which corresponds to a negative shift in the home investment curve.

Likewise, one finds for foreign capital:

$$\left.\frac{\mathrm{d}K_1^*}{\mathrm{d}\theta}\right|_{\frac{p_0^*}{p_1^*} R_1^* \text{ constant}} = -\frac{1}{\alpha(1-\alpha)} \left(\frac{L^*}{K_1^*}\right)^{\alpha-2} L^* \left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*} R_{K,1}^*\right)}{\mathrm{d}\theta}\right|_{\frac{p_0^*}{p_1^*} R_1^* \text{ constant}}, \tag{E.3}$$

where

$$\left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*} R_{K,1}^*\right)}{\mathrm{d}\theta}\right|_{\frac{p_0^*}{p_1^*} R_1^* \text{ constant}} = \left.\frac{\mathrm{d}}{\mathrm{d}\theta}\left(\frac{p_0^*}{p_1^*} R_1^* \left(1 + \frac{\lambda^*}{1+\lambda^*}\theta^*\right)\right)\right|_{\frac{p_0^*}{p_1^*} R_1^* \text{ constant}} = 0. \tag{E.4}$$

Hence, given a foreign real interest rate $\frac{p_0^*}{p_1^*} R_1^*$, an increase in θ has no effect on foreign investment (remember that $\theta^* = 0$).

- **Effect of an increase in θ^*:**

  By the same reasoning as above, one finds:

$$\left.\frac{\mathrm{d}K_1}{\mathrm{d}\theta^*}\right|_{\frac{p_0}{p_1} R_1 \text{ constant}} = -\frac{1}{\alpha(1-\alpha)} \left(\frac{L}{K_1}\right)^{\alpha-2} L \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1} R_{k,1}\right)}{\mathrm{d}\theta^*}\right|_{\frac{p_0}{p_1} R_1 \text{ constant}}, \tag{E.5}$$

where

$$\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta^*}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = \frac{\mathrm{d}}{\mathrm{d}\theta^*}\left.\left(\frac{p_0}{p_1}R_1\left(1+\frac{\lambda}{1+\lambda}\theta\right)\right)\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = 0. \quad \text{(E.6)}$$

Hence, given a real interest rate $\frac{p_0}{p_1}R_1$, an increase in θ^* has no effect on home investment (remember that $\theta = 0$).

For the foreign investment curve, one finds:

$$\left.\frac{\mathrm{d}K_1^*}{\mathrm{d}\theta^*}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = -\frac{1}{\alpha(1-\alpha)}\left(\frac{L^*}{K_1^*}\right)^{\alpha-2} L^* \left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*}R_{K,1}^*\right)}{\mathrm{d}\theta^*}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}}, \quad \text{(E.7)}$$

where

$$\left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*}R_{K,1}^*\right)}{\mathrm{d}\theta^*}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = \frac{p_0^*}{p_1^*}R_1^*\overline{\theta^*}\left.\frac{\mathrm{d}\lambda^*}{\mathrm{d}\theta^*}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} > 0. \quad \text{(E.8)}$$

Hence, given a foreign real interest rate $\frac{p_0^*}{p_1^*}R_1^*$, an increase in θ^* leads to a decrease in the level of foreign capital, which corresponds to a negative shift in the foreign investment curve.

- **Effect of an increase in Δ:**
  For the home investment curve, one finds:

$$\left.\frac{\mathrm{d}K_1}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = -\frac{1}{\alpha(1-\alpha)}\left(\frac{L}{K_1}\right)^{\alpha-2} L \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}, \tag{E.9}$$

where (remember that $\theta = 0$):

$$\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = \left.\frac{\mathrm{d}}{\mathrm{d}\Delta}\left(\frac{p_0}{p_1}R_1\left(1+\frac{\lambda}{1+\lambda}\theta\right)\right)\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = 0. \tag{E.10}$$

Likewise, one finds for foreign capital:

$$\left.\frac{\mathrm{d}K_1^*}{\mathrm{d}\Delta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = -\frac{1}{\alpha(1-\alpha)}\left(\frac{L^*}{K_1^*}\right)^{\alpha-2} L^* \left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*}R_{K,1}^*\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}}, \tag{E.11}$$

where (remember that $\theta^* = 0$):

$$\left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*}R_{K,1}^*\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = \left.\frac{\mathrm{d}}{\mathrm{d}\Delta}\left(\frac{p_0^*}{p_1^*}R_1^*\left(1+\frac{\lambda^*}{1+\lambda^*}\theta^*\right)\right)\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = 0. \tag{E.12}$$

Hence, given a real interest rate, an increase in Δ does not have any effect the level of investment of either country.

E.2 Savings Schedules

The home country's saving schedule is described by the Euler equation (expressed in terms of traded goods, see Equations (5) and (7)), where the households' intertemporal budget constraint (given by $C_{NT,1} + p_1 C_{T,1} = R_1(p_0 Y_{T,0} - p_0 N_0 + Y_{NT,0} - C_{NT,0} - p_0 C_{T,0}) + w_1 L + p_1 N_1 + Y_{NT,1}$) is used to eliminate $p_1 C_{T,1}$:

$$\begin{aligned} p_0 C_{T,0} &= \frac{p_1 C_{T,1}}{\beta R_1} \\ \Leftrightarrow p_0 C_{T,0} &= \frac{1}{\beta R_1}(R_1(p_0 Y_{T,0} - p_0 N_0 + Y_{NT,0} - C_{NT,0} - p_0 C_{T,0}) \\ &\quad + w_1 L + p_1 N_1 + Y_{NT,1} - C_{NT,1}). \end{aligned}$$

Market clearing in the non-traded goods' sector implies that in equilibrium demand and endowment for non-traded goods always have to cancel each other out, and the non-financial firms' technology and optimization behavior ensure that labor income is a constant share of (nominal) output ($w_1 L = p_1(1-\alpha)Y_{T,1}$). Finally, the value of the equity capital in period 1, $p_1 N_1$, is given by Equation (B.4), where by market clearing $S_{p,0} = K_1$ and $e_0 A_{p,0} = p_0 NX_0 = p_0(Y_{T,0} - K_1 - C_{T,0})$. Thus, the home households' savings schedule (SS) is defined by

$$\begin{aligned} p_0 C_{T,0} &= \frac{1}{\beta R_1}\Bigg(R_1(p_0 Y_{T,0} - p_0 C_{T,0}) + p_1(1-\alpha)Y_{T,1} \\ &\quad + (R_{k,1} - R_1)\, p_0 K_1 + \left(R_1^* \frac{e_1}{e_0} - R_1\right) p_0 NX_0\Bigg) \\ \Leftrightarrow C_{T,0} &= \frac{1}{(1+\beta)\frac{p_0}{p_1} R_1}\Bigg(\frac{p_0}{p_1} R_1 Y_{T,0} + (1-\alpha)Y_{T,1} \\ &\quad + \frac{p_0}{p_1}(R_{k,1} - R_1)\, K_1 + \frac{p_0}{p_1}\left(R_1^* \frac{e_1}{e_0} - R_1\right) NX_0\Bigg). \end{aligned} \tag{E.13}$$

Likewise, the foreign country's saving schedule (SS^*) is implicitly given by

$$C^*_{T,0} = \frac{1}{(1+\beta^*)\,\frac{p^*_0}{p^*_1}R^*_1}\left(\frac{p^*_0}{p^*_1}R^*_1 Y^*_{T,0} + (1-\alpha)Y^*_{T,1} + \frac{p^*_0}{p^*_1}\left(R^*_{k,1} - R^*_1\right)K^*_1\right). \tag{E.14}$$

- **Effect of an increase in θ:**

  The effect of an increase in θ on the home country's saving curve can be found by differentiating Equation (E.13), holding $\frac{p_0}{p_1}R_1$ constant:

$$\begin{aligned}
&\left.\frac{\mathrm{d}C_{T,0}}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&= \frac{1}{(1+\beta)\,\frac{p_0}{p_1}R_1}\Bigg((1-\alpha)\frac{\partial Y_{T,1}}{\partial K_1}\left.\frac{\mathrm{d}K_1}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&\quad + \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&\quad K_1 + \frac{p_0}{p_1}(R_{k,1} - R_1)\left.\frac{\mathrm{d}K_1}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&\quad + \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R^*_1\frac{e_1}{e_0}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&\quad NX_0 + \frac{p_0}{p_1}\left(R^*_1\frac{e_1}{e_0} - R_1\right)\left.\frac{\mathrm{d}NX_0}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}\Bigg).
\end{aligned}$$

  As we consider the case where the countries are initially in the frictionless state, excess returns are zero and the two respective terms disappear. Furthermore, by using Equations (5) and (17), we find that $\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R^*_1\frac{e_1}{e_0}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = \frac{\mathrm{d}}{\mathrm{d}\theta}\left.\left(\frac{p_0}{p_1}R_1\left(1+\frac{\lambda}{1+\lambda}\Delta\right)\right)\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = 0$ (remember that

$\Delta = 0$). The two terms that are then still left represent the effect of the decrease in the households' second-period labor income (due to the lower level of capital) and the higher return on equity capital due to the increase in the excess return on home investment securities. Using Equation (E.1) to replace $\left.\frac{\mathrm{d}K_1}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}$ and substituting the marginal product to capital $\left(\frac{\partial Y_{T,1}}{\partial K_1} = \alpha\left(\frac{L}{K_1}\right)^{1-\alpha}\right)$, one finds that they just cancel each other out:

$$
\begin{aligned}
&\left.\frac{\mathrm{d}C_{T,0}}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&= \frac{1}{(1+\beta)\frac{p_0}{p_1}R_1}\left((1-\alpha)\frac{\partial Y_{T,1}}{\partial K_1}\left.\frac{\mathrm{d}K_1}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}\right. \\
&\quad \left. + \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} K_1\right) \\
&= \frac{1}{(1+\beta)\frac{p_0}{p_1}R_1}\left((1-\alpha)\alpha\left(\frac{L}{K_1}\right)^{1-\alpha}\right. \\
&\quad \frac{-1}{\alpha(1-\alpha)}\left(\frac{L}{K_1}\right)^{\alpha-2} L \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} \\
&\quad \left. + \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} K_1\right) \\
&= \frac{1}{(1+\beta)\frac{p_0}{p_1}R_1}\left(-K_1\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}\right. \\
&\quad \left. + \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\theta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} K_1\right) \\
&= 0.
\end{aligned}
$$

Hence, for a given real interest rate, an increase in θ has no effect on the level of consumption in the home country.

Likewise, by differentiating Equation (E.14), holding $\frac{p_0^*}{p_1^*}R_1^*$ constant, one finds the effect of an increase in θ on the foreign country's saving curve:

$$\begin{aligned}
&\left.\frac{\mathrm{d}C_{T,0}^*}{\mathrm{d}\theta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} \\
&= \frac{1}{(1+\beta^*)\,\frac{p_0^*}{p_1^*}R_1^*}\Bigg((1-\alpha)\frac{\partial Y_{T,1}^*}{\partial K_1^*}\left.\frac{\mathrm{d}K_1^*}{\mathrm{d}\theta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} \\
&\quad + \left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*}R_{k,1}^*\right)}{\mathrm{d}\theta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} K_1^* \\
&\quad + \frac{p_0^*}{p_1^*}(R_{k,1}^* - R_1^*)\left.\frac{\mathrm{d}K_1^*}{\mathrm{d}\theta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}}\Bigg) \\
&= 0.
\end{aligned}$$

We know from the analysis of the foreign investment curves that the first two expressions in the big brackets are equal to zero (see Equations (E.3) and (E.4)), and given that the economy is initially in a frictionless state, excess returns are zero as well. Hence, for a given real interest rate, an increase in θ has no effect on the foreign country's consumption.

- **Effect of an increase in θ^*:**

  Following the same reasoning as in the case of an increase in θ, one finds that $\left.\frac{\mathrm{d}C_{T,0}}{\mathrm{d}\theta^*}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = \left.\frac{\mathrm{d}C_{T,0}^*}{\mathrm{d}\theta^*}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = 0$. Thus, an increase in θ^* does not lead to a shift in the two countries' saving schedules.

- **Effect of an increase in Δ:**

  The effect of an increase in Δ on the home country's consumption, holding $\frac{p_0^*}{p_1^*}R_1^*$ constant:

$$\left.\frac{\mathrm{d}C_{T,0}}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}$$

$$= \frac{1}{(1+\beta)\,\frac{p_0}{p_1}R_1}\Bigg((1-\alpha)\frac{\partial Y_{T,1}}{\partial K_1}\left.\frac{\mathrm{d}K_1}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}$$

$$+ \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}$$

$$K_1 + \frac{p_0}{p_1}(R_{k,1}-R_1)\left.\frac{\mathrm{d}K_1}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}$$

$$+ \left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_1^*\frac{e_1}{e_0}\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}$$

$$NX_0 + \frac{p_0}{p_1}\left(R_1^*\frac{e_1}{e_0} - R_1\right)\left.\frac{\mathrm{d}NX_0}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}\Bigg).$$

Again, excess returns are zero and the two respective terms disappear. From the analysis above, we know that $\left.\frac{\mathrm{d}K_1}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = 0$ (see Equations (E.9) and (E.10)). Furthermore, by using Equations (5) and (16), we find that $\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_{k,1}\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = \left.\frac{\mathrm{d}}{\mathrm{d}\Delta}\left(\frac{p_0}{p_1}R_1\left(1+\frac{\lambda}{1+\lambda}\theta\right)\right)\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} = 0$ (remember that $\theta = 0$). The one term that is still left represents the effect of the higher return on equity capital due to the increase in excess returns on foreign assets and equals (using Equation (17)):

$$\left.\frac{\mathrm{d}C_{T,0}}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}}$$

$$= \frac{1}{(1+\beta)\,\frac{p_0}{p_1}R_1}\left(\left.\frac{\mathrm{d}\left(\frac{p_0}{p_1}R_1^*\frac{e_1}{e_0}\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0}{p_1}R_1 \text{ constant}} NX_0\right)$$

$$= \frac{1}{(1+\beta)\,\frac{p_0}{p_1}R_1}\left(\frac{\mathrm{d}}{\mathrm{d}\Delta}\left(\frac{p_0}{p_1}R_1\left(1+\frac{\lambda}{1+\lambda}\Delta\right)\right)\Big|_{\frac{p_0}{p_1}R_1 \text{ constant}} NX_0\right)$$

$$= \frac{1}{(1+\beta)\,\frac{p_0}{p_1}R_1}\left(\frac{p_0}{p_1}R_1\,\overline{\Delta}\frac{\mathrm{d}\lambda}{\mathrm{d}\Delta}\Big|_{\frac{p_0}{p_1}R_1 \text{ constant}} NX_0\right)$$

$$= \frac{1}{(1+\beta)}\left(\overline{\Delta}\frac{\mathrm{d}\lambda}{\mathrm{d}\Delta}\Big|_{\frac{p_0}{p_1}R_1 \text{ constant}} NX_0\right)$$

$$= \frac{1}{(1+\beta)}\left(\overline{\Delta}\frac{\mathrm{d}\lambda}{\mathrm{d}\Delta}\Big|_{\frac{p_0}{p_1}R_1 \text{ constant}} (Y_{T,0}-K_1-C_{T,0})\right) > 0.$$

As we look at the effect where the friction parameter passes from being non-binding to being binding, we know that $\frac{\mathrm{d}\lambda}{\mathrm{d}\Delta}\Big|_{\frac{p_0}{p_1}R_1 \text{ constant}} > 0$. Hence, given a real interest rate, an increase in Δ leads to an increase in consumption due to the higher return on home equity capital, which in turn corresponds to a negative shift in the home country's saving curve.

On the other hand, as the foreign banks have no international portfolio and their equity capital is independent of the excess return on foreign international transactions, there is no shift in the foreign country's saving curve:

$$\frac{\mathrm{d}C^*_{T,0}}{\mathrm{d}\Delta}\Big|_{\frac{p^*_0}{p^*_1}R^*_1 \text{ constant}}$$

$$= \frac{1}{(1+\beta^*)\,\frac{p^*_0}{p^*_1}R^*_1}\left((1-\alpha)\frac{\partial Y^*_{T,1}}{\partial K^*_1}\,\frac{\mathrm{d}K^*_1}{\mathrm{d}\Delta}\Big|_{\frac{p^*_0}{p^*_1}R^*_1 \text{ constant}}\right.$$

$$+\frac{\mathrm{d}\left(\frac{p^*_0}{p^*_1}R^*_{k,1}\right)}{\mathrm{d}\Delta}\Big|_{\frac{p^*_0}{p^*_1}R^*_1 \text{ constant}} K^*_1 + \frac{p^*_0}{p^*_1}(R^*_{k,1}-R^*_1)$$

$$\left.\frac{\mathrm{d}K^*_1}{\mathrm{d}\Delta}\Big|_{\frac{p^*_0}{p^*_1}R^*_1 \text{ constant}}\right) = 0.$$

Again, excess returns in the initial frictionless state are zero, and the first two terms drop out as well: $\left.\frac{\mathrm{d}\left(\frac{p_0^*}{p_1^*}R_{K,1}^*\right)}{\mathrm{d}\Delta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = \frac{\mathrm{d}}{\mathrm{d}\Delta}\left(\frac{p_0^*}{p_1^*}R_1^*\left(1+\frac{\lambda^*}{1+\lambda^*}\theta^*\right)\right)\Big|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}} = 0$ (remember that $\theta^* = 0$) and therefore $\left.\frac{\mathrm{d}K_1^*}{\mathrm{d}\Delta}\right|_{\frac{p_0^*}{p_1^*}R_1^* \text{ constant}}$ (see Equation (E.11)).

Appendix F. CPI-Based Real Exchange Rate

The exchange rate e_t is equal to the relative price of the non-traded goods, i.e., the numéraires, respectively, in our model. To find the CPI-based real exchange rate, we first need to derive the price indices. This is done by replacing $C_{NT,t}$ and $C_{T,t}$ in the consumption index by the demand functions resulting from the intratemporal optimization problem (see Equations (6) and (7)):

$$\begin{aligned} C_t &= \left(C_{NT,t}^{\chi} C_{T,t}\right)^{\frac{1}{1+\chi}} \\ &= \left(\left(\frac{\chi}{1+\chi}\left(\frac{1}{P_t}\right)^{-1} C_t\right)^{\chi}\left(\frac{1}{1+\chi}\left(\frac{p_t}{P_t}\right)^{-1} C_t\right)\right)^{\frac{1}{1+\chi}} \\ &= \frac{\chi^{\frac{\chi}{1+\chi}}}{1+\chi}\left(\frac{1}{p_t}\right)^{\frac{1}{1+\chi}} P_t C_t. \end{aligned}$$

Hence,

$$P_t = \frac{1+\chi}{\chi^{\frac{\chi}{1+\chi}}}\, p_t^{\frac{1}{1+\chi}}. \tag{F.1}$$

Similarly, the foreign price index is found to be

$$P_t^* = \frac{1+\chi}{\chi^{\frac{\chi}{1+\chi}}}\, p_t^{*\frac{1}{1+\chi}}. \tag{F.2}$$

It follows that the CPI-based real exchange rate, defined as the ratio of the price indices multiplied by the relative price of the two numéraires, is given by[35]

$$\mathcal{E}_t \equiv \frac{P_t^*}{P_t} e_t = \left(\frac{p_t^*}{p_t}\right)^{\frac{1}{1+\chi}} e_t = \left(\frac{1}{e_t}\right)^{\frac{1}{1+\chi}} e_t$$
$$= e_t^{\frac{\chi}{1+\chi}}. \tag{F.3}$$

Thus, the exchange rate as we define it in our model is very closely related to the CPI-based real exchange rate. Whenever e_t is larger (smaller) than 1, this also holds for $\mathcal{E}_t$.

References

Adrian, T., C. Erceg, J. Lindé, P. Zabczyk, and J. Zhou. 2020. "A Quantitative Model for the Integrated Policy Framework." IMF Working Paper.

Andrés, J., J. D. López-Salido, and E. Nelson. 2004. "Tobin's Imperfect Asset Substitution in Optimizing General Equilibrium." *Journal of Money, Credit and Banking* 36 (4): 665–90.

Bernanke, B. S. 2009. "The Crisis and the Policy Response." Stamp Lecture, London School of Economics, London, January 13.

Borio, Claudio E. V., R. N. McCauley, P. McGuire, and V. Sushko. 2016. "Covered Interest Parity Lost: Understanding the Cross-Currency Basis." *BIS Quarterly Review* (September): 45–64.

Branson, W. H., and D. W. Henderson. 1985. "The Specification and Influence of Asset Markets." In *Handbook of International Economics* Vol. 2, ed. R. W. Jones and P. B. Kenen, 749–805 (chapter 15). Elsevier.

Cavallino, P. 2019. "Capital Flows and Foreign Exchange Intervention." *American Economic Journal: Macroeconomics* 11 (2): 127–70.

Christensen, J. H. E., and S. Krogstrup. 2019. "Transmission of Quantitative Easing: The Role of Central Bank Reserves." *Economic Journal* 129 (617): 249–72.

[35] $\mathcal{E}_t$ is defined to be the price of a foreign consumption bundle expressed in terms of home consumption bundles. Thus, if $\mathcal{E}_t < 1$, one consumption bundle in the home country gives more than one consumption bundle in the foreign country.

Coenen, G., P. Karadi, S. Schmidt, and A. Warne. 2018. "The New Area-Wide Model II: An Extended Version of the ECB's Micro-founded Model for Forecasting and Policy Analysis with a Financial Sector." ECB Working Paper No. 2200.

Dedola, L., G. Georgiadis, J. Grab, and A. Mehl. 2021. "Does a Big Bazooka Matter? Quantitative Easing Policies and Exchange Rates." *Journal of Monetary Economics* 117 (January): 489–506.

Dedola, L., P. Karadi, and G. Lombardo. 2013. "Global Implications of National Unconventional Policies." *Journal of Monetary Economics* 60 (1): 66–85.

Dornbusch, R. 1976. "Expectations and Exchange Rate Dynamics." *Journal of Political Economy* 84 (6): 1161–76.

Du, W., A. Tepper, and A. Verdelhan. 2018. "Deviations from Covered Interest Rate Parity." *Journal of Finance* 73 (3): 915–57.

Eggertsson, G. B., and M. Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* (1): 139–233

Engel, C. 2014. "Exchange Rates and Interest Parity." In *Handbook of International Economics*, Vol. 4, ed. G. Gopinath, E. Helpman, and K. Rogoff, 453–48 (chapter 8). Elsevier.

Fuhrer, L. M., T. Nitschka, and D. Wunderli. 2021. "Central Bank Reserves and Bank Lending Spreads." *Applied Economics Letters* 28 (15): 1301–05.

Gabaix, X., and M. Maggiori. 2015. "International Liquidity and Exchange Rate Dynamics." *Quarterly Journal of Economics* 130 (3): 1369–1420.

Gertler, M., and P. Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58 (1): 17–34.

———. 2013. "QE 1 vs. 2 vs. 3. . . : A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool." *International Journal of Central Banking* 9 (1): 5–53.

Gertler, M., and N. Kiyotaki. 2011. "Financial Intermediation and Credit Policy in Business Cycle Analysis." In *Handbook of Monetary Economics*, Vol. 3, ed. B. M. Friedman and M. Woodford, 547–99 (chapter 11). Elsevier.

Henderson, D. W., and K. S. Rogoff. 1982. "Negative Net Foreign Asset Positions and Stability in a World Portfolio Balance Model." *Journal of International Economics* 13 (1–2): 85–104.

Jordan, T. 2011. "Does the Swiss National Bank Need Equity?" Speech at Statistisch-Volkswirtschaftliche Gesellschaft, Basel, Switzerland, September 28.

Kollmann, R. 2016. "Discussion of 'Financial Intermediation in a Global Environment'." *International Journal of Central Banking* 12 (3): 345–51.

Kouri, P. J. K. 1976. "The Exchange Rate and the Balance of Payments in the Short Run and in the Long Run: A Monetary Approach." *Scandinavian Journal of Economics* 78 (2): 280–304.

Kumhof, M. 2010. "On the Theory of Sterilized Foreign Exchange Intervention." *Journal of Economic Dynamics and Control* 34 (8): 1403–20.

Kwon, D. 2020. "Risk Shocks and Credit Spreads." *Journal of Macroeconomics* 64 (June): Article 103208.

Leutert, J. 2018. "The Swiss Franc Safety Premium." *Swiss Journal of Economics and Statistics* 154 (1): 1–21.

Maggiori, M. 2013. "The U.S. Dollar Safety Premium."

Menkhoff, L. 2010. "High-Frequency Analysis of Foreign Exchange Interventions: What Do We Learn?" *Journal of Economic Surveys* 24 (1): 85–112.

Neely, C. J. 2005. "An Analysis of Recent Studies of the Effect of Foreign Exchange Intervention." *Review* (Federal Reserve Bank of St. Louis) 87 (6): 685–717.

Nuguer, V. 2016. "Financial Intermediation in a Global Environment." *International Journal of Central Banking* 12 (3, September): 291–344.

Obstfeld, M., and K. Rogoff. 1995. "Exchange Rate Dynamics Redux." *Journal of Political Economy* 103 (3): 624–60.

Saadi Sedik, T., L. I. Jacome H, and A. Ziegenbein. 2018. "Is Credit Easing Viable in Emerging and Developing Economies? An Empirical Approach." IMF Working Paper No. 18/43.

Sarno, L., and M. P. Taylor. 2001. "Official Intervention in the Foreign Exchange Market: Is It Effective and, If So, How Does It Work?" *Journal of Economic Literature* 39 (3): 839–68.

Shleifer, A., and R. W. Vishny. 1997. "The Limits of Arbitrage." *Journal of Finance* 52 (1): 35–55.

Tobin, J. 1958. "Liquidity Preference as Behavior Towards Risk." *Review of Economic Studies* 25 (2): 65–86.

———. 1969. "A General Equilibrium Approach To Monetary Theory." *Journal of Money, Credit and Banking* 1 (1): 15–29.

Vayanos, D., and J.-L. Vila. 2009. "A Preferred-Habitat Model of the Term Structure of Interest Rates." NBER Working Paper No. 15487.

Funding Behavior of Debt Management Offices and the ECB's Public Sector Purchase Program*

Katharina Plessen-Mátyás, Christoph Kaufmann, and Julian von Landesberger
European Central Bank

This paper investigates whether the funding behavior of euro area debt management offices (DMOs) changed with the start of the ECB's Public Sector Purchase Program (PSPP). Our results show that (i) lower yield levels and (ii) PSPP purchases supported higher maturities at issuance. The former indicates a behavior of "locking in low rates for longer," while the latter suggests the existence of an additional "demand effect" of the PSPP on DMO strategies beyond the PSPP's effect via yields. The combined impact of the PSPP via these channels amounts to maturity extensions at issuance of about one year in our estimation.

JEL Codes: E51, E58, E63, H63.

*We wish to thank Christophe Beuve, Ulrich Bindseil, Alexander Duering, Björn Fischer, Benjamin Friedman, Bryan Gurhy, Wolfgang Lemke, Annette Kamps, Imène Rahmouni-Rousseau, Isabel Schnabel, Quentin Vandeweyer, and Ralph Weidenfeller as well as seminar and conference participants at the European Central Bank, the European Economic Association Annual Congress 2020, and the Dutch Ministry of Finance for helpful comments and suggestions. We also thank Dagmar Hartwig Lojsch and Niki Theodosiou for help with accessing and understanding the data on government debt issuance. The views expressed in this paper are those of the authors only and do not necessarily reflect the views of the European Central Bank or the Eurosystem. Author contact: Katharina Plessen-Mátyás, Directorate General Market Operations, European Central Bank, Sonnemannstrasse 20, 60314 Frankfurt am Main, Germany; e-mail: katharina.von_plessen@ecb.europa.eu. Christoph Kaufmann, Directorate General Macroprudential Policy and Financial Stability, European Central Bank; christoph.kaufmann@ecb.europa.eu. Julian von Landesberger, Directorate General Market Operations, European Central Bank; julian.von_landesberger@ecb.europa.eu.

1. Introduction

The conduct of large-scale asset purchase programs by central banks can affect government bond supply. This paper investigates how the funding behavior of public debt management offices (DMOs) has been affected by the European Central Bank's (ECB) Public Sector Purchase Program (PSPP).[1] We analyze, in particular, the maturity of newly issued securities before the start of the PSPP and over the first entire phase of net asset purchases in a panel of seven euro area countries between December 2009 and April 2019.

In response to changing funding conditions and subject to adequate demand across maturities, DMOs regularly optimize the maturity structure of debt with regard to a trade-off between debt servicing costs, which usually increase with debt maturity, and refinancing risks, which decrease with debt maturity. This optimization has been formalized, e.g., by Greenwood, Hanson, and Stein (2015). By alleviating funding conditions, asset purchase programs are likely to improve the trade-off faced by DMOs, which should have observable impacts on their financing in terms of cost of funding and/or maturity at issuance.

It is well established that asset purchase programs led to a decrease in sovereign bond yields over the last years.[2] Estimates from term structure models for the euro area imply that the PSPP compressed sovereign bond term premia via the duration channel significantly (Eser et al. 2019). The yield compression implies that DMOs can, ceteris paribus, fund cheaper in particular at longer maturities.

Beyond the direct effect resulting from the compression of the yield curve, an additional "demand" effect of the PSPP on funding behavior could arise when DMOs expect that a larger amount of eligible longer-dated securities could be absorbed by the market and in view of the ECB's lower price sensitivity relative to private-sector market participants.

[1]Debt management offices are entities that are operationally responsible for public debt management. They can either be part of the ministry of finance or the ministry can delegate operational responsibility to them. See, for example, Wolswijk and de Haan (2005) for further details.

[2]See, for example, D'Amico et al. (2012), D'Amico and King (2013), Li and Wei (2013), Altavilla, Carboni, and Motto (2015), Andrade et al. (2016), and De Santis and Holm-Hadulla (2020).

The Italian DMO, for example, notes in its 2016 annual report that "[the] Treasury was able to issue large volumes of debt with maturities of more than 10 years while securing a much lower extra cost than historical average. These issuance choices became possible mainly thanks to two factors: on the one hand—obviously—PSPP that can absorb a substantial quantity of bonds of all maturities including the longer term ones (up to 31 years of residual maturity); on the other hand, a relatively high number of investors shifting to very long maturities, as a result of the strong reduction of yields on shorter maturities traditionally chosen by these investors" (Dipartimento del Tesoro 2016, p. 28).

This excerpt suggests that a demand effect was indeed perceived by the Italian DMO during 2016, whereby stronger issuance of longer-dated securities became possible due to the additional demand for such bonds in the secondary market, both from the ECB and from other investors.[3]

This paper quantifies the importance of the direct yield effect and the demand effect and addresses the following research questions: Did the PSPP lead to an extension of maturities issued by euro area DMOs via the lowering of government bond yields, as DMOs wanted to lock in low rates for longer? Are there additional demand effects arising from the PSPP, which affected the issuance maturities targeted by DMOs beyond the PSPP's effect on yields? To address these questions, we empirically investigate the relationship between the weighted average maturity (WAM) at issuance and (i) the cost of issuance, measured by government bond yields and (ii) a demand variable, measured by PSPP gross purchases as a share of total issuance. We estimate our empirical models using the common correlated effects estimator by Pesaran (2006) to account for the strong cross-sectional dependencies that we find in our data. We also estimate local projections (see Jordà 2005) of WAM at issuance to high-frequency yield shocks taken from Altavilla et al. (2019) that allow for a causal interpretation of our results. While some empirical evidence has started emerging on the link between yields and DMO issuance behavior, the relationship of the latter to central bank

[3]Under the PSPP, the ECB conducts purchases in the secondary market only and hence does not buy securities from the issuers directly in the primary market.

asset purchases has to the best of our knowledge not been analyzed systematically to date.

Our results show that a 1 percentage point decrease in 10-year government bond yields leads to an increase of the WAM at issuance by about five months. A 10 percentage points higher ratio of PSPP gross purchases to issuance volume contributes to an increase of the WAM at issuance by about one month. The statistical significance and order of magnitude of these findings is robust over sub-samples, in the presence of several DMO-specific and macro control variables, as well as to different estimation methods. Spanning from December 2009 to April 2019, our sample covers a period of high policy relevance. Specifically, the period covers the euro area sovereign debt crisis from 2010 to 2013 as well as the implementation of the full first period of net asset purchases under the PSPP that was announced in January 2015 and ended in December 2018.

We gauge the economic significance of our results by assessing estimates for the PSPP's effect on term premia and data on PSPP gross purchases as a share of total issuance through the lens of our model. For the euro area "Big 4" countries—France, Germany, Italy, and Spain—that represent three-quarters of the euro area GDP, our estimates suggest the following impact of the PSPP on public funding maturities. (i) The reduced yield level over the PSPP episode led to a lengthening of issuance maturities by seven months on average. (ii) The increased demand for PSPP-eligible bonds led to a lengthening of issuance maturities by six months on average. The overall monthly average effect of the PSPP on issuance maturities is, accordingly, an increase of about one year, which compares to the average maturity of all debt outstanding of Germany, France, Italy, and Spain before the PSPP of about six years.

Furthermore, we present a segmentation of our sample into countries that were more and less vulnerable to fiscal stress during the European sovereign debt crisis, respectively. Given the higher uncertainty and potentially higher rollover risks, more vulnerable countries may have higher incentives to make use of a favorable market environment. Indeed, we find that these countries increase the maturity of their issuance relatively stronger in response to yield changes and PSPP purchases than less vulnerable countries.

The results of this paper are a basis for further work on the economic impact of maturity extension by DMOs during the PSPP

and its potential relevance for the transmission of monetary policy. A key element in the transmission of accommodative monetary policy to the real economy is how lower market rates—in particular, for longer-dated debt—improve financing conditions of borrowers. This alleviation in funding constraints contributes to an increase in aggregate demand and to a reduction of refinancing risk for the borrower.

A decline in public debt servicing costs and reductions in its refinancing risk enables governments to increase general public spending, alleviates financing restrictions for longer-term projects, and/or reduces the overall tax burden and tax variability. Additionally, a possible improvement in debt sustainability may lead to a reduction of credit risk and thereby of yields, thus further propagating the beneficial effect.

Similar transmission effects could potentially be present across all sectors of the economy. In fact, non-financial corporations can also be found to adjust their maturity structure in response to central bank purchases of public debt. Following the "gap-filling" theory by Greenwood, Hanson, and Stein (2010), firms issue longer when central bank asset purchases reduce the effective supply of long-term debt in the markets, given the inelastic demand of preferred-habitat investors, such as insurance corporations.[4] This maturity extension of private-sector debt can also be understood as an intended consequence of central banks' asset purchase programs, as it alleviates funding constraints of the private sector.

Our findings can also have implications for research that quantifies the duration channel of quantitative easing (QE) that works through a crowding-out of price-sensitive investors, who rebalance their portfolios towards riskier assets. When governments issue longer maturities, there is a crowding-in of price-sensitive investors due to the increased supply of longer-dated bonds, which are not purchased by the central bank, with a positive overall impact on government bond yields. This is in line with the arguments presented by Greenwood et al. (2014), who show that the issuance of longer-dated securities by the U.S. Treasury following the start of QE counteracted between one-third and two-thirds of the impact that QE had

[4]Badoer and James (2016) and Foley-Fisher, Ramcharan, and Yu (2016) provide empirical evidence in support of this theory.

on yield levels. Performing this analysis based on our data and estimation results, we find that the maturity lengthening by euro area DMOs due to the yield and demand effect may have offset one-third of the PSPP's effect on term premia. Estimations of the duration channel, which treat government funding maturity as unresponsive to yield and demand effects of QE (i.e., as an exogenous variable) may, therefore, overstate the total impact of the QE on yields since a portion may have been counteracted through maturity lengthening by governments.

The rest of the paper is structured as follows. After giving a review of the existing literature in Section 2, we provide a description of DMOs' objectives and behavior in Section 3. Section 4 describes the data set and its properties as well as our econometric model to analyze DMO behavior. All regression and local projection results as well as an illustration of their economic significance are provided in Section 5. We conclude the analysis in Section 6.

2. Review of the Existing Literature

This section offers a summary of the related literature focusing on three strands relevant for this work. The first strand of literature analyzes characteristics as well as monetary policy implications of the maturity structure of government debt. The second strand argues that the maturity structure of debt issuance can be used as a tool of macroeconomic stabilization policy itself. The third strand of literature analyzes the behavior of DMOs empirically.

Related to the first strand of literature, Vayanos and Vila (2021) formalize a term structure model of preferred-habitat investors, where risk-averse arbitrageurs conduct substitution across maturities. In their model, scarcity of securities that have a preferred-habitat investor base can drive up bond prices. The total supply of bonds is thereby a determinant of their yields and affects the market price for duration risk.[5] Building on this model, Greenwood and Vayanos (2014) investigate how the supply and maturity distribution of public debt affects bond yields and expected returns. They

[5]Duration risk measures the sensitivity of the value of a fixed-income asset or portfolio to a change in interest rates.

find a positive relationship between maturity-weighted debt-to-GDP and longer-dated bond yields.

Krishnamurthy and Vissing-Jorgensen (2012) conduct an empirical analysis of the aggregate demand for treasury debt showing that changes in treasury supply have large effects on a variety of yield spreads, such as for safety and liquidity, due to preferred-habitat investors. Greenwood et al. (2014) argue that some types of, in particular, short-term government debt securities are cash-like due to safe-haven/liquidity characteristics and that the marginal holder of long-term government debt is a specialized fixed-income investor, who demands compensation for bearing interest rate risk. These two papers underpin the notion of scarcity and supply effects, according to which changes in the supply behavior of DMOs can have an effect on government bond yields.

Potential interactions between central banks' asset purchase programs (APP) and government debt issuance have gained attention in the literature since the start of such programs in the United States in 2008. Li and Wei (2013) and Eser et al. (2019) develop term structure models to estimate the effects of central banks' asset purchases in the United States and the euro area, respectively. Both papers consider bond supply and duration factors in their models. The results by Li and Wei (2013) imply that the Federal Reserve's QE programs until 2011 in sum reduced the 10-year U.S. Treasury yield by about 100 basis points. Eser et al. (2019) estimate that the ECB's APP has reduced the 10-year term premium in the euro area by 95 basis points. Greenwood et al. (2014) argue that while the Federal Reserve's QE program led to a sizable reduction in 10-year U.S. Treasury yields of 137 basis points, the simultaneous impact of maturity extension by the U.S. Treasury counteracted this effect by 48 basis points. The paper does not assess the relationship between large-scale asset purchases and the funding behavior of DMOs systematically using econometric methods, though. The quantification in this work is based on the growth in the total maturity of outstanding debt by the U.S. Treasury during the QE implementation and point estimates for the 10-year yield impact of QE in other academic papers. Our paper shows that the weighted average maturity of newly issued government debt in the euro area reacts significantly to yield changes and central bank asset purchases. Using the approach by Greenwood et al. (2014), we find that in the euro area about

33 percent of the PSPP's term premium effect may have been counteracted by longer-maturity issuance. The results from our paper could therefore be used to serve as an input to term structure models that estimate the effect and persistence of QE programs.

A second strand of literature analyzes the role of the maturity structure of debt as a tool for macroeconomic stabilization policy. Leong (1999) and Wolswijk and de Haan (2005) note a discrepancy between the academic debate and practice when it comes to public debt management. While much of the scientific literature focuses on macroeconomic stabilization goals, DMO practitioners take a more microeconomic approach by focusing on the cost-risk tradeoff inherent in an upward-sloping yield curve. The macroeconomic literature often views public debt management from the perspective of a government optimization problem, where DMO and government are one single entity. This approach ignores principal-agent problems that could arise, for example, due to potentially different planning horizons between governments searching for re-election and DMOs having a long-term perspective.

Tobin (1963) argues that governments should follow a countercyclical debt maturity policy for macroeconomic stabilization purposes. He argues that governments should issue longer-dated maturities during economic expansions to drive up long-term interest rates, while the minimization of financing costs is considered a secondary priority and risk minimization is not considered. Friedman (1992) studies the proposition of Tobin (1963) empirically to quantify the impact of debt-management policies on both interest rates and real economic activity. The simulations suggest, for a given budget deficit and therefore a given amount of debt to be issued, that long-term bond yields fall if a government issued short- rather than long-term securities. This in turn stimulates business investment, residential construction, and other interest-sensitive elements of aggregate spending. Angeletos (2002) studies the optimal maturity structure of public debt in a general equilibrium model. He shows that a broad range of Arrow-Debreu allocations are implementable when the government has the possibility of issuing debt at different maturities. Optimal policy consists of issuing long-term debt, which is used to invest into short-term debt as a reserve fund. The government can draw from this fund in bad times to stabilize the economy. Relatedly, Bhandari et al. (2017) derive prescriptions for optimal debt maturity

in a dynamic macro model. They show that the government's optimal target debt level is negative when a Ramsey planner can control the maturity structure of a non-state-contingent debt portfolio.

Krause and Moyen (2016) formalize in a New Keynesian model that the capability of a central bank to reduce real debt levels by setting higher inflation targets increases with the average maturity of government debt. Missale and Blanchard (1994) delve into the same issue. They document that higher debt levels are related to lower average maturity of debt. They rationalize this finding in a reputation model, where the government decreases the maturity with rising debt levels, in order to keep its commitment to low inflation credible.

A further strand of literature analyzes the behavior of DMOs themselves, mainly using empirical methods. Greenwood et al. (2014) formalize an objective function of DMOs regarding the issuance of short- and long-term debt. This function captures the trade-off between the liquidity premium of short-term debt versus its higher refinancing risk compared with long-term debt, also considering the costliness of budget variability. Hoogduin, Öztürk, and Wierts (2011) investigate DMOs' reaction functions in terms of long- versus short-dated issuance. They find that higher term spreads and higher long-term yield levels translate into a higher proportion of short-term debt in a panel of 11 euro area countries between 1990 and 2009. They document an increase in the proportion of short-term debt after 1999 and after the start of the global financial crisis in 2008, but the sample period does not allow them to study effects of QE.

Abbas et al. (2014) study the structure of public debt in a panel of 13 advanced economies between 1900 and 2011. Their results suggest that changes in the debt composition that increase exposure to crisis risk, such as a maturity shortening, can be related with subsequent financial crises. De Broeck and Guscina (2011) investigate crisis-related changes in government debt issuance in a panel of 16 European countries between 2007 and 2009. They find a shift away from fixed-interest rate instruments with longer maturities towards shorter-dated debt during the financial crisis. These works do not, however, consider the DMOs' response to changes in yields.

Beetsma et al. (2021) construct a theoretical model, where the public debt maturity choice depends on the liquidity services of

short-term debt, rollover risk, and credit risk. Using data for six euro area countries between 1999 and 2017 in a panel vector autoregression framework, they find that higher risk aversion, credit risk, and demand for short-term liquid assets have negative effects on the maturity of newly issued debt. The paper does not discuss any separate effects of QE policies. Wolswijk (2020) finds that higher interest rate spreads are related to a rising share of short-term debt issuance in a panel of 10 euro area countries between 1992 and 2017. This effect is found to be stronger in more vulnerable countries with higher debt levels. The paper also finds that growing debt levels imply more short-term financing, but that this effect vanishes after 2015 when the PSPP was introduced.

3. An Illustrative Description of DMO Behavior

Practitioners generally frame the government debt-management problem in terms of a trade-off between cost and risk, as formalized, for example, by Greenwood, Hanson, and Stein (2015). Former U.S. Treasury Secretary Lawrence Summers summarized the considerations as follows: "I think the right theory is that one tries to [borrow] short to save money but not [so much as] to be imprudent with respect to rollover risk. Hence there is certain tolerance for [short-term] debt but marginal debt once [total] debt goes up has to be more long term" (cited after Greenwood, Hanson, and Stein 2015).

Accordingly, DMOs optimize the maturity of its debt issuance intertemporally with regard to the funding cost (or interest expense) and to the refinancing risk, equating the marginal benefit of reducing refinancing risk with the marginal expense of higher funding costs. Figure 1 presents a stylized illustration of the cost-risk trade-off faced by DMOs in the spirit of the model by Greenwood, Hanson, and Stein (2015).[6]

We assume that a DMO faces an objective to minimize both funding costs and refinancing risk, subject to their funding need

[6]In practice DMOs employ a variety of indicators to measure portfolio risk, such as duration targets, average interest rate re-fixing period of the debt portfolio, and other risk-adjusted cash-flow-based targets. See Jonasson and Papaioannou (2018) for a detailed report of such targets.

Figure 1. Illustrative Maturity Funding Trade-Off for DMOs

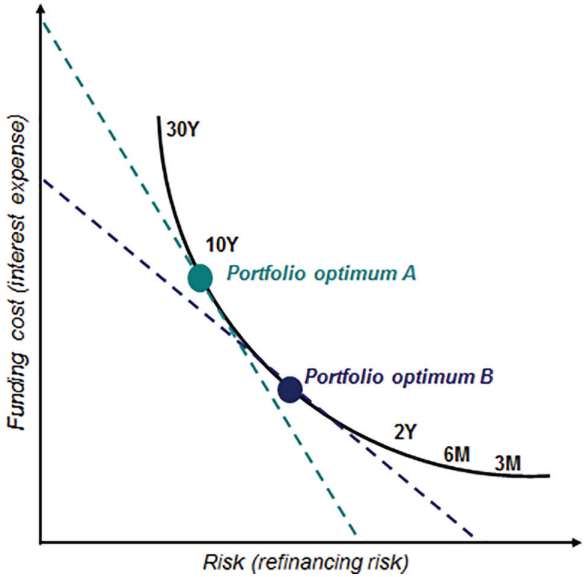

Note: Dashed lines represent indifference curves for DMOs A and B in the cost-risk space. Solid line depicts a market curve of debt at different maturities.

and current market conditions.[7] The dashed lines in the figure represent indifference curves, which represent combinations of costs and refinancing risk that yield the same level to the objective function for the DMO.[8] Figure 1 shows indifference curves for two different

[7]For example, the Italian DMO states, "Italy has focused on two principal risks: that posed by the interest rate [...] and that of refinancing, in order to distribute the maturities uniformly over time so that new debt may be placed with greater ease [...]. It is therefore crucial for Italy to set up an approach to debt management that places at the centre of its strategy risk control, and particularly those risks posed by rates and refinancing" Dipartimento del Tesoro (2015, pp. 5–6).

[8]For simplicity of the exposition, we assume in the figure that costs and refinancing risk are perfectly substitutable against each other, resulting in linear curves. In a more general case with imperfect substitutability, the indifference curves would be concave, as both arguments of the objective, i.e., funding cost and refinancing risk, are "bads" that the DMO seeks to minimize.

DMOs, denoted as A and B. Reducing refinancing risk by funding via longer-term debt is associated with an acceptable increase in the funding cost captured by the slope of the DMO indifference curve. Depending on whether the DMO requires a large (or small) reduction in refinancing risk in order to accept an increase in funding costs, the slope of the line is flatter (or steeper). The steepness of the indifference curves may vary across DMOs and possibly also over time. For example, some DMOs may have become more tilted towards risk reduction in response to the European sovereign debt crisis.

A DMO can issue debt instruments with different tenors for funding its government's financing needs. The debt instruments issued will carry varying interest expenses, as the level of interest rates fluctuates in the market. The instruments will be associated with changing refinancing risk, since future demand for such bonds may shift. At each point in time, we assume that the government faces a market frontier in the cost-risk space reflecting market conditions, such as interest rates and term spreads, for a given funding need of the DMO. This market curve is depicted by the solid line in Figure 1. As the y-axis is given in percent and the risk depicted on the x-axis is inversely related with debt maturity, this market curve can also be thought of as a mirrored yield curve. A downward-sloping market curve would then correspond to an upward-sloping government bond yield curve: In an environment with an upward-sloping yield curve, short-dated instruments (located at the lower-right part of the market curve) will be cheaper to issue but carry higher refinancing risk, while longer-dated instruments (located at the upper-left part of the market curve) will be more expensive but postpone refinancing risk further into the future.

The stylized representation in Figure 1 illustrates the optimal average maturity composition of government debt as the point where the slope of the DMO's indifference curve is equal to the slope of the market frontier. A steeper cost-risk trade-off translates into a portfolio optimum with a higher average maturity. This is reflected in the different portfolio optima for DMOs A and B in Figure 1. Through its liability management (e.g., in the form of new issuance, buybacks, and exchanges) and by entering derivative agreements (such as interest rate swaps), a DMO can transform its average portfolio maturity toward the optimum level. A DMO may to some extent be able to

Figure 2. Effect of the PSPP on DMOs' Funding Trade-Off

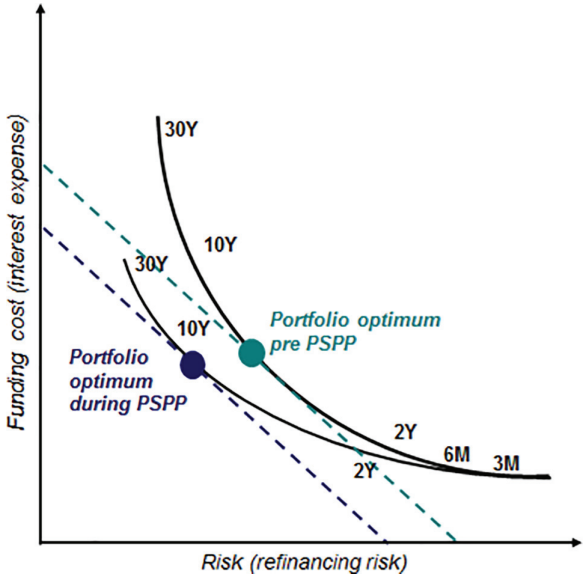

Note: Dashed lines represent CMO indifference curves in the cost-risk space. Solid lines depict market curves of debt at different maturities before and during the PSPP.

influence the position of the market curve by consistently following a specific strategy across time, thus giving investors predictability and reducing the risk premia embedded in government yields.

The potential effect of the PSPP on DMOs' maturity funding behavior is illustrated in Figure 2. Assuming that the yield curve for government debt declines and flattens due to central bank asset purchases, the market curve (black solid line) shifts down and inwards. The market curve also becomes shorter at its upper-left end, as the cost for any long-term tenor is now lower than before. In this new environment, the DMO extends the average maturity of its debt from the portfolio "pre-PSPP" to the portfolio "during PSPP," where it faces a generally better combination of cost and risk. Based on this simplified representation, we hypothesize that the PSPP, in addition to its effect of lowering yields and term spreads, could lead to an overall reduction in refinancing risk for DMOs. In particular,

Table 1. Summary of DMO Objectives, Tools, and Drivers

Objectives	Tools	Drivers
• Minimize funding costs • Reduce risks, including market risk, refinancing risk, liquidity risk, credit risk, settlement risk, and operational risk • Develop and maintain an efficient market for government securities (depth, liquidity) • Reduce uncertainty for investors • Maintain a diverse investor base.	• Issuance means: auction, syndication, buyback, exchange, retention, private placements • Security types: e.g., floating vs. fixed, currency denomination, bond characteristics, green, derivatives (in particular, swaps) • Maturity and size of issuance • Number/size of benchmarks • External communication: length of funding plans/calendars, pre-auction announcements, communication of portfolio composition targets • Number of primary dealers	• Yield environment • Market access/ investor base • Market depth/ liquidity • Political considerations • Macroeconomic and financial sector policies • Perception of market cycle • Risk aversion/ targets set by finance ministry and cost of budget variability • Borrowing requirement • Sustainability of debt/ creditworthiness

Note: Market risk represents the risk of cost variability due to changes in market variables (such as interest or exchange rates), whereas refinancing risk (also known as rollover risk) represents the risk that debt has to be refinanced at an unusually high cost, or cannot be rolled over at all. See Jonasson and Papaioannou (2018) for a comprehensive description of different types of risk faced by DMOs, including how they are managed and measured.

it may have provided assurance that longer-dated bonds would be purchased by the market at an acceptable price for the duration of PSPP implementation.

As with any model, there are limitations to the practical applicability of the illustrative funding trade-off depicted in Figures 1 and 2. In particular, it is not a comprehensive representation of funding choices. In practice, DMOs decide on the instruments issued, the issuance means, and overall transparency, depending on the size of their funding requirements, the liquidity of their markets, their investor base, their risk tolerance, and other internal and external drivers. A non-exhaustive overview of their decision space is provided in Table 1.

One limitation is that the illustrative model abstracts from the intertemporal nature of the DMO decision problem in the following sense. Only in rare circumstances will a DMO have to refinance the entire government debt in one period. In general, DMOs refinance around 10 to 20 percent of the outstanding government debt within one year. The overwhelming share of government debt has been "locked in" by funding decisions made in previous years. The relevant metric for capturing the outcome of a DMO's optimization at a given point in time with regard to the maturity composition of the debt is, therefore, the weighted average maturity of the newly issued debt. This indicator will, accordingly, be the main variable of interest in our empirical analysis. The relation between WAM at issuance compared with the WAM of the overall debt portfolio outstanding (hereafter also denoted as WAM outstanding) is further described in Appendix A.

4. Data and Estimation Procedure

This section has three parts. The first part describes the data set. The second part analyzes data properties, such as cross-sectional dependence and non-stationarity. The third part explains the econometric model employed in this paper and discusses issues of identification.

4.1 Data

The empirical analysis is based on a newly constructed panel data set of seven euro area countries over a monthly sample period of just under 10 years from December 2009 to April 2019, thereby covering euro area sovereign debt crisis from 2010 to 2013 as well as the implementation of the full first period of net asset purchases under the PSPP from 2015 to 2018. The size of both panel dimensions is determined by data availability. This results in a balanced panel with 113 periods and 791 observations. The countries covered in our sample and the name of their respective DMO are summarized in Table 2. We collect data on DMO portfolios, government bond yields, and a set of control variables to account for the macroeconomic environment.

Table 2. Sample Overview of Countries and DMOs

Country	DMO
Belgium (BE)	Agence Federale de la Dette/Federaal Agentschap van de Schuld
France (FR)	Agence France Trésor
Germany (DE)	Bundesrepublik Deutschland — Finanzagentur GmbH
Italy (IT)	Dipartimento del Tesoro
The Netherlands (NL)	Agentschap van de Generale Thesaurie
Portugal (PT)	Agencia de Gestão da Tesouraria e da Dívida Pública
Spain (ES)	Tesoro Público

Note: 791 usable observations between December 2009 and April 2019.

All data on DMO portfolios are taken from the European System of Central Banks' (ESCB) Centralised Securities Database (CSDB), which is the most comprehensive database for euro area sovereign debt securities. It consolidates security-level data from both ESCB-internal and commercial sources as of December 2009. The data undergo extensive data quality testing within the ECB (European Central Bank 2010). Monthly WAM outstanding, monthly nominal issuance, and monthly nominal redemptions for euro area general government securities are publicly available on the ECB's website for Government Finance Statistics (GFS). For the purpose of this paper it is, however, important to obtain data specific to central, rather than general, government issuers, i.e., excluding regional government issuers. These regional governments often have individual issuance strategies that may differ from the central government DMO strategy. Aggregating diverging strategies may impede the clear identification of DMO behavior. In addition, monthly data on WAM issuance as opposed to WAM outstanding is required to directly measure changes in DMO behavior as explained in Section 3. The data on central government debt are provided to us by ECB statistics.

We obtain data on monthly PSPP purchases by jurisdiction from an internal database maintained by the ECB. This database contains security-level information, enabling the classification of issuer type,

e.g., central government versus regional government and agencies.[9] The publicly available data do not allow for such a distinction. Additionally, the published data only disclose net as opposed to gross purchase values.

We use government bond yield data that are publicly available on the ECB's website. The choice of countries is restricted by the availability of monthly yield data in the 5- and 10-year maturity segments.[10] Furthermore, Ireland was excluded due to a series of floating rate bonds with a nominal value of EUR 25 billion, issued by the Irish government in connection with the Irish Bank Resolution Corporation Act 2013, with original maturities ranging from 25 to 40 years, of which more than 60 percent have been canceled to date. This issuance leads to massive structural breaks in the time series for WAM indicators of Irish government bonds (National Treasury Management Agency 2013). Austria is excluded for precautionary reasons, as the presence of outliers could distort the results. If we include data on Austria in the sample, our results become even stronger.

We use HICP inflation and industrial production excluding construction from the ECB's website as additional macroeconomic control variables. Summary statistics for all variables over different subsamples are provided in Tables B.1 and B.2 in Appendix B. Table B.3 provides bivariate correlations between all variables used in the analysis. The generally low correlations indicate that multicollinearity does not pose an issue in our regressions.

4.2 Cross-Sectional Dependence and Non-stationarity

To determine the appropriate estimation method, we test the data for cross-sectional correlation and non-stationarity.

Table 3 presents average (absolute) cross-sectional correlation coefficients and results for the CD-test for cross-sectional dependence (Pesaran 2004). The test statistic follows a standard normal

[9] EU supranational bonds are excluded from our analysis as well.

[10] Ten euro area jurisdictions did not have 5- and 10-year benchmark bonds outstanding throughout the sample period: Cyprus, Estonia, Finland, Greece, Latvia, Lithuania, Luxembourg, Malta, Slovakia, and Slovenia.

Table 3. Cross-Sectional Dependence Tests

	CD_P	*avg.* (r_{ij})	*avg.* $(\lvert r_{ij}\rvert)$
WAM Issuance	8.61***	0.177	0.179
PSPP/Issuance	34.96***	0.770	0.770
5-Year Yield	37.52***	0.836	0.836
10-Year Yield	42.09***	0.864	0.864
Redemptions	6.65***	0.137	0.170
WAM Outstanding	21.01***	0.431	0.512
Δ Industrial Production	11.15***	0.231	0.231
Inflation	23.93***	0.491	0.491

Note: CD_P denotes Pesaran (2004) cross-sectional dependence test statistic. Asterisks indicate rejection of the null hypothesis of cross-sectional independence at 10 percent (*), 5 percent (**), and 1 percent (***). *avg.* (r_{ij}) and *avg.* $(\lvert r_{ij}\rvert)$ denote average and average absolute cross-section correlation coefficients.

distribution under the null hypothesis of cross-sectional independence. It is shown to be efficient even when the time dimension is relatively small. According to the CD-test, cross-sectional independence is rejected for all variables at the 1 percent level. Also the cross-sectional correlation coefficients indicate strong dependencies for several variables in the data set. Ignoring cross-sectional dependence would lead to inefficient standard errors and even biased coefficient estimates. Hence, we address it in our model.

Given the presence of strong cross-sectional correlation in the sample, we apply the CIPS test by Pesaran (2007), a second-generation panel unit-root test, to analyze the stationarity properties of the data set.[11] The test is based on standard augmented Dickey–Fuller regressions, extended with the cross-section averages of lagged levels and first-differences of the individual series. Results for two versions of the CIPS test are presented in Table 4. The null hypothesis of the test assumes that the variable tested features a unit root. Autoregressive lags are included to control for autocorrelation, where the appropriate number of lags is determined by the

[11]First-generation unit-root tests, such as those by Maddala and Wu (1999), Levin, Lin, and Chu (2002), and Im, Pesaran, and Shin (2003), assume that variables are cross-sectionally independent and are therefore not appropriate for this data set.

Table 4. Panel Unit-Root Tests

	CIPS without Trend	CIPS with Trend
WAM Issuance	–12.935***	–12.942***
PSPP/Issuance	–2.410***	–2.056**
5-Year Yield	–1.235*	–2.037**
10-Year Yield	–1.749**	–1.278*
Redemptions	–12.586***	–12.625***
WAM Outstanding	–2.299***	–1.774**
Δ Industrial Production	–2.544***	–2.063***
Inflation	–3.597***	–2.648***

Note: Results of CIPS panel unit-root test statistics (Pesaran 2007). Asterisks indicate rejection of the null hypothesis of a unit root at 10 percent (*), 5 percent (**), and 1 percent (***). Optimal lag length determined by Akaike and Bayesian information criteria searching between zero and four lags.

Bayesian information criterion, searching between zero and four lags. The presence of a unit root is rejected for all variables in the data set. Industrial production in levels features a unit root. Throughout the paper, it is therefore used as the 12-month difference (denoted by Δ), for which non-stationarity can be rejected. Accordingly, it is not necessary to analyze the stationarity of regression residuals or to test for co-integration to rule out spurious regression results.

4.3 Econometric Model

To analyze the behavior of DMOs before and during the PSPP period, we estimate the following econometric model:

$$wam_{it}^{iss} = \alpha_i + \beta_1 yield_{it-1} + \beta_2 \left[\frac{PSPP}{issuance}\right]_{it-1} + \beta_3' X_{it-1} + \beta_4' M_{it-1} + u_{it}, \quad (1)$$

where the subindices denote DMO/country i and month t, α_i are DMO/country-fixed effects, and u_{it} is an error term. The dependent variable wam_{it}^{iss} is a measure of DMO funding behavior and denotes the weighted average maturity in years of securities issued by DMO i in month t.

To analyze the effect of the PSPP on DMOs' issuance behavior, we regress WAM at issuance on the lagged 10-year sovereign bond yield, denoted by $yield_{it-1}$, and on a measure of lagged monthly central government bond purchases of country i by the ECB, scaled by monthly government debt issuance in that country, denoted $[PSPP/issuance]_{it-1}$. The coefficient β_1 captures how strongly DMOs adjust the maturity of their new issuance to changes in yields. One important driver of sovereign bond yields and term spreads in the euro area over the last years was the ECB's APP, which is therefore partially reflected in this coefficient.[12] The coefficient β_2 measures whether there are additional demand effects of the PSPP that drives DMO behavior in addition to its impact on yields, which is captured by β_1.

When estimating these effects, some issues of endogeneity may potentially arise. A higher net supply of bonds (not included as a variable in our model) in a given tenor will ceteris paribus have a positive effect on the respective yield. Our variable of focus, WAM at issuance, is independent of the total amount issued. Nevertheless, it is still possible that a higher WAM of a *given* amount of newly issued debt can increase yield curve steepness and, thus, also the 10-year yield used in the model. We approach the issue of potential reversed causality in the following ways. In our regressions, we use lagged yields which are by definition exogenous to the current WAM at issuance. As yields are, however, highly correlated over time, this approach may not resolve the problem completely in practice. We, therefore, also consider a special case of exogenous variations in yields. To this end, we proxy monthly yield changes by monetary policy shocks that are identified using recent high-frequency identification methods (see, e.g., Gertler and Karadi 2015 and Altavilla et al. 2019). We describe this alternative estimation approach in

[12]Naturally, yields are also directly affected by changes in private investor demand and in the funding needs of sovereigns themselves. Moreover, the PSPP effect on yields would also endogenously interact with the demand of private investors. For example, Koijen et al. (2021) find that foreign investors had the most elastic demand for euro-denominated securities and sold considerable amounts of securities to the ECB after the PSPP's introduction. Boermans and Vermeulen (2018) find evidence that euro area investors instead acted as preferred-habitat investors with no significant change in the coefficients of their bond demand function after 2015.

Section 5.2, where we find that the results of the regression approach described here are fully robust to this treatment of potential yield endogeneity.

Problems of reverse causality are less likely to occur between the monthly volume of PSPP purchases and DMO's WAM choices. In its press releases regarding the PSPP, the ECB announced the overall monthly purchase targets well in advance. Volumes by country are then determined according to the countries' share in the ECB capital. Publicly available data of the monthly purchases make clear that any deviations from these announced purchase plans are not linked to yield movements or to the WAM of newly issued government debt.[13] From the perspective of the econometric model, the PSPP purchases can therefore be treated as an exogenous variable. As we scale the PSPP purchases by the debt issued in a given month, we also lag this ratio to rule out further endogeneity concerns.[14]

The regression model is augmented by a set of further control variables for DMO behavior, summarized in the vector X_{it-1}. This set includes variables that control for portfolio redemption effects and the WAM of the total portfolio outstanding of DMO i. Since we want to quantify changes in DMO behavior that occur as a response to changes in the external funding environment and the PSPP in particular, we control for deterministic portfolio effects that could affect wam_{it}^{iss} and are unrelated to the funding environment in month t.

The redemption variable is calculated as the logarithm of the total nominal value of all redemptions in the current and the previous month. Redemptions can affect the WAM at issuance because a high volume of redemptions may need to be replaced with issuances at relatively short maturities, as market demand and thus liquidity at the shorter end of the yield curve is often higher. We control for redemptions in the current and the previous month, since

[13]See Hammermann et al. (2019) for a detailed description of the practical implementation of the PSPP.

[14]We also test for Granger causality between the variables of main interest, i.e., WAM at issuance, yields, and PSPP/issuance. Although Granger causality does not necessarily imply a "true" causal relationship, it can be used to analyze the explanatory power of these variables for each other. In line with our argumentation, we find evidence that yields and the scaled asset purchases Granger-cause WAM at issuance, but not the other way around.

redemptions occur on a fixed date, while gross issuance is mainly implemented via the auction system in gradual steps.[15]

While controlling for redemptions captures most of the short-term portfolio legacy effects, we additionally control for WAM outstanding in levels to ensure that roll-down of debt is also captured across multiple periods. For example, DMOs that previously had a relatively low WAM outstanding could be more likely to issue at shorter maturities on average, independent of changes in external funding conditions.

The set of macro controls M_{it-1} includes lagged inflation and the lagged annual change of industrial production. These variables capture potential effects of the state of the business cycle on WAM at issuance, as governments may have incentives to borrow short term during a recession to reduce funding costs. Moreover, larger funding requirements during a downturn could be easier to place at shorter maturities, where liquidity is typically higher.

The data set used is a macro panel with a relatively large time dimension and a small cross-sectional dimension ("large-T-small-N"), where cross-sectional dependence is found to be an issue. The common correlated effects pooled (CCEP) estimator by Pesaran (2006) is designed specifically for this type of data. Compared with, for example, the standard two-way fixed effects estimator well suited for "large-N-small-T" micro panels, the CCEP estimator has several advantages. By assuming a multi-factorial structure as the data-generating process, the estimator allows that each country in the panel can respond differently to common time effects in each variable of the model, while allowing for arbitrary degrees of auto- and cross-correlation among all variables.

These properties are particularly useful for the macro-financial variables in this data set. For example, euro area government bond yields can be considered to be driven by a number of common factors. These may include the common monetary policy as well as the global trends behind the low interest rate environment, e.g., demographics or the productivity slowdown. At the same time, the yields of each country may depend on these factors to a varying extent, besides further country-idiosyncratic factors. In comparison, the time fixed

[15]The results are robust to using redemptions of the current month only.

effects in the conventional two-way fixed effects estimator merely allow all countries to depend homogeneously on one single common time factor.[16]

The CCEP estimator is practically computed as an ordinary least squares regression, augmented with cross-sectional averages of the dependent and independent variables, which are interacted with the country dummies as additional regressors. As a robustness check, we also provide results using the two-way fixed effects estimator and the panel-corrected standard errors (PCSE) estimator by Beck and Katz (1995) that can account for heteroskedasticity, autocorrelation, and cross-sectional correlation in the regression errors.

5. Results

In the following, Section 5.1 presents and discusses all our main regression results. Section 5.2 shows the local projections to high-frequency yield shocks to explore potential endogeneity concerns about our main regressions. Section 5.3 provides a quantification of our findings' economic significance.

5.1 Yield and Demand Effect

Table 5 presents the main regression results for the whole sample of seven euro area countries from December 2009 to April 2019. The regressions are all based on the model in Equation (1), where WAM at issuance, measured in years, is regressed on different combinations of the independent variables. All regressions include DMO/country-fixed effects and use the CCEP estimator. Our two regressors of

[16] Although non-stationarity could be rejected for the variables in our data set, Kapetanios, Pesaran, and Yamagata (2011) show that the CCEP estimator even remains consistent if the data is driven by unit-root processes. This is a further advantage of this estimator in applications using macro-financial data. In a recent contribution, Juodis, Karabiyik, and Westerlund (2021) confirm the consistency of the CCEP estimator under very general conditions for the data-generating process. Acknowledging the estimator's good small-sample performance, they show, however, that asymptotic normality no longer holds when the number of underlying common factors in the data is larger than the number of regressors in the model. As we will show in Section 5, the coefficient size of our main regressors of interest, β_1 and β_2, remains fairly stable when further control variables are added.

Table 5. The Effect of Yields and the PSPP on WAM at Issuance: Euro Area

Dependent Variable: WAM Issuance	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
PSPP/Issuance						0.014***	0.011**	0.011**	0.011**	0.011**
						(0.0032)	(0.0040)	(0.0044)	(0.0041)	(0.0044)
10-Year Yield	–0.34***	–0.35***	–0.38**	–0.38***	–0.49***	–0.26**	–0.28**	–0.31*	–0.29**	–0.39**
	(0.075)	(0.075)	(0.10)	(0.089)	(0.12)	(0.097)	(0.10)	(0.14)	(0.098)	(0.11)
Redemptions		–0.54*	–0.58*	–0.51	–0.63**		–0.49	–0.49	–0.42	–0.54*
		(0.24)	(0.24)	(0.27)	(0.25)		(0.26)	(0.26)	(0.28)	(0.24)
WAM Outstanding			–0.25	–0.48	–0.70*			–0.10	–0.41	–0.62
			(0.21)	(0.33)	(0.29)			(0.29)	(0.42)	(0.37)
Δ Industrial Production				0.021	0.017				0.031	0.027
				(0.040)	(0.038)				(0.034)	(0.033)
Inflation					0.42					0.38
					(0.31)					(0.30)
Observations	784	784	784	784	784	784	784	784	784	784
Adjusted R^2	0.366	0.371	0.369	0.372	0.373	0.371	0.373	0.371	0.374	0.374

Note: Robust standard errors presented in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. CCEP estimator is used for all regressions. All models include country fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. All independent variables except for redemptions are lagged by one month. The sample includes BE, DE, FR, ES, IT, NL, and PT over the period December 2009 to April 2019.

interest are yields and the demand variable, PSPP/issuance. The models in columns 1 to 5 focus on the effect of the 10-year yield on WAM at issuance, while the models in columns 6 to 10 additionally consider the effect of PSPP/issuance.

We find throughout all regressions that yields have a significant negative relationship with WAM at issuance. For instance, a 1 percentage point decrease in 10-year yields is related to an increase in the WAM of securities issued by 0.49 years in column 5. The size of the coefficients remains relatively stable, ranging from −0.34 to −0.49, when different controls are added or removed. All effects are statistically significant at the 5 percent and in most cases even at the 1 percent level. This result indicates that DMOs change their weighted average maturity at issuance in response to changes in the yield environment. This finding is generally consistent with results in the literature. When regressing WAM at issuance on current, instead of lagged, yields, Beetsma et al. (2021) also find a negative sign with somewhat larger coefficients than we do. Hoogduin, Öztürk, and Wierts (2011) and Wolswijk (2020) use the share of short-term debt issuance as dependent variable. Consistent with us, they find that higher yields and term spreads are related to a higher share of short-term debt issuance.

In a next step, we add PSPP/issuance to the regression model, in order to test for additional effects of the PSPP on DMOs' behavior due to the higher and stable demand by the ECB. We find that PSPP/issuance has a significant positive relationship with WAM at issuance in all our regressions. Specifically, a 10 percentage point higher ratio of PSPP/issuance is related to an increase of WAM at issuance by 0.11 years in column 10. The coefficients of the 10-year yield variable remain statistically significant and in the same order of magnitude with values ranging from −0.26 to −0.39 when PSPP/issuance is added.

The fact that both yields and PSPP/issuance enter the regressions significantly at the same time indicates the existence of an additional demand effect of the PSPP on WAM at issuance, which is not explained through the PSPP's effect on yields. An explanation is that the PSPP purchases by the ECB, as a relatively price-insensitive investor, enabled DMOs to issue additional longer-dated securities. When setting auction prices for new debt issuances, DMOs have to consider that primary dealer demand tends to be

lower for high-duration debt. Liquidity is typically higher at the shorter end of the curve, meaning that dealers have a lower risk of not being able to offload their positions in the secondary market. Some of this risk is removed through the presence of the PSPP, as dealers can expect that the ECB will exert a significant secondary market demand for longer-dated maturities. The PSPP eligibility criteria prohibit purchases of securities with a residual maturity below one year and allow purchases at a yield to maturity below the deposit facility rate only to the extent necessary, thereby limiting shorter-maturity purchases. To the best of our knowledge, such QE-related demand effects have not been documented empirically before.[17]

The control variables in Table 5 generally have the expected sign or are insignificant. The effect of redemptions is negative, indicating that DMOs decrease their WAM at issuance in the presence of higher redemption volumes. The redemptions variable is statistically significant in three out of eight cases. WAM outstanding also has a negative coefficient, but is found to be insignificant in all but one regression. The two macroeconomic controls, inflation and industrial production, enter the regressions with positive signs, which is in line with the notion that governments issue shorter in a downturn. None of the coefficients are statistically different from zero, though. This generally supports the notion that DMOs focus on funding costs and risks, with alternative fiscal objectives playing a secondary role, if any.

In principle, yield levels and the amount of asset purchases are expected to correlate negatively. A high correlation would complicate the identification of separate yield and demand effects on WAM at issuance in our analysis. As Table B.3 in Appendix B shows, however, the correlation between the 10-year yield and the PSPP/issuance variable that we employ in the regressions is rather moderate, with a value of -0.40. Notably, the coefficients of PSPP/issuance and yields are robust to the exclusion of the respective other term compared with when they are added jointly. This can be seen by comparing Table 5 with columns 6 to 10 of Table C.1 in Appendix C, indicating that the regressions can identify separate yield and demand effects on *WAM issuance*.

[17]Wolswijk (2020) finds that DMOs generally tend to issue more short term when government debt is rising. The effect disappears after 2015, though. The paper hypothesizes that this may be related to the presence of the ECB, as a predictable and relatively price-insensitive buyer, in sovereign bond markets.

As a robustness check, we rerun columns 1 to 5 of Table 5 with 5-year instead of 10-year yields. The results are displayed in Table C.1 of Appendix C. The results are generally very close to each other, although the 5-year yields tend to have moderately smaller coefficients than 10-year yields. Adding yields at different maturities to the model simultaneously does not generate meaningful results given their very high levels of correlation. As one of the main objectives of the PSPP is term premium compression, we also test whether term spreads (i.e., the steepness of the curve) can play a role in determining WAM at issuance in addition to yield levels. Adding the 10 minus 2-year term spread to the regressions of Table 5 leaves all results unchanged.[18] The coefficient of the term spread itself is not found to be statistically significant. Yield levels and not curve steepness, therefore, appear to be the main driver of issuance maturity.

Our results also remain robust when using the two-way fixed effects estimator or an estimator with panel-corrected standard errors, as shown in Tables C.2 to C.5 in Appendix C. The signs of all coefficients as well as the patterns of significance remain broadly unchanged. With the two-way fixed effects estimator, the yield coefficients are somewhat larger, while the effects of PSPP/issuance are smaller and turn insignificant. As shown in Table 3, PSPP/issuance is among the variables in the data set with the largest degree of cross-sectional dependence. The insignificance of its coefficient is therefore likely a result of the more limited treatment of this issue under the two-way fixed effects estimator. Notably, when using the PCSE estimator, where standard errors are corrected for cross-sectional correlations, the effect of PSPP/issuance is again found to be significant. Also, the effect of the yield variable becomes a bit smaller and the effects of PSPP/issuance get larger when using this estimator.

Table 6 shows results of the regression model in Equation (1) for different sub-samples in the columns indicated with a superscript a. The table also analyzes whether the responsiveness of DMOs to yields changed after the onset of the PSPP. The results for this are given in the columns with superscript b. We consider the following sub-samples: We study effects for the "Big 4" group, which consists of DE, FR, IT, and ES. We also analyze whether effects are different

[18]Results are available on request.

Table 6. The Effect of 10-Year Yields and the PSPP on WAM at Issuance over Different Sub-samples

	EA		Big 4		Stressed		Non-stressed	
Dependent Variable: WAM Issuance	**(1^a)**	**(1^b)**	**(2^a)**	**(2^b)**	**(3^a)**	**(3^b)**	**(4^a)**	**(4^b)**
PSPP/Issuance	0.011**	0.011**	0.019**	0.019**	0.018*	0.018*	0.014	0.014
	(0.0044)	(0.0039)	(0.0051)	(0.0047)	(0.0053)	(0.0047)	(0.0062)	(0.0064)
10-Year Yield	–0.39**	–0.32**‡	–0.69**	–0.66**‡	–0.58*	–0.52‡	–1.15	–1.15
	(0.11)	(0.11)	(0.21)	(0.17)	(0.16)	(0.22)	(1.05)	(1.03)
10-Year Yield × PSPP-Dummy		–0.43‡		–0.20‡		–0.15‡		0.064
		(0.32)		(0.33)		(0.15)		(0.41)
Redemption Effect	–0.54*	–0.56**	–1.01**	–1.02**	–0.27*	–0.30**	–1.70**	–1.70**
	(0.24)	(0.22)	(0.28)	(0.26)	(0.089)	(0.065)	(0.23)	(0.23)
WAM Outstanding	–0.62	–0.57	–1.54	–1.55	–0.18	–0.15	–0.86	–0.87
	(0.37)	(0.39)	(1.40)	(1.38)	(0.19)	(0.26)	(0.59)	(0.61)
Δ Industrial Production	0.027	0.030	0.0041	0.0047	0.036	0.038	–0.072	–0.072
	(0.033)	(0.027)	(0.049)	(0.052)	(0.031)	(0.040)	(0.030)	(0.030)
Inflation	0.38	0.35	0.14	0.12	0.74**	0.72***	–0.092	–0.091
	(0.30)	(0.28)	(0.24)	(0.26)	(0.087)	(0.064)	(0.26)	(0.27)
Observations	784	784	448	448	336	336	336	336
Adjusted R^2	0.374	0.376	0.522	0.521	0.596	0.596	0.453	0.451
F(Yield^b, Interaction = 0)		6.95**		11.17**		294.3***		8.15
F(Yield^b = Yield^a)		0.42		0.03		0.06		0.00

Note: Robust standard errors presented in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$. CCEP estimator is used for all regressions. All models include country fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. PSPP-dummy is 1 as of March 2015, otherwise 0. All independent variables except for redemptions are lagged by one month. ‡ denotes joint significance of the yield and the interaction term of yield*PSPP-dummy in the first F-test presented below the table. When the null hypothesis of the second F-test, Yield^b = Yield^a, cannot be rejected, the effect of 10-year yield on WAM issuance is not statistically different over the whole sample period (December 2009 to April 2019) and before the PSPP (until February 2015). Big 4 includes DE, FR, IT, and ES. Stressed includes ES, IT, and PT. Non-stressed includes DE, FR, and NL.

for countries that were more and less affected during the European sovereign debt crisis of 2010–12. Our "stressed" sample includes IT, ES, and PT, while our "non-stressed" sample includes DE, FR, and NL. The columns with superscript a in Table 6 compare the effects of yields and PSPP/issuance in the euro area sample (1^a), which is repeated here from Table 5 for convenience, with the different sub-samples. Overall, the negative effects of higher yields and the positive effect of higher PSPP/issuance prevails over all sub-samples considered. Notably, the effects of both variables are larger in the "Big 4" group (2^a) than in the overall sample.

The effects in the "stressed" group (3^a) are found to be larger than those in the full sample and in the "non-stressed" group (4^a). For example, the effect of yields reads -0.58 for the stressed DMOs, while it is -0.39 in the full sample and insignificant for the non-stressed DMOs. The same holds true for the effect of PSPP/issuance with coefficients of 0.018 versus 0.011 and 0.014. Notably, the coefficients in the "non-stressed" group are not statistically different from zero. This can be interpreted that these DMOs are less reactive to changes in their financing environment, but it may also be due to the relatively small sample size. Overall, these findings imply that DMOs that are more vulnerable to fiscal stress increase the maturity of their issuance relatively stronger in response to yield changes and PSPP purchases. Given the higher uncertainty and potentially higher rollover risks, these countries may have higher incentives to make use of a favorable market environment.[19]

Columns 1^b, 2^b, 3^b, and 4^b in Table 6 are augmented with an interaction term of the 10-yield yield and a binary PSPP-dummy variable that takes a value of one after the onset of the PSPP (as of March 2015) and is zero otherwise. When this interaction term is included to the model, the coefficient of the plain yield variable (hereafter denoted as Yield^b) captures the effect of yields on WAM at issuance before the PSPP. Meanwhile, the interaction coefficients represent any additional effect of 10-year yields on WAM at issuance

[19] Using a similar country-split, Wolswijk (2020) finds that the share of short-term debt issuance by vulnerable countries is more sensitive to yield spreads than it is for strong countries. Instead, Beetsma et al. (2021) find that the more vulnerable countries (Italy and Spain in their sample) react less, which would indicate that these countries favor lower costs over a reduction in rollover risks.

during the PSPP, on top of the effect of 10-year yields before the start of the PSPP. Given the (by construction) high degree of correlation between the interaction term and the 10-year yield variable, their standard errors increase in some of the sub-samples, influencing their individual significance but not leading to biased estimates. The interaction term and the yield variable are, however, jointly significant for all groups except "non-stressed," as presented in the first F-test at the bottom of the table, which tests the null hypothesis of joint significance. For convenience, we indicate joint significance of the yield and its interaction with a ‡ in the table.

The coefficients of the interaction term are insignificant in all of the sub-samples considered. Accordingly, DMOs did not change their responsiveness to yield changes after the onset of the PSPP. To analyze this further, we test whether there is a statistically significant difference in the effect of 10-year yields on WAM at issuance before and during the PSPP by means of another F-test. The null hypothesis of this test is that the pre-PSPP yield impact (denoted by Yield^b and taken from the columns with superscript b) is equal to the yield impact of the full sample period (denoted by Yield^a and taken from the columns with superscript a), i.e., that there is no additional yield effect during the PSPP ($\text{Yield}^b = \text{Yield}^a$). This null hypothesis cannot be rejected for any of the groups, which indicates that the effect of the 10-year yield on WAM issuance is not statistically different in the period before the PSPP (December 2009 to February 2015) and the full sample (December 2009 to April 2019). This indicates a continuation of the existing DMO behavior before and after the PSPP. DMOs neither became more nor less responsive to yields than can be expected if they continued to act in line with their mandates. It is the intention of the PSPP to alleviate financing conditions, and DMOs acted accordingly and endogenously in response to the changed conditions.

As a robustness check of Table 6, we repeat all regressions using five-year yields in Table C.6. All results remain fully robust.

5.2 Yield Shocks: A Special Case

To explore whether our results from the previous section are subject to issues of reverse causality between WAM at issuance and yields,

we now focus on a special case of exogenous variation in government bond yields: monetary policy shocks. Arguably and given the statutory independence of the ECB's monetary policy from national fiscal policies, it can be assumed that monetary policy decisions of the ECB do not react in a systematic manner to changes in the WAM at issuance of euro area DMOs. To the contrary, monetary policy shocks can be expected to drive government bond yields and, thus, also exert effects on DMO issuance behavior.

In this section, we therefore assess the direct effect of monetary policy-induced yield shocks on WAM at issuance. To this end, we estimate a dynamic version of Equation (1) using local projections methods as introduced by Jordà (2005). Impulse response functions (IRFs) are obtained following

$$Y_{i,t+h} = \alpha_{i,h} + \beta_h \sum_{p=1}^{P} Z_{i,t-p} + \gamma_h shock_{i,t}^{yld} + u_{i,t+h} \ , \qquad (2)$$

where subindex h denotes the IRF horizon, while p gives the number of lags in the matrix of independent variables $Z_{i,t-p}$, and β_h is the corresponding matrix of coefficients. The vector of dependent variables is given by $Y_{i,t+h}$, and $shock_{i,t}^{yld}$ indicates an exogenous country-specific shock to the 10-year yield with coefficient γ_h.

The dependent variables are WAM at issuance and the 10-year yield. Yields are also used as dependent variable here to make sure that the monetary-policy-induced yield shock has the expected effect on monthly yields. Finding a significant response here can be seen as a precondition for finding effects of the yield shock on the WAM at issuance. The list of independent variables can include all variables introduced in Section 4.3 as well as autoregressive terms $Y_{i,t-p}$.

To obtain an exogenous measure for the yield shocks, we use a high-frequency identification approach.[20] This approach employs the change of financial variables, e.g., government bond yields, in a small time window around monetary policy announcements. This allows to obtain the surprise effect of a monetary policy change on financial

[20]See, among many other papers, in particular Kuttner (2001), Gürkaynak, Sack, and Swanson (2005), Gertler and Karadi (2015), and Jarociński and Karadi (2020). For a recent application in a macroeconomic panel analysis, see Holm-Hadulla and Thürwächter (2021).

markets relative to market expectations before the announcement. The chosen time windows need to be sufficiently small to rule out confounding effects from other market-relevant news.

We use high-frequency intraday data from the Euro Area Monetary Policy Event-Study Database by Altavilla et al. (2019). This data set contains the changes of a broad set of financial market variables in a narrow time window of monetary policy events on all monetary policy meetings of the ECB's Governing Council since January 1999. We use changes over the whole monetary event window that is calculated as the difference between the median quote from the time window 13:25–13:35 before the press release and the median quote from the time window 15:40–15:50 after the press conference. For our application, we use the intraday changes in the 10-year government bond yields of France, Italy, Germany, and Spain. For the other countries of our sample—Belgium, the Netherlands, and Portugal—no data are available in this database. We, therefore, use the one-day changes of the respective 10-year yields around the Governing Council dates that we obtain from Reuters. These high-frequency yield shocks correspond as closely as possible to the monthly 10-year yield series that we use in the regressions and the local projections. The shock measure includes the combined yield effect from all monetary policy instruments discussed at the respective dates, such as conventional interest rate changes, as well as forward guidance and asset purchases.

Figure 3 shows our 10-year yield shock measure over time for all seven countries in the sample. While the mean value of the shock is close to zero, it features considerable variation within the cross-section and over time. The standard deviation reads 7 basis points (bps), while the min-max range spans from −44 to 54 bps.

Figure 4 depicts IRFs of WAM at issuance and the 10-year yield to a 1 basis point high-frequency yield shock. The IRFs are based on Equation (2) and include all possible independent variables, i.e., an autoregressive term, the 10-year yield, PSPP/issuance, redemptions, WAM outstanding, the change of industrial production, and inflation. The model features one lag for each variable. We find that a 1 basis point yield shock translates to an increase of the monthly 10-year yield by 0.6 bps after one month, growing further to slightly beyond 1 basis point after four months. The effect is highly statistically significant and persistent. This finding is decisive if we want

Figure 3. High-Frequency Yield Shocks over Time

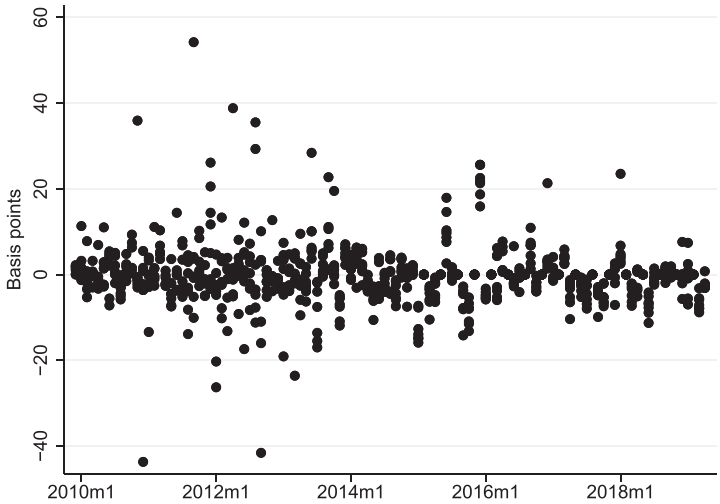

Note: Dots show intraday/end-of-day changes of the 10-year yield of all countries in the sample on all ECB Governing Council meeting dates between December 2009 and April 2019.

to find a meaningful effect of the shock on WAM at issuance transmitted via the monthly yields. Consistent with the previous panel regression results, WAM at issuance is also found to react to the shock significantly. In the IRF the effect starts being statistically significant with a negative sign three months after the shock before reaching a peak response of −0.03 years five months after the shock.

While the panel regressions in Section 5.1 can be understood to provide long-run estimates on the effect of yields on WAM at issuance, the local projections also allow to study dynamic effects. The finding that WAM at issuance falls gradually corresponds to the gradual rise of the monthly yield to the shock. At the same time, this could also be rationalized economically by the fact that most DMOs plan and schedule the modalities of their issuance some months in advance.

In order to compare the local projection results with the long-run effects from Table 5, we calculate the average coefficient size over the first 12 months after the 1 basis point shock, which reads −0.0042 years. Scaling this up to a 1 percent yield shock, the average effect of

Figure 4. Impulse Responses to 1 Basis Point Yield Shock

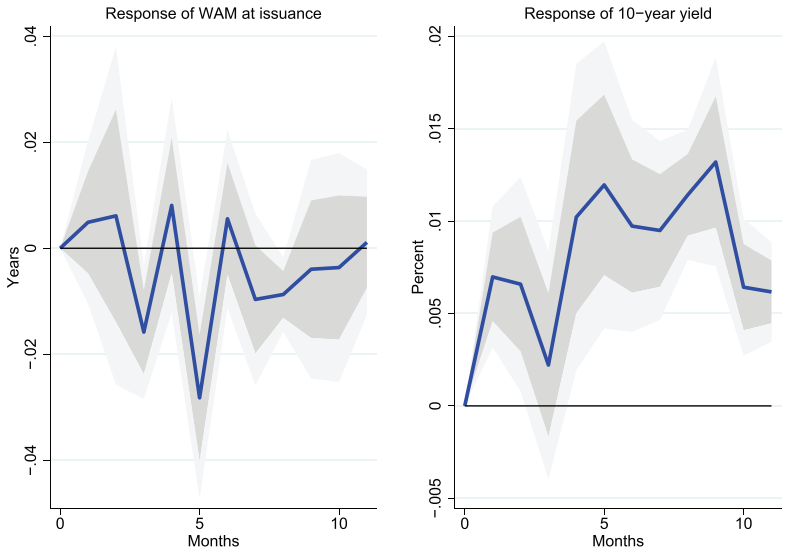

Note: Impulse response functions to a 1 basis point high-frequency yield shock based on local projections as defined in Equation (2). Control variables include an autoregressive term, the 10-year yield, PSPP/issuance, redemptions, WAM outstanding, the change of industrial production, and inflation. Models include one lag for each variable. Robust standard errors are clustered by country. Dark (light) gray-shaded areas indicate 68 percent and 90 percent confidence intervals.

−0.42 years lies very well in the range of yield coefficient estimates in Table 5, which is between −0.26 and −0.49 years. As the results of the panel regressions and the local projections are consistent with each other, we do not find an indication that the yield estimates in Section 5.1 are subject to reverse causality.

While Figure 4 shows results for the full model, the results continue to hold in more parsimonious specifications including yields and autoregressive terms only. Results are also robust to using more than one lag and when no lagged dependent variable is employed (see Figures C.1, C.2, and C.3 in Appendix C).

5.3 Economic Significance of Maturity Extension and Duration Extraction

Our results support the hypothesis that the PSPP led to an overall lengthening of issuance maturities through its impact on euro area

Figure 5. PSPP "Yield" Impact on WAM Issuance

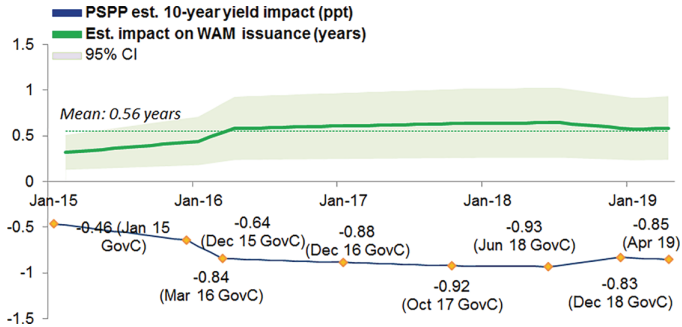

Note: The figure shows the estimated effect of PSPP-induced changes in 10-year yields on WAM at issuance for DE, FR, ES, IT, based on regression results in column 2^a, Table 6. Point estimates for the PSPP's 10-year yield term premium compression are taken from Eser et al. (2019). The point estimates correspond to the initial announcement of the AOO by the ECB Governing Council (GovC) in January 2015 with net purchases of EUR 60 billion per month from March 2015 to at least September 2016, and to subsequent changes to the purchase horizon and/or net purchase volumes. Estimates shown together with 95 percent confidence interval.

yields. We can now quantify this impact on the WAM at issuance based on results of Eser et al. (2019), who provide point estimates for the PSPP's term premium compression in the 10-year segment for the Big 4 countries following APP announcements by the ECB Governing Council.[21] We use these point estimates in our model for the Big 4 countries from column 2^a in Table 6. The reaction of WAM at issuance to the PSPP-induced yield changes together with a 95 percent confidence band are shown in Figure 5. Our quantification shows that the mean monthly PSPP yield impact on WAM at issuance in the Big 4 countries after March 2015 is estimated to be 0.56 years (seven months).

In a next step, we quantify the economic effect of the additional demand effect of the PSPP. Column 2^a in Table 6 indicates that for the Big 4 countries a 1 percentage point increase in PSPP/issuance

[21] The point estimates correspond to the initial announcement of the APP in January 2015, with net purchases of EUR 60 billion per month from March 2015 to at least September 2016, and to subsequent changes to the purchase horizon and/or net purchase volumes.

Table 7. PSPP "Demand" Effect Impact Quantification

(EUR Million Unless Stated Differently)	
μ Issuance (per Country) Implied Value of a 1 Percent Change in PSPP/Issuance	30,661 307
μ WAM Issuance Impact of a EUR 1 Billion Increase in Monthly PSPP (Years): 1000/307 × 0.019 (Coefficient)	0.062
Note: Calculation based on data for DE, FR, ES, IT from March 2015 to April 2019.	

Figure 6. PSPP "Demand" Impact on WAM Issuance

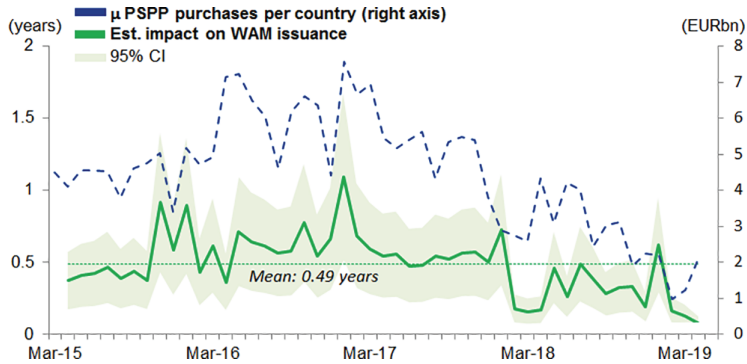

Note: The figure shows the estimated effect of PSPP/issuance on WAM at issuance for DE, FR, ES, IT, based on regression results in column 2^a, Table 6. Estimates shown together with 95 percent confidence interval. Gross PSPP purchases per country for central government issuers in EUR billion shown on right axis.

coincides with a 0.02 year increase in WAM at issuance. To gauge the economic significance of this effect, we convert this into an effect of nominal PSPP purchases in euro on WAM at issuance in Table 7. The average monthly issuance volume of the Big 4 countries across the full sample is EUR 30 billion per country. A 1 percent change therefore equals on average EUR 300 million. We can then calculate that a EUR 1 billion increase in monthly PSPP purchases results in a rise of WAM at issuance by 0.062 years (about one month) thereafter.

Figure 6 shows the PSPP "demand" effect implied by column 2^a in Table 6 on WAM at issuance over time. The average monthly

PSPP impact on WAM at issuance amounts to 0.49 years (six months) per country. As an illustration, we plot the average monthly PSPP purchases per country, i.e., the numerator of our demand variable, on the right axis of the figure.

The findings suggest that the PSPP did have a significant positive impact on WAM at issuance through yields and demand, leading to an extension of issuance maturities by around 1.1 years on average. This is an economically meaningful number given that the WAM outstanding of the Big 4 countries at the onset of the PSPP was 6.4 years on average. During the PSPP implementation phase, this average increased by almost one year to 7.2[22] and the calculations imply that the PSPP may explain a significant portion of the overall increase.

In a next step, we use our data and estimation results to calculate the change in the WAM outstanding that can be attributed to the yield and the demand effect as well as the offsetting effect these maturity extensions had on the term premium compression effect of the ECB's asset purchases. To this end, we follow the approach by Greenwood et al. (2014), who find that—depending on the reference period chosen—between 35 percent and 63 percent of the Federal Reserve QE's term premium compression effect was canceled by higher-maturity issuance of the U.S. Treasury.

In the same vein, Table 8 compares the duration extraction through the ECB's PSPP with the injection of duration by DMOs for the euro area, the Big 4 countries, as well as our "stressed" country sample.

To account for different maturities, we convert the total amount of the ECB's euro area sovereign debt holdings under the PSPP at the end of our sample in April 2019 into 10-year duration equivalents following

$$Debt_t^{10ye} = \frac{Debt_t \cdot Dur_t}{Dur_t^{10y}}, \tag{3}$$

where $Debt_t$ and Dur_t denote nominal debt amounts and average portfolio duration, respectively, while the superscript "$10ye$" indicates the 10-year equivalent. For example, using the nominal amount

[22]See Figure B.1, Appendix B for the time series of the WAM of all debt outstanding for each country in our sample.

Table 8. Term Premium Compression Offset Due to DMO Maturity Extension

	EA		Big 4		Stressed	
	Total	**Est.**	**Total**	**Est.**	**Total**	**Est.**
ECB						
PSPP Holdings Apr. 2019 (EUR bn)	1,929		1,374		590	
WAM Outstanding Apr. 2019 (Years)	7.27		7.13		7.72	
Term Premium Compression (bps)	50.49		35.25		16.39	
DMOs						
Sovereign Debt Feb. 2015	7,580		6,084		2,845	
Sovereign Debt Apr. 2019	8,193		6,604		3,225	
WAM Outstanding Feb. 2015	6.52	6.52	6.43	6.43	6.29	6.29
WAM Outstanding Apr. 2019	7.37	6.80	7.23	6.69	6.98	6.62
Δ WAM Outstanding	0.85	0.57	0.80	0.54	0.69	0.35
Maturity Extension Effect (EUR bn)	692	463	528	357	222	114
Term Premium Offset (bps)	24.93	16.67	19.02	12.87	7.98	4.09
– in Percent	49%	33%	54%	37%	49%	25%

Note: Table shows the offset of the PSPP term premium compression effect due to maturity extension of euro area DMOs for two cases: (i) Columns labeled "Total": The offset effect under the total change of WAM outstanding actually observed between the start of the PSPP (Feb. 2015) and the end of the sample (Apr. 2019); (ii) Columns labeled "Est.": The offset effect due to changes in WAM outstanding that can be attributed to the yield and demand effect from the regression models (Table 5, column 10) Sensitivity of term premia to asset purchases taken from Eser et al. (2019). Maturity extension effect as defined in Equation (4). PSPP holdings and sovereign debt outstanding are given in nominal amounts.

of the ECB's PSPP holdings of EUR 1929 billion and the average WAM outstanding of the ECB's portfolio of 7.27 years, we obtain ECB holdings in 10-year equivalents of EUR 1402 billion at the end of our sample period. Results by Eser et al. (2019) indicate a term premium compression of about 3.6 bps per EUR 100 billion of ECB asset purchases. The gross term premium compression until April 2019, accordingly, amounted to 51 bps.[23]

Greenwood et al. (2014) then decompose the change of total government debt outstanding in 10-year equivalents into pure debt expansion and maturity extension according to

$$\Delta\left(\frac{Debt_t \cdot Dur_t}{Dur_t^{10y}}\right) = \frac{1}{Dur_t^{10y}}\left(\underbrace{\Delta Debt_t \cdot Dur_{t-1}}_{\text{Debt Expansion}} + \underbrace{\Delta Dur_t \cdot Debt_t}_{\text{Maturity Extension}}\right). \quad (4)$$

We then calculate the maturity extension effect and the resulting term premium offset for two cases. In the first case (shown in columns labeled "Total" of Table 8), we use the total change of the WAM outstanding observed in the data between the start of the PSPP in February 2015 and the end of our sample period. In the second case (shown in columns labeled "Est."), we calculate an estimated change in WAM outstanding that can be directly attributed to the yield and demand effect. For this, we calculate the effect of the yield and demand effect on WAM at issuance using the regression model in column 10 of Table 5. The estimated WAM at issuance is then transformed into WAM outstanding as described in Appendix A.

While the total observed change of WAM outstanding amounted to 0.85 years for the euro area, our estimation states that 0.57 years of this change are due to the response of the DMOs to the lower yields and the additional demand effect of the PSPP purchases. Calculating the maturity extension effect, we find that 33 percent

[23]This yield elasticity abstracts from the announcement effect of the PSPP on term premia, which Eser et al. (2019) estimate to be around 50 bps, with the total effect, thus, being around 100 bps at the end of our sample period.

of the PSPP's effect on term premia (17 out of 51 bps) may have been offset due the yield and the demand effect. As the total change of WAM outstanding was somewhat higher, even up to 49 percent of the PSPP's effect may have been canceled due to behavior of euro area DMOs. We find similar effect sizes for the Big 4 and the stressed countries. These numbers are within the range found for the United States by Greenwood et al. (2014). This finding underlines the potentially large significance of our paper's empirical results on the effects and the transmission of QE policies.

6. Conclusion

The findings of this paper suggest that the impact of the PSPP on public funding maturities in the euro area is twofold. (i) The reduced yield level led to a lengthening of issuance maturities by about seven months, while (ii) increased demand for PSPP-eligible bonds led to a lengthening of issuance maturities by about six months on average. The overall monthly average effect of the PSPP on issuance maturities is, hence, estimated to be 1.1 years, which compares with the average maturity outstanding of Germany, France, Italy, and Spain before the PSPP of about six years. We argue that the maturity extension by euro area governments represents a rational response to the altered cost-risk trade-off faced by DMOs, whereby overall funding costs, term spreads, as well as refinancing risks are reduced. It is the intention of the PSPP to alleviate financing conditions for the whole economy, and DMOs acted accordingly and endogenously in response to the changed conditions.

This paper represents a first assessment of an interaction between asset purchase programs and DMO funding behavior in the euro area. It contributes to the literature that examines such a relationship in the United States (see in particular Greenwood et al. 2014).

The empirical literature to date does not investigate the real economic consequences of longer-dated maturity structures of public debt. This link is also not addressed in our paper. We do, however, provide empirical evidence of a link between QE effects and longer-dated public debt. The results of this paper are thereby a basis for further work on the economic impact of maturity extension by DMOs

during the PSPP and its potential relevance for the transmission of monetary policy (see also Friedman 1992).

Finally, our results imply that DMO reaction functions should be internalized where relevant in monetary policy research and not treated as an exogenous variable. This paper illustrates that DMO behavior in response to yield changes and demand factors can to some extent be predicted. This opens the possibility to also treat DMO funding maturities as an endogenous factor in impact estimations of central bank purchase programs.

Appendix A. Decomposition of WAM Outstanding

The total outstanding amount of a DMO debt portfolio, denoted by out_t, evolves according to

$$out_t = out_{t-1} + iss_t - red_t, \tag{A.1}$$

where iss_t denotes the nominal amount of newly issued debt and red_t is the nominal amount of outstanding debt that is redeemed in period t (called redemptions hereafter), consisting both of regularly maturing securities and of active buybacks by the DMO.

The weighted average maturity of the outstanding debt portfolio, WAM_t^{out}, changes over time according to the following identity:

$$WAM_t^{out} = (WAM_{t-1}^{out} - a) \tag{A.2a}$$

$$+ \frac{iss_t}{out_t}\left[WAM_t^{iss} - (WAM_{t-1}^{out} - a)\right] \tag{A.2b}$$

$$+ \frac{red_t}{out_t}\left[WAM_t^{red} - (WAM_{t-1}^{out} - a)\right], \tag{A.2c}$$

where WAM_t^{iss} represents the weighted average (residual) maturity of newly issued securities and WAM_t^{red} is the weighted average (residual) maturity of redemptions at the time of their redemption. In the absence of any DMO buybacks and subsequent cancellations, $WAM_t^{red} = 0$.

The first component of the identity, given by a on the right-hand side of (A.2a), captures the roll-down of all outstanding maturities by one period every period. This deterministic reduction is termed the *aging effect*. Component (A.2b), termed the *issuance*

Table A.1. Decomposed Cumulative Change of WAM Outstanding

Cumulative Total Jan. 10 to Apr. 19 (Years)	DE	FR	IT	ES
Aging Effect	–9.3	–9.3	–9.3	–9.3
Redemption Effect	15.4	23.1	13.7	11.1
Issuance Effect	–5.7	–12.5	–4.7	–0.4
Other (Currency Conversion/Accounting)	0.2	0.2	0.1	–0.3
WAM Outstanding Change Period	0.6	1.5	–0.3	1
WAM Outstanding Start	5.9	6.6	7	6.5
WAM Outstanding End	6.6	8	6.8	7.6
Average WAM Outstanding	6	7.2	6.7	6.5
Average WAM Issuance	4	3.8	4.7	5.6

Note: All numbers are measured in years. The cumulative change in each component of WAM_t^{out} as shown in (A.2a) to (A.2c) is presented in the upper part of this table.

effect, covers the effect of newly issued debt on WAM_t^{out}. Whenever the WAM of the newly issued debt is higher than last period's WAM of the outstanding debt minus aging, the WAM outstanding will increase in the current period. The closer WAM_t^{iss} is to WAM_t^{out}, the smaller the overall impact of new issuances will be. Similarly, Component (A.2c), termed the *redemption effect*, covers the effect of debt redemptions on WAM_t^{out}. As redemptions generally have low or zero maturities, the weighted average maturity of the outstanding portfolio increases after redemptions. While the issuance effect depends on the DMOs' funding decisions in the given period, the aging and redemption effects on WAM outstanding are a consequence of historical portfolio legacy effects (except for buyback events).

In order to illustrate the working of Equation (A.2), we quantify the relative contribution of issuance compared with aging and redemptions in determining changes in WAM_t^{out} in our data set for the largest four euro area countries (DE, FR, IT, ES). To this end, Table A.1 summarizes the cumulative sum of the *aging effect*, the *issuance effect,* and the *redemption effect* over the whole sample period of 9.3 years.

In line with the description given above, the WAM of the outstanding portfolio increases significantly following redemptions. New

issuances contribute less than both redemptions and aging to the total change in WAM outstanding. In fact, new issuances are found to have a negative cumulative impact although they inject duration into the market. All four jurisdictions have active bill markets (with maturities of less than one year), which make up a relatively high contribution of total bonds outstanding and which roll over on a regular basis due to their relatively low maturity at issuance. The negative effect of new issuance on the WAM outstanding is comparably large for France, which can be explained by the relatively high volume of bills in the AFT's portfolio.[24] Overall, the effect of redemptions net of aging is larger than the effect of newly issued debt, implying that portfolio legacy effects contribute more to changes in WAM_t^{out} than current portfolio decisions. The WAM of the outstanding debt is therefore a poor behavioral indicator for current DMO funding decisions.

Figure A.1 shows the components of WAM outstanding across the sample period for the euro area, where the redemption and the constant aging effect have been summed. The component indicators fluctuate and broadly counterbalance each other, thereby generally stabilizing the WAM of outstanding debt. It can be seen that WAM_t^{out} remained relatively stable until the middle of 2014, after which it increased by more than one year.

[24]See, for example, the ECB Government Finance Statistics on this. France has a relatively large number of money market funds, which contribute to the active bill market.

Figure A.1. WAM Components: Euro Area

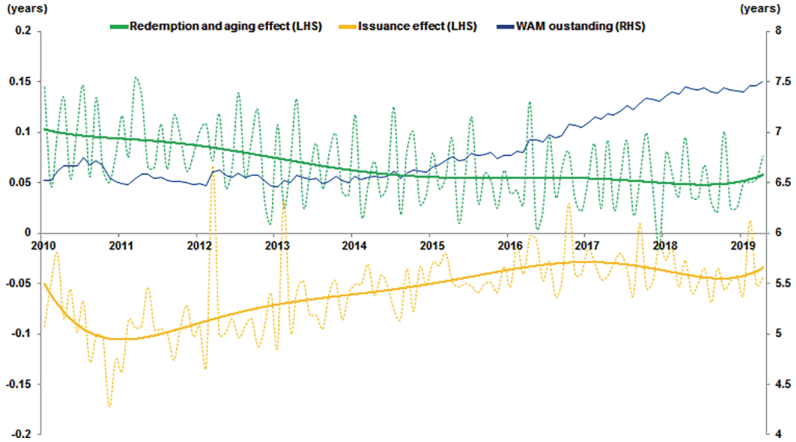

Note: Euro area changing composition based on CSDB data. WAM outstanding denotes the weighted average maturity of the outstanding debt portfolio. The issuance, redemption, and aging effects are calculated as shown in Equations (A.2a) to (A.2c). The aging effect represents the roll-down of outstanding maturities by one period every period. Redemption and aging effects are summed. Trend lines are added for illustrative purposes and are derived from a sixth-order polynomial for all indicators.

Appendix B. Summary and Descriptive Statistics

Table B.1. Summary Statistics for the Euro Area

	N	Unit	Mean	SD	Min.	Max.
EA: During PSPP (March 2015–April 2019)						
WAM Issuance	350	Years	4.77	2.52	0.14	18.27
PSPP/Issuance	350	Pct.	24.1	17.25	0.68	134.62
10-Year Yield	350	Pct.	1.19	0.88	–0.09	4.06
5-Year Yield	350	Pct.	0.29	0.65	–0.55	2.65
Redemptions	350	EURm.	19,048	15,355	0	67,651
WAM Outstanding	350	Years	7.13	0.9	6.05	10
Δ Industrial Production	350	Index	1.51	2.73	–7.2	10.5
Inflation	350	Pct. Change	1.1	0.89	–1.2	3.3
EA: Before PSPP (December 2009–February 2015)						
WAM Issuance	441	Years	3.69	2.39	0.16	13.85
PSPP/Issuance	441	Pct.	0	0	0	0
10-Year Yield	441	Pct.	3.63	2.23	0.35	14.09
5-Year Yield	441	Pct.	2.72	2.7	–0.06	17.5
Redemptions	441	EURm.	22,837	16,949	9	73,505
WAM Outstanding	441	Years	6.36	0.61	4.97	7.8
Δ Industrial Production	441	Index	0.53	4.49	–12.3	14
Inflation	441	Pct. Change	1.67	1.19	–1.5	4
EA: Full Sample (December 2009–April 2019)						
WAM Issuance	791	Years	4.17	2.51	0.14	18.27
PSPP/Issuance	791	Pct.	10.66	16.58	0	134.62
10-Year Yield	791	Pct.	2.55	2.14	–0.09	14.09
5-Year Yield	791	Pct.	1.64	2.39	–0.55	17.5
Redemptions	791	EURm.	21,160	16,362	0	73,505
WAM Outstanding	791	Years	6.7	0.84	4.97	10
Δ Industrial Production	791	Index	0.96	3.84	–12.3	14
Inflation	791	Pct. Change	1.41	1.1	–1.5	4

Note: Euro area includes BE, DE, FR, ES, IT, NL, and PT. Issuance and PSPP are monthly nominal values. PSPP is based on gross purchases. Industrial production index is excluding construction and calculated as annual rate of change. Inflation is based on annual rate of change of Eurostat HICP Index, neither seasonally nor working-day adjusted.

Table B.2. Summary Statistics for the Big 4 Countries

	N	Unit	Mean	SD	Min.	Max.
Big 4: During PSPP (March 2015–April 2019)						
WAM Issuance	200	Years	5.17	1.89	0.66	12.76
PSPP/Issuance	200	Pct.	25.53	17.24	2.18	134.62
10-Year Yield	200	Pct.	1.15	0.75	–0.09	3.48
5-Year Yield	200	Pct.	0.28	0.57	–0.55	2.65
Redemptions	200	EURm.	27,974	14,335	2,068	67,651
WAM Outstanding	200	Years	6.84	0.55	6.07	8.04
Δ Industrial Production	200	Index	1.56	2.34	–6.4	7.2
Inflation	200	Pct. Change	0.96	0.92	–1.2	3
Big 4: Before PSPP (December 2009–February 2015)						
WAM Issuance	252	Years	4.01	2	0.37	11.9
PSPP/Issuance	252	Pct.	0	0	0	0
10-Year Yield	252	Pct.	3.31	1.5	0.35	6.86
5-Year Yield	252	Pct.	2.28	1.53	–0.06	6.88
Redemptions	252	EURm.	32,606	15,552	5,856	73,505
WAM Outstanding	252	Years	6.42	0.55	5.24	7.21
Δ Industrial Production	252	Index	0.37	4.43	–9.1	13.8
Inflation	252	Pct. Change	1.63	1.08	–1.5	3.8
Big 4: Full Sample (December 2009–April 2019)						
WAM Issuance	452	Years	4.52	2.03	0.37	12.76
PSPP/Issuance	452	Pct.	11.3	17.09	0	134.62
10-Year Yield	452	Pct.	2.36	1.63	–0.09	6.86
5-Year Yield	452	Pct.	1.4	1.56	–0.55	6.88
Redemptions	452	EURm.	30,556	15,185	2,068	73,505
WAM Outstanding	452	Years	6.61	0.58	5.24	8.04
Δ Industrial Production	452	Index	0.9	3.7	–9.1	13.8
Inflation	452	Pct. Change	1.33	1.07	–1.5	3.8

Note: Big 4 countries includes DE, FR, IT, and ES. Issuance and PSPP are monthly nominal values. PSPP is based on gross purchases. Industrial production index is excluding construction and calculated as annual rate of change. Inflation is based on annual rate of change of Eurostat HICP Index, neither seasonally nor working-day adjusted.

Figure B.1. WAM Outstanding by Country

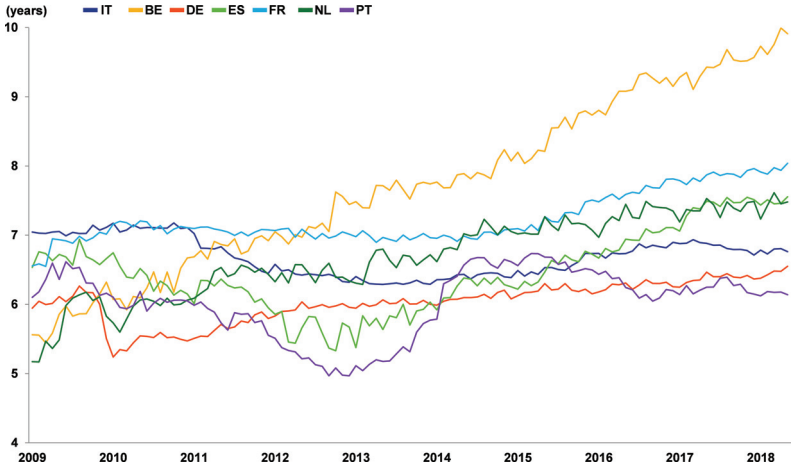

Figure B.2. WAM at Issuance by Country

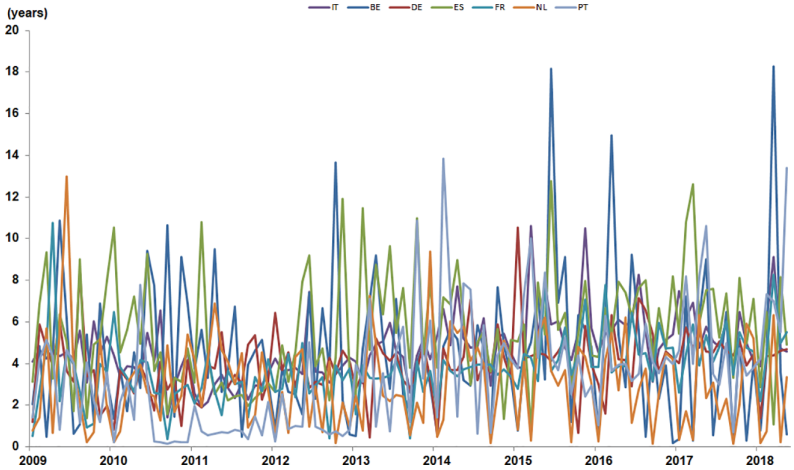

Table B.3. Bivariate Correlation Coefficients

	wam^{iss}	PSPP/Iss.	$yield^{5y}$	$yield^{10y}$	Redem.	wam^{out}	ΔIP	Infl.
WAM Issuance	1.00							
PSPP/Issuance	0.08	1.00						
5-Year Yield	–0.22	–0.36	1.00					
10-Year Yield	–0.20	–0.40	0.98	1.00				
Redemptions	–0.05	–0.11	–0.25	–0.26	1.00			
WAM Outstanding	0.13	0.16	–0.38	–0.42	0.07	1.00		
Δ Ind. Prod.	0.12	0.13	–0.24	–0.22	0.05	0.02	1.00	
Inflation	–0.17	–0.30	0.42	0.41	0.01	0.08	–0.12	1.00

Appendix C. Robustness Checks

Table C.1. The Effect of 5-Year Yields and the PSPP on WAM at Issuance: Euro Area

Dependent Variable: WAM Issuance	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
PSPP/Issuance						0.012**	0.010**	0.012**	0.012*	0.012**
						(0.0034)	(0.0033)	(0.0044)	(0.0052)	(0.0045)
5-Year Yield	–0.23**	–0.24**	–0.28**	–0.30***	–0.46***					
	(0.076)	(0.081)	(0.096)	(0.078)	(0.100)					
Redemptions		–0.55*	–0.57*	–0.48	–0.63**		–0.37	–0.38	–0.41	–0.51*
		(0.25)	(0.26)	(0.28)	(0.23)		(0.23)	(0.27)	(0.31)	(0.26)
WAM Outstanding			–0.16	–0.37	–0.70**			–0.60**	–0.96**	–0.63*
			(0.16)	(0.29)	(0.24)			(0.18)	(0.37)	(0.31)
Δ Industrial Production				0.029	0.026				0.064*	0.040
				(0.035)	(0.035)				(0.029)	(0.031)
Inflation					0.35					0.23
					(0.28)					(0.22)
Observations	784	784	784	784	784	784	784	784	784	784
Adjusted R^2	0.366	0.371	0.370	0.372	0.376	0.308	0.311	0.332	0.348	0.373

Note: Robust standard errors presented in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$. CCEP estimator is used for all regressions. All models include country fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. All independent variables except for redemptions are lagged by one month. The sample includes BE, DE, FR, ES, IT, NL, and PT over the period December 2009 to April 2019.

Table C.2. The Effect of 10-Year Yields and the PSPP on WAM at Issuance: Euro Area, Two-Way Fixed Effects Estimator

Dependent Variable: WAM Issuance	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
PSPP/Issuance						0.0074	0.0063	0.0035	0.0035	0.0050
						(0.0082)	(0.0085)	(0.0085)	(0.0085)	(0.0083)
10-Year Yield	–0.49***	–0.50***	–0.50***	–0.48***	–0.53***	–0.49***	–0.50***	–0.50***	–0.48***	–0.53***
	(0.065)	(0.065)	(0.065)	(0.074)	(0.079)	(0.065)	(0.065)	(0.065)	(0.074)	(0.079)
Redemptions		–0.24	–0.35	–0.36	–0.37		–0.21	–0.34	–0.34	–0.35
		(0.28)	(0.29)	(0.29)	(0.29)		(0.29)	(0.29)	(0.29)	(0.29)
WAM Outstanding			–0.65***	–0.63***	–0.73***			–0.64***	–0.62***	–0.73***
			(0.23)	(0.23)	(0.24)			(0.22)	(0.23)	(0.24)
Δ Industrial Production				0.021	0.024				0.021	0.024
				(0.037)	(0.037)				(0.037)	(0.037)
Inflation					0.36**					0.36**
					(0.18)					(0.18)
Observations	784	784	784	784	784	784	784	784	784	784
R^2	0.371	0.371	0.385	0.386	0.390	0.371	0.372	0.385	0.386	0.391

Note: Robust standard errors presented in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$. Fixed effects estimator is used for all regressions. All models include country and time fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. All independent variables except for redemptions are lagged by one month. The sample includes BE, DE, FR, ES, IT, NL, and PT over the period December 2009 to April 2019.

Table C.3. The Effect of 5-Year Yields and the PSPP on WAM at Issuance: Euro Area, Two-Way Fixed Effects Estimator

Dependent Variable: WAM Issuance	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
PSPP/Issuance						0.0078	0.0072	0.0044	0.0045	0.0052
						(0.0085)	(0.0087)	(0.0088)	(0.0087)	(0.0086)
5-Year Yield	–0.34***	–0.34***	–0.33***	–0.32***	–0.35***					
	(0.045)	(0.045)	(0.045)	(0.051)	(0.054)					
Redemptions		–0.19	–0.29	–0.30	–0.31		–0.11	–0.23	–0.27	–0.27
		(0.28)	(0.29)	(0.29)	(0.29)		(0.30)	(0.30)	(0.30)	(0.30)
WAM Outstanding			–0.60***	–0.58**	–0.68***			–0.64***	–0.56**	–0.60**
			(0.22)	(0.23)	(0.24)			(0.23)	(0.23)	(0.24)
Δ Industrial Production				0.024	0.026				0.072**	0.075**
				(0.036)	(0.036)				(0.034)	(0.034)
Inflation					0.35*					0.15
					(0.18)					(0.18)
Observations	784	784	784	784	784	784	784	784	784	784
R^2	0.369	0.370	0.382	0.382	0.387	0.338	0.339	0.352	0.357	0.358

Note: Robust standard errors presented in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Fixed effects estimator is used for all regressions. All models include country and time fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. All independent variables except for redemptions are lagged by one month. The sample includes BE, DE, FR, ES, IT, NL, and PT over the period December 2009 to April 2019.

Table C.4. The Effect of 10-Year Yields and the PSPP on WAM at Issuance: Euro Area, PCSE Estimator

Dependent Variable: WAM Issuance	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
PSPP/Issuance						0.018***	0.018***	0.018***	0.017***	0.017**
						(0.0064)	(0.0067)	(0.0067)	(0.0067)	(0.0066)
10-Year Yield	–0.27***	–0.28***	–0.32***	–0.29***	–0.24***	–0.21***	–0.21***	–0.24***	–0.22***	–0.18***
	(0.051)	(0.056)	(0.057)	(0.058)	(0.062)	(0.054)	(0.061)	(0.062)	(0.063)	(0.065)
Redemptions		–0.093	–0.10	–0.090	–0.041		–0.0065	–0.018	–0.0068	0.029
		(0.10)	(0.10)	(0.10)	(0.10)		(0.11)	(0.11)	(0.11)	(0.11)
WAM Outstanding			–0.17	–0.15	–0.072			–0.15	–0.13	–0.063
			(0.15)	(0.15)	(0.15)			(0.15)	(0.15)	(0.15)
Δ Industrial Production				0.055*	0.056*				0.053*	0.054*
				(0.030)	(0.029)				(0.029)	(0.028)
Inflation					–0.14					–0.12
					(0.12)					(0.12)
Observations	784	784	784	784	784	784	784	784	784	784
R^2	0.079	0.080	0.082	0.087	0.078	0.081	0.079	0.081	0.086	0.082

Note: Panel-corrected standard errors presented in parentheses, using panel-specific AR(1) autocorrelation structure. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. PCSE estimator is used for all regressions. All models include country fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. All independent variables except for redemptions are lagged by one month. The sample includes BE, DE, FR, ES, IT, NL, and PT over the period December 2009 to April 2019.

Table C.5. The Effect of 5-Year Yields and the PSPP on WAM at Issuance: Euro Area, PCSE Estimator

Dependent Variable: WAM Issuance	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
PSPP/Issuance						0.026***	0.027***	0.026***	0.025***	0.022***
						(0.0061)	(0.0062)	(0.0061)	(0.0062)	(0.0062)
5-Year Yield	–0.26***	–0.26***	–0.28***	–0.26***	–0.23***					
	(0.043)	(0.047)	(0.047)	(0.049)	(0.052)					
Redemptions		–0.083	–0.086	–0.074	–0.040		0.11	0.11	0.11	0.14
		(0.10)	(0.10)	(0.10)	(0.10)		(0.11)	(0.11)	(0.10)	(0.100)
WAM Outstanding			–0.12	–0.11	–0.058			0.052	0.055	0.10
			(0.14)	(0.14)	(0.15)			(0.14)	(0.14)	(0.14)
Δ Industrial Production				0.051*	0.052*				0.062**	0.061**
				(0.030)	(0.029)				(0.029)	(0.028)
Inflation					–0.12					–0.23**
					(0.12)					(0.11)
Observations	784	784	784	784	784	784	784	784	784	784
R^2	0.086	0.087	0.088	0.092	0.085	0.076	0.066	0.068	0.070	0.075

Note: Panel-corrected standard errors presented in parentheses, using panel-specific AR(1) autocorrelation structure. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. PCSE estimator is used for all regressions. All models include country fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. All independent variables except for redemptions are lagged by one month. The sample includes BE, DE, FR, ES, IT, NL, and PT over the period December 2009 to April 2019.

Figure C.1. Impulse Responses to 1 Basis Point Yield Shock: Parsimonious Specification

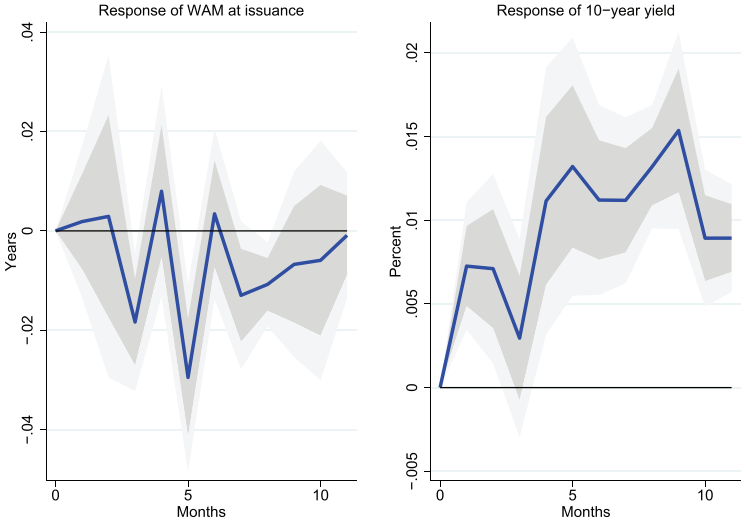

Note: Impulse response functions to a 1 basis point high-frequency yield shock based on local projections as defined in Equation (2). Control variables in regressions for WAM at issuance include an autoregressive term and the 10-year yield. The yield regression includes an autoregressive term only. Models include one lag for each variable. Robust standard errors are clustered by country. Dark (light) gray-shaded areas indicate 68 percent and 90 percent confidence intervals.

Table C.6. The Effect of 5-Year Yields and the PSPP on WAM at Issuance over Different Sub-samples

	EA		Big 4		Stressed		Non-stressed	
Dependent Variable: WAM Issuance	**(1^a)**	**(1^b)**	**(2^a)**	**(2^b)**	**(3^a)**	**(3^b)**	**(4^a)**	**(4^b)**
PSPP/Issuance	0.010*	0.0094*	0.018**	0.018**	0.020*	0.018**	0.013	0.013
	(0.0044)	(0.0043)	(0.0049)	(0.0045)	(0.0047)	(0.0039)	(0.0062)	(0.0070)
5-Year Yield	–0.38***	–0.30**‡	–0.55**	–0.53**‡	–0.27	–0.18‡	0.100	0.054
	(0.088)	(0.11)	(0.17)	(0.14)	(0.12)	(0.18)	(1.17)	(1.04)
5-Year Yield × PSPP-Dummy		–0.68‡		–0.18‡		–0.49‡		0.42
		(0.56)		(0.36)		(0.43)		(1.11)
Redemption Effect	–0.55*	–0.57**	–0.99**	–1.01**	–0.23	–0.29**	–1.73**	–1.73**
	(0.24)	(0.22)	(0.29)	(0.28)	(0.11)	(0.066)	(0.23)	(0.23)
WAM Outstanding	–0.64*	–0.56	–1.66	–1.69	–0.39	–0.26	–0.71	–0.71
	(0.33)	(0.36)	(1.51)	(1.54)	(0.31)	(0.50)	(0.52)	(0.58)
Δ Industrial Production	0.035	0.036	0.011	0.011	0.037	0.037	–0.071	–0.071
	(0.032)	(0.028)	(0.052)	(0.053)	(0.024)	(0.038)	(0.032)	(0.030)
Inflation	0.31	0.29	0.12	0.098	0.66**	0.64**	–0.18	–0.17
	(0.28)	(0.27)	(0.25)	(0.26)	(0.14)	(0.093)	(0.25)	(0.25)
Observations	784	784	448	448	336	336	336	336
Adjusted R^2	0.377	0.379	0.519	0.518	0.589	0.591	0.446	0.445
F(Yield^b, Interaction = 0)		10.07***		13.01**		173.5***		5.94
F(Yield^b = Yield^a)		0.45		0.02		0.21		0.00

Note: Robust standard errors presented in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$. CCEP estimator is used for all regressions. All models include country fixed effects. Dependent variable: WAM at issuance in years. PSPP/Issuance are monthly PSPP purchases per country divided by the monthly issuance volume of debt securities in the country. PSPP-dummy is 1 as of March 2015, otherwise 0. All independent variables except for redemptions are lagged by one month. ‡ denotes joint significance of the yield and the interaction term of yield*PSPP-dummy in the first F-test presented below the table. When the null hypothesis of the second F-test, Yield^b = Yield^a, cannot be rejected, the effect of 10-year yield on WAM issuance is not statistically different over the whole sample period (December 2009 to April 2019) and before the PSPP (until February 2015). Big 4 includes DE, FR, IT, and ES. Stressed includes ES, IT, and PT. Non-stressed includes DE, FR, and NL.

Figure C.2. Impulse Responses to 1 Basis Point Yield Shock: Model with Two Lags

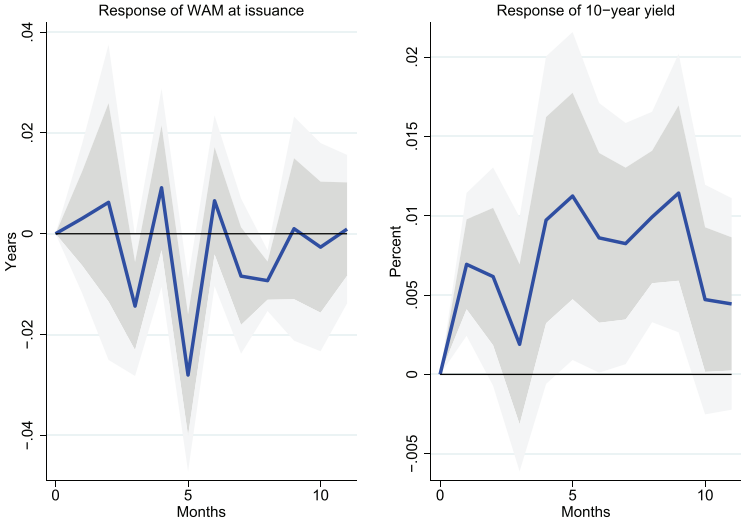

Note: Impulse response functions to a 1 basis point high-frequency yield shock based on local projections as defined in Equation (2). Control variables include an autoregressive term, the 10-year yield, PSPP/issuance, redemptions, WAM outstanding, the change of industrial production, and inflation. Models include two lags for each variable. Robust standard errors are clustered by country. Dark (light) gray-shaded areas indicate 68 percent and 90 percent confidence intervals.

Figure C.3. Impulse Responses to 1 Basis Point Yield Shock: No Lagged Dependent Variables

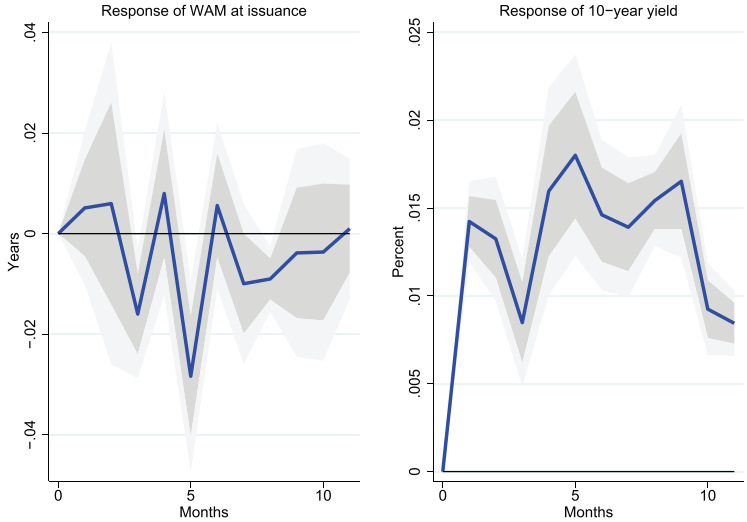

Note: Impulse response functions to a 1 basis point high-frequency yield shock based on local projections as defined in Equation (2). Control variables include the 10-year yield, PSPP/issuance, redemptions, WAM outstanding, the change of industrial production, and inflation, but no autoregressive lagged dependent variables. Models include one lag for each variable. Robust standard errors are clustered by country. Dark (light) gray-shaded areas indicate 68 percent and 90 percent confidence intervals.

References

Abbas, S. M. A., L. Blattner, A. E.-G. De Broeck, and M. Hu. 2014. "Sovereign Debt Composition in Advanced Economies: A Historical Perspective." Working Paper No. 14/162, International Monetary Fund.

Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa. 2019. "Measuring Euro Area Monetary Policy." *Journal of Monetary Economics* 108 (December): 162–79.

Altavilla, C., G. Carboni, and R. Motto. 2015. "Asset Purchase Programmes and Financial Markets: Lessons from the Euro Area." Working Paper No. 1864, European Central Bank.

Andrade, P., J. Breckenfelder, F. De Fiore, P. Karadi, and O. Tristani. 2016. "The ECB's Asset Purchase Programme: An Early Assessment." Working Paper No. 1956, European Central Bank.

Angeletos, G.-M. 2002. "Fiscal Policy with Noncontingent Debt and the Optimal Maturity Structure." *Quarterly Journal of Economics* 117 (3): 1105–31.

Badoer, D. C., and C. M. James. 2016. "The Determinants of Long-Term Corporate Debt Issuances." *Journal of Finance* 71 (1): 457–92.

Beck, N., and J. N. Katz. 1995. "What to Do (and Not to Do) with Time-Series Cross-Section Data." *American Political Science Review* 89 (3): 634–47.

Beetsma, R., M. Giuliodori, J. Hanson, and F. de Jong. 2021. "The Maturity of Sovereign Debt Issuance in the Euro Area." *Journal of International Money and Finance* 110 (February): Article 102293.

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent. 2017. "Fiscal Policy and Debt Management with Incomplete Markets." *Quarterly Journal of Economics* 132 (2): 617–63.

Boermans, M., and R. Vermeulen. 2018. "Quantitative Easing and Preferred Habitat Investors in the Euro Area Bond Market." Working Paper No. 586, De Nederlandsche Bank.

D'Amico, S., W. English, D. López-Salido, and E. Nelson. 2012. "The Federal Reserve's Large-scale Asset Purchase Programmes: Rationale and Effects." *Economic Journal* 122 (564): F415–F446.

D'Amico, S., and T. B. King. 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *Journal of Financial Economics* 108 (2): 425–48.

De Broeck, M., and A. Guscina. 2011. "Government Debt Issuance in the Euro Area: The Impact of the Financial Crisis." Working Paper No. 11/21, International Monetary Fund.

De Santis, R., and F. Holm-Hadulla. 2020. "Flow Effects of Central Bank Asset Purchases on Sovereign Bond Prices: Evidence from a Natural Experiment." *Journal of Money, Credit and Banking* 52 (6): 1467–91.

Dipartimento del Tesoro. 2015. "Annual Report 2015." http://www.dt.tesoro.it/modules/documenti_en/debito_pubblico/presentazioni_studi_relazioni/Public_Debt_Report_2015 (last accessed September 15, 2019).

———. 2016. "Public Debt Report 2016." www.dt.tesoro.it/export/sites/sitodt/modules/documenti_en/debito_pubblico/presentazioni_studi_relazioni/Public_Debt_Report_2016.pdf (last accessed October 28, 2019).

Eser, F., W. Lemke, K. Nyholm, S. Radde, and A. Vladu. 2019. "Tracing the Impact of the ECB's Asset Purchase Programme on the Yield Curve." Working Paper No. 2293, European Central Bank.

European Central Bank. 2010. "The 'Centralised Securities Database' in Brief." https://www.ecb.europa.eu/pub/pdf/other/centralisedsecuritiesdatabase201002en.pdf (last accessed October 25, 2019).

Foley-Fisher, N., R. Ramcharan, and E. Yu. 2016. "The Impact of Unconventional Monetary Policy on Firm Financing Constraints: Evidence from the Maturity Extension Program." *Journal of Financial Economics* 122 (2): 409–29.

Friedman, B. 1992. "Debt Management Policy, Interest Rates, and Economic Activity." In *Does Debt Management Matter?* ed. J. Agell, M. Persson, and B. Friedman. FIEF Studies in Labour Markets and Economic Policy. Oxford: Oxford University Press.

Gertler, M., and P. Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76.

Greenwood, R., S. Hanson, J. S. Rudolph, and L. H. Summers. 2014. "Government Debt Management at the Zero Lower Bound." Working Paper No. 5, Hutchins Center on Fiscal and Monetary Policy at Brookings.

Greenwood, R., S. Hanson, and J. C. Stein. 2010. "A Gap-Filling Theory of Corporate Debt Maturity Choice." *Journal of Finance* 65 (3): 993–1028.

———. 2015. "A Comparative-Advantage Approach to Government Debt Maturity." *Journal of Finance* 70 (4): 1683–1722.

Greenwood, R., and D. Vayanos. 2014. "Bond Supply and Excess Bond Returns." *Review of Financial Studies* 27 (3): 663–713.

Gürkaynak, R. S., B. Sack, and E. T. Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1 (1): 55–93.

Hammermann, F., K. Leonard, S. Nardelli, and J. von Landesberger. 2019. "Taking Stock of the Eurosystem's Asset Purchase Programme after the End of Net Asset Purchases." *ECB Economic Bulletin* (2/2019): 69–92.

Holm-Hadulla, F., and C. Thürwächter. 2021. "Heterogeneity in Corporate Debt Structures and the Transmission of Monetary Policy." *European Economic Review* 136 (July): Article 103743.

Hoogduin, L., B. Öztürk, and P. Wierts. 2011. "Public Debt Managers' Behaviour: Interactions with Macro Policies." *Revue Économique* 62 (6): 1105–22.

Im, K. S., M. H. Pesaran, and Y. Shin. 2003. "Testing for Unit Roots in Heterogeneous Panels." *Journal of Econometrics* 115 (1): 53–74.

Jarociński, M., and P. Karadi. 2020. "Deconstructing Monetary Policy Surprises—The Role of Information Shocks." *American Economic Journal: Macroeconomics* 12 (2): 1–43.

Jonasson, T., and M. Papaioannou. 2018. "A Primer on Managing Sovereign Debt-Portfolio Risks." Working Paper No. 18/74, International Monetary Fund.

Jordà, O. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.

Juodis, A., H. Karabiyik, and J. Westerlund. 2021. "On the Robustness of the Pooled CCE Estimator." *Journal of Econometrics* 220 (2): 325–48.

Kapetanios, G., M. H. Pesaran, and T. Yamagata. 2011. "Panels with Non-Stationary Multifactor Error Structures." *Journal of Econometrics* 160 (2): 326–48.

Koijen, R. S., F. Koulischer, B. Nguyen, and M. Yogo. 2021. "Inspecting the Mechanism of Quantitative Easing in the Euro Area." *Journal of Financial Economics* 140 (1): 1–20.

Krause, M. U., and S. Moyen. 2016. "Public Debt and Changing Inflation Targets." *American Economic Journal: Macroeconomics* 8 (4): 142–76.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120 (2): 233–67.

Kuttner, K. N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics* 47 (3): 523–44.

Leong, D. 1999. "Debt Management — Theory and Practice." Occasional Paper No. 18, HM Treasury.

Levin, A., C.-F. Lin, and C.-S. J. Chu. 2002. "Unit Root Tests in Panel Data: Asymptotic and Finite-Sample Properties." *Journal of Econometrics* 108 (1): 1–24.

Li, C., and M. Wei. 2013. "Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs." *International Journal of Central Banking* 9 (1): 3–39.

Maddala, G. S., and S. Wu. 1999. "A Comparative Study of Unit Root Tests with Panel Data and a New Simple Tests." *Oxford Bulletin of Economics and Statistics* 61 (S1): 631–52.

Missale, A., and O. J. Blanchard. 1994. "The Debt Burden and Debt Maturity." *American Economic Review* 84 (1): 309–19.

National Treasury Management Agency. 2013. "Promissory Note Transaction." https://www.ntma.ie/download/investor_presentation/PromissoryNoteTransaction070213.pdf (last accessed October 25, 2019).

Pesaran, M. H. 2004. "General Diagnostic Tests for Cross Section Dependence in Panels." Discussion Paper No. 1240, Institute of Labor Economics (IZA).

———. 2006. "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure." *Econometrica* 74 (4): 967–1012.

———. 2007. "A Simple Panel Unit Root Test in the Presence of Cross-section Dependence." *Journal of Applied Econometrics* 22 (2): 265–312.

Tobin, J. 1963. "An Essay on the Principles of Debt Management." In *Fiscal and Debt Management Policies*, Vol. 1, ed. W. Fellner and B. Fox. Englewood Cliffs.

Vayanos, D., and J.-L. Vila. 2021. "A Preferred-Habitat Model of the Term Structure of Interest Rates." *Econometrica* 89 (1): 77–112.

Wolswijk, G. 2020. "Drivers of European Public Debt Management." Working Paper No. 2437, European Central Bank.

Wolswijk, G., and J. de Haan. 2005. "Government Debt Management in the Euro Area: Recent Theoretical Developments and Changes in Practices." Occasional Paper No. 25, European Central Bank.

Optimal Central Bank Forward Guidance*

Eunmi Ko
Rochester Institute of Technology

When the future state of the economy is uncertain, yet a central bank has more information about the possible scenarios, how should the central bank communicate its private information to the public? This paper analyzes the optimal *tone* of central bank Delphic forward guidance using the Bayesian persuasion model (Kamenica and Gentzkow 2011). Assuming that monetary policy is an exogenously given function over the states of the economy and that the central bank is precommitted to a forward-guidance policy function, under certain conditions, the optimal tone of communication about the uncertain future is overly pessimistic.

JEL Codes: D82, D83, E58.

1. Introduction

Suppose central bank governors are required to offer an impromptu prediction of the state of the economy in the next one or two years. In the near future, they know that the economy could either experience a mild recession or stay expansionary. They think that both scenarios are possible, and that they need further information to offer a specific prediction. Then, what should they say to the public? What kind of tone should they use? Being vague by describing all the scenarios, including the above two, is one option. Or should they be optimistic, focusing only on the leading indicators that provide

*I am deeply indebted to Narayana Kocherlakota for his invaluable guidance and support throughout this project. I thank the editor, Klaus Adam, and an anonymous referee for helpful comments and suggestions. I also thank (including but not limited to) Mark Bils, Anmol Bhandari, George Alessandria, Yu Awaya, Yan Bai, Paulo Barelli, Asen Kochov, and participants at the department student seminars at the University of Rochester in 2018, the EEA-ESEM (Cologne, Germany, August, 2018), and the Midwest Macro Meetings (Vanderbilt University, November, 2018) for useful comments and questions. All remaining errors are my own. Author contact: Rochester Institute of Technology, 92 Lomb Memorial Dr, Rochester, NY 14623; e-mail: emkgse@rit.edu.

them with an expansionary forecast? Should they be pessimistic, emphasizing the negative indicators that predict a looming economic recession? Otherwise, should they simply point to the mixed signs from the economic indicators and the fact that they need further information?

Inadequate communication can mislead private agents' beliefs and risk evaluation about the future, thereby encouraging them to make premature decisions that may lead to either overly borrow to make inappropriately aggressive equipment investments or miss opportunities to grow faster. Moreover, it is not helpful to be always optimistic or always pessimistic because such central bank's communication may be regarded as uninformative or even incredible. Then, what is the optimal way of communicating the uncertain future?

Motivated by the above thought experiment, this paper describes a simple framework based on the Bayesian persuasion model (Kamenica and Gentzkow 2011) for analyzing optimal communication of a central bank's economic forecasts. The model consists of spaces with two signals, two states of the economy, and two messages. The state of the economy is about supply shocks. The monetary policy is exogenously determined as a response to the state of the economy: in the weak economy with an adverse cost shock, inflation is allowed to be high, while in the strong economy with an advantageous cost shock, inflation is controlled to be low. I have three results conditional on the characteristic of the exogenously given monetary policy regime (relatively unemployment fighting versus relatively inflation fighting) and the economic fundamental (the prior of an economic forecast of a strong or weak economy). First, given a relatively unemployment-fighting monetary policy regime and an economic environment with an a priori robust fundamental, it is optimal for the central bank to be overly pessimistic in communication. Second, under a relatively unemployment-fighting monetary policy regime and an economic environment with an a priori weak fundamental, central bank communication should be truthful. Third, given a relatively inflation-fighting monetary policy regime, it is optimal for the central bank to be uninformative.

The most interesting result is that the optimal central bank communication should be overly pessimistic if monetary policy regime is relatively unemployment fighting and the economic fundamental is a priori robust. The overly pessimistic communication means

that, first, when the central bank receives a signal forecasting the weak economy more likely, it should send a message associated with the weak economy ("pessimistic message") with probability 1. Second, when the central bank receives a signal forecasting the strong economy more likely, it should mix the pessimistic message and an optimistic message, which is associated with the strong economy. By the overly pessimistic forward guidance, the central bank can mitigate the welfare damage caused by a large deviation from inflation target under a relatively unemployment-fighting regime.

In the model, the central bank is able to manage private sector's expected inflation by communicating the uncertain future. The private sector sets expected inflation with anticipating central bank's monetary policy action, while the monetary policy action is assigned by an exogenously preprogrammed function of the state of the economy. The state of the economy is realized stochastically, and the monetary policy action is confirmed only after the state is realized. Before the state is realized, the central bank obtains a better probability distribution over the states of the economy as its private information. Hence, communication of the central bank's economic forecast ("forward guidance") helps the private sector set better inflation expectations.

Before it observes the private information, the central bank designs forward-guidance policy function to maximize social welfare. The social welfare function reflects both inflation stability and output gap stability. Central bank communication changes the private sector's beliefs about the likelihood of future states and affects the expected inflation. In this way, inflation surprises can be managed by the central bank even though the realized inflation will be exogenously bounded by the preprogrammed monetary policy function. Once the optimal forward-guidance policy function is determined, the central bank is pre-committed to it, and the private sector is informed of the forward-guidance policy function. The model is one-shot. However, the result does not change under a dynamic model: as forward-guidance policy function is pre-committed, there is no room for the private sector's learning.

The novel feature of this study is that it analyzes the optimal *tone* for the central bank communication to convey information about the uncertain future. The Bayesian persuasion model allows to assign a probability distribution over the message space (see Table 3). For

example, suppose the central bank privately forecasts a weak economy more likely, while there is a small but positive probability that a strong economy is realized. In this environment, the tone of the communication is formulated by a probability distribution over the message space. The tone is referred to as *truthful* if probability 1 is assigned to the pessimistic message in the above scenario, which means the truthful communication of the central bank's private information. However, the tone is referred to as *overly optimistic* if positive probabilities are assigned to both the optimistic message as well as the pessimistic message. By contrast, suppose the central bank privately forecasts the strong economy more likely, while there is a small but positive probability of the forecast proven wrong. The tone is referred to as *overly pessimistic* if a positive probability is also assigned to the pessimistic message along with the optimistic message.

In the sense that this paper deals with a central bank's communication about the future, it follows the line of literature on forward guidance. Gürkaynak, Sack, and Swanson (2005) empirically show the effect of verbal communication. Rudebusch and Williams (2008), Campbell et al. (2012), Del Negro, Giannoni, and Patterson (2012), and Andrade et al. (2019), among others, admit that empirically central bank forward guidance includes the communication of the future uncertainty, so-called Delphic forward guidance. Given reality, this paper normatively approaches how to optimally communicate the future uncertainty.

There has been numerous research on optimal central bank communication, especially using cheap talk by Crawford and Sobel (1982) or global game by Morris and Shin (2002). Morris and Shin (2018) describe optimal communication when the central bank has private noisy signal on the realized state, focusing on the reflection problem between the market and the central bank. Bassetto (2019) shows that when the central bank receives a noisy signal about the realized state, the central bank's cheap talk can increase social welfare.[1] This paper, however, provides a new framework to analyze the

[1]The research area of central bank communication about the future uncertainty is nascent, so most papers are at the working paper stage. Notably, Cieslak, Malamud, and Schrimpf (2019) suggest that the central bank's optimal communication of the state of the economy should be clustered. Fujiwara and Waki (2019)

optimal *tone* of central bank communication of the uncertain future: given a private signal forecasting either a strong economy more likely or a weak economy more likely, how optimistic or pessimistic the central bank communication should be.

The remainder of the chapter proceeds as follows. The simple two-signal, two-state, two-message model is introduced in Section 2. Sections 3 presents the threefold optimal central bank forward guidance. Section 4 discusses how the overly pessimistic communication can be better than the truthful communication in certain economic environments. In the appendix, I present that the threefold results of forward guidance hold in a generalized economic environment that incorporates possibly self-fulfilling effect of the private sector's expected inflation and allows for endogenous asymmetric inflation-targeting bandwidth.

2. The Model

2.1 Environment

In this section, I describe the primitives: the players, their payoff functions, their actions, and the private information of the central bank.

There are Nature, a central bank, and the private sector. It is a one-shot problem where the central bank pre-commits on a communication function as well as a monetary policy function. After the pre-commitment, the central bank executes actions according to the policy functions without reoptimization. The private sector determines the expected inflation after communication but before the monetary policy action.

The private sector's payoff function is given by a quadratic loss function from inflation surprises:

$$-\left(\pi^e - \pi\right)^2, \tag{1}$$

incorporate Bayesian persuasion into the New Keynesian model to suggest that central bank's Delphic communication may not increase social welfare. Related to communication *tone* in forecasting, Beaudry and Willems (2022) analyze International Monetary Fund (IMF) data to find that an overly optimistic tone in forecasting is correlated with economic contractions with time lags.

where $\pi^e \in \mathbb{R}$ is expected inflation and $\pi \in \mathbb{R}$ is realized inflation. The expected inflation is the private sector's action.

The ex post payoff of the central bank is given by a linear combination of the quadratic losses from both the inflation gap and the output gap. The inflation gap is measured by the deviation of realized inflation from the target, and the output gap is measured by cyclical unemployment. Let u^{NAR} and π^T denote the natural unemployment rate and inflation target, respectively, where u and π stand for the unemployment and inflation rates. Then, the payoff function is given by

$$W(u, \pi) = -(u - u^{NAR})^2 - \alpha \cdot (\pi - \pi^T)^2, \tag{2}$$

where u^{NAR} and π^T are parameters. The unemployment u is determined by a classic Phillips curve:

$$u - u^{NAR} = \theta + \gamma \cdot (\pi^e - \pi), \tag{3}$$

where γ is a parameter of unemployment for inflation surprises and θ represents the state of the economy. In addition, the realized inflation π is determined by an exogenously given monetary policy function $\pi(\cdot)$ and the realized state of the economy θ.

$$\text{(monetary policy) } \pi : \ \Theta \ \rightarrow \mathbb{R} \tag{4}$$

Nature informs θ to the central bank only as the private information. Nature draws the state of the economy from an exogenously given probability distribution Φ. The timing when the state of the economy is realized is *after* the communication occurs, and this formulation makes the communication be about the future uncertainty. The disposition of the state of the economy is adverse or advantageous cost-push shocks.

There are two possible states of the economy: $\Theta = \{\theta_S, \theta_W\}$. Here, θ_S and θ_W are real numbers related to the unemployment rate where $\theta_W > \theta_S$. When the economy is hit by an adverse cost-push shock, the state is θ_W (the "weak economy"); then, cyclical unemployment soars, assuming all other factors remained the same. When the economy is hit by an advantageous cost shock, the state is θ_S (the "strong economy"); then, cyclical unemployment is relatively low.

Table 1. Prior Distribution for Signals

	s_S	s_W
Prob.	ψ_0	$(1-\psi_0)$

Table 2. Stochastic Evolution of the State of the Economy

	θ_S	θ_W
s_S	ψ_S	$(1-\psi_S)$
s_W	$(1-\psi_W)$	ψ_W

In addition to the state of the economy, θ, Nature draws one more kind of private information for the central bank with a time lead: a private signal, s, which corresponds to the central bank's economic forecasts to the state of the economy. There are two possible private signals: $\{s_S,\ s_W\}$. Here, s_S denotes a signal to forecast the future *strong* economy with higher probability, whereas s_W denotes a signal to forecast the future *weak* economy with higher probability. The probability $Pr(s_S)$ of a signal to forecast future strong economy with higher probability is denoted by $\psi_0 \in [0,1]$. The private sector does not know the realized signal; however, they know the prior probability distribution over signals as summarized in Table 1.

The communication of the central bank concerns the private signal rather than the state of the economy. When its communication occurs, the central bank does not know which state of the economy will be realized but knows better about the stochastic evolution of the state of the economy than the private sector. Each private signal can be recast as a probability distribution over the state of the economy. The stochastic evolution of the state of the economy conditional on each signal is summarized in Table 2, and the Markov chain is common knowledge to the private sector.

Assume that $\psi_S \in (1/2, 1]$ as well as $\psi_W \in (1/2, 1]$. Hence, $\psi_S > 1 - \psi_W$. If the central bank observes the signal s_S, it is more likely that θ_S is realized; however, there is a small but positive probability that θ_W is realized. Similarly, if the central bank observes the signal s_W, it is more likely that θ_W is realized; nevertheless, there

Figure 1. Timeline

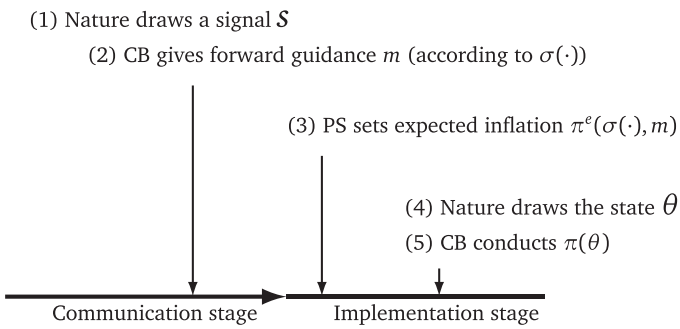

is a small but positive probability that θ_S is realized.[2] Owing to an empirical condition, I assume that $\left\{\theta_W \psi_W + \theta_S \cdot (1 - \psi_W)\right\} > 0$.[3] This condition sets a lower bound for θ_S, especially when $\theta_S < 0$.

With two types of private information, the timeline (Figure 1) is divided into two stages: the communication stage and the monetary policy implementation stage. In the communication stage, the central bank and private sector play in turn. The stage begins with the central bank observing the private signal. Then, the central bank sends a message according to a pre-committed communication policy function. There are two possible messages: an optimistic message and a pessimistic message. Here, the message space is denoted by $M = \{m_{opt},\ m_{pes}\}$. The communication policy function prescribes the central bank a probability distribution over the message space conditional on the realized state of the economy.

Reflecting the fact that the central bank communicates the uncertain future, the central bank's communication policy function

[2]Intuitively, this formation between the economic forecast and realized future state of the economy reflects the following observation: even if the central bank privately forecasts the strong economy, it is still just a forecast. There is a chance that the economy deteriorates rapidly. Conversely, even if the central bank (privately) forecasts the weak economy, there is a chance that economic conditions improve.

[3]The empirical conditions stems from the Phillips curve: Cogley and Sargent (2005) have a lagged term in their equation (1). The lagged term plays a role as a positive bias for cyclical unemployment. The theoretical results of this paper hold without the empirical condition.

is referred to as a *forward-guidance* policy function.[4] It is the objective of policy design as a functional of which domain is the set of signals. Let $\sigma(\cdot)$ denote the policy function:

$$\text{(forward-guidance policy function) } \sigma : \ S \ \rightarrow \Delta(M), \qquad (5)$$

where $\Delta(M)$ means the set of probability distributions over the message space M. Let Σ denote the set of all feasible communication policy functions.

The choice of optimal forward-guidance policy function $\sigma(\cdot)$ is of interest in this paper. Once the optimal forward-guidance policy function is imposed, the central bank is pre-committed to the policy function, and it sends messages according to the policy prescription. Forward guidance does not affect the stochastic evolution of the state of the economy.

Observe that in comparison with the Kamenica and Gentzkow (2011), the communication stage corresponds to the conventional Bayesian persuasion model. I refer to the "state of the world" in Kamenica and Gentzkow (2011) as a "private signal," the "signal" in Kamenica and Gentzkow (2011) as a "forward-guidance policy function," and the "signal realization" in Kamenica and Gentzkow (2011) as a "message." The term "state" in this paper is reserved for the basis of monetary policy in the implementation stage. Each private signal can be recast as a probability distribution over the state space.

The private sector knows that the central bank communicates the future uncertainty, and the central bank's message does not guarantee a certain state of the economy. Also, the private sector is aware that the central bank's pre-committed communication policy is designed strategically to maximize the social welfare. The private sector updates their posterior after the central bank's message is received. The private sector decides, first, whether to "listen to" (Bergemann and Morris 2016 refer to this action as "obey") the announcement of the central bank and, second, how to formulate expected inflation. Once the private sector listens to the message, then it updates its beliefs about the true state of the economy. Each

[4]More specifically, it is Delphic forward guidance according to Campbell et al. (2012).

posterior is determined by the realized message as well as the pre-committed forward-guidance policy function from which the message is drawn.

After receiving the message m, the private sector's posterior is $\mathcal{P}(\cdot \mid \sigma(\cdot), m)$. The private sector's maximizer is given by

$$\pi^e\Big(\sigma(\cdot), m\Big) = \int_\theta \pi(\theta)\ \mathcal{P}\left(d\theta \mid \sigma(\cdot), m\right). \tag{6}$$

Given the monetary policy function, the expected inflation is a function of the forward-guidance policy function and a realized message:[5]

$$\text{(the private sector's action) } \pi^e :\ \Sigma \times M \to \mathbb{R}. \tag{7}$$

I abuse the notation for π^e so that it represents both a realized value and a measurable function.

After the private sector's expected inflation is chosen, the monetary policy implementation stage begins. The central bank undertakes its monetary policy action $\pi(\theta)$ according to an exogenously given monetary policy function $\pi(\cdot)$ and the realized state of the economy θ. Then, the ex post payoff of the central bank can be rewritten as

$$\begin{aligned} &W(\pi^e\Big(\sigma(\cdot),\ m\Big), \pi(\theta), \theta) \\ &\quad = -\left[\theta + \gamma \cdot \left\{\pi^e\left(\sigma(\cdot), m\right) - \pi(\theta)\right\}\right]^2 - \alpha \cdot (\pi(\theta) - \pi^T)^2, \end{aligned} \tag{8}$$

where $\gamma > 0$, and $\alpha > 0$. With this model, the goal of this study is to design the optimal *forward-guidance* policy function to maximize social welfare, which is ex ante expected payoff of the central bank. Social welfare is measured even before any signal is drawn by Nature.

2.2 Parameterization and Monetary Policy Regime

In this subsection, the forward-guidance policy function and monetary policy function are parameterized. Observe that the co-domain

[5]The choice of expected inflation π^e does not affect the stochastic evolution of the state of the economy. Hence, there are no self-fulfilling effects.

Table 3. Forward-Guidance Policy Function $\boldsymbol{\sigma : S \to \Delta(M)}$

$\boldsymbol{\sigma(s) \in \Delta(M)}$	$\boldsymbol{m_{opt}}$	$\boldsymbol{m_{pes}}$
s_S	ρ_{opt}	$(1-\rho_{opt})$
s_W	$(1-\rho_{pes})$	ρ_{pes}

of forward-guidance policy function σ is the set of probability distribution over the message space. The parameterization of policy function allows us to solve the optimization problem analytically. The abstract form of forward-guidance policy function is given by

$$\text{(forward-guidance policy function) } \sigma : \ S \ \to \Delta(M). \tag{9}$$

Since there are two possible messages, a probability distribution over the message space can be parameterized by a probability mass on a specific message. That is, first, when the central bank observes the signal s_S, a probability distribution over message space can be represented by a probability mass $Pr(m_{opt} \mid s_S) \equiv \rho_{opt}$ on the optimistic message m_{opt}. Then, the probability mass on the pessimistic message m_{pes} is the complementary probability $(1-\rho_{opt})$. Second, symmetrically, when the central bank observes the signal s_W, a probability distribution can be represented by a probability mass $Pr(m_{pes} \mid s_W) \equiv \rho_{pes}$ on the pessimistic message m_{pes}. Assume that the unconditional probabilities are not zero: $Pr(m_{opt}) = \psi_0 \cdot \rho_{opt} + (1-\psi_0) \cdot (1-\rho_{pes}) \neq 0$ and $Pr(m_{pes}) = \psi_0 \cdot (1-\rho_{opt}) + (1-\psi_0) \cdot \rho_{pes} \neq 0$. Table 3 summarizes the parameterization.

The optimal values of parameters ρ_{opt} and ρ_{pes} are of interest in this paper. Observe that the case in which $\rho_{opt} = (1-\rho_{pes})$ means that the central bank's message is *uninformative*: given any message, the private sector's updated beliefs will be exactly the same to the prior.

In addition, for analytical simplicity, monetary policy function $\pi : \Theta \to \mathbb{R}$ takes a linear functional form as follows:

$$\pi(\theta_S) = \pi^T - v, \qquad \pi(\theta_W) = \pi^T + v, \tag{10}$$

where $v \geq 0$. The monetary policy function reflects a simplified flexible inflation targeting under the Qvigstad rule (Walsh 2014).[6]

There are 62 countries that use inflation targeting with an explicitly stated inflation target (Foerster and Davig 2017). The majority of them (40 countries) use flexible inflation targeting with a range for target inflation rate as opposed to inflation targeting with pointwise target. For instance, the two pioneers of inflation targeting, New Zealand and Canada, adopted flexible inflation targeting while the United Kingdom, the third, adopted pointwise inflation targeting.

The optimal monetary policy under flexible inflation targeting is suggested by Qvigstad (2006), and coined as the "Qvigstad rule" by Walsh (2014). With flexible inflation targeting, the optimal monetary policy is characterized by the opposite sign of output gap and inflation deviation from the numerical target inflation point (also known as the inflation gap). It is an empirically supposed relationship based on the following evaluation: if the output gap is positive and inflation is above the target point, monetary policy is too accommodative. If the output gap is negative and inflation is below the target, monetary policy is too tight.

Equation (10) follows the Qvigstad rule. When the economy is hit by an adverse cost-push shock, monetary policy prescribes accommodative policy action so that inflation rate can rise above the target. When the economy is hit by advantageous cost shock, monetary policy prescribes tightening monetary policy action so that inflation rate decreases below the target. The degree of deviation from the inflation target v is a parameter. While the optimal values for ρ_{opt} and ρ_{pes} are to be solved, the value of v is assumed as exogenously given.

The monetary policy regime is defined by the magnitude of v. The threshold is $\left(\dfrac{\theta_W - \theta_S}{\gamma}\right)$. If monetary policy regime is "relatively unemployment fighting," i.e., $v > \left(\frac{\theta_W - \theta_S}{\gamma}\right)$, the central bank bears with sufficiently high inflation when unemployment is high under the adverse cost-push shock by the exogenously imposed

[6]More generalized formulations give the same results.

policy prescription. Otherwise, monetary policy is referred to as "relatively inflation fighting." Monetary policy regime is public information; the private sector knows whether it is relatively unemployment fighting or relatively inflation fighting before it chooses expected inflation.

3. Optimal Central Bank Forward Guidance

The main theorem of this paper is as follows:

THEOREM 1. *Optimal central bank forward-guidance policy function should be*

(i) overly pessimistic if monetary policy is relatively unemployment fighting and if it is unlikely to forecast a future weak economy;

(ii) truthful if monetary policy is relatively unemployment fighting and if it is highly likely to forecast a future weak economy;

(iii) uninformative if monetary policy is relatively inflation fighting.

Sketch of Proof. The policy designer wants to maximize the ex ante expected payoff of the central bank, subject to the private sector's best response. To obtain the optimal policy, the idea of backward induction is applied: For any pre-committed forward-guidance policy function and any realized message, the private sector best responds to it. The private sector updates posterior and takes its action by setting the expected inflation.

The private sector updates its beliefs over the future state of the economy as follows: given a message m and a pre-committed forward-guidance policy function $\sigma : S \rightarrow \Delta(M)$,

$$Pr(\theta \mid \sigma(\cdot), m) = \sum_{s \in S} Pr(\theta \mid s) \cdot Pr(s \mid \sigma(\cdot), m). \tag{11}$$

Table 4. Posterior Using $\sigma : S \to \Delta(M)$ and $Pr(s_S) = \psi_0, Pr(s_W) = 1 - \psi_0$

$Pr(s\|\sigma(\cdot), m)$	s_S	s_W
m_{opt}	$\frac{\psi_0\rho_{opt}}{(1-\psi_0)(1-\rho_{pes})+\psi 0\rho_{opt}}$	$\frac{(1-\psi_0)(1-\rho_{pes})}{(1-\psi_0)(1-\rho_{pes})+\psi_0\rho_{opt}}$
m_{pes}	$\frac{\psi_0(1-\rho_{opt})}{(1-\psi_0)\rho_{pes}+\psi_0(1-\rho_{opt})}$	$\frac{(1-\psi_0)\rho_{pes}}{(1-\psi_0)\rho_{pes}+\psi_0(1-\rho_{opt})}$

For an exogenously given monetary policy function $\pi : \Theta \to \{\pi^T - v,\ \pi^T + v\}$, the expected inflation is set as

$$\pi^e(\sigma(\cdot),\ m) = \left(\pi^T - v\right) \cdot \left[\sum_{s\in S} Pr(\theta_S \mid s) \cdot Pr(s \mid \sigma(\cdot), m)\right] \tag{12}$$
$$+ \left(\pi^T + v\right) \cdot \left[\sum_{s\in S} Pr(\theta_W \mid s) \cdot Pr(s \mid \sigma(\cdot), m)\right].$$

For arbitrary ρ_{opt} and ρ_{pes}, for each $m \in M$, the posterior is as in Table 4. Then, the expected inflation conditional on each message is given by, respectively,

$$\pi^e(\rho_{opt},\ \rho_{pes},\ m_{opt}) \tag{13}$$
$$= \pi^T + v\left[\frac{(2\psi_W - 1)(1-\psi_0)(1-\rho_{pes}) - (2\psi_S - 1)\psi_0\rho_{opt}}{(1-\psi_0)(1-\rho_{pes}) + \psi_0\rho_{opt}}\right],$$

$$\pi^e(\rho_{opt},\ \rho_{pes},\ m_{pes}) \tag{14}$$
$$= \pi^T + v\left[\frac{(2\psi_W - 1)(1-\psi_0)\rho_{pes} - (2\psi_S - 1)\psi_0(1-\rho_{opt})}{(1-\psi_0)\rho_{pes} + \psi_0(1-\rho_{opt})}\right].$$

The efficacy of forward-guidance policy function relies on the private sector participating in the communication by updating their posterior. It is at issue to ensure that the private sector's expected payoff does not deteriorate by updating the posterior in order to give the private sector the incentive to adjust their belief after

the central bank's forward guidance. In conventional Bayesian persuasion models, the incentive compatibility of the private sector participating in a communication becomes a restriction for the communication design (see the obedience condition in Bergemann and Morris 2016, p. 587). However, in this game structure of a central bank and the private sector, the incentive compatibility does not restrict the central bank's choice of forward-guidance policy function. It is due to the private sector's action space to be richer than the message space. In particular, for a given triple of $(\rho_{opt},\ \rho_{pes},\ m_{pes})$, the expected inflation can potentially be any number in between the upper bound and the lower bound of the inflation,

$$\pi^e(\rho_{opt},\ \rho_{pes},\ m_{pes}) \in \left[\pi^T - v,\ \ \pi^T + v,\right], \tag{15}$$

whereas the message space has only two elements, $M = \{m_{opt},\ m_{pes}\}$. Then, the private sector's expected payoff under the posterior is not less than the expected payoff under the prior. The policy designer can optimize the forward-guidance policy function relaxed from the private sector's incentive compatibility in communication.

Then, the policy designer solves the optimization problem for the forward guidance subject to the private sector's expected inflation only. The results are threefold conditional on the monetary policy regime, i.e., the value of $(\theta_W - \theta_S - \gamma v)$, and the economic fundamental, i.e., the probability of observing the signal that predicts the strong economy, $Pr(s_S) = \psi_0$. The mathematical derivation is elaborated in the online math appendix (Ko 2022).

Optimal Uninformative Forward-Guidance Policy Function. For the case where monetary policy regime is relatively inflation fighting, i.e., $\theta_W - \theta_S - \gamma v > 0$, the optimal forward-guidance policy function is as follows: irrespective of the signal that the central bank observes, it should send both the optimistic message and the pessimistic message with probability $\frac{1}{2}$. The message is completely uninformative because the posterior of the private sector remains equivalent to the prior. It can be summarized as in Table 5.

Effects of Optimal Uninformative Forward Guidance. When the central bank pre-commits to implement the above forward

Table 5. Optimal Uninformative Forward-Guidance Policy Function if $\theta_W - \theta_S - \gamma v > 0$ (hawkish)

$\sigma : S \rightarrow \Delta(M)$	m_{opt}	m_{pes}
s_S	1/2	1/2
s_W	1/2	1/2

guidance, i.e., $\rho_{opt} = 1/2,\ \rho_{pes} = 1/2$ to Equation (13) and Equation (14), the private sector forms the following expected inflation after receiving the optimistic and pessimistic message:

$$\pi^e(\sigma(\cdot), m_{opt}) = \pi^T + v \cdot \left[(1 - 2\psi_S)\psi_0 + (2\psi_W - 1)(1 - \psi_0)\right]$$
$$\pi^e(\sigma(\cdot), m_{pes}) = \pi^T + v \cdot \left[(1 - 2\psi_S)\psi_0 + (2\psi_W - 1)(1 - \psi_0)\right].$$

Under the optimal uninformative forward guidance, however, irrespective of the received message, the private sector sets the same expected inflation based only on the prior. It is because the private sector is given no additional information through the communication of the central bank. The inflation surprise in each state of the economy is, respectively,

$$\pi(\theta_S) - \pi^e(m_.) = -v \cdot \left[2\psi_0(1 - \psi_S) + 2(1 - \psi_0)\psi_W\right] < 0,$$
$$\pi(\theta_W) - \pi^e(m_.) = v \cdot \left[2\psi_0\psi_S + 2(1 - \psi_0)(1 - \psi_W)\right] > 0,$$

for any $m_. \in \{m_{opt}, m_{pes}\}$. Since $\psi_S, \psi_W \in (\frac{1}{2}, 1]$, $(2 - 2\psi_S) > 0$, and $(2 - 2\psi_W) > 0$, the sign of inflation surprise in the strong economy is negative, i.e., $\{\pi(\theta_S) - \pi^e(m_.)\} < 0$. By contrast, the sign of inflation surprise in the weak economy is positive, i.e., $\{\pi(\theta_W) - \pi^e(m_.)\} > 0$. Therefore, the cyclical unemployment will be lower than θ_W if the weak economy is realized, while it will rise higher than θ_S when the strong economy is realized. The inflation surprises contribute to a fall in cyclical unemployment from θ_W in the weak economy and to a rise in cyclical unemployment from θ_S in the strong economy. By doing so, inflation surprises reduce the variance of cyclical unemployment across the states of the economy. The underlying economic environment $v < \frac{\theta_W - \theta_S}{\gamma}$ helps that the

Table 6. Optimal Overly Pessimistic Forward-Guidance Plan if $\theta_W - \theta_S - \gamma v < 0$ (relatively unemployment fighting) and $\psi_0 \in [\frac{1}{2}, 1]$

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	$\frac{1}{2\psi_0}$	$(1 - \frac{1}{2\psi_0})$
s_W	0	1

overall ex ante social welfare is to be maximized with such amounts of inflation surprises under the uninformative forward guidance.[7]

Optimal Overly Pessimistic Forward Guidance. For the case where monetary policy regime is relatively unemployment fighting ($\theta_W - \theta_S - \gamma v < 0$) and the economy is considered robust so that it is more likely to have the strong economy ($\psi_0 \in \left[\frac{1}{2}, 1\right]$), the optimal forward guidance is as follows: when the central bank observes the signal s_W, it should send the pessimistic message m_{pes} with probability 1. When it observes the signal s_S, it should send the optimistic message and the pessimistic message with mixed probabilities such that $Pr(m_{opt} \mid s_S) = \frac{1}{2\psi_0}$ and $Pr(m_{pes} \mid s_S) = \left(1 - \frac{1}{2\psi_0}\right)$. This forward guidance is pre-committed, so the private sector knows the function, albeit the private sector does not know the realized signal. It can be summarized as in Table 6.

Effects of Optimal Overly Pessimistic Forward Guidance. When the central bank pre-commits to implement the above forward guidance, i.e., $\rho_{opt} = \frac{1}{2\psi_0}$, $\rho_{pes} = 1$, the private sector forms the expected inflation as follows after receiving optimistic or pessimistic message, respectively:

[7]The losses from the inflation gap $-\alpha \cdot v^2$ are not affected by the forward-guidance policy under the assumptions that the inflation-targeting bandwidth v is exogenously given and that the central bank has complete control over the inflation, without the influence of the private sector's expected inflation. The complete control assumption can be relaxed so that the ex post realized inflation is a convex combination of the policy implementation of the central bank and the expected inflation of the private sector ("incomplete control"). The results for the forward-guidance policy remain the same, while the welfare benefits from reduction in cyclical unemployment gain more weights in the welfare function. See Appendix Section C for more detail.

$$\pi^e(\sigma(\cdot), m_{opt}) = \pi^T + v \cdot (1 - 2\psi_S)$$
$$\pi^e(\sigma(\cdot), m_{pes}) = \pi^T + v \cdot [(1 - 2\psi_S)(2\psi_0 - 1) + (2\psi_W - 1)(2 - 2\psi_0)].$$

Since the expected inflation is formed before the state of the economy is realized, it only depends on the received message. Observe that the private sector will have a truthful signal when the optimistic message m_{opt} is realized. It is because when the central bank observes the signal s_W, it never sends the optimistic message. Therefore, the optimistic message means that the central bank observes the signal s_S. By contrast, the pessimistic message m_{pes} generates obfuscation: under the overly pessimistic forward guidance, the pessimistic message can be sent when the central bank observes either the signal s_S or s_W (with different probabilities). If the pessimistic message is received, the private sector will set the expected inflation considering that both cases are possible. That helps the expected inflation become lower than the truthful communication.

Given any message, either the strong economy or the weak economy can be realized as the state of the economy. Hence, there are four scenarios for inflation surprises as follows:

$$\pi(\theta_S) - \pi^e(m_{opt}) = -v \cdot (2 - 2\psi_S) < 0$$
$$\pi(\theta_S) - \pi^e(m_{pes}) = -v \cdot [(2 - 2\psi_S)(2\psi_0 - 1) + 2\psi_W(2 - 2\psi_0)] < 0$$
$$\pi(\theta_W) - \pi^e(m_{opt}) = v \cdot 2\psi_S > 0$$
$$\pi(\theta_W) - \pi^e(m_{pes}) = v \cdot [2\psi_S(2\psi_0 - 1) + (2 - 2\psi_W)(2 - 2\psi_0)] > 0.$$

The direction of influence on cyclical unemployment is the same as the previous case of uninformative forward-guidance policy: Since $\psi_S, \psi_W \in (\frac{1}{2}, 1]$ and $\psi_0 \in \left(\frac{1}{2}, 1\right]$, the sign of inflation surprises is negative when the strong economy is realized, i.e., $\{\pi(\theta_S) - \pi^e(m_.)\} < 0$, while the sign of inflation surprise is positive in the weak economy, i.e., $\{\pi(\theta_W) - \pi^e(m_.)\} > 0$. Therefore, the cyclical unemployment will be lower than θ_W if the weak economy is realized, while it will rise higher than θ_S when the strong economy is realized.

Recall that in the previous case, the exogenously given bandwidth $v < \frac{\theta_W - \theta_S}{\gamma}$ while the overly pessimistic forward guidance is applied to the environment $v > \frac{\theta_W - \theta_S}{\gamma}$. Therefore, it is unfair to compare the inflation surprises between the previous uninformative

forward guidance and the overly pessimistic forward guidance. However, it can give us insight on the underlying mechanism to compare the coefficients of inflation surprises, except for the bandwidth v.

The average surprise in inflation of the weak economy is given by

$$v\psi_S(1-\psi_S) + v\left[\psi_S(2\psi_0-1) + (1-\psi_W)(2-2\psi_0)\right] \\ \left[(1-\psi_S)(2\psi_0-1) + \psi_W(2-2\psi_0)\right].$$

Then, the coefficient of v can be written as $\big[\psi_S(1-\psi_S) + \mathbb{A}_w (1-\mathbb{A}_w)\big]$, where $\mathbb{A}_w \equiv \psi_S(2\psi_0-1) + (1-\psi_W)(2-2\psi_0)$. Recall that under the uninformative forward guidance, the corresponding coefficient of v is $[2\psi_0\psi_S + 2(1-\psi_0)(1-\psi_W)]$, which is $\psi_S + \mathbb{A}_w$. To compare these two coefficients, as $(1-\psi_S) < 1$ and $(1-\mathbb{A}_w) < 1$,

$$\psi_S(1-\psi_S) + \mathbb{A}_w(1-\mathbb{A}_w) < \psi_S + \mathbb{A}_w.$$

Therefore, when the weak economy is realized, the coefficient of inflation surprise in uninformative forward guidance is greater than the one in overly pessimistic forward guidance. If the same inflation-targeting bandwidth v is applied, the reduction in cyclical unemployment in the weak economy will be greater under the uninformative forward-guidance policy. Then, why does the overly pessimistic forward guidance become optimal in certain economic environments? It is because such forward guidance reduces the variance of cyclical unemployment more effectively in those economic environments. Although there is a rise in cyclical unemployment in the strong economy, such a cost is less than the rise under the uninformative forward guidance. To compare the costs induced by the overly pessimistic forward guidance and the uninformative forward guidance more concretely, recall that the coefficient part of inflation surprise under uninformative forward guidance is given by $-[2\psi_0(1-\psi_S) + 2(1-\psi_0)\psi_W]$, which is $-(1-\psi_S) - (1-\mathbb{A}_w)$. The average surprise in inflation of the strong economy under overly pessimistic forward guidance is given by

$$-v\psi_S(1-\psi_S) - v\big[(1-\psi_S)(2\psi_0-1) + \psi_W(2-2\psi_0)\big] \\ \big[\psi_S(2\psi_0-1) + (1-\psi_W)(2-2\psi_0)\big].$$

Table 7. Truthful Forward-Guidance Policy Function if $\theta_W - \theta_S - \gamma v < 0$ (relatively unemployment fighting) and $\psi_0 \in [0, \frac{1}{2}]$

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	1	0
s_W	0	1

The coefficient of v can be written by $-[\psi_S(1-\psi_S) + \mathbb{A}_w(1-\mathbb{A}_w)]$. Then, the cost is less under the overly pessimistic forward guidance than the one under the uninformative forward guidance because

$$\psi_S(1-\psi_S) + \mathbb{A}_w(1-\mathbb{A}_w) < (1-\psi_S) + (1-\mathbb{A}_w).$$

Therefore, the overly pessimistic forward guidance can be advantageous in certain economic environments with relatively larger v, i.e., $v > \frac{\theta_W - \theta_S}{\gamma}$, while the uninformative forward guidance can be beneficial in the other economic environments, i.e., $v < \frac{\theta_W - \theta_S}{\gamma}$.

Intuitively, inflation surprises under overly pessimistic forward guidance contribute to reducing cyclical unemployment less than θ_W in the weak economy while it causes a rise in cyclical unemployment from θ_S in the strong economy: such inflation surprises reduce the variance of cyclical unemployment.

Optimal Truthful Forward Guidance. For the case where monetary policy regime is relatively unemployment fighting ($\theta_W - \theta_S - \gamma v < 0$) and the economy is considered strained so that it is more likely to have the weak economy ($\psi_0 \in \left[0, \frac{1}{2}\right]$), the optimization problem has a corner solution. The case requires direct comparison of ex ante expected payoffs among policy functions. The resulting optimal forward-guidance policy function prescribes the central bank to be truthful. Then, the private sector can infer exactly which signal the central bank observes. It can be summarized as in Table 7.

Effects of Truthful Forward Guidance. As the central bank pre-commits to implement the above forward guidance, plugging in

$\rho_{opt} = 1$, $\rho_{pes} = 1$, the private sector forms the expected inflation as follows:

$$\pi^e(\sigma(\cdot), m_{opt}) = \pi^T + v \cdot (1 - 2\psi_S)$$
$$\pi^e(\sigma(\cdot), m_{pes}) = \pi^T + v \cdot (2\psi_W - 1).$$

There are four scenarios for inflation surprises given by the following:

$$\pi(\theta_S) - \pi^e(m_{opt}) = -v \cdot (2 - 2\psi_S) < 0$$
$$\pi(\theta_S) - \pi^e(m_{pes}) = -v \cdot 2\psi_W < 0$$
$$\pi(\theta_W) - \pi^e(m_{opt}) = v \cdot 2\psi_S > 0$$
$$\pi(\theta_W) - \pi^e(m_{pes}) = v \cdot (2 - 2\psi_W) > 0.$$

The direction of influence on cyclical unemployment is the same to the previous cases of the uninformative forward guidance and the overly pessimistic forward guidance: When the weak economy is realized, the inflation surprises lower the cyclical unemployment than θ_W, whereas the inflation surprises raise cyclical unemployment higher than θ_S. This helps lower variance in cyclical unemployment across the states of the economy.

The inflation surprises under the optimistic messages are the same in both cases under the overly pessimistic forward guidance and the truthful forward guidance. However, the inflation surprises under the pessimistic message are different: given $\psi_0 \in [0, \frac{1}{2})$, the effects of inflation surprises under the truthful forward guidance are greater than the effects under the overly pessimistic forward guidance, with a smaller cost and a greater benefit: $2\psi_W < [(2 - 2\psi_S)(2\psi_0 - 1) + 2\psi_W(2 - 2\psi_0)]$ and $(2 - 2\psi_W) > [2\psi_S(2\psi_0 - 1) + (2 - 2\psi_W)(2 - 2\psi_0)]$ where $\psi_0 \in [0, \frac{1}{2}]$. By contrast, if $\psi_0 \in (\frac{1}{2}, 1]$, the inequality reverses, and the effects of the overly pessimistic forward guidance becomes greater, with a smaller cost and a greater benefit.

The expected benefit and expected cost of the overly pessimistic policy function can be computed with respect to the truthful

policy function. First, the expected cost incurs when the signal s_S forecasting the future strong economy is realized:

$$\begin{aligned}
&(\text{after } s_S)\\
&EW_{\text{overly pes}}(s_S) - EW_{\text{truthful}}(s_S) \equiv C\\
&\quad = -\left(1 - \frac{1}{2\cdot\psi_0}\right)\cdot(2 - 2\cdot\psi_0)\cdot(2\psi_W + 2\psi_W - 2)\cdot\gamma\cdot v\\
&\qquad \times\Big\{2\psi_S\cdot(2\cdot\gamma\cdot v + \theta_S - \theta_W) + (2\theta_W - 4\cdot\gamma\cdot v)\\
&\qquad - 2\cdot\psi_0\cdot(2\psi_W + 2\psi_S - 2)\Big\} < 0. \qquad (16)
\end{aligned}$$

The expected benefit is earned when the signal s_W forecasting the future weak economy is realized:

$$\begin{aligned}
&(\text{after } s_W)\\
&EW_{\text{overly pes}}(s_W) - EW_{\text{truthful}}(s_W) \equiv B\\
&\quad = \Big[2\cdot\theta_S + \{(2\cdot\psi_0 - 1)\cdot(2 - 2\cdot\psi_S) + (3 - 2\cdot\psi_0)\cdot(2\cdot\psi_W)\}\\
&\qquad \cdot\gamma\cdot v + 2\psi_W\cdot(2\cdot\gamma\cdot v + \theta_S - \theta_W)\Big]\\
&\qquad \times(2\cdot\psi_0 - 1)\cdot(2\cdot\psi_W + 2\cdot\psi_S - 2)\cdot\gamma\cdot v > 0. \qquad (17)
\end{aligned}$$

As long as the monetary policy regime is relatively unemployment fighting ($\theta_W - \theta_S < \gamma\cdot v$) and it is more likely to have the signal forecasting the strong economy a priori ($\psi_0 \in [\frac{1}{2}, 1]$), the overly pessimistic communication's expected benefit outweighs the expected cost multiplied by the probability of each signal being realized: $B \times (1 - \psi_0) + C \times \psi_0 \geq 0$. In such economic environment, the overly pessimistic communication function is better than the truthful communication function.

4. Discussion

In this paper, the monetary policy function of the central bank is given exogenously, and it is clearly understood by the private sector.

Hence, the monetary policy can be said to be transparent,[8] and it is well known to the private sector how high or low the realized inflation rate will be as soon as the state of the economy is realized; for instance, if the weak economy θ_W is realized, then the inflation rate will be realized as $\pi^T + v$. Given the monetary policy transparency, however, at issue is the transparency in communication or information symmetry between the central bank and the private sector. In this model, before the state of the economy is realized, the central bank is assumed to have only a probabilistic signal. Then, should the central bank communicate the information truthfully? The result that the uninformative forward guidance is optimal for the relatively inflation-fighting central bank may appear counterintuitive considering the emphasis on transparency by the central banks around the world.

The benefits from opaqueness in communication as opposed to transparency have long been acknowledged by the literature on central bank communication (See Geraats 2002, p. F547 for benefits from economic uncertainty). To see the benefits from inflation surprises under the uninformative forward guidance in the main model,[9] a central bank with strict inflation targeting can be compared with the relatively inflation-fighting central bank with flexible inflation targeting. In strict inflation targeting, the ex post realized inflation is π^T regardless of the realized state of the economy, and thus, the expected inflation is π^T irrespective of forward-guidance policy. As a result, there are zero surprises in inflation. It means there is no buffer for the supply shock: unemployment will fully swing between θ_W and θ_S, whereas the inflation rate is always at π^T. This case

[8] According to the classification of Geraats (2002, p. F540), the environment satisfies political (objective) transparency, procedural transparency, policy transparency, and operational transparency.

[9] In this model, it is assumed that the central bank completely controls realized inflation, and there is no role for the expected inflation in ex post realized inflation. However, in reality, the expected inflation may be self-fulfilling, and central banks use communication extensively to gravitate the expected inflation near the inflation target. In Appendix Section C, the model is extended to incorporate such reality that the central bank may have only incomplete control over the realized inflation while the private sector's expected inflation may have self-fulfilling effect. In addition, a central bank is allowed to choose its optimal inflation-targeting bandwidth endogenously. The inflation-targeting bandwidth can be asymmetric.

is suboptimal. Rather, the positive inflation-targeting bandwidth and uninformative forward guidance that generate inflation surprises contribute welfare improvement through decreases in the variance of cyclical unemployment across the states of the economy.[10]

The next question raised will be why overly pessimistic communication is better than truthful communication in certain unemployment-fighting environments (where parameters satisfy $\theta_W - \theta_S - \gamma v < 0$ and $\psi_0 \in [\frac{1}{2}, 1]$). The economic intuition stems from the idea that it is more useful to have a lower variance in cyclical unemployment across the states of the economy in order to maximize social welfare. To achieve such a lower variance in cyclical unemployment, it is better to lower the expected inflation to induce a greater inflation surprise in the weak economy. Assume the environment is given as $\theta_W - \theta_S - \gamma v < 0$, where $\psi_0 \in [\frac{1}{2}, 1]$. Compared with truthful communication, the overly pessimistic communication lowers the expected inflation more effectively when the central bank received the signal forecasting the future weak economy. If the weak economy is actually realized following the signal, then the inflation surprise will be larger than the truthful communication. The larger inflation surprise contributes to lower cyclical unemployment more effectively. This benefit, however, comes at a cost. To have the larger inflation surprise in the weak economy, the central bank will endure a smaller inflation surprise in the strong economy than the truthful communication. Both the benefit and the cost stem from the strategic part of the overly pessimistic communication that the central bank may send the pessimistic message strategically under both signals forecasting the future strong economy or forecasting the future weak economy.

This paper shows optimal tone of central bank forward guidance when a central bank has been confined with an exogenous monetary policy and tries to use forward guidance about future uncertainty as an additional tool. I also show in the appendix that when the inflation bandwidth is endogenously determined, global optimal is the relatively inflation-fighting monetary policy and the uninformative forward guidance.[11] Considering Woodford (2005) and the current emphasis on central bank communication in practice, the result that

[10]Please see the Appendix Section B for more detail.
[11]See Appendix Sections B and C for more detail.

global optimal forward guidance being uninformative one when the monetary policy is endogenous determined jointly with forward guidance may seem puzzling. As future research, the role of central bank forward guidance on the monetary policy transmission mechanism (for instance, see Leombroni et al. 2021) and other factors that affect the private sector's formation of expected inflation (see Binder and Makridis 2022, Coibion, Gorodnichenko, and Weber 2022, and others) may require more attention to further investigate the global optimal pair of endogenous monetary policy and forward guidance.

Appendix

In this appendix, I show the expected inflation of the private sector under each forward guidance in Section A. I derive the optimal symmetric inflation-targeting bandwidth paired under each forward guidance (Section B.1–3). I present generalized formulation incorporating possibly self-fulfilling effect of the private sector's expected inflation and allowing for asymmetric inflation-targeting bandwidth (Section C.1–5). Finally, I do welfare comparison among the pairs of forward guidance and monetary policy (Section B.4 and C.6, respectively).[12] All the mathematical derivations are elaborated in the online math appendix (Ko 2022).

A. Expected Inflation under Each Communication Policy

In this section, I demonstrate how the expected-inflation is computed in each forward-guidance regime. For the economic intuition, please see Section 3 of the main text.

A.1 $\psi_0 \in [\frac{1}{2}, 1]$ *and* $(\theta_W - \theta_S - \gamma v) < 0$

The optimal forward-guidance policy plan in this environment is shown in Table A.1.

Plugging in the values shown in Table A.1, the expected inflations of the private sector are given as follows: after the private sector receives the optimistic message from the central bank,

[12] I thank the anonymous referee and the editor for the suggestion to obtain forward guidance and inflation-targeting bandwidth jointly and compare welfare levels.

Table A.1. Optimal Overly Pessimistic Forward-Guidance Plan if $(\theta_W - \theta_S - \gamma v) < 0$ (relatively unemployment fighting) and $\psi_0 \in [\frac{1}{2}, 1]$

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	$\frac{1}{2\psi_0}$	$(1 - \frac{1}{2\psi_0})$
s_W	0	1

$$
\begin{aligned}
&\pi^e(\sigma(\cdot), m_{opt}) \\
&\quad = \pi^T + v \cdot \left[\frac{(1-2\psi_S)\psi_0\rho_{opt} + (2\psi_W - 1)(1-\psi_0)(1-\rho_{pes})}{(1-\psi_0)(1-\rho_{pes}) + \psi_0\rho_{opt}} \right] \\
&\qquad \text{plugging in } \rho_{opt} = \frac{1}{2\psi_0},\ 1 - \rho_{pes} = 0 \qquad \text{(A.1)} \\
&\quad = \pi^T + v \cdot (1 - 2\psi_S).
\end{aligned}
$$

That is, the private sector recognizes that m_{opt} is only sent when the central bank observes s_S with probability 1. By contrast, after the private sector receives the pessimistic message from the central bank, the private sector is not sure if the central bank observes s_S or s_W.

$$
\begin{aligned}
&\pi^e(\sigma(\cdot), m_{pes}) \\
&\quad = \pi^T + v \cdot \left[\frac{(1-2\psi_S)\psi_0(1-\rho_{opt}) + (2\psi_W - 1)(1-\psi_0)\rho_{pes}}{(1-\psi_0)\rho_{pes} + \psi_0(1-\rho_{opt})} \right] \\
&\qquad \text{plugging in } 1 - \rho_{opt} = \left(1 - \frac{1}{2\psi_0}\right),\ \rho_{pes} = 1 \qquad \text{(A.2)} \\
&\quad = \pi^T + v \cdot \{(1-2\psi_S)(2\psi_0 - 1) + (2\psi_W - 1)(2 - 2\psi_0)\}
\end{aligned}
$$

A.2 $\psi_0 \in [0, 1/2)$ *and* $(\theta_W - \theta_S - \gamma v) < 0$

The optimal forward-guidance policy plan in this environment is shown in Table A.2.

Under the truthful forward-guidance policy, after the private sector receives the optimistic message m_{opt}, the private sector is

Table A.2. Truthful Forward-Guidance Policy Function if $\theta_W - \theta_S - \gamma v < 0$ (relatively unemployment fighting) and $\psi_0 \in [0, \frac{1}{2}]$

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	1	0
s_W	0	1

assured that the central bank observes the signals for the future strong economy s_S. The expected inflation is updated as follows:

$$\begin{aligned}
&\pi^e(\sigma(\cdot), m_{opt}) \\
&\quad = \pi^T + v \cdot \left[\frac{(1-2\psi_S)\psi_0\rho_{opt} + (2\psi_W - 1)(1-\psi_0)(1-\rho_{pes})}{(1-\psi_0)(1-\rho_{pes}) + \psi_0\rho_{opt}} \right] \\
&\qquad \text{plugging in } \rho_{opt} = 1,\ 1-\rho_{pes} = 0 \qquad \text{(A.3)} \\
&\quad = \pi^T + v \cdot (1 - 2\psi_S).
\end{aligned}$$

Notice that this is also the case under the above optimal overly pessimistic forward-guidance policy: receiving m_{opt} means the central bank observing s_S for sure. Therefore, the expected inflation $\pi^e(\sigma(\cdot), m_{opt})$ becomes equivalent.

When the private sector receives the pessimistic message m_{pes} from the central bank, the private sector is assured that the central bank observes the signals for the future weak economy s_W. The expected inflation is given by the following:

$$\begin{aligned}
&\pi^e(\sigma(\cdot), m_{pes}) \\
&\quad = \pi^T + v \cdot \left[\frac{(1-2\psi_S)\psi_0(1-\rho_{opt}) + (2\psi_W - 1)(1-\psi_0)\rho_{pes}}{(1-\psi_0)\rho_{pes} + \psi_0(1-\rho_{opt})} \right] \\
&\qquad \text{plugging in } 1-\rho_{opt} = 0,\ \rho_{pes} = 1 \qquad \text{(A.4)} \\
&\quad = \pi^T + v \cdot (2\psi_W - 1).
\end{aligned}$$

A.3 $(\theta_W - \theta_S - \gamma v) > 0$

In the case of uninformative communication policy, the central bank announcement does not induce any updates in the expected inflation.

Table A.3. Optimal Uninformative Forward-Guidance Policy Function if $\theta_W - \theta_S - \gamma v > 0$ (relatively inflation fighting)

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	1/2	1/2
s_W	1/2	1/2

Hence, the expected inflation remains the same as under the prior irrespective of the messages.

$$\begin{aligned}\pi^e(\sigma(\cdot), m_{opt}) &= \pi^e(\sigma(\cdot), m_{pes}) \\ &= \pi^T + v \cdot [(1-2\psi_S)\psi_0 + (2\psi_W - 1)(1-\psi_0)]\end{aligned}$$

B. Optimal Symmetric Inflation-Targeting Bandwidth v^*

In this section, the optimal pairs of forward guidance and inflation-targeting bandwidth are presented. For each optimal communication policy, a paired optimal inflation-targeting bandwidth can be computed.

B.1 $\psi_0 \in [1/2, 1]$ *and* $(\theta_W - \theta_S - \gamma v) < 0$

In this environment, by plugging in the optimal overly pessimistic forward-guidance policy, the ex ante social welfare is given as follows:

$$\begin{aligned}\max_v \quad & \mathbb{E}[SW] \\ &= \left\{-\left(\theta_S + \gamma \cdot v(2-2\psi_S)\right)^2 - \alpha v^2\right\} \cdot \frac{\psi_S}{2} \\ &\quad + \left\{-\Big(\theta_S + \gamma \cdot v(2-2\psi_S) + 4 \cdot v\gamma \cdot (1-\psi_0)\right. \\ &\quad \left. \cdot (\psi_S + \psi_W - 1)\Big)^2 - \alpha v^2\right\} \times \left(\phi - \frac{\psi_S}{2}\right) \\ &\quad + \left\{-\left(\theta_W + \gamma \cdot v(-2\psi_S)\right)^2 - \alpha v^2\right\} \cdot \frac{1-\psi_S}{2} \\ &\quad + \left\{-\Big(\theta_W + \gamma \cdot v(-2\psi_S) + 4 \cdot v\gamma \cdot (1-\psi_0)\right. \\ &\quad \left. \cdot (\psi_S + \psi_W - 1)\Big)^2 - \alpha v^2\right\} \left(1 - \phi - \frac{1-\psi_S}{2}\right), \qquad \text{(B.1)}\end{aligned}$$

where $\phi = \psi_0\psi_S + (1-\psi_0)(1-\psi_W)$ and $(1-\phi_0) = \psi_0(1-\psi_S) + (1-\psi_0)\psi_W$.

Then, by taking derivatives, we have the welfare-maximizing inflation-targeting bandwidth v:

$$v = \frac{\gamma(\theta_W - \theta_S)\left[\psi_S(1-\psi_S) + (2\phi - \psi_S)\left\{2(1-\phi) - (1-\psi_S)\right\}\right]}{2\cdot\gamma^2\left[\psi_S(1-\psi_S) + (2\phi-\psi_S)\left\{2(1-\phi)-(1-\psi_S)\right\}\right] + \alpha}. \tag{B.2}$$

However, this value of v will not satisfy the condition $(\theta_W - \theta_S - \gamma v) < 0$. Therefore, the infimum of this bound $v_{threshold} = \frac{\theta_W - \theta_S}{\gamma}$ is to be paired as optimal inflation-targeting bandwidth with the optimally overly pessimistic forward-guidance policy.

B.2 $\psi_0 \in [0, \frac{1}{2})$ *and* $(\theta_W - \theta_S - \gamma v) < 0$

In this environment, by plugging in the optimal truthful forward-guidance policy, the ex ante social welfare is given as follows:

$$\begin{aligned}
\max_v \quad & \mathbb{E}[SW] \\
&= -\left[(\theta_S + \gamma\cdot v\cdot(2-2\psi_S))^2 + \alpha v^2\right]\psi_0\psi_S \\
&\quad - \left[\left(\theta_S + \gamma\cdot v\cdot(2\psi_W)\right)^2 + \alpha v^2\right](1-\psi_0)(1-\psi_W) \\
&\quad - \left[(\theta_W + \gamma\cdot v\cdot(-2\psi_S))^2 + \alpha v^2\right]\psi_0(1-\psi_S) \\
&\quad - \left[(\theta_W + \gamma\cdot v\cdot(2\psi_W - 2))^2 + \alpha v^2\right](1-\psi_0)\psi_W.
\end{aligned} \tag{B.3}$$

Then, from the optimization condition, we have

$$v = \frac{2\gamma\cdot(\theta_W - \theta_S)\left(\psi_0(1-\psi_S)\psi_S + (1-\psi_0)\psi_W(1-\psi_W)\right)}{4\gamma^2\cdot\left(\psi_0(1-\psi_S)\psi_S + (1-\psi_0)\psi_W(1-\psi_W)\right) + \alpha}. \tag{B.4}$$

This value of v also fails to satisfy the threshold condition $(\theta_W - \theta_S - \gamma v) < 0$. Therefore, the infimum of this bound $v_{threshold} = \frac{\theta_W - \theta_S}{\gamma}$ is to be paired as optimal inflation-targeting bandwidth with the optimal truthful forward-guidance policy.

B.3 $(\theta_W - \theta_S - \gamma v) > 0$

In this inflation-fighting environment, by plugging in the optimal uninformative forward-guidance policy, the ex ante social welfare is given as follows:

$$\begin{aligned}\max_{v} \quad & \mathbb{E}[SW] \\ &= -\left[(\theta_S + \gamma \cdot v \cdot (2\psi_W - Z\psi_0))^2 + \alpha v^2\right]\phi \\ &\quad - \left[(\theta_W + \gamma \cdot v \cdot (2\psi_W - Z\psi_0 - 2))^2 + \alpha v^2\right](1-\phi), \qquad \text{(B.5)}\end{aligned}$$

where $\phi = \big(\psi_0\psi_S + (1-\psi_0)(1-\psi_W)\big)$, $(1-\phi) = \big(\psi_0(1-\psi_S) + (1-\psi_0)\psi_W\big)$. Then, the optimal v^* is given as follows:

$$v^* = \frac{2\gamma\,(\theta_W - \theta_S)\,\phi(1-\phi)}{4 \cdot \gamma^2\phi(1-\phi) + \alpha}. \qquad \text{(B.6)}$$

Notice that as the weight on quadratic loss from inflation gap, α, increases, the optimal inflation-targeting bandwidth v^* will decrease. Also, optimal v^* is not zero for the relatively inflation-fighting central bank.

To understand why strict inflation targeting is not optimal, consider the scenario where v is set zero. Then ex post realized inflation is always π^T, irrespective of the state of the economy in this model. That means expected inflation of the private sector will be fixed at π^T for any communication policy. As a result, the economy will experience zero losses from inflation gap, whereas the economy shall take the welfare losses from cyclical unemployment shocks θ_W or θ_S, without any mitigation. By contrast, if inflation bandwidth v^* is set as positive, the central bank will be able to generate surprises in inflation which can mitigate the impact from the cyclical unemployment shocks.

B.4 Global Optimal among Pairs of Forward Guidance and Symmetric Inflation-Targeting Bandwidth

I have three local optimal pairs of forward-guidance policy and inflation-targeting bandwidth in different economic environments as in Table B.1: first, overly pessimistic forward-guidance policy

Table B.1. Optimal Pairs of Forward Guidance and Symmetric Inflation-Targeting Bandwidth (no self-fulfilling effect of private sector expected inflation)

	$\psi_0 \geq 1/2$	$\psi_0 < 1/2$	$\forall \psi_0$
Forward Guidance	Overly Pessimistic	Truthful	Uninformative
Inflation Bandwidth	$v_{threshold} = \frac{\theta_W - \theta_S}{\gamma}$		$v^* = \frac{2\gamma(\theta_W - \theta_S)\varphi(1-\varphi)}{4\cdot\gamma^2\varphi(1-\varphi)+\alpha}$
Welfare Comparison	Suboptimal	Suboptimal	Global Optimal Pair

and inflation-targeting bandwidth of $v_{threshold} = \frac{\theta_W - \theta_S}{\gamma}$; second, truthful forward-guidance policy and inflation-targeting bandwidth $v_{threshold} = \frac{\theta_W - \theta_S}{\gamma}$; third, uninformative forward guidance and optimal inflation-targeting bandwidth

$$v^* = \frac{(\theta_W - \theta_S)}{2\gamma + \frac{\alpha}{2\gamma(1-\phi)\phi}}, \tag{B.7}$$

where $\phi(1-\phi) = \{\psi_0\psi_S + (1-\psi_0)(1-\psi_W)\} \cdot \{\psi_0(1-\psi_S) + (1-\psi_0)\psi_W\}$.

Then, which one would give the highest ex ante expected social welfare? To examine the levels of ex ante expected social welfare among these threefold results, first the overly pessimistic forward-guidance policy and $\gamma \cdot v_{threshold} = \theta_W - \theta_S$ are plugged in:

$$\begin{aligned}
&\mathbb{E}[SW_{pess}] \\
&\quad = -\left(\theta_S + \gamma \cdot v_{threshold} 4p_3\right)^2 \cdot p_1 - \left[\theta_S + \gamma \cdot v_{threshold} 4p_4\right]^2 \cdot p_2 \\
&\qquad - \left(\theta_W + \gamma \cdot v_{threshold}(-4p_1)\right)^2 \cdot p_3 \\
&\qquad - \left[\theta_W + \gamma \cdot v_{threshold}(-4p_2)\right]^2 \cdot p_4 - \alpha(v_{threshold})^2,
\end{aligned} \tag{B.8}$$

where $p_1 = \frac{\psi_S}{2}$, $p_2 = \left(\phi - \frac{\psi_S}{2}\right)$, $p_3 = \frac{1-\psi_S}{2}$, $p_4 = \left(1 - \phi - \frac{1-\psi_S}{2}\right)$. Observe that $p_1 + p_2 = \phi$ and $p_3 + p_4 = 1 - \phi$. Also, $p_1 + p_3 = \frac{1}{2} = p_2 + p_4$.

Second, the truthful forward-guidance policy and $\gamma \cdot v_{threshold} = \theta_W - \theta_S$ are plugged in:

$$\begin{aligned}
&\mathbb{E}[SW_{truth}] \\
&= -\left(\theta_S + \gamma \cdot v_{threshold} \cdot (2 - 2\psi_S)\right)^2 \psi_0 \psi_S \\
&\quad - \left(\theta_S + \gamma \cdot v_{threshold} \cdot (2\psi_W)\right)^2 (1 - \psi_0)(1 - \psi_W) \\
&\quad - \left(\theta_W + \gamma \cdot v_{threshold} \cdot (-2\psi_S)\right)^2 \psi_0 (1 - \psi_S) \\
&\quad - \left(\theta_W + \gamma \cdot v_{threshold} \cdot (2\psi_W - 2)\right)^2 (1 - \psi_0)\psi_W - \alpha (v_{threshold})^2.
\end{aligned} \tag{B.9}$$

Comparing $\mathbb{E}[SW_{pess}]$ and $\mathbb{E}[SW_{truth}]$, the inequality depends on the value of ψ_0. If $\psi_0 \geq 1/2$, then $\mathbb{E}[SW_{pess}]$ gives greater ex ante expected social welfare; otherwise, $\mathbb{E}[SW_{truth}]$ is greater.

Third, the uninformative forward-guidance policy and paired optimal inflation-targeting bandwidth $v^*_{uninfo} = \frac{(\theta_W - \theta_S)}{2\gamma + \frac{\alpha}{2\gamma(1-\phi)\phi}}$ are plugged in. Then, $\gamma \cdot v^*_{uninfo}$ part is replaced by $\left(\frac{\theta_W - \theta_S}{2} - \epsilon\right)$ while $\epsilon = \frac{\alpha(\theta_W - \theta_S)}{2\alpha + 8\gamma^2\phi(1-\phi)}$.

$$\begin{aligned}
&\mathbb{E}[SW_{uninfo}] \\
&\quad = -\left(\theta_S + \gamma \cdot v^*_{uninfo} \cdot 2(1 - \phi)\right)^2 \cdot \phi \\
&\qquad - \left(\theta_W + \gamma \cdot v^*_{uninfo} \cdot (-2)\phi\right)^2 \cdot (1 - \phi) - \alpha (v^*_{uninfo})^2
\end{aligned} \tag{B.10}$$

Assume that $\psi_0 > \frac{1}{2}$. To compare Equation (B.10) of $\mathbb{E}[SW_{uninfo}]$ with Equation (B.8) of $\mathbb{E}[SW_{pess}]$, first, a fixed inflation-targeting bandwidth v is plugged in, and next, the fact that (B.10) of $\mathbb{E}[SW_{uninfo}]$ is local optimal is utilized.

Fix $v = v_{threshold}$ and observe that $p_1 > p_2$ and $p_3 < p_4$.

$$\begin{aligned}
&\mathbb{E}[SW_{pess};\ v = v_{threshold}] - \mathbb{E}[SW_{uninfo};\ v = v_{threshold}] \\
&\quad = -(\theta_W - \theta_S - \gamma \cdot v_{threshold}) 4\gamma \cdot v_{threshold} (p_1 - p_2)(p_4 - p_3) \\
&\qquad + (p_4 - p_3 - p_1 + p_2) 4\gamma^2 (v_{threshold})^2 \{2p_1p_3 - 2p_2p_4\}
\end{aligned} \tag{B.11}$$

Since $(p_4 - p_3 - p_1 + p_2) = 0$, it is clear that $\mathbb{E}[SW_{pess};\ v = v_{threshold}] = \mathbb{E}[SW_{uninfo};\ v = v_{threshold}]$ by plugging in

$\gamma \cdot v_{threshold} = (\theta_W - \theta_S)$. Then, as local optimal pair for the uninformative forward-guidance policy is given by $v = v^*_{uninfo}$, it is obvious that $\mathbb{E}[SW_{uninfo};\ v = v^*_{uninfo}] > \mathbb{E}[SW_{uninfo};\ v = v_{threshold}]$. Therefore, $\mathbb{E}[SW_{pess};\ v = v_{threshold}] < \mathbb{E}[SW_{uninfo};\ v = v^*_{uninfo}]$. In conclusion, if $\psi_0 \geq 1/2$, the global optimal pair of forward-guidance policy and symmetric inflation-targeting policy will be the uninformative forward-guidance and inflation bandwidth v^*_{uninfo}.

Now assume that $\psi_0 < \frac{1}{2}$. To compare Equation (B.10) of $\mathbb{E}[SW_{uninfo}]$ with Equation (B.9) of $\mathbb{E}[SW_{truth}]$, first, a fixed inflation-targeting bandwidth v is plugged in, and next, the fact that (B.10) of $\mathbb{E}[SW_{uninfo}]$ is local optimal is utilized.

Fix $v = v_{threshold}$ and observe that $\psi_S - \phi = \psi - \psi_0\psi_S - (1-\psi_0)(1-\psi_W) = (1-\psi_0)(\psi_S + \psi_W - 1)$ and $1 - \phi - \psi_W = \psi_0(1-\psi_S) + (1-\psi_0)\psi_W - \psi_W = -\psi_0(\psi_S + \psi_W - 1)$.

$$
\begin{aligned}
&\mathbb{E}[SW_{truth}] - \mathbb{E}[SW_{uninfo}] \\
&\quad = -\left(\theta_S + \gamma \cdot v_{threshold} \cdot (2 - 2\psi_S)\right)^2 \psi_0\psi_S \\
&\qquad + \left(\theta_S + \gamma \cdot v_{threshold} \cdot 2(1-\phi)\right)^2 \psi_0\psi_S \\
&\qquad - \left(\theta_S + \gamma \cdot v_{threshold} \cdot (2\psi_W)\right)^2 (1-\psi_0)(1-\psi_W) \\
&\qquad + \left(\theta_S + \gamma \cdot v_{threshold} \cdot 2(1-\phi)\right)^2 (1-\psi_0)(1-\psi_W) \\
&\qquad - \left(\theta_W + \gamma \cdot v_{threshold} \cdot (-2\psi_S)\right)^2 \psi_0(1-\psi_S) \\
&\qquad + \left(\theta_W + \gamma \cdot v_{threshold} \cdot (-2)\phi\right)^2 \psi_0(1-\psi_S) \\
&\qquad - \left(\theta_W + \gamma \cdot v_{threshold} \cdot (2\psi_W - 2)\right)^2 (1-\psi_0)\psi_W \\
&\qquad + \left(\theta_W + \gamma \cdot v_{threshold} \cdot (-2)\phi\right)^2 (1-\psi_0)\psi_W \\
&\qquad - \alpha(v_{threshold})^2 + \alpha(v_{threshold})^2 \\
&\quad = -4\gamma \cdot v_{threshold} \cdot \psi_0(1-\psi_0)(\psi_S + \psi_W - 1)^2 \\
&\qquad (\theta_W - \theta_S - \gamma \cdot v_{threshold}).
\end{aligned} \tag{B.12}
$$

It is clear that $\mathbb{E}[SW_{truth};\ v = v_{threshold}] = \mathbb{E}[SW_{uninfo};\ v = v_{threshold}] < \mathbb{E}[SW_{uninfo};\ v = v_{threshold}]$. Hence, if $\psi_0 < 1/2$, the global optimal pair of forward guidance and inflation-targeting policy will be the uninformative forward guidance with the bandwidth v^*_{uninfo}.

The welfare comparison results imply that if a central bank is allowed to choose a pair of forward-guidance policy and a symmetric inflation-targeting policy, in order to maximize the ex ante expected social welfare, the central bank should follow the uninformative forward guidance and inflation targeting with a relatively narrow inflation bandwidth for all ψ_0. Additionally, notice that $v^*_{uninfo} > 0$: the strict inflation targeting is not optimal.

C. Generalization: Asymmetric Inflation-Targeting Bandwidth under Incomplete Control of Central Bank over Inflation

C.1 Incomplete Control of Central Bank over Inflation

In the main model, the central bank has complete control over the realized inflation. The monetary policy function $\pi : \Theta \to \mathbb{R}$ is given as the actually realized inflation function. That is, when the state of the economy is θ_S, the central bank implements monetary policy toward $\pi(\theta_S) = \pi^T - v$, and the realized inflation is actually $\pi^T - v$. Similarly, when the state of the economy is θ_W, the central bank implements its monetary policy toward $\pi(\theta_W) = \pi^T + v$, and the realized inflation is as the central bank desires. In this sense, inflation is under complete control of the central bank in the main model, and the expected inflation of the private sector does not affect the realized inflation at all. Therefore, inflation gap between realized inflation and the target π^T is either $\pm v$.

Although the assumption that the central bank completely control the ex post realized inflation is not uncommon in the literature (see Geraats 2002, p. F537), a more realistic assumptions would be that the central bank has incomplete control over the realized inflation. The expected inflation of the private sector can affect the realized inflation, such as

$$\pi^{ex.post}(\theta_., m_.) = \xi \cdot \pi^{policy}(\theta_.) + (1 - \xi) \cdot \pi^e(\sigma(\cdot), m_.),$$

where ξ is the degree of control of the central bank over the ex post realized inflation. If $\xi = 1$, the central bank has complete control over inflation, which is the main model. As ξ declines, the private sector's expected inflation begins to affect the realized inflation, and if ξ approaches zero, it is so called self-fulfilling expectation. Observe

that expected inflation is set as the private sector forecasts the future inflation as follows:

$$\begin{aligned}\pi^e(\sigma(\cdot), m_.) &= \widehat{E}\left[\pi^{ex.post}(\theta_., m_.) \mid m_.\right] \\ &= \xi \cdot \widehat{E}\left[\pi^{policy}(\theta_.) \mid m_.\right] + (1-\xi) \cdot \pi^e(\sigma(\cdot), m_.).\end{aligned}$$

Hence, expected inflation is the solution of the above fixed-point problem, and the expected inflation remains the same to the formulation under the case where the central bank has complete control over realized inflation.

$$\pi^e(\sigma(\cdot), m_.) = \widehat{E}\left[\pi^{policy}(\theta_.) \mid m_.\right]$$

As a result, the threefold results of the optimal forward-guidance policy under exogenously given inflation targeting do not change. However, the optimal inflation-targeting bandwidth paired with each optimal forward guidance needs modification because the losses from inflation gap contributes less to the ex ante expected social welfare as much as $\alpha\xi$. In the model of incomplete control of the central bank over inflation, the social welfare function is defined by

$$\begin{aligned}&W(\pi^e(\sigma(\cdot), m_.),\ \pi^{ex.post}(\theta, m), \theta) \\ &\quad = -(u - u^{NAR})^2 - \alpha \cdot (\pi^{ex.post}(\theta_., m_.) - \pi^T)^2 \\ &\quad = -[\theta_. + \gamma \cdot (\pi^e(\sigma(\cdot), m_.) - \pi^{ex.post}(\theta_., m_.))]^2 \\ &\qquad - \alpha \cdot (\pi^{ex.post}(\theta_., m_.) - \pi^T)^2 \\ &\quad = -[\theta_. + \gamma \cdot \xi \cdot (\pi^e(\sigma(\cdot), m_.) - \pi^{policy}(\theta_., m_.))]^2 \\ &\qquad - \alpha \cdot \left[(1-\xi)\left\{\pi^e(\sigma(\cdot), m_.) - \pi^T\right\} + \xi(\pi^{policy}(\theta_.) - \pi^T)\right]^2.\end{aligned} \tag{C.1}$$

Observe that inflation gap is measured by the difference between ex post realized inflation and the inflation target π^T. The quadratic loss from inflation gap is given by $-\alpha(\pi^{ex.post}(\theta_., m_.) - \pi^T)^2$, weighed by $\alpha > 0$.

C.2 Optimal Forward Guidance under Exogenous Asymmetric Inflation-Targeting Bandwidth

How would the results change if the inflation-targeting bandwidth is allowed to be asymmetric?[13] The monetary policy function is now set to $\pi^{policy}(\theta_S) = \pi^T - v_S$ and $\pi^{policy}(\theta_W) = \pi^T + v_W$. Assuming the possible incomplete control of the central bank over realized inflation, the private sector's expected inflation under each message is given as follows:

$$\begin{aligned}
&\pi^e(\sigma(\cdot), m_{opt}) \\
&\quad = \pi^T + \left[\frac{\begin{array}{c}\{(1-\psi_S)v_W - \psi_S v_S\}\psi_0\rho_{opt} \\ +\{(\psi_W - 1)v_S + \psi_W v_W\}(1-\psi_0)(1-\rho_{pes})\end{array}}{(1-\psi_0)(1-\rho_{pes}) + \psi_0\rho_{opt}} \right] \\
&\pi^e(\sigma(\cdot), m_{pes}) \\
&\quad = \pi^T + \left[\frac{\begin{array}{c}\{(1-\psi_S)v_W - \psi_S v_S\}\psi_0(1-\rho_{opt}) \\ +\{(\psi_W - 1)v_S + \psi_W v_W\}(1-\psi_0)\rho_{pes}\end{array}}{(1-\psi_0)\rho_{pes} + \psi_0(1-\rho_{opt})} \right]. \qquad \text{(C.2)}
\end{aligned}$$

Then the forward-guidance policy is derived from the following first-order conditions (FOCs) with respect to ρ_{opt} and d:

$$\begin{aligned}
&\frac{\partial \mathbb{E}_{s_\cdot, m_\cdot, \theta_\cdot}(SW)}{\partial \rho_{opt}} \\
&\quad = \left(\frac{(v_S + v_W)(\psi_S + \psi_W - 1)^2 \psi_0^2 (1-\psi_0)^2 d^2 (1-2X)}{X^2(1-X)^2} \right) \\
&\qquad \times \left[2\gamma\xi(\theta_W - \theta_S) + \left\{ \alpha(1-\xi^2) - \gamma^2\xi^2 \right\} (v_S + v_W) \right] \qquad \text{(C.3)} \\
&\frac{\partial \mathbb{E}_{s_\cdot, m_\cdot, \theta_\cdot}(SW)}{\partial d} = \left\{ \frac{(v_S + v_W)(\psi_S + \psi_W - 1)^2 \psi_0^2 (1-\psi_0)^2 \cdot d}{X^2(1-X)^2} \right. \\
&\quad \left. [X(1-\rho_{opt}) + (1-X)\rho_{opt}] \right\} \\
&\quad \times \left[-2\gamma\xi(\theta_W - \theta_S) - \left\{ \alpha(1-\xi^2) - \gamma^2\xi^2 \right\} (v_S + v_W) \right], \qquad \text{(C.4)}
\end{aligned}$$

[13] I thank the anonymous referee for encouraging me to develop this section.

Table C.1. Optimal Overly Pessimistic Forward-Guidance Plan if $[2\gamma\xi(\theta_W - \theta_S) + \{\alpha(1-\xi^2) - \gamma^2\xi^2\}(v_S + v_W)] < 0$ and $\psi_0 \in [\frac{1}{2}, 1]$

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	$\frac{1}{2\psi_0}$	$(1 - \frac{1}{2\psi_0})$
s_W	0	1

where $X \overset{def}{=} (1-\psi_0)(\rho_{opt} - d) + \psi_0\rho_{opt} = \rho_{opt} - d(1-\psi_0)$, $(1 - X) \overset{def}{=} (1-\psi_0)(1 - \rho_{opt} + d) + \psi_0(1 - \rho_{opt}) = (1 - \rho_{opt}) + d(1-\psi_0)$, $Z \overset{def}{=} 2\psi_W + 2\psi_S - 2$. The threefold results of the optimal forward-guidance policy function remain unchanged, but the threshold condition changes. The threshold condition is now $\left[2\gamma\xi(\theta_W - \theta_S) + \left\{\alpha(1-\xi^2) - \gamma^2\xi^2\right\}(v_S + v_W)\right] \gtrless 0$. For the paired optimal inflation-targeting bandwidth, again, plug in the optimal forward-guidance policy for each environment, take the first-derivative w.r.t. v_S and v_W, and find the values that maximize the ex ante expected social welfare: $\frac{\partial \mathbb{E}_{s_\cdot, m_\cdot, \theta_\cdot}(SW)}{\partial v_\cdot} = 0$. The optimal pairs of forward guidance and asymmetric inflation-targeting bandwidth are presented below by the types of forward-guidance policy.

C.3 Optimal Asymmetric Inflation Bandwidth Paired with Overly Pessimistic Forward Guidance

When the parameters about the economic fundamental satisfy $\psi_0 \in [\frac{1}{2}, 1]$ and $\left[2\gamma\xi(\theta_W - \theta_S) + \left\{\alpha(1-\xi^2) - \gamma^2\xi^2\right\}(v_S + v_W)\right] < 0$, the optimal forward-guidance policy is as the optimal overly pessimistic forward-guidance plan (see Table C.1).

By plugging the forward-guidance policy function from Table C.1 into the ex ante expected social welfare function (for computational details, please see Ko 2022), the paired optimal inflation-targeting bandwidth is derived as follows:

$$v_S^* = \frac{(\theta_W - \theta_S)\,\gamma\xi\,(C_1 + (1-\psi_S))\left\{(1-\psi_S)\psi_S + 2A_1\left(1 - 2A_1\right)\right\}}{B_1},$$

$$v_W^* = \frac{(\theta_W - \theta_S)\,\gamma\xi \cdot C_1\left\{(1-\psi_S)\psi_S + 2A_1\,(1 - 2A_1)\right\}}{B_1},$$

where $2A_1 \overset{def}{=} 2(1-\psi_0)(\psi_S+\psi_W-1)+(1-\psi_S) = 2\left(\frac{2\psi_0-1}{2}\cdot(1-\psi_S)+(1-\psi_0)\cdot\psi_W\right)$, $(1-2A_1) \overset{def}{=} -2(1-\psi_0)(\psi_S+\psi_W-1)+\psi_S = 2\left(\frac{2\psi_0-1}{2}\cdot\psi_S+(1-\psi_0)\cdot(1-\psi_W)\right)$, $B_1 \overset{def}{=} 2\alpha(1-\psi_0)(\psi_S+\psi_W-1)\Big\{(1-\psi_0)(\psi_S+\psi_W-1)+(1-\psi_S)\Big\}+\Big\{\gamma^2\xi^2-\alpha(1-\xi^2)\Big\}2A_1\Big\{(1-\psi_S)\psi_S+2A_1\,(1-2A_1)\Big\}$, $C_1 \overset{def}{=} (1-\psi_0)(\psi_S+\psi_W-1)$.

However, the value above $(v_W^*+v_S^*)$ does not satisfy the threshold condition about $\left[2\gamma\xi(\theta_W-\theta_S)+\left\{\alpha(1-\xi^2)-\gamma^2\xi^2\right\}(v_W^*+v_S^*)\right]$. Then, the optimal asymmetric inflation-targeting bandwidth paired with the overly pessimistic forward guidance is determined by the threshold:

$$v_S'+v_W' = \frac{2\gamma\xi(\theta_W-\theta_S)}{\gamma^2\xi^2-\alpha(1-\xi^2)}$$
$$v_S' = \frac{\{C_1+(1-\psi_S)\}\,2\gamma\xi(\theta_W-\theta_S)}{A_1\left\{\gamma^2\xi^2-\alpha(1-\xi^2)\right\}},\quad v_W' = \frac{C_1\cdot 2\gamma\xi(\theta_W-\theta_S)}{A_1\left\{\gamma^2\xi^2-\alpha(1-\xi^2)\right\}}. \tag{C.5}$$

Clearly, $v_W' < v_S'$, where $v_W' = \frac{C_1}{C_1+(1-\psi_S)}\cdot v_S'$ following the optimal relationship given by $v_W^* = \frac{C_1}{C_1+(1-\psi_S)}\cdot v_S^*$. As the monetary policy is constructed as $\pi^{policy}(\theta_W) = \pi + v_W$ and $\pi^{policy}(\theta_S) = \pi - v_S$, the bandwidth above the inflation target is relatively narrow while the bandwidth beneath the inflation target is relatively wide.

Recall that α is the weight which the central bank puts on the inflation gap. Clearly, as α increases, the threshold increases that bifurcates whether the overly pessimistic forward guidance being optimal or the uninformative forward guidance being optimal $\left(\frac{\partial(v_S'+v_W')}{\partial\alpha} > 0\right)$ when $\psi_0 \geq 1/2$. Concurrently, the constrained optimal asymmetric inflation-targeting bandwidth (v_S', v_W') paired with the overly pessimistic forward guidance increases.

It is worth noting that as ξ decreases, i.e., the private sector plays a key role in realization of ex post inflation, the denominator decreases rapidly since $\xi^2 < \xi \leq 1$. Hence, if the central bank has weaker influence over ex post realized inflation and the private sector's influence becomes more significant, i.e., when ξ decreases, then the optimal inflation-targeting bandwidth $(v_S' + v_W')$ increases

Table C.2. Truth Forward-Guidance Policy Function if $\psi_0 \in [0, \frac{1}{2}]$ and $[2\gamma\xi(\theta_W - \theta_S) + \{\alpha(1-\xi^2) - \gamma^2\xi^2\}(v_S + v_W)] < 0$

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	1	0
s_W	0	1

$\left(\frac{\partial(v'_S+v'_W)}{\partial\xi} < 0\right)$. This result seems intuitive that when the private sector influences more significantly, the role of the expected inflation of the private sector to the social welfare function becomes more significant: the economy is better off when the central bank has more room to adjust cyclical unemployment through the expected inflation.

C.4 Optimal Asymmetric Inflation Bandwidth Paired with Truthful Forward Guidance

When the parameters about the economic fundamental satisfies $\psi_0 \in [0, \frac{1}{2})$ and $\left[2\gamma\xi(\theta_W - \theta_S) + \left\{\alpha(1-\xi^2) - \gamma^2\xi^2\right\}(v_S + v_W)\right] < 0$, the optimal forward-guidance policy plan is as the truthful policy (see Table C.2).

By plugging in the truthful policy from Table C.2, the optimal asymmetric inflation-targeting bandwidth is derived as follows:

$$v_S^{**} = \frac{(\theta_W - \theta_S)\,\gamma \cdot \xi \cdot (1-\phi)}{\gamma^2\xi^2 - \alpha\xi(1-\xi) + \alpha\xi(1-\phi)},$$
$$v_W^{**} = \frac{(\theta_W - \theta_S)\,\gamma \cdot \xi \cdot \phi}{\gamma^2\xi^2 - \alpha\xi(1-\xi) + \alpha\xi(1-\phi)}, \tag{C.6}$$

where $\phi = \psi_0\psi_S + (1-\psi_0)(1-\psi_W)$, $(1-\phi) = \psi_0(1-\psi_S) + (1-\psi_0)\psi_W$.

Unfortunately, the values above (v_S^{**}, v_W^{**}) do not satisfy the threshold condition: $\left[2\gamma\xi(\theta_W - \theta_S) + \left\{\alpha(1-\xi^2) - \gamma^2\xi^2\right\}(v_S^{**} + v_W^{**})\right] < 0$.

Therefore, the optimal bandwidth paired with the truthful forward guidance is determined by the threshold:

$$v_S'' + v_W'' = \frac{2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)} \tag{C.7}$$
$$v_S'' = \frac{(1-\phi)2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}, \quad v_W'' = \frac{\phi 2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}.$$

The above values are determined by $v_S'' + v_W'' = \frac{1}{\phi}v_W'' = \frac{1}{1-\phi}v_S''$. This relationship follows the optimal condition that $v_S'' + v_W'' = \frac{1}{\phi}v_W'' = \frac{1}{1-\phi}v_S''$. Observe that, since $\psi_0 < 1/2$, it is more likely that $\phi < (1-\phi)$, and thus $v_W'' = \frac{\phi}{1-\phi}v_S'' < v_S''$. However, the other case $\phi > (1-\phi)$ is not impossible, and $v_W'' = \frac{\phi}{1-\phi}v_S'' > v_S''$ holds if $\phi > (1-\phi)$. Observe that if $\psi_0 = 1/3$, $\psi_S = 8/9$, $\psi_W = 5/9$, then $\phi > (1-\phi)$ holds even if $\psi_0 = 1/3 < 1/2$.

For changes in α and ξ, observe that $\frac{\partial(v_S''+v_W'')}{\partial\alpha} > 0$ and $\frac{\partial(v_S''+v_W'')}{\partial\xi} < 0$. Similarly to the overly pessimistic case, as α increases, the threshold for $(v_S + v_W)$ rises that bifurcates whether the truthful forward guidance being optimal or the uninformative forward guidance being optimal when $\psi_0 < 1/2$. Also, when the private sector influences more significantly, i.e., $(1-\xi)$ rises, the role of the expected inflation of the private sector to the social welfare function becomes more significant and the central bank should have more room to adjust the expected inflation by increasing the values of (v_W'', v_S'') accordingly.

C.5 Optimal Asymmetric Inflation Bandwidth Paired with Uninformative Forward Guidance

When the parameters satisfy $\left[2\gamma\xi(\theta_W - \theta_S) + \left\{\alpha(1-\xi^2) - \gamma^2\xi^2\right\}(v_S + v_W)\right] > 0$, the optimal forward-guidance policy is the uninformative one (see Table C.3).

Given the forward-guidance policy shown in Table C.3 as optimal, the paired optimal asymmetric inflation-targeting bandwidth is as follows:

Table C.3. Optimal Uninformative Forward-Guidance Policy Function if $[2\gamma\xi(\theta_W - \theta_S) + \{\alpha(1-\xi^2) - \gamma^2\xi^2\}(v_S + v_W)] > 0$

$\sigma : S \to \Delta(M)$	m_{opt}	m_{pes}
s_S	1/2	1/2
s_W	1/2	1/2

Table C.4. Pairs of Forward Guidance and Asymmetric Inflation-Targeting Bandwidth

	$\psi_0 \geq 1/2$	$\psi_0 < 1/2$	$\forall \psi_0$
Threefold Forward Guidance	OverlyPessimistic	Truthful	Uninformative
	$v'_S + v'_W = v''_S + v''_W = \frac{2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}$		$v_S^{***} + v_S^{***} = \frac{(\theta_W - \theta_S)\gamma}{\xi(\gamma^2+\alpha)}$

$$v_S^{***} = \frac{(\theta_W - \theta_S)\gamma\,(1-\phi)}{\xi(\gamma^2+\alpha)}, \quad v_W^{***} = \frac{(\theta_W - \theta_S)\gamma\phi}{\xi(\gamma^2+\alpha)} \tag{C.8}$$

where $\phi = (\psi_0\psi_S + (1-\psi_0)(1-\psi_W))$, $(1-\phi) = \big(\psi_0(1-\psi_S) + (1-\psi_0)\psi_W\big)$.

The sum $v_S^{***} + v_W^{***} = \frac{(\theta_W - \theta_S)\gamma}{\xi(\gamma^2+\alpha)}$ satisfies the threshold condition. If α is sufficiently large, the values of v_S^{***} and v_W^{***} are decreasing in α $\left(\frac{\partial(v_S^{***}+v_W^{***})}{\partial\alpha} < 0\right)$. Clearly, if the private sector's expectation plays a bigger role in the realization of ex post inflation, i.e., ξ declines, then the values of v_S^{***}, v_W^{***} increases $\left(\frac{\partial(v_S^{***}+v_W^{***})}{\partial\xi} < 0\right)$.

C.6 Welfare Comparison

So far three distinct local optimal pairs of forward-guidance policy and asymmetric inflation-targeting band are derived as in Table C.4. In this section, I compare the welfare levels of these three distinct local optimal pairs. The notations are introduced as $L_S \overset{def}{=} (1-\psi_S)v_W - \psi_S v_S$, $L_W = \psi_W v_W - (1-\psi_W)v_S$, $\phi = \psi_0\psi_S + (1-\psi_0)(1-\psi_W)$, $(1-\phi) = \psi_0(1-\psi_S) + (1-\psi_0)\psi_W$, $p_1 = \frac{\psi_S}{2}$, $p_2 = \left(\phi - \frac{\psi_S}{2}\right)$, $p_3 = \frac{1-\psi_S}{2}$, $p_4 = \left(1 - \phi - \frac{1-\psi_S}{2}\right)$. Observe

that $p_1 + p_2 = \phi$ and $p_3 + p_4 = 1 - \phi$. Also, $p_1 + p_3 = \frac{1}{2} = p_2 + p_4$. Finally, $\overline{v} \equiv (\psi_0(1-\psi_S) + (1-\psi_0)\psi_W)\, v_W - (\psi_0\psi_S + (1-\psi_0)(1-\psi_W))\, v_S$.

First, the difference between the ex ante expected social welfare associated with the overly pessimistic forward-guidance policy and the one for the uninformative forward-guidance policy is given as follows (when $\psi_0 \in [1/2, 1]$):

$$
\begin{aligned}
&\mathbb{E}\left[SW_{pess};\ (v_W + v_S) = \frac{2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}\right] \\
&\quad - \mathbb{E}\left[SW_{uninfo};\ (v_W + v_S) = \frac{2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}\right] \\
&= -\left(\theta_S + \gamma\xi(v_S + v_W)\cdot 2p_3\right)^2 \cdot p_1 \\
&\quad + \left\{\theta_S + \gamma\cdot\xi\cdot(p_3 + p_4)(v_W + v_S)\right\}^2 p_1 \\
&\quad - \left[\theta_S + \gamma\xi(v_S + v_W)\cdot 2p_4\right]^2 \cdot p_2 \\
&\quad + \left\{\theta_S + \gamma\cdot\xi\cdot(p_3 + p_4)(v_W + v_S)\right\}^2 p_2 \\
&\quad - \left(\theta_W + \gamma\xi(v_S + v_W)\cdot(-2p_1)\right)^2 \cdot p_3 \\
&\quad + \left\{\theta_W + \gamma\cdot\xi\cdot(-p_1 - p_2)(v_W + v_S)\right\}^2 p_3 \\
&\quad - \left[\theta_W + \gamma\xi(v_S + v_W)\cdot(-2p_2)\right]^2 \cdot p_4 \\
&\quad + \left\{\theta_W + \gamma\cdot\xi\cdot(-p_1 - p_2)(v_W + v_S)\right\}^2 p_4 \\
&\quad - \alpha\left\{(1-\xi)L_S + \xi(-v_S)\right\}^2 p_1 \\
&\quad + \alpha\left\{(1-\xi)(1-\psi_0)(L_W - L_S) + (1-\xi)\cdot L_S + \xi(-v_S)\right\}^2 p_1 \\
&\quad - \alpha\left\{(1-\xi)2(1-\psi_0)(L_W - L_S) + (1-\xi)\cdot L_S + \xi(-v_S)\right\}^2 p_2 \\
&\quad + \alpha\left\{(1-\xi)(1-\psi_0)(L_W - L_S) + (1-\xi)\cdot L_S + \xi(-v_S)\right\}^2 p_2 \\
&\quad - \alpha\left\{(1-\xi)L_S + \xi v_W\right\}^2 p_3 \\
&\quad + \alpha\left\{(1-\xi)(1-\psi_0)(L_W - L_S) + (1-\xi)\cdot L_S + \xi v_W\right\}^2 p_3 \\
&\quad - \alpha\left\{(1-\xi)2(1-\psi_0)(L_W - L_S) + (1-\xi)\cdot L_S + \xi v_W\right\}^2 p_4 \\
&\quad + \alpha\left\{(1-\xi)(1-\psi_0)(L_W - L_S) + (1-\xi)\cdot L_S + \xi v_W\right\}^2 p_4 \\
&= -\left\{2\gamma\xi(\theta_W - \theta_S) - \gamma^2\xi^2(v_W + v_S) + \alpha(1-\xi^2)(v_S + v_W)\right\} \\
&\quad (v_S + v_W)(1-\psi_0)^2(\psi_W + \psi_S - 1)^2 \\
&= 0.
\end{aligned}
$$

That is, $\mathbb{E}\left[SW_{pess};\ (v_W + v_S) = \frac{2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}\right] = \mathbb{E}\left[SW_{uninfo};\ (v_W + v_S) = \frac{2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}\right]$. For computation details, see the mathematical appendix (Ko 2022). If $\alpha > \frac{\xi^2(1-2\xi)}{1-\xi^2+2\xi^3}\gamma^2$, then the uninformative forward guidance and its paired asymmetric inflation-targeting bandwidth achieves the global optimum, because by construction of $(v_W^{***} + v_S^{***})$, $\mathbb{E}\left[SW_{uninfo};\ (v_W + v_S) = \frac{2\gamma\xi(\theta_W - \theta_S)}{\gamma^2\xi^2 - \alpha(1-\xi^2)}\right] < \mathbb{E}\left[SW_{uninfo};\ (v_W^{***} + v_S^{***}) = \frac{\gamma(\theta_W - \theta_S)}{\xi(\gamma^2+\alpha)}\right]$.

Second, assuming $\psi_0 < 1/2$, the ex ante expected social welfare associated with the truthful forward guidance is now compared with the one associated with the uninformative forward guidance.

$$
\begin{aligned}
&\mathbb{E}[SW_{truth};\ (v_W + v_S)] - \mathbb{E}[SW_{uninfo};\ (v_W + v_S)] \\
&\quad = -\left\{\theta_S + \gamma\cdot\xi\cdot(L_S + v_S)\right\}^2 \psi_0\psi_S \\
&\qquad + \left\{\theta_S + \gamma\cdot\xi\cdot(\overline{v} + v_S)\right\}^2 \psi_0\psi_S \\
&\qquad - \alpha\left\{(1-\xi)\cdot L_S + \xi(-v_S)\right\}^2 \psi_0\psi_S \\
&\qquad + \alpha\left\{(1-\xi)\cdot \overline{v} + \xi(-v_S)\right\}^2 \psi_0\psi_S \\
&\qquad - \left\{\theta_S + \gamma\cdot\xi(L_W + v_S)\right\}^2 (1-\psi_0)(1-\psi_W) \\
&\qquad + \left\{\theta_S + \gamma\cdot\xi\cdot(\overline{v} + v_S)\right\}^2 (1-\psi_0)(1-\psi_W) \\
&\qquad - \alpha\left\{(1-\xi)\cdot L_W + \xi(-v_S)\right\}^2 (1-\psi_0)(1-\psi_W) \\
&\qquad + \alpha\left\{(1-\xi)\cdot \overline{v} + \xi(-v_S)\right\}^2 (1-\psi_0)(1-\psi_W) \\
&\qquad - \left\{\theta_W + \gamma\cdot\xi\cdot(L_S - v_W)\right\}^2 \psi_0(1-\psi_S) \\
&\qquad + \left\{\theta_W + \gamma\cdot\xi\cdot(\overline{v} - v_W)\right\}^2 \psi_0(1-\psi_S) \\
&\qquad - \alpha\left\{(1-\xi)\cdot L_S + \xi v_W\right\}^2 \psi_0(1-\psi_S) \\
&\qquad + \alpha\left\{(1-\xi)\cdot \overline{v} + \xi(+v_W)\right\}^2 \psi_0(1-\psi_S) \\
&\qquad - \left\{\theta_W + \gamma\cdot\xi\cdot(L_W - v_W)\right\}^2 (1-\psi_0)\psi_W \\
&\qquad + \left\{\theta_W + \gamma\cdot\xi\cdot(\overline{v} - v_W)\right\}^2 (1-\psi_0)\psi_W \\
&\qquad - \alpha\left\{(1-\xi)\cdot L_W + \xi v_W\right\}^2 (1-\psi_0)\psi_W \\
&\qquad + \alpha\left\{(1-\xi)\cdot \overline{v} + \xi(+v_W)\right\}^2 (1-\psi_0)\psi_W
\end{aligned}
$$

$$= - \psi_0(1-\psi_0)(\psi_S+\psi_W-1)^2(v_W+v_S)\big[2\gamma\xi(\theta_W-\theta_S) - \left\{\gamma^2\xi^2-\alpha(1-\xi^2)\right\}(v_W+v_S)\big] \tag{C.9}$$

For any given value of $(v_W + v_S)$, the above holds. If $(v_W + v_S) = \frac{2\gamma\xi(\theta_W-\theta_S)}{\gamma^2\xi^2-\alpha(1-\xi^2)}$ is plugged in, then it is easy to show that $\mathbb{E}\left[SW_{truth};\ (v''_W+v''_S)=\frac{2\gamma\xi(\theta_W-\theta_S)}{\gamma^2\xi^2-\alpha(1-\xi^2)}\right] = \mathbb{E}\left[SW_{uninfo};\ (v''_W+v''_S)=\frac{2\gamma\xi(\theta_W-\theta_S)}{\gamma^2\xi^2-\alpha(1-\xi^2)}\right]$. When $\alpha > \frac{\xi^2(1-2\xi)}{1-\xi^2+2\xi^3}\gamma^2$, the $(v_W^{***}+v_S^{***})$ gives a better level of the ex ante expected social welfare than the one achieved by $(v''_W+v''_S)=\frac{2\gamma\xi(\theta_W-\theta_S)}{\gamma^2\xi^2-\alpha(1-\xi^2)}$. Therefore, the uninformative forward guidance and its paired asymmetric inflation-targeting bandwidth achieves the global optimum under the aforementioned condition.

C.7 Discussion

When the ex post realized inflation is influenced by private expectation with a weight of $(1-\xi)$ and inflation-targeting bandwidth is allowed to be asymmetric by $\pi^{policy}(\theta_S)=\pi^T-v_S$ and $\pi^{policy}(\theta_W)=\pi^T+v_W$, the forward-guidance threshold is determined by the sum of the bandwidth around the inflation target, (v_W+v_S) or $2v$. That is, irrespective of asymmetric inflation targeting or symmetric one, the threshold for forward-guidance is at the point where the sum of positive and negative side bandwidth, (v_W+v_S) or $2v$ is given as $\frac{2\gamma\xi(\theta_W-\theta_S)}{\gamma^2\xi^2-\alpha(1-\xi^2)}$. Also, the threefold optimal forward-guidance results remain the same. However, the paired inflation-targeting bandwidth becomes different from the symmetric case. It is because, for example, when the optimality condition for v_W is formed via the first-order derivative with respect to v_W, the terms with v_S and its associated probabilities are treated as constant. Similarly, for the optimality condition for v_S, the terms with v_W and the associated probability are treated as constant. Then, the optimal conditions for v_S and v_W are given differently, i.e., $\frac{\partial\mathbb{E}[SW]}{\partial v_W}\neq\frac{\partial\mathbb{E}[SW]}{\partial v_S}$. In the symmetric inflation-targeting bandwidth case, by contrast, there is only one FOC $\frac{\partial\mathbb{E}[SW]}{\partial v}$ w.r.t. v. Due to this distinction between symmetric and asymmetric inflation targeting, the optimal values of v_W and v_S of the asymmetric case do not converge to the optimal value of v

even if the private sector's self-fulfilling effect is ignored ($\xi \to 1$). And the asymmetric inflation-targeting bandwidth achieves higher social welfare level than the symmetric one: otherwise, the results of asymmetric inflation targeting would have converged to the symmetric inflation-targeting bandwidth.

References

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon. 2019. "Forward Guidance and Heterogeneous Beliefs." *American Economic Journal: Macroeconomics* 11 (3): 1–29.

Bassetto, M. 2019. "Forward Guidance: Communication, Commitment, or Both?" *Journal of Monetary Economics* 108 (December): 69–86.

Beaudry, P., and T. Willems. 2022. "On the Macroeconomic Consequences of Over-Optimism." *American Economic Journal: Macroeconomics* 14 (1): 38–59.

Bergemann, D., and S. Morris. 2016. "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium." *American Economic Review* 106 (5): 586–91.

Binder, C., and C. Makridis. 2022. "Stuck in the Seventies: Gas Prices and Consumer Sentiment." *Review of Economics and Statistics* 104 (2): 293–305.

Campbell, J. R., C. L. Evans, J. D. M. Fisher, and A. Justiniano. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity* (Spring): 1–80.

Cieslak, A., S. Malamud, and A. Schrimpf. 2019. "Policy Announcement Design." Research Paper No. 20-17, Swiss Finance Institute.

Cogley, T., and T. J. Sargent. 2005. "The Conquest of US Inflation: Learning and Robustness to Model Uncertainty." *Review of Economic Dynamics* 8 (2): 528–63.

Coibion, O., Y. Gorodnichenko, and M. Weber. 2022. "Monetary Policy Communications and Their Effects on Household Inflation Expectations." *Journal of Political Economy* 130 (6): 1537–84.

Crawford, V. P., and J. Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.

Del Negro, M., M. Giannoni, and C. Patterson. 2012. "The Forward Guidance Puzzle." Staff Report No. 574, Federal Reserve Bank of New York.

Foerster, A., and T. Davig. 2017. "Communicating Monetary Policy Rules." Working Paper No. 17-04, Federal Reserve Bank of Kansas City.

Fujiwara, I., and Y. Waki. 2019. "Private News and Monetary Policy–Forward Guidance as Bayesian Persuasion." CAMA Working Paper No. 2019-91, Crawford School of Public Policy, Australian National University.

Geraats, P. M. 2002. "Central Bank Transparency." *Economic Journal* 112 (483): F532–F565.

Gürkaynak, R. S., B. Sack, and E. T. Swanson. 2005. "Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1 (1): 55–94.

Kamenica, E., and M. Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6): 2590–2615.

Ko, E. 2022. "Online Math Appendix to 'Optimal Central Bank Forward Guidance'." Unpublished Manuscript. https://www.dropbox.com/s/w2lzrmll0za62oo/EK_2022_may_IJCB_online_appendix.pdf?dl=0.

Leombroni, M., A. Vedolin, G. Venter, and P. Whelan. 2021. "Central Bank Communication and the Yield Curve." *Journal of Financial Economics* 141 (3): 860–80.

Morris, S., and H. S. Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.

———. 2018. "Central Bank Forward Guidance and the Signal Value of Market Prices." *AEA Papers and Proceedings* 108 (May): 572–77.

Qvigstad, J. F. 2006. "When Does an Interest Rate Path Look Good? Criteria for an Appropriate Future Interest Rate Path." Monetary Policy Working Paper ANO 2006/5, Norges Bank.

Rudebusch, G. D., and J. C. Williams. 2008. "Revealing the Secrets of the Temple: the Value of Publishing Central Bank Interest Rate Projections." In *Asset Prices and Monetary Policy*, ed. J. Y. Campbell, 247–89. University of Chicago Press.

Walsh, C. E. 2014. "Multiple Objectives and Central Bank Trade-offs Under Flexible Inflation Targeting." CESifo Working Paper No. 5097.

Woodford, M. 2005. "Central Bank Communication and Policy Effectiveness." In *The Greenspan Era: Lessons for the Future.* Proceedings of the Annual Economic Policy Symposium held in Jackson Hole, Wyoming. Federal Reserve Bank of Kansas City.

An Interpretable Machine Learning Workflow with an Application to Economic Forecasting*

Marcus Buckmann and Andreas Joseph
Bank of England

We propose a generic workflow for the use of machine learning models to inform decisionmaking and to communicate modeling results with stakeholders. It involves three steps: (i) a comparative model evaluation, (ii) a feature importance analysis, and (iii) statistical inference based on Shapley value decompositions. We discuss the different steps of the workflow in detail and demonstrate each by forecasting changes in U.S. unemployment one year ahead using the well-established FRED-MD data set. We find that universal function approximators from the machine learning literature, including gradient boosting and artificial neural networks, outperform more conventional linear models. This better performance is associated with greater flexibility, allowing the machine learning models to account for time-varying and non-linear relationships in the data-generating process. The Shapley value decomposition identifies economically meaningful non-linearities learned by the models. Shapley regressions for statistical inference on machine learning models enable us to assess and communicate variable importance akin to conventional econometric approaches. While we also explore high-dimensional models, our findings suggest that the best trade-off between interpretability and performance of the models is achieved when a small set of variables is selected by domain experts.

JEL Codes: C14, C45, C53, E27.

*The content and views presented in this article do not represent the views of the Bank of England (BoE). Many thanks to David Bholat, Philippe Bracke, and Helena Robertson for support of the project and helpful comments on the manuscript. The code for implementing and replicating the paper's results can be found at https://github.com/bank-of-england/InterpretableMLWorkflow.

1. Introduction

Predictive machine learning models are increasingly being used at decisionmaking institutions, such as central banks, governments, and international institutions (Doerr, Gambacorta, and Garralda 2021). Major appeals of these models are that they often give more accurate predictions than conventional approaches and can handle high-dimensional data (Haldane 2018).

On the downside, many machine learning methods suffer from the black box critique. It is not straightforward to assess the factors driving predictions and therefore to understand the relations between the inputs and output of the model. However, this understanding of a model is crucial, especially for decisionmaking processes, for several reasons. First, both decisionmakers and their audiences naturally have a desire to understand the inputs leading to decisions and legitimize them. Second, decisionmaking processes often involve multiple models.[1] The information derived from different models should be compatible, leading to a coherent picture. The understanding of all models involved is needed for this. Third, models can "misfire" for several reasons—for example, by picking up spurious relations in the data. This often can only be detected and prevented if one has a good understanding of a model.

Prediction models whose accuracy is a key motivation behind their deployment—which often holds for machine learning methods—should also help to inform the narrative approach behind any economic policy decision rather than providing mere black box predictions (George 1999; Burgess et al. 2013; Independent Evaluation Office 2015). Machine learning models also can provide a richer set of information compared with more conventional statistical models, like linear regression models. In particular, they can implicitly learn non-linear functional forms and interaction from the data without the need to specify them a priori.

In this paper, we lay out a multi-step workflow for the use of machine learning models, which we deem suitable to inform

[1]Without any stringent assumption on the data-generating process, machine learning models can be labeled "non-structural," describing correlations between inputs and targets. Other "structural" models make richer assumptions about the data-generating process. Comparing the two requires analyses of the assumptions on the data-generating process, estimation techniques, and results.

decisionmaking processes. It consists of three steps which can be directly applied to other contexts as well than those presented in the accompanying case study. First, a model comparison is conducted between conventional statistical methods and machine learning models to provide prima facie evidence of whether a machine learning approach is likely to deliver benefits. If the primary objective is model accuracy, e.g., for forecasting, this would be a model horse race to minimize the forecasting error. Second, the machine learning predictions are decomposed into the contributions of the individual model variables. This allows us to uncover the relative importance of variables and understand the functional forms learned by the different machine learning models. By a comparison across models, one can gauge how robust feature decompositions are to the choice of the algorithm. Third, statistical inference is conducted to understand which variables make a statistically significant contribution to the accuracy of a model, providing a level of confidence for our interpretations and any narrative attached to them. This inference uses a parametric regression analysis, allowing for a standardized communication of statistical model results. A rich set of robustness checks provides guidance for frequently encountered challenges, especially when using machine learning. These include variable selection in a high-dimensional setting, model stability, and computational requirements.

Throughout the paper, we apply the proposed procedures to a macroeconomic case study, where we forecast changes in unemployment—an important input for fiscal and monetary policy decisions (Burgess et al. 2013). Along the steps of the workflow, we contrast the use of machine learning models with a simpler but less flexible linear model.

Most of the presented techniques generalize to other settings in a straightforward manner, such that this paper provides a one-stop reference for practitioners for the use of machine learning models in situations where the understanding and communication of model results is crucial. This is especially the case at policy institutions. Inevitably this ties to many technicalities. We discuss the most relevant ones in intuitive terms and provide references to the related literature for further guidance.

There are several levels of communication involved in a data-driven decision process, from technical modeling experts performing

the analysis, over communications to management, up to the decisionmaking bodies. The ability to communicate modeling results with varying levels of complexity is crucial in this setting. However, effective communication is highly contextual. The technical knowledge and experience can vary greatly within decisionmaking bodies and their audiences. Thus, there is no one-size-fits-all mapping between workflow outputs and target audiences. Instead we provide some broad guidance and suggestions on matching individual outputs with target audiences as follows. We layer target audiences by how close they are to the technical details of the analysis, going from analysts, who perform the analysis, to management, who aggregate and filter information from different sources (e.g., different teams of analysts) and distill information for decisionmaking bodies, and finally decisionmakers and their audiences. This gives three levels, where guidance is to be understood as a "smaller or equal," meaning that if the target audience is the management, it also includes analysts, and if it is decisionmakers, it may serve for all. Labels for the target audience are mostly attached to table and figure captions which summarize the outputs of our workflow.

The present paper connects different fields, ranging from machine learning and model interpretability to statistical inference and economic forecasting. There is a growing literature that suggests that machine learning methods can outperform more conventional models in economic prediction problems including forecasting. For example, machine learning methods have been shown to be better at predicting bond risk premia (Bianchi, Büchner, and Tamoni 2019), forecasting macroeconomic variables such as unemployment and inflation (Sermpinis et al. 2014; Chen et al. 2022), recessions (Döpke, Fritsche, and Pierdzioch 2017), and financial crises (Bluwstein et al. 2020).[2] However, other papers do not observe consistently improved performance by using machine learning, instead finding that it is state or horizon dependent (Kock and Teräsvirta 2014). This mixed

[2]In these problems, several variables are used to forecast the outcome variable. In the univariate case, when only the lagged outcome is used for prediction, evidence suggests that statistical methods or hybrid models combining statistical and machine learning approaches outperform pure machine learning methods, on average (Makridakis, Spiliotis, and Assimkopoulos 2018a, 2018b; Parmezan et al. 2019).

evidence validates our horse race as an important first step for the workflow.

Predicting macroeconomic dynamics is challenging. Relationships between variables may not hold over time, and shocks such as recessions or financial crises might lead to a breakdown of previously observed relationships (Elliott and Timmermann 2008; Ng and Wright 2013). In line with the literature, we suggest that it is the inherent non-linearity of non-parametric models that allows them to learn and exploit complex relationships for prediction (Wang and Manning 2013). Coulombe et al. (2020) show that this advantage of machine learning models to exploit non-linearities in macroforecasting is enhanced at longer horizons. However, the non-linear relationships learned are not directly observable, which has led to the aforementioned black box critique of these models as a major challenge to their applicability to inform decisions.

Approaches to interpretable machine learning come from different directions: epistemic discussions about what it means for a model to be interpretable (Miller 2019), technical approaches in machine learning research (Doshi-Velez and Kim 2017), and methodology in econometrics and statistics (Chernozhukov et al. 2018).

Miller (2019) analyzes the psychology of explanations and suggests that humans expect explanations that are based on a limited number of causes rather than an exhaustive account of all factors—acknowledging that the simplification of the problem risks introducing bias. Relatedly, Lipton (2016) argues that a high-dimensional linear model is not necessarily more interpretable than a compact artificial neural network that learns from only few features. Also, if the linear model is trained on abstract features, for instance, obtained by principal component analysis, its parameters may not provide an obvious economic interpretation.

In the machine learning literature, approaches to interpretability usually focus on measuring how important input variables are for prediction. *Variable attributions* can be either global, by assessing the variable importance across the whole data set, or local, by measuring the importance of the variables at the level of individual observations in the form of a decomposition. Such local attributions can always be summarized in a global variable importance measure by averaging local attributions across all observations. Popular global methods are permutation importance or Gini

importance for tree-based models (Breiman 2001a). Popular local decomposition methods are LIME (Ribeiro, Singh, and Guestrin 2016), DeepLIFT (Shrikumar, Greenside, and Anshul 2017), and Shapley values (Štrumbelj and Kononenko 2010; Lundberg and Lee 2017). Lundberg and Lee (2017) demonstrate that Shapley values offer a unified framework of LIME (local interpretable model-agnostic explanations) and DeepLIFT with appealing properties. Most importantly, Shapley values guarantee *consistency*, where a consistent measure of variable importance preserves the relative importance between variables across situations where such a ranking is imposed. We therefore focus on Shapley values when describing the workflow and presenting the case study. For illustrative purposes, we contrast the use of Shapley values with permutation importance.

These global and local attribution methods are only descriptive—they explain the drivers of model predictions and performance, but they do not assess the predictors' statistical significance, i.e., how certain one can be that a variable is actually important to describe a specific outcome. We extend our interpretation of machine learning models for forecasting by statistically testing the predictors in a Shapley regression framework (Joseph 2019). Shapley values and inference based on them is arguably the most general and rigorous approach to address the issues of machine learning interpretability and model communication. In this way, we close the gap between two traditional modeling approaches, the maximization of predictive performance using "black box" machine learning methods and the application of statistical techniques to make inferences about the data-generating process (Breiman 2001b).

The remainder of this paper is structured as follows. Section 2 describes the proposed workflow. The data and the methodology used for the macroeconomic forecasting study used throughout this paper is introduced in Section 3. Section 4 presents the outputs of the workflow for our baseline scenario. This includes model performances, the analysis of feature importances and learned functional forms, and statistical inference. Section 5 discusses a rich set of robustness checks and how they relate to different aspects of the proposed workflow. We conclude with a short discussion in Section 6. The technical appendix (Appendix B) discusses the computation of model interpretability measures used in the case study.

2. A Machine Learning Workflow

Our proposed workflow for the use of machine learning models is geared towards situations where model interpretability and the communication of modeling results is important. It consists of three steps: a model comparison, the assessment of variable importances, and statistical inference on model components. The latter two steps are required due to the opaque nature of machine learning models.

We keep the notation deliberately simple and general, with a more detailed and specific description used in the next section and Appendix B. We say that a model $f(x;\theta) = \hat{y}$ takes inputs x, consisting of N variables indexed k, and has fitted parameters θ. The model predicts the target variable $y = \hat{y} + \epsilon$, with ϵ being the error term of which some form is minimized during model training (fitting), e.g., the squared error ϵ^2. The Greek letter ϕ_k denotes variable components of f, i.e., $f(x) = \sum_{k=0}^{N} \phi_k(x)$, with ϕ_0 being an intercept. Additionally, let us denote $\mathcal{P}$ as a performance metric to evaluate the goodness of a model. For example, we may define $\mathcal{P}$ as the prediction error, which we aim to minimize.

2.1 Step 1: Model Comparison

Machine learning methods require additional effort from the modeler compared with conventionally used econometric models (steps 2 and 3). Thus, the first step is to decide whether to proceed with any further analysis of the machine learning models by comparing their performance with that of a benchmark model. The performance metric $\mathcal{P}$ can, for example, be the absolute forecasting error when predicting a continuous variable in a time series. However, $\mathcal{P}$ can also have a more complex forms. For example, it may be a function describing the trade-off between type 1 and 2 errors when predicting a binary variable, or when describing treatment heterogeneity in an experimental setting. Crucial questions regarding this horse race are what models (not) to compare, their stability, and what or how much data to use.

2.1.1 Model Selection

The choice of models can be both specific and general—specific in the sense that there are established models for certain tasks, which can

serve as benchmarks, and general in the sense that it is usually not possible to know for machine learning models which models will predict best in a given situation (Fernández-Delgado et al. 2014). A reasonable start are popular general-purpose machine learning models, such as random forests, gradient boosting, support vector machines, and artificial neural networks (see Friedman, Hastie, and Tibshirani 2009 for an introduction to different models). Some authors include penalized regressions in the machine learning toolbox. These models arguably lie at the boundary between traditional econometric and machine learning techniques. We do not include them in the latter, as they do not have the universal approximator property (see, for example, Cybenko 1989). Universal function appoximation means that a model will, under the right circumstances, learn to approximate any functional form given enough training data. A further difference between the two classes of models is that the parameter vector θ is clearly defined for (penalized) regression models, while it is generally undefined in terms of its shape and values for machine learning models. These are set during the cross-validation and training process, respectively.

Some machine learning models may not be suited for a certain prediction task. For example, support vector machines have substantial computational costs when the data set is large. Random forests can be memory intensive, especially when allowed to grow many large trees on a large data set. Further, random forests are not suited for extrapolation beyond the training set, i.e., they cannot make predictions that exceed the observed values in the training set. On the other hand, random forests can deal well with high-dimensional data and a limited number of observations. Compared with other methods, they also deal well with extreme values and correlated variables.

Artificial neural networks are both computation and data intensive. They have a wide range of architectures, some adapted to certain data types (see Goodfellow, Bengio, and Courville 2016 for an overview), but finding the appropriate architecture and other hyperparameters can be a challenging task. In contrast, support vector machines only have a few relevant hyperparameters, and the random forest often performs well without tuning the hyperparameters at all.

2.1.2 Model Stability

Another aspect to consider is model stability. A linear or support vector regression (SVR) will always produce a deterministic optimal solution, while the training process of a random forest or artificial neural network is not deterministic, leading to different solutions when trained repeatedly on the same data with different random seeds. D'Amour et al. (2020) showed for complex neural networks that these different solutions can produce substantially different predictions on new data that differs from the distribution on which the model was trained. We may encounter this situation in economic forecasting when there is some unknown drift in the data-generating process. Further, the hyperparameter search can introduce randomness into the training of any machine learning method. For complex models, such as gradient boosting or artificial neural networks, an exhaustive search for hyperparameters often is infeasible. In practice one only tests a few values for selected key hyperparameters. Alternatively, one employs a random search, testing only a subset of all possible combinations in a larger hyperparameter space. Ideally, a machine learning model is insensitive to changes in its hyperparameters, which makes a model comparison more robust and increases the replicability of the modeling.

A remedy for low model stability is averaging several models trained based on different random seeds or slightly different training samples. However, this comes at the cost of increased computational requirements—for training the ensemble of models, storing them, and explaining their predictions.

2.1.3 Data and Variables Selection

We assume that the data set is structured, i.e., it can be represented well in a table.[3] A modeler might be tempted to use all available variables as predictors. However, including more variables increases the likelihood that a model exploits spurious correlations that might

[3]The complement to this is unstructured data such as text, images, video, etc. On these kind of data, artificial neural networks generally perform better than other machine learning approaches, while data always need to be brought into some potentially very high-dimensional tabular representation.

not hold outside the training sample, increasing model variance and lowering performance.

Accordingly, one strand of the forecasting literature recommends to hand-pick a few predictors based on prior causal knowledge (Einhorn and Hogarth 1985; Armstrong, Green, and Graefe 2015). Another important consideration is the number of observations per variable. For example, a standard least squares regression model cannot find a model if there are more variables than observations in the data. Penalized regression methods, like lasso and ridge, can produce a solution in that case (Chetverikov, Liao, and Chernozhukov 2020). However, the rate of convergence of a non-parametric estimator depends on the dimension of the input space (Stone 1982). Convergence rates for machine learning models can be low and the theory-based estimates of convergence rates is mostly not practical in real-world situations.[4] The potential problem with slow convergence, especially in high-dimensional settings, is that a model may show high variance and thus perform poorly.

On the other hand, some studies have reported successes in forecasting with large sets of variables (Medeiros et al. 2021; Chen et al. 2022)—but at the cost of interpretability: The more features are used in a model, the more difficult it is to understand and communicate the model independent of the type of model used.

A common approach to increase the predictive performance and simplify the prediction model is to calibrate it on common factors that provide a lower-rank representation of the large set input variables (Stock and Watson 2002; Kim and Swanson 2018). However, this approach also makes the interpretation of the resulting model challenging, as the data-driven factors do not necessarily have a clear interpretation. In contrast, using a small hand-picked set of diverse predictors allows us to interpret their relationship with the response variable as learned by the prediction models. But this might lead to a decrease in performance in some data sets. Giannone, Lenza, and Primiceri (2021) use Bayesian modeling on a handful of data sets to show that selecting a small set of predictors from a large set of variables is often not feasible without trading off the performance of a linear model.

[4]They may, nevertheless, be estimated empirically if enough observations are available.

Two other aspects that need to be considered are data revisions and a reporting lag. Macroeconomic data are often substantially revised (Runkle 1998). Using the most recent vintage in pseudo out-of-sample forecasting removes the data uncertainty, but revised data cannot be used when the forecasting model is used in real time to make predictions about the future. Furthermore, data are often only reported with a delay, e.g., GDP growth for this month might only be published the following month. In this case, a real-time forecast 12 months ahead on this variable is actually a 13-month-ahead forecast.

Finally, in a forecasting setting, the modeler needs to determine how many observations should be used to train the model. There is a trade-off between using all past data, which improves the convergence of the model, or only recent data, which avoids that the model is trained on observations that do not reflect the present and future due to structural shifts over time.

2.2 Step 2: Variable Importance

Variable importance measures usually answer one of two questions. How important is a variable for a model's performance $\mathcal{P}$? Or, how important is a variable to generate a predicted value $\hat{y}$? The two questions are related: the more accurate a model is, the closer $\hat{y}$ is to y and thus the more similar the two metrics of importance will be.

Measures related to $\mathcal{P}$ often are *global*, which means they provide a single number for each variable and model across the test set.[5] This is practical for communication, as one obtains a simple variable ranking. Global measures, however, can obscure many nuances of a model. For instance, machine learning models are non-linear (often non-monotonic) and, as such, a global measure risks oversimplification or producing inconclusive results when evaluated across differing domains of the input space.

[5] Variable importance can be evaluated on any fraction of the test set. Evaluation of the training set has to be interpreted with caution because of overfitting. However, comparisons across training and test sets can help to identify problems of model generalization, such as overfitting.

Local importance measures decompose individual predictions $f(x_t)$ of observation t into attributions of the individual features:

$$f(x_t) = \sum_{k=0}^{N} \phi_k(x_t)\,, \tag{1}$$

with ϕ_0 being a model baseline value (intercept). Equation (1) defines an additive feature attribution. The advantage of feature importance measures of this form is that it provides more detailed information. For instance, comparing inputs x_k with attributions ϕ_k provides the functional form of feature k learned by this model. Furthermore, any local measure can provide global information via aggregation.

We employ two feature importance measures that are model agnostic, unlike other approaches, such as Gini impurity (Friedman, Hastie, and Tibshirani 2009; Kazemitabar et al. 2017), that are only compatible with specific machine learning methods. We argue for the use of *Shapley values* (Shapley 1953; Štrumbelj and Kononenko 2010; Lundberg and Lee 2017), which are local and of the form (1), and contrast them with *permutation importance* (Breiman 2001a; Fisher, Rudin, and Dominici 2019), a simple global measure. A concise technical description of both measures is provided in Appendix B.

The idea of these and other feature importance measures is to either remove the information of the variables of interest or that of the other variables in the model, and then to observe how model outputs change. Permutation importance does this by randomly shuffling the values of a variable to observe how much the performance of the model deteriorates over the test set. Shapley values, on the other hand, explain individual predictions by measuring the contribution that a variable makes on top of others in the model. Shapley values have the advantage that they come with a set of appealing mathematical properties inherited from their game-theoretic origins (Young 1985; Lundberg and Lee 2017). In particular, Shapley values are the only variable attribution scheme which provides accurate local, linear attributions (Equation (1)), respects null contributions, and is consistent. Consistency is a monotonicity property, that is, if a variable is more important in a model compared with another model, then it should also have a larger importance attributed to it. Most popular feature attribution metrics do violate

consistency (Lundberg et al. 2020), which makes Shapley values the preferred importance measure, especially for local attribution. Computing Shapley values is computationally more demanding than computing permutation importance. However, there exist accurate approximations that substantially reduce the computation time of Shapley values which we will investigate as well.

While the primary goal of a model comparison (Step 1 of the workflow) is to identify the most accurate model, it is also informative for the modeler to compare different machine learning methods in their variable importance. Strong disagreements in the importance or functional forms learned by the models can be an indication that the modeling needs to be refined. The better aligned the models are in how they use the predictive variables, the more confident the modeler can be that the models generalize well to the data-generating process.

The information derived from global and local feature importance measures is descriptive. They do not by themselves provide measures of certainty, i.e., an estimate on how certain one can be that a variable is actually important to describe or predict the outcome. This is the realm of statistical testing, e.g., in the form of hypothesis testing.

2.3 Step 3: Shapley Regressions

Linear regression-based models are the workhorse in many applied settings, as they allow for standardized and well-established communication. They achieve that by means of regression coefficients and statistical tests that show whether coefficients are different from zero. The canonical test against the null that there is no effect means that we test if there is a significant alignment between a variable of interest and the target (reject the null). We can ask the same about local attributions coming from the variables components ϕ_k in Equation (1) within a linear regression setting (Joseph 2019),[6]

$$y_t = \phi_0^S + \sum_{k=1}^{N} \phi_k^S(x_t)\beta_k^S + \epsilon^S \quad \text{with} \quad \mathcal{H}_k^0 : \{\beta_k^S \leq 0 \,|\, x_t \in \Omega\}. \quad (2)$$

[6]This approach differs from the previous use of Shapley values in econometrics to analyze multicollinearity (Lipovetsky and Conklin 2001).

Equation (2) is almost identical to a standard linear regression, with two differences.[7] First, the null hypothesis $\mathcal{H}_k^0$ includes negative values. This is because Shapley values absorb the sign of a contribution, such that only significant positive values for β^S mean alignment. Second, inference from Equation (2) is only valid within a region Ω, usually the test set. This is because non-linear machine learning models may show alignment with the target only in bounded regions of the input space. Non-linearity also means that we cannot summarize a variable's importance by a single coefficient universally. We can, however, define something akin to a linear regression coefficient within a region Ω. Let *sign* be the sign of the coefficients when regressing y on x, and let $\psi_k^t = \frac{|\phi_k^S(x_t))|}{\sum_{l=1}^N |\phi_l^S(x_t)|}$ be the share of absolute Shapley values of observation t attributed to variable k. The average share of Shapley values across all observations in Ω is denoted by $\bar{\psi}_k = \frac{1}{|\Omega|}\sum_{x_t \in \Omega} \psi_k^t$.[8] Further, let $(*)$ indicate the confidence level with which we can reject $\mathcal{H}_k^0$ according to Equation (2); then we define the following:

$$\textit{Shapley share coefficient:} \qquad \Gamma_k \equiv sign\ \bar{\psi}_k^{(*)} \quad \in \quad [-1, 1]\,. \qquad (3)$$

Shapley share coefficients can be communicated as commonly used regression coefficients in a well-known table form. The interpretation of Γ_k is also similar to that of a regression coefficient, as it measures strength and confidence in alignment with the target variable. However, it cannot be interpreted as a marginal effect, unless the model is linear. In this case, Shapley share coefficients are aligned with the actual linear regression coefficients.

3. Data and Experimental Setup

We describe the notation, data, and experimental procedure for the macroeconomic forecasting exercise which we use to demonstrate the proposed machine learning workflow.

[7]Equation (2) is based on generated regressors (Pagan 1984). The validity of inference and asymptotic properties of estimating the β^S are discussed in detail in Joseph (2019).

[8]Shapley values do not have a natural scale on which to represent them and they can change alongside the region Ω being considered. This motivates the normalizing denominator in the definition of ψ_k^t.

We first introduce the necessary notation. Let y and $\hat{y} \in \mathbb{R}^T$ be the observed and predicted outcome, respectively, where T is the number of observations in the time series. The feature matrix[9] is denoted by $x \in \mathbb{R}^{T \times N}$, where N is the number of features in the data set. The feature vector of observation t is denoted by x_t. Generally, we use t to index the point in time of the observation and k to index features. The forecasting horizon in months is denoted by h. The forecasting horizon is a crucial aspect regarding the purpose of a forecast. One does not necessarily expect models to perform equally well or to pick up the same information across horizons. The default discussed in this paper is the one-year forecast ($h = 12$), sitting between short- and medium-term projections.

3.1 Data

We use the FRED-MD macroeconomic database (McCracken and Ng 2016), which contains monthly macroeconomic indicators for the United States. Our vintage of the data goes from 1959 to 2019. Our forecast target is changes in unemployment, and we hand-pick nine variables as predictors in our baseline approach, each capturing a different macroeconomic channel. We use the stationarity transformations suggested by the authors of the data set that include first differences ($\Delta^l(x) = x_t - x_{t-l}$), log differences ($\Delta^l log(x)$), and second-order log differences ($\Delta^l log(x_t) - \Delta^l log(x_{t-l})$). Given that we predict the yearly change of unemployment, we set transformation span l to 12 for the outcome and lagged outcome (predictor) variables. For the remaining predictors, we set $l = 3$ in our baseline setup. Table 1 shows the variables, with the respective transformations and the series names in the original database. The augmented Dickey-Fuller test confirms that all transformed series are stationary ($p < 0.01$). Different choices for handling the data, like choosing l as well as the aspects discussed in Section 2.1, are investigated in detail in Section 5.

[9]Features or predictors in the machine learning literature correspond to independent variables, or just variables. The observed response, outcome, or dependent variable is often referred to as the target.

Table 1. Series Used in Our Baseline Forecasting Experiment

Variable	Transformation	Name in the FRED-MD Database
Unemployment	Changes	UNRATE
Three-Month Treasury Bill	Changes	TB3MS
Real Personal Income	Log Changes	RPI
Industrial Production	Log Changes	INDPRO
Consumption	Log Changes	DPCERA3M086SBEA
S&P 500	Log Changes	S&P 500
Business Loans	Second-Order Log Changes	BUSLOANS
CPI	Second-Order Log Changes	CPIAUCSL
Oil Price	Second-Order Log Changes	OILPRICEx
M2 Money	Second-Order Log Changes	M2SL

Note: The middle column shows the transformations suggested by the authors of the FRED-MD database, the right column shows how the series are named in that database. Target audience: analysts.

3.2 Models

We test two types of models, a simple autoregressive model and several full-information models containing lags of the response variable and the features. The latter group is further split into linear regression models, with and without penalization, and non-linear machine learning models.

- The *autoregressive model (AR)* uses lagged values of the response variable as predictors: $\hat{y}_t = \alpha + \sum_{l=1}^{L} \theta_l y_{t-l}$. We test AR models of lag lengths $1 \leq L \leq 12$, where we chose L using the Akaike information criterion in the training set. We also test a simple AR1 model by setting $L = 1$. The model's fitted coefficients are given by $\theta \in \mathbb{R}^L$. Forecasts over a horizon h are obtained iteratively from $\hat{y}_{t+h} = \alpha + \sum_{l=1}^{L} \theta_l \hat{y}_{t+h-l}$.
- The *full-information models* use the h-month lag of the outcome variable and the other features as independent variables: $\hat{y}_t = f(y_{t-h}, x_{t-h}; \theta)$, where f is any given predictive model. For example, if f is a linear model, a horizon-h projection takes the form $\hat{y}_t = \alpha + \theta_0 y_{t-h} + \sum_{k=1}^{N} \theta_k x_{t-h,k} + \epsilon_t$, with ϵ_t being the error term. To simplify notation in what follows, we include the lagged outcome in the feature matrix x. We test seven full-information models: ordinary least squares (OLS) regression, regularized regression with ridge and lasso penality, and four machine learning models: random forests (Breiman 2001a), gradient boosting (Friedman 2001), support vector regression (Drucker et al. 1997), and artificial neural networks (Goodfellow, Bengio, and Courville 2016). Table B.1 in Appendix B provides details on the implementation of the models.

3.3 Experimental Procedure

We evaluate how all models predict $l = 12$ month changes in unemployment $h = 12$ months ahead in a pseudo out-of-sample setting with an expanding horizon. All methods are evaluated on the 359 data points of the forecasts between January 1990 and November 2019 using an expanding window approach. We choose the absolute error as our performance metric. It is easy to interpret and less

sensitive to outliers than the squared error. Accordingly, we pick the hyperparameter values that minimize absolute error.[10]

We fit, i.e., train, the AR models every month. The full-information models are trained every 12 months such that each model makes 12 predictions before it is updated. Compared with updating the models every month, this reduces the computational cost considerably, while only minimally affecting model performance in normal times. However, one may refit a particular model more frequently during operation. Especially machine learning models can be quick in picking up different or new (economic) regimes, as we will see below.

As the models predict changes h months ahead, we have to create an initial gap between training and test set when making predictions to avoid a look-ahead bias. For a model trained on observations $1 \dots t$, the earliest observation in the test set that provides a pseudo real-time h-month forecast is $t + h$. For observations $t + 1, \dots, t + h - 1$, the time difference from the last observation in the training set t is less than one year.

Most of the models we use can be affected by outliers. We therefore test how winsorization of the features at the 1st and 99th percentile affects the predictive performance. We do not winsorize the response variable and the lagged response that is used as a feature.

All machine learning models that we test have hyperparameters which need calibration. We use two types of cross-validation for the hyperparameter tuning. First, we employ ordinary fivefold cross-validation (see Chakraborty and Joseph 2017), which does not consider temporal dependencies in the data, but randomly assigns the observations in the training set to five folds. Second, we use fivefold block cross-validation (Snijders 1988; Bergmeir and Benítez 2012), where the folds are assigned to five consecutive blocks. This approach respects the temporal dependency of the training and test data. More concretely, we use *hv-block cross-validation* (Racine 2000), which additionally introduces a gap of 12 months between blocks of the training and test set. We employ a random search across 100 hyperparameter combinations and pick the hyperparameters that minimize the mean absolute error.

[10]Note however, that different models have different loss functions. Minimizing these is not necessarily equivalent to minimizing the absolute error.

As the hyperparameter optimization is computationally expensive, we conduct it only every 36 months. Even during operation, it is unlikely that hyperparameters need to be updated with a higher frequency unless one expects dramatic model changes, e.g., due to a large change in the data-generating process. Smaller changes will be reflected in the changes in the model parameters (e.g., the weights of the neural network) rather than hyperparameters (e.g., the architecture of the neural network).

To increase the stability of the full-information models, we train each model 30 times on different bootstrapped samples of the training set and average their predictions. This bootstrap aggregation approach is also referred to as *bagging* in the literature (Breiman 1996). Each of the 30 models uses the same hyperparameters, which are calibrated on the full training set.

To estimate how stable our models are, we repeat each experiment—including the training of the 30 bootstrapped models every 12 months and the hyperparmaeter search every 36 months—10 times for all methods, each time with a different random seed. Sources of randomness that can lead to performance differences between these 10 iterations are the chosen hyperparameters,[11] the random bootstrap samples, and random initializations of the random forest, gradient boosting, and neural networks.

4. Workflow Output

This section presents the results of our workflow when applying it to the baseline setting for forecasting changes in the U.S. unemployment rate on a one-year horizon as described in the previous two sections. Not all results presented here are meant to be communicated during operation. Rather, we also present additional analyses that are only relevant for the technical expert that is developing a forecasting model or that are shown for illustration purposes to help the reader better understand the technicalities of the workflow.

[11]Both the randomly selected hyperparameter combinations and the random assignments to folds when using k-fold cross-validation (folds in hv-block cross-validation are not randomly assigned) can induce randomness.

4.1 Step 1: Model Performance

Table 2 summarizes the empirical performance of the different forecasting models. For this table and the following analyses, we applied winsorization to all models and used hv-block cross-validation for the hyperparameter search. Further, to obtain more stable predictions, we average the predictions of 10 models, each trained with a different random seed. In the table, the models are ordered by decreasing mean absolute error over the whole test period between 1990 and 2019.

The table also breaks down the performance in three periods: the 1990s and the periods before and after the global financial crisis (GFC, September 2008). The best model in the individual periods is highlighted in bold. We statistically compare the error of the best model in each period against all other models using a Diebold-Mariano test.[12]

All machine learning models outperform the linear models on the whole sample. In the 1990s and the periods before the global financial crisis, the difference in performance between the models is rather small compared with the period after the crisis. This is indicative that machine learning models may be particularly suited for detecting regime shifts or the modeling of non-linearities, both aspects we will investigate in more detail.

The simple AR_1 model performs better than the AR_{12} model and the linear full-information models. Ridge and lasso regression perform very similarly, both outperforming the OLS regression. In the following analyses we will only consider one AR model, the AR_1, and will focus on one regression model, the ridge regression.

Apart from the aggregated performance across the test period, it is informative to look at the models' individual predictions. Figure 1 (top panel) shows the observed response variable and the predictions of gradient boosting, ridge regression, and the AR_1 model. The bottom panel shows the prediction error. The vertical lines indicate

[12]The horizon of the Diebold-Mariano test is set to 1 for all tests. Note however, that the horizon of the AR models is 12 so that the p-values for this comparison are biased. Setting the horizon of the Diebold-Mariano test to 12, we do not observe significant differences between the absolute error of the best model and the AR.

Table 2. Forecasting Performance for the Different Prediction Models in the Baseline Setup

Time Period	**01/1990–11/2019**	**01/1990–12/1999**	**01/2000–08/2008**	**09/2008–11/2019**
Gradient Boosting	**0.559** —	**0.460** —	**0.466** —	0.718 (0.353)
SVR	0.565 (0.323)	0.470 (0.328)	0.489 (0.219)	**0.709** —
Forest	0.581 (0.018)	0.472 (0.240)	0.471 (0.413)	0.762 (0.005)
Neural Network	0.589 (0.009)	0.468 (0.336)	0.503 (0.070)	0.762 (0.001)
AR_1	0.608 (0.063)	0.472 (0.382)	0.503 (0.216)	0.811 (0.064)
AR_{12}	0.626 (0.001)	0.543 (0.011)	0.482 (0.356)	0.810 (0.001)
Lasso Regression	0.637 (0.000)	0.498 (0.061)	0.474 (0.378)	0.886 (0.000)
Ridge Regression	0.639 (0.000)	0.497 (0.065)	0.481 (0.272)	0.886 (0.000)
OLS Regressions	0.648 (0.000)	0.516 (0.016)	0.508 (0.053)	0.872 (0.000)

Note: The models are ordered by decreasing MAE on the whole sample. The best-performing model in each time period is highlighted in bold. The p-values in parentheses indicate the statistical significance (one-sided) of the Diebold-Mariano test comparing the best model in each column with the other models. Target audience: management to decisionmakers.

Figure 1. Comparison of Observed and Predicted Outcome

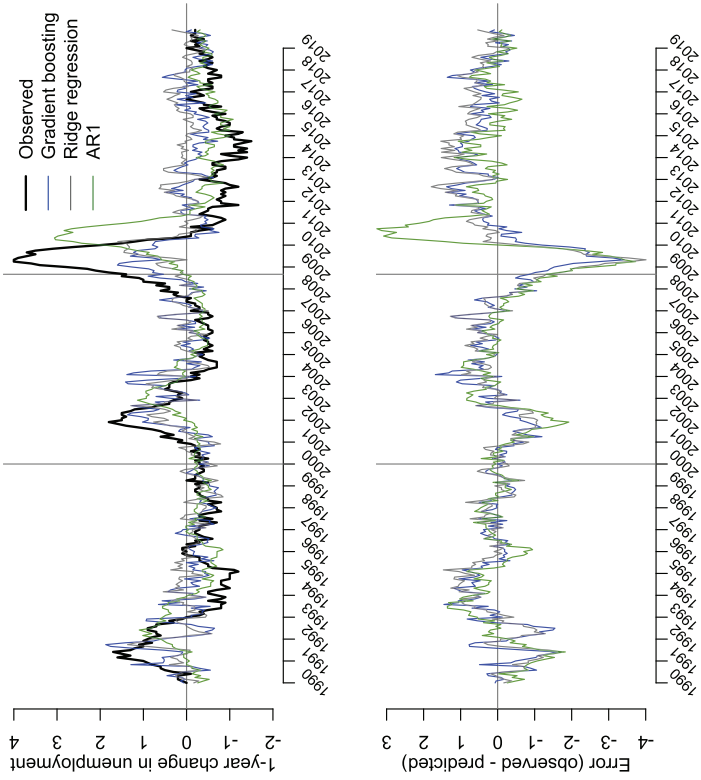

Note: The top panel shows the observed one-year change in unemployment and the predictions by the gradient boosting model, ridge regression, and the AR_1. The bottom panel shows the error of these two methods. Target audience: management or decisionmakers.

the different time periods distinguished in Table 2. All three models underestimate unemployment growth during the global financial crisis and overestimate it during the recovery. However, the gradient boosting model is least biased in those periods and forecasts the increase in unemployment earlier during the crisis. A similar observation can be made after the burst of the dot-com bubble in the early 2000s. Such a chart can be presented to the policymaker to convey the model's performance in a clear and detailed way.

4.2 Step 2: Feature Importance

We explain the predictions of the machine learning models and the linear regression as calibrated in our baseline setup. Our focus is largely on explaining forecast predictions in a pseudo real-time setting. However, in some cases it can be instructive to explain the predictions of a model that was trained on observations across the whole time period. For that, we exploit the fact that we trained the models on 30 different bootstrapped samples across the whole time series. Each of these models can make predictions on those observations not in the bootstrapped training sample. In this way we obtain several predictions for each observation in the time series, which are then averaged. This *out-of-bag* analysis is subject to look-ahead bias, as we use future data to predict the past.

We first analyze our two methods of model interpretation at a global level. Figure 2 compares Shapley shares $|\Gamma^S|$ (left panel) with permutation importance (middle panel). The variables are sorted by average Shapley shares of the four machine learning models. Vertical lines connect the lowest and highest share across models for each feature to highlight the disagreement between models.

Shapley values and permutation importance do not agree in their ranking of feature importance. For instance, using a random forest model, the three-month Treasury bill seems to be a more important indicator according to permutation importance than according to Shapley calculations.

The permutation importance is a measure of a feature's influence on the accuracy of the model and is affected by how the relationship between outcome and features changes over time. In contrast, Shapley values reflect a variable's influence on the predicted value, independent of that value's accuracy. Arguably this

Figure 2. Variable Importance According to Different Measures

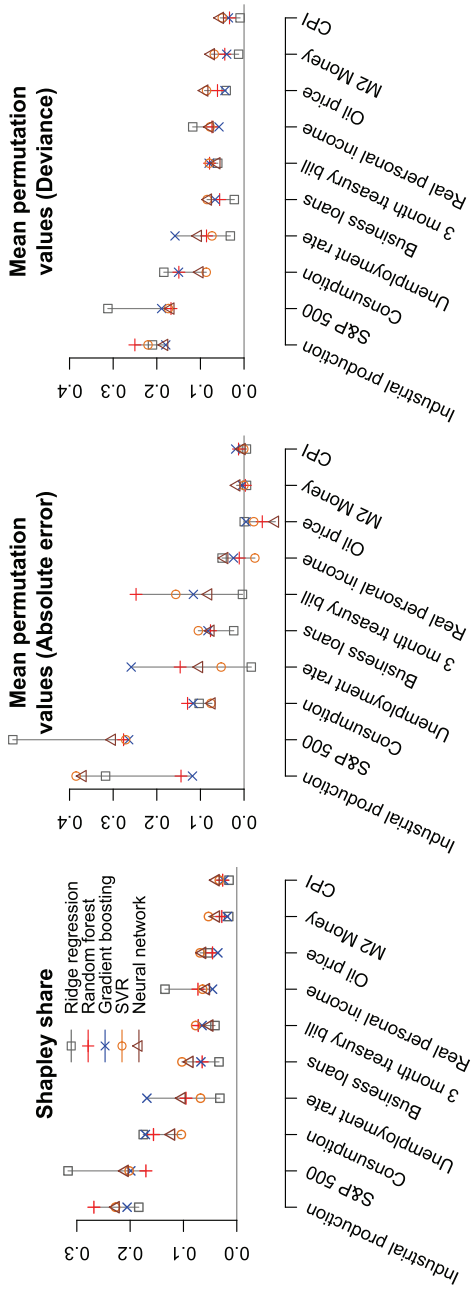

Note: The left panel shows the importance according to the Shapley shares $|\Gamma^S|$ and the middle panel shows the variable importance according to permutation importance. The right panel shows an altered metric of permutation importance that measures the effect of permutation on the predicted value rather than prediction error. Target audience for the left panel: decisionmakers. Permutation importance is shown for illustrative purposes.

measure of importance is more useful in a forecasting setting when the variable importance should be computed for data points for which the true outcome has not been observed yet, which means that permutation importance is not computable. The right panel of Figure 2 shows an altered measure of permutation importance. Instead of measuring the change in the error due to permutations, we measure the change in the predicted value.[13] We see that this importance measure is more closely aligned with Shapley values. Further, when we evaluate the error-based permutation importance metric using predictions based on the out-of-bag analysis, we find a strong alignment with Shapley values (not shown), as the relationship between variables is not affected by the changes between the training and test set.[14]

Overall, the different prediction models have a similar importance ranking of the features according to the Shapley share. There are, however some notable differences—especially the ridge regression model often differs substantially from the other models in the Shapley shares. Even the different machine learning models do not completely agree on the relative importance of features. For example, gradient boosting gives more importance to the lagged unemployment indicator than the other methods.

While the computation of Shapley values is technically rather complex and difficult to communicate to a non-technical audience, we believe that the intuition behind Shapley values as the contribution to the model's predictions is easy to understand. Thus, a chart such as the left panel of Figure 2—but only showing the best-performing model—can be communicated to decisionmakers.

This global analysis only conveys which variables are important across all observations in the test set. Local attributions will often be more useful in a pratical setting, as they allow to assess individual, e.g., the latest, predictions.

[13]This metric computes the mean absolute difference between the observed predicted values and the predicted values after permuting feature k m times: $\frac{1}{m}\sum_{i=1}^{m}|\hat{y}_i - \hat{y}_{i(k)}^{perm}|$. The higher this difference, the higher the importance of the feature k (see Lemaire, Féraud, and Voisine 2008 and Robnik-Šikonja and Kononenko 2008 for similar approaches to measure variable importance).

[14]Comparing out-of-sample and out-of-bag measures allows to evaluate model drift and look-ahead bias more generally.

Local attributions also reveal the functional form learned by a model. To illustrate this, we consider the out-of-bag predictions, abstracting away from model drift which we discuss in a moment. Here, the most accurate models are the gradient boosting model (absolute error of 0.431) followed by the random forest (0.450), the SVR (0.452), the neural network (0.452), and ridge regression (0.584).[15]

Figure 3 shows the functional forms that the machine learning models have learned from the most important features according to Shapley shares shown in Figure 2 (left panel). It depicts local Shapley values against the observed input values (horizontal axis), with rows showing the variables and columns the different models. The approximate functional form learned by each model for each feature is traced out by a best-fit third-degree polynomial. Although the four machine learning methods use very different learning mechanisms and even do not agree perfectly on the global importance of features, the functional forms learned by all of them are highly consistent for all variables shown. This gives us confidence that the functions learned are meaningful and robust.

For example, consider the S&P 500. The ridge regression learns a steep negative slope with higher stock market values being associated with lower unemployment one year ahead. This makes economic sense. However, we can make more nuanced statements when looking at the other models. There is an asymmetry between market increases and decreases. While large decreases suggest large increases in unemployment down the line, there is a saturation effect for high market valuations with only a small expected decrease in unemployment.

For unemployment, all machine learning models learn a quadratic function. A high increase in unemployment makes future increases in unemployment less likely compared with a medium increase. For business loans we also observe a quadratic function, where very low and high loans lead to a positive predicted change in unemployment. In contrast, the linear model cannot model quadratic trends, so it is

[15] As in the forecasting setting, winsorization is applied, as it helps the performance of the SVR (see Section 5), while it does not substantially affect the other models' performance. In the out-of-bag analysis, we use k-fold cross-validation rather than blocked cross-validation, as this generally improves the performance.

Figure 3. Functional Forms Learned by Different Models for Five Features with the Highest Average Shapley Share

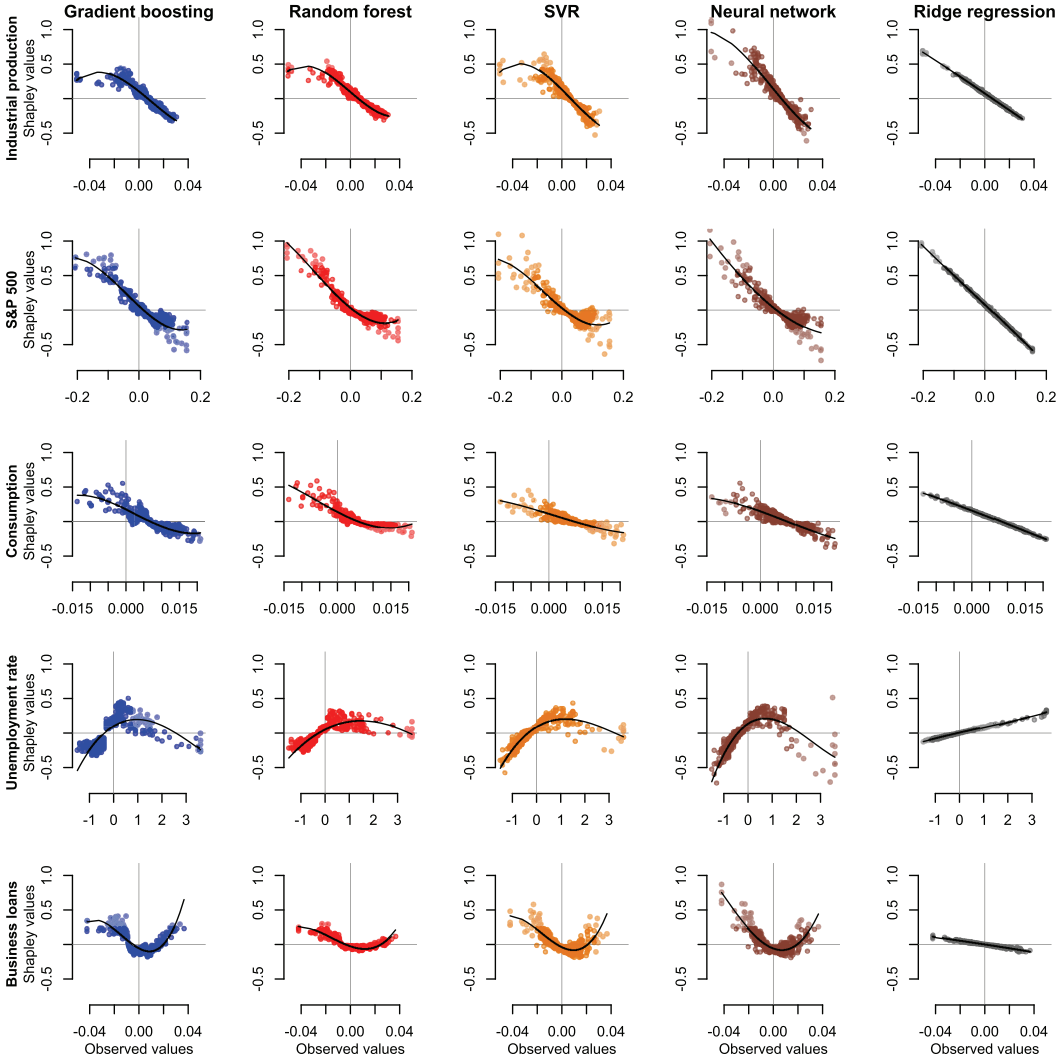

Note: The lines show a third-degree polynomial fitted to the data. The Shapley values are computed on the out-of-bag predictions and are therefore subject to look-ahead bias. Target audience: analysts (comparison); decisionmakers (single model if robust).

not surprising that the Shapley share of these two variables (Figure 2, left panel) are substantially smaller according to the linear model compared with the machine learning models.

Figure 3 serves several purposes. The functional forms learned from the data, although not causal, might provide new economic insights about the underlying processes. These can then be further investigated by, for example, using structural models. Further, by providing information about the inner workings of different models, these charts can be used as a diagnostic tool for the technical expert training and tuning the model. For instance, if the functional forms learned by an SVR are mostly linear whereas those of the other machine learning models are not, this might suggest a problem constraining the flexibility of the SVR. Finally, by evaluating the functional forms learned at different points in time, model drift or structural breaks can be detected.

For example, we consider the out-of-bag predictions of the models trained up to three different points in time. Figure 4 shows the functional form for the lagged unemployment change variable. The ridge regression model (left panel) trained up to the periods 2000 and 2008 finds no predictive power for lagged unemployment. It is only after the onset of the GFC that the regression learns a positive relationship—an increase of unemployment increases the predicted unemployment change one year ahead. However, this is simply reflective of the trend—the one-year unemployment change was high for a prolonged period following the financial crisis: it was persistently greater than or equal to 1 percentage point for 23 consecutive months (May 2008–March 2010). In contrast, the functional form of the gradient boosting model (right panel) is rather stable. Across the three time periods it learns a non-monotonic relationship where high absolute values in the unemployment make future increases in unemployment less likely compared with small changes. The scale of this learned functional form increased after the GFC in line with larger movements in the unemployment rate during this time.

To better understand the non-monotonic function of lagged unemployment change learned by the gradient boosting model, we look into the role of recessions within the model.[16] Figure 5 (left)

[16]We use the definition of recessions provided by the Federal Reserve Bank of St. Louis (Federal Reserve Bank of St. Louis 2020).

Figure 4. Functional Form of Lagged Unemployment Change Learned by Ridge Regression (left panel) and Gradient Boosting (right panel) for Three Models Trained Up to Different Points in Time

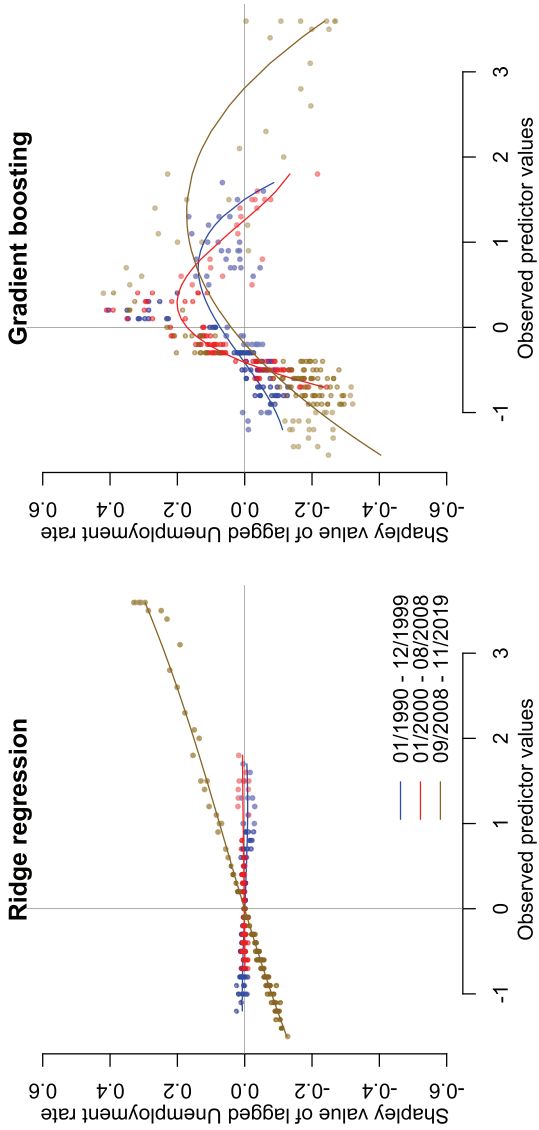

Note: The lines show third-degree polynomials fitted to the data. Target audience: analysts.

Figure 5. Interaction between Unemployment Changes and Recessions as Learned by Gradient Boosting Model

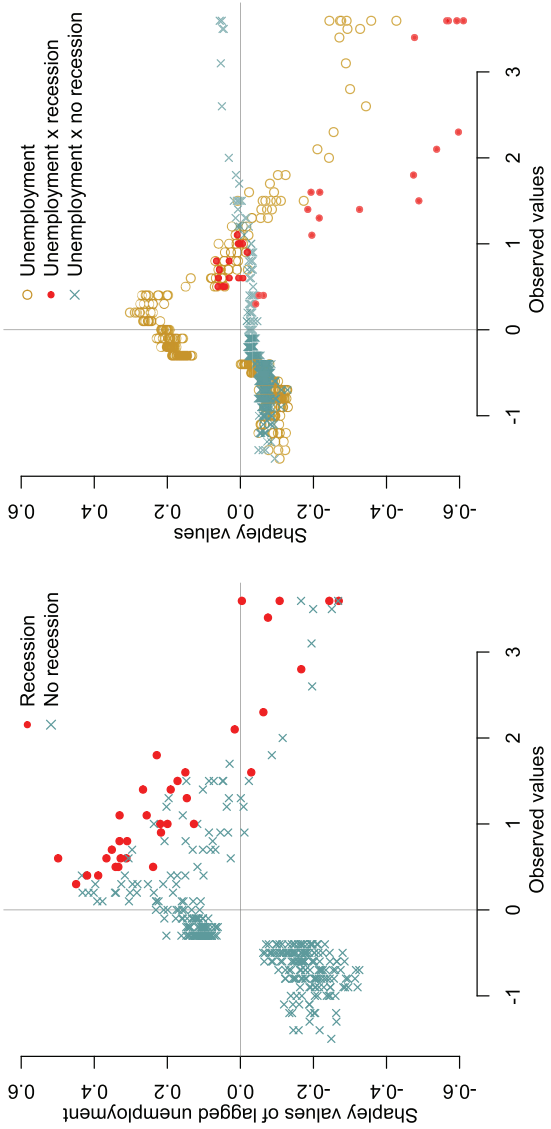

Note: The left panel shows the functional form of lagged unemployment changes when the model is trained on the baseline features without a recession dummy (as in Figure 4). The right panel includes the Shapley values of the interactions when the model was trained with a recession indicator. Target audience: analysts (comparison); decisionmakers (recession effect only, right-hand side).

again shows the functional form of lagged unemployment as learned by the gradient boosting model in the out-of-bag setup. But now recession observations (also lagged by a year) in the input space are highlighted in red. Even though we did not include recessions explicitly as an indicator the model could learn from, these periods account for a large share of the downward-sloping part on the right-hand side. This makes economic sense as larger current increases of unemployment during recessions make future decreases of unemployment during the recovery period more likely.

We further elaborate on this observation by including a lagged recession dummy in our models and compute the Shapley-Taylor index (Agarwal, Dhamdhere, and Sundararajan 2019) to decompose the predictions into the main effects from past unemployment, the recession dummy, and their interaction.[17]

The Shapley values of this interaction as well as the main effect of lagged unemployment are shown in the right panel of Figure 5. The main effect of lagged unemployment still shows the inverted U-shaped form—even after controlling for interactions with all other variables. The Shapley values of the interaction show that during a recession thereis an additional strong negative effect of lagged unemployment on the prediction for larger input values (red), which is in line with the above "reversal-to-the-mean" explanation.

While including the recession indicator improves the interpretation of the results and the interaction with unemployment has high Shapley values that contribute substantially to the prediction, the predictive accuracy of the gradient boosting model does not increase meaningfully. Adding the recession indicator, the forecast error only slightly decreases from 0.559 to 0.554. This suggests that the model learned the role of recession periods implicitly, incorporating two different regimes, normal times and recessions.

4.3 Step 3: Statistical Inference with Shapley Regressions

Shapley value-based inference (Equation (2)) allows us to communicate machine learning models analogously to a linear regression analysis. In Table 3, we present the Shapley regression for the full

[17]We follow Joseph (2019) in his empirical approach and group all remaining variables into a single "other" variable to reduce computation time. We compute the Shapley Taylor expansion to the third order (see Section 6) such that two-way interaction terms are unbiased.

Table 3. Shapley Regression of Gradient Boosting Mode (left) and the Ridge Regression (right) for the Forecasting Predictions between 1990 and 2019

	Gradient Boosting			Ridge Regression		
	β^S	p-value	Γ^S	β^S	p-value	Γ^S
Industrial Production	1.132	0.000	–0.217***	2.280	0.000	–0.185***
S&P 500	0.942	0.000	–0.191***	0.907	0.000	–0.317***
Consumption	1.103	0.000	–0.177***	0.966	0.012	–0.173**
Unemployment	1.443	0.000	+0.175***	9.789	0.000	+0.031***
Business Loans	3.086	0.000	–0.066***	5.615	0.006	–0.035***
Three-Month Treasury Bill	4.273	0.000	–0.062***	–6.816	1.000	–0.042
Personal Income	–0.394	0.682	+0.04	–0.658	0.870	+0.138
Oil Price	0.298	0.387	–0.035	–2.256	0.973	–0.055
CPI	0.272	0.438	+0.021	–4.294	0.875	+0.014
M2 Money	–8.468	1.000	–0.016	–18.545	0.994	–0.009

Note: Significance levels: *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$. Target audience: analysts (comparison); decisionmaker (left-hand side).

out-of-sample forecasting period between 1990 and 2019 based on the predictions of the gradient boosting model. For illustrative purposes, the table also shows the Shapley regression for the ridge regression.

As mentioned above, the coefficients β^S measure the alignment of a variable with the target. Values close to one indicate perfect alignment and convergence of the learning process. Values larger than one indicate that a model underestimates the effect of a variable on the outcome. And the opposite is the case for values smaller than one. This can intuitively be understood as the model hyperplane of the Shapley regression either tilting more towards a Shapley component from a variable (underestimation, $\beta_k^S > 1$) or away from it (overestimation, $\beta_k^S < 1$). Significance decreases as the β_k^S approaches zero.

Variables with higher Shapley shares $\left|\Gamma^S\right|$ (same as in Figure 2) tend to have lower p-values. This is intuitive, demonstrating that the model learns to rely more on features that are important for predicting the target variable. However, this does not hold by construction, especially not in a forecasting setting where the relationships between variables change over time. Any statistical significance may disappear in the test set—even for features with high Shapley shares.

One more variable is statistically significant for the gradient boosting method than for the linear model. This is expected given the greater flexibility of machine learning models. It also provides further evidence of how non-parametric methods, like gradient boosting forests, exploit non-linear relationships that linear regression cannot account for (as in Figure 3).

A Shapley regression table can provide meaningful insights for decisionmakers that are acquainted with standard statistical inference for regression. Further, it can help the technical expert to refine the model—for example, by removing variables with negative coefficients—or adjust the period of analysis until the coefficients align better with the target.

5. Robustness

We consider a wide array of alternative choices made during our baseline analysis and how these affect the outputs from the first two steps of our workflow, model performance and Shapley feature

importance.[18] We first propose and run a set of analyses that vary key parameters of our experimental setup to test whether our results that machine learning models outperform linear models is robust. Next, we replicate our results using real-time data and show how our workflow can be applied to substantially larger sets of predictors. Finally, we investigate the robustness of the estimation of Shapley values, discussing computationally cheap approximations.

5.1 Experimental Setup

In our main analysis we used a bagging ensemble of 30 models for each of the full-information methods. We show in Figure A.1 (Appendix A) that only the gradient boosting model and the neural network improve using bagging. The figure shows the mean absolute error across the 10 iterations (based on different seeds) as a function of the bagging ensemble size. The confidence intervals ($\pm$ 1.96 standard errors of the mean) of the gradient boosting and the neural network decrease visibly when increasing the size of the bagging ensembles, suggesting that bagging makes these models more stable and thus less sensitive to the random seed. In the following, we use bagging only for these two methods to save computation time.

Further, we used blocked cross-validation for the hyperparameter search and have averaged the predictions of 10 models, each trained on a different random seed. Figure 6 (left panel) investigates the impact of these choices on model performance. It shows the mean absolute error (MAE) across the whole test period between 1990 and 2019 and conveys several findings.

First, the machine learning models, especially the neural network and the SVR, show a substantial variance in performance (smaller transparent dots) for the 10 different iterations based on different random seeds. Averaging the predictions across these iterations (bigger non-transparent points) tends to produce more accurate predictions than the average individual model. This variance in the error across the 10 iterations reflects substantial differences in the predictions on individual data points.

[18]Investigating changes in statistical alignment of feature components (step 3 of the workflow) can be interesting during practical applications, but does not add much value to the discussion here, we believe.

Figure 6. Robustness of Predictions

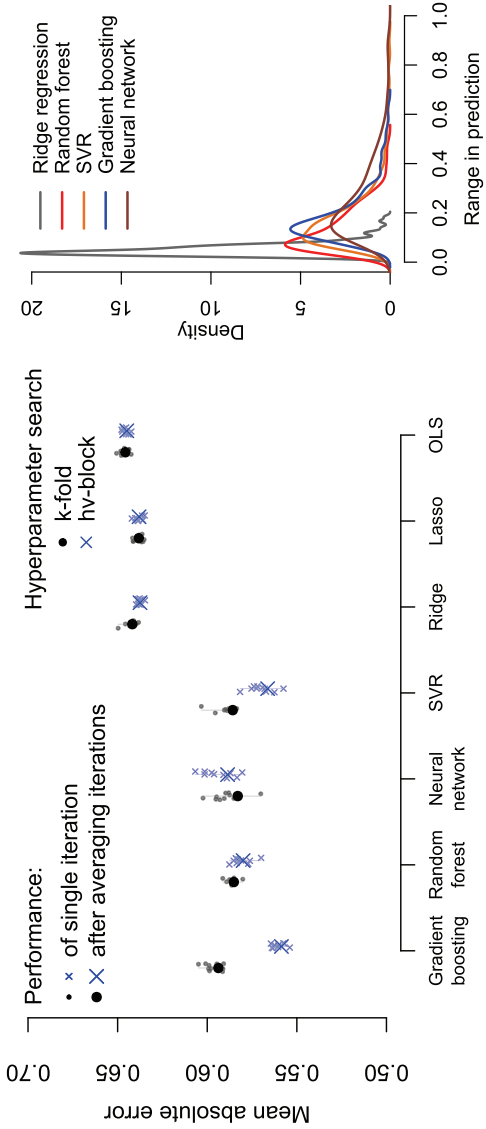

Note: The left panel shows the MAE of our main prediction models. The small markers show the performance of the individual interations, each based on a different random seed. The larger markers show the average performance across these 10 iterations. For each model we test two types of hyperparameter searches (k-fold versus hv-block cross-validation). The OLS does not have hyperparameters and is not affected by this test. The right panel shows the variation in the predicted values across the 10 iterations. For each observation in the test period (1990–2019), we measure the range of predicted values and show the distribution of this measure in the chart. Target audience: analysts.

To investigate this further we measure, for each observation in the test set, the range of predicted values across the 10 iterations. The right panel of Figure 6 plots the distribution of this range across observations. The 10 models are very similar for the ridge regression, with a mean range of 0.05 (90 percent percentile: $P_{90} = 0.08$). The random forest is less stable, with a mean range of 0.14 ($P_{90} = 0.26$) but a factor of two more stable than the neural network with a mean range of 0.27 ($P_{90} = 0.5$). This is—given a mean absolute error of less than 0.6—a substantial variation in the prediction of the models.

In a practical forecasting setting, the modeler might decide to slightly trade off predictive performance against model stability and choose, for example, the random forest over the SVR. We believe that the repeated training of the same model with different random seeds is crucial to get a sense of the stability of their performance. To stabilize the models and make them less susceptible to the random seeds, we suggest averaging them.

Second, the type of cross-validation employed in the hyperparameter search matters for the performance of some of the methods. The linear models, the random forest, and the neural network do not differ markedly in their performance for the two types of cross-validation. However, for the gradient boosting model and the support vector regression, we observe a substantially better performance when using the blocked cross-validation approach.

Even rather small design factors such as the type of cross-validation can change the conclusion about which model performs best. The fact that this and other factors (see below) affect the performance of the models in different ways suggests that the modeler should conduct a extensive set of experiments before identifying the best prediction model, and also assure its stability.

We next alter several parameters with respect to our baseline setup. The results are shown in Table 4, with the best model in each row highlighted in bold.

- *Prediction Horizon:* In the baseline setup, we have predicted unemployment changes $h = 12$ months ahead. Here, we alter the prediction horizon between 1 and 36 months. We observe that the AR_1 models competes well with the full-information models at prediction horizons 1, 3, and 6 months but falls

Table 4. Performance for Different Parameter Specifications

	Gradient Boosting	SVR	Random Forest	Neural Network	Ridge Regression	AR_1
Prediction Horizon h (lag between response and predictors in months)						
1	0.20	0.19	**0.17**	0.18	0.18	**0.17**
3	0.28	0.28	**0.27**	**0.27**	**0.27**	**0.27**
6	0.41	0.41	**0.39**	0.42	0.43	0.41
12 (Baseline)	**0.56**	0.57	0.58	0.59	0.64	0.61
24	0.68	0.67	**0.62**	0.69	0.73	0.79
36	0.64	0.63	**0.61**	0.72	0.72	0.80
Training Set Size (in months)						
60	0.83	0.87	**0.79**	0.84	0.87	0.95
120	0.63	0.67	**0.57**	0.66	0.66	0.71
240	0.58	**0.56**	0.57	0.58	0.61	0.67
360	**0.57**	0.58	0.58	0.60	0.61	0.64
480	**0.56**	0.57	0.57	0.57	0.63	0.62
Max (Baseline)	**0.56**	0.57	0.58	0.59	0.64	0.61
Transformation Span l (in months)						
1	0.57	0.60	**0.55**	0.59	0.64	—
3 (Baseline)	**0.56**	0.57	0.58	0.59	0.64	—
6	**0.60**	**0.60**	**0.60**	0.67	0.66	—
9	**0.65**	0.68	0.67	0.70	0.70	—
12	0.68	0.74	0.70	0.71	0.74	**0.61**
Winsorization at 1 Percent and 99 Percent						
Yes (Baseline	**0.56**	0.57	0.58	0.59	0.64	0.61
No	**0.56**	0.59	0.58	0.60	0.64	0.61

Note: The shown metric is mean absolute error. The best model(s) in each row are highlighted in bold. Target audience: management.

behind when increasing the horizon. This is not surprising, as the autocorrelation in the response variables decreases with increasing h. The table shows that the machine learning models can provide meaningful signals for the unemployment changes at longer horizons, even three years ahead. The good performance of the random forest is notable: For all horizons different from 12, it performs as well as or better than the other models.

- *Window Size:* In the baseline setup, the training set grows over time (expanding window). This can potentially improve the performance, as more observations may facilitate a better approximation of the data-generating process. On the other hand, it may make the model sluggish and prevents quick adoption to structural changes. To differentiate between these two cases, we test sliding windows of 60 to 480 months. All methods perform worst on the smallest horizons of 60, and only the random forest performs well on a sample of just 120 months. Gradient boosting consistently improves its performance with a growing sample size. This is not surprising for machine learning models, as they can learn different regimes for different time periods due to their flexibility and exploit them for prediction. For instance, different paths down a tree model, or different trees in a forest, are all different submodels. By contrast, the ridge regression, like all linear models, cannot adjust in this way and needs to fit the best hyperplane to the current situation. This can explain why its performance declines for sample sizes larger than 360.
- *Transformation Span:* We use $l = 3$ months in the baseline, when calculating first differences, log differences, and second-order log differences of the predictors (see Table 1). Testing lag lengths of 1, 6, 9, and 12 months, we find that shorter horizons of 1 or 3 months generally lead to better performance than longer ones. This is useful from a practical point of view, as quarterly changes are commonly used for short-term economic projections.
- *Winsorization:* Winsorization only helps the SVR and the neural network. It does not have a visible impact on the performance of the other models. As the response variable is not

winsorized, there is, by design, no effect on the performance of the AR_1 model.

Testing different training set sizes, transformation horizons, and winsorization of predictors are crucial to refine and improve the prediction models. The choice of the prediction horizons will be informed by the needs of the decisionmakers. But testing different horizons can help to assess the change in predictability of the response and by explaining predictions (see Section 4.2), one can detect differential signals provided by the predictors at different horizons.

5.2 Real-Time Data

Our pseudo out-of-sample forecasting approach does not reflect how forecasts are made in the real world. When training and testing our models in Section 4.1, we used revised, macroeconomic data. In a practical setting, we have to rely on early vintages that are likely to be revised. We investigate whether the results change substantially when replicating the horse race using real-time data.

There exist monthly vintages of the FRED-MD database[19] starting from August 1999, each providing estimates for the indicators lagging one month behind.[20]

As before, we predict the change in unemployment one year ahead, this time for the period for which we can produce real-time forecasts (August 2000–November 2019). As the real-time data are delayed by a month, an actual one-year-ahead forecast requires lagging the variables by 13 months. More formally, we use the data (features and response) of the vintage at time t, which contains the measurements up to date $t-1$. We train the models with response

[19]https://research.stlouisfed.org/econ/mccracken/fred-databases/.

[20]The *consumption* variable is not included in the vintages before 2004. When using these vintages for training, we use the revised time series from our baseline data set. Further, the variables *business loans* and *real personal income* have missing values in some of the vintages. Again, we replace these missing values with the revised series. Some variables (e.g., real personal income and industrial production), have been re-indexed for the different vintages. This does not affect our modeling, as we use variable transformations such that level differences do not matter.

Table 5. Comparison of the Forecasting Performance When Using Real-Time vs. Revised Data

	Real Time (13 Months)	Real Time (12 Months)	Revised Data (12 Months)
Gradient Boosting	0.63 —	0.62 —	0.62 —
SVR	0.64 (0.33)	0.62 (0.48)	0.63 (0.37)
Forest	0.66 (0.04)	0.64 (0.02)	0.65 (0.02)
Neural Network	0.67 (0.01)	0.64 (0.05)	0.66 (0.01)
AR_1	0.72 (0.04)	0.69 (0.05)	0.69 (0.06)
Lasso Regression	0.73 (0.00)	0.71 (0.00)	0.72 (0.00)
Ridge Regression	0.75 (0.00)	0.72 (0.00)	0.73 (0.00)
Linear Regressions	0.75 (0.00)	0.72 (0.00)	0.73 (0.00)

Note: The performance metric is mean absolute error. The models are tested in the period between August 2000 and November 2019. Target audience: analysts.

$y_{t'}$ and lagged predictors $x_{t'-13}$, where $t' \leq t-1$. To make a prediction one year ahead ($\hat{y}_{t+12}$), we use the latest feature values of the same vintage (x_{t-1}) and compare the prediction against the revised response variable y^r_{t+12}. As in the previous experiments, we update the machine learning models every 12 months, winsorize the features, and use hv-block cross-validation to calibrate the hyperparameters.

Table 5 compares the model performances using real-time (left two columns) and revised data (right column). As in our main empirical analysis (see Table 2), the machine learning methods outperform the linear models, with gradient boosting being the best model. The p-values in parentheses indicate the statistical significance (one sided) of the Diebold-Mariano test, estimating whether the gradient boosting model significantly outperforms the other models.

The predictions based on the revised data are slightly more accurate than those based on the real-time data. This is driven by the fact that the real-time prediction is a 13-month forecast rather than a one-year-ahead forecast because of the reporting lag of one month. The middle column of the table shows the performance when this reporting lag would not exist. Here the real-time data is used to make predictions 12 months ahead of the latest available data, which effectively is a forecast 11 months ahead. The performance differences between these real-time predictions and the predictions based on

revised data are small and do not suggest that the models improve when using revised data. This is not surprising, given that the real-time and revised series most often only differ by a small degree, as shown in Figure A.2 in Appendix A. We therefore do not investigate real-time data in detail but focus on the revised data in this study, which allows us to investigate the models over a longer time period.

5.3 Extending the Set of Features

So far, we have used nine hand-picked key features (see Table 1) to predict unemployment changes. However, the FRED-MD database (McCracken and Ng 2016) offers a much richer set of variables—97 of which do not have any missing values between 1959 and 2019. Can we improve the forecasting performance by exploiting all of these? We make the variables stationary by applying the transformations suggested by the authors of the database using a transformation span of $l = 3$ for all variables.

Table 6 compares the performance when using the key features (first column) versus all features (second column). Using all features, the performance of the best models, gradient boosting, as well as the OLS regression, declines, whereas the performance of the other models improves or does not change. The random forest based on all features performs even slightly better than gradient boosting based on the key features. However, the Diebold-Mariano test shows that the difference is not significant ($p = 0.72$, two sided).

An analysis of the Shapley values shows that the machine learning models do not learn a sparse model when trained on all features. For the three best models—random forest, SVR, and neural network—the Shapley share of the top 10 features, respectively, only accounts for 41 percent, 32 percent, and 34 percent of the variance in the predictions. To account for at least 80 percent, we need to consider at least the top 39, 53, and 47 features, respectively. The large number of variables also increases the disagreement between models. While the agreement in the Shapley share is high between the SVR and the neural network (correlation of 0.93), it is lower between the forest and the other two methods (0.69, 0.70) (see Figure A.3 in Appendix A). Further, unlike the models trained on the key features

Table 6. Comparison of the Forecasting Performance When Using Different Input Data

	Key Features	**All Features**	**PCA$_1$**	**PCA$_2$**	**PCA$_3$**	**PCA$_5$**	**PCA$_7$**
Gradient Boosting	0.56	0.58	0.67	0.53	**0.52**	0.54	0.57
SVR	0.57	0.57	0.61	**0.52**	**0.52**	0.55	0.59
Random Forest	0.58	0.55	0.62	**0.52**	0.53	0.55	0.61
Neural Network	0.59	0.57	0.69	**0.52**	0.53	0.55	0.55
Lasso	0.64	0.63	0.65	0.56	**0.54**	0.56	0.59
Ridge	0.64	0.58	0.65	0.56	**0.54**	0.56	0.58
OLS	0.65	0.80	0.65	0.56	**0.54**	0.56	0.59

Note: The models are trained on the 10 key features, all 97 features, or the k top components of a principal component analysis (PCA$_k$), which was calibrated on all features. The performance metric is the mean absolute error. The best input data for each model (rows) is highlighted in bold. Target audience: analysts.

only (Figure 3), the functional forms do not align well between methods when trained on all features. Figure A.4 in Appendix A shows this for some of the key features.

This is not surprising given the rather small number of observations in our data and the fact that non-parametric convergence often is slow when the number of features is high (Stone 1982).

5.3.1 Dimensionality Reduction

In the literature on economic forecasting, a standard approach to exploit the predictive power of many features is to calibrate a dimensionality reduction model (e.g., principal components analysis, or PCA) and train the prediction models on the most important components (Stock and Watson 2002; Kim and Swanson 2018). Aggregating redundant variables in the same component allows models to learn more effectively from a lower-dimensional feature representation.

We follow this approach and use a PCA to summarize all 97 features.[21] Table 6 shows the performance for the forecast error when the machine learning models are trained on one, two, three, five, and seven components of the PCA. The best performance is achieved when using only two components, with the SVR, neural network, and random forest all performing equally well. Comparing these three models with the gradient boosting model trained on the key features, the Diebold-Mariano test estimates the following p-values (two sided), respectively: 0.054, 0.028, and 0.064. This suggests that using the PCA leads to a superior performance.

A model based on just two components may seem easy to interpret at first sight. However, as shown in Figure 7, the loadings of the 97 variables on these components are not sparse. We show to which group a variable belongs, where the groups have been defined by the authors of the data set (see also Ludvigson and Ng 2009). Most variables with high loadings on the first component belong to the labor market and output and income variable groups, but other variables have substantial loadings as well. Similarly, on the second component, the variables with the highest loadings belong to the

[21] We calibrate the PCA model on the training set only and update it every year as we do for the machine learning models.

Figure 7. Absolute Loadings of the 97 Features on the First (left panel) and Second (right panel) PCA Component

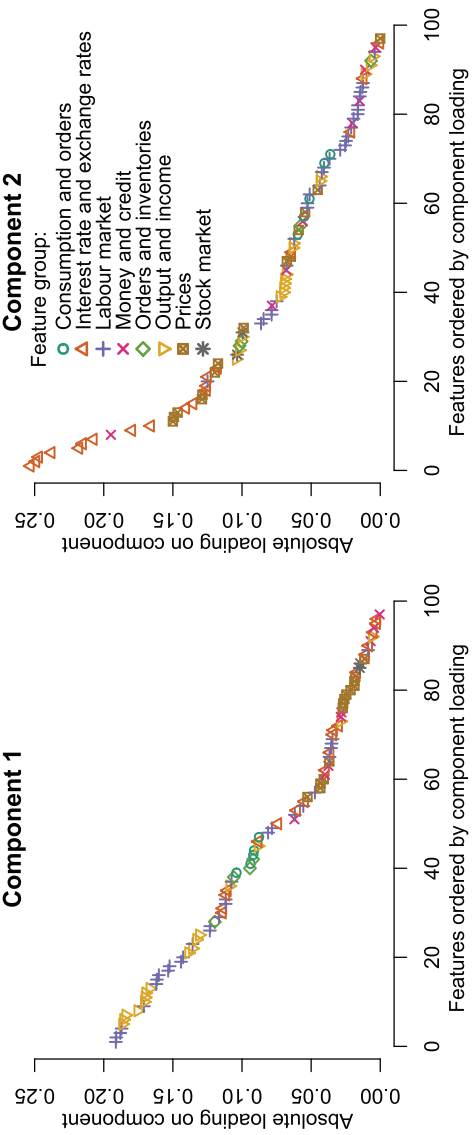

Note: The loadings shown are averages based on 30 PCA models trained on bootstrapped samples of the complete time series. Target audience: analysts.

interest rate and exchange rates group, but other variables also contribute substantially.[22] This suggests that the components do not have a simple economic interpretation.

At the same time, using the PCA components also limits the insights we can draw from the analysis of Shapley values. The first two components only explain 24 percent and 9 percent of the total variance in the data, respectively. Thus, most of the variation in the variables is not accounted for by the first components of the PCA. Further, making a machine learning model learn from only a few components will confine its ability to learn idiosyncratic functions of the individual features underlying that components. Rather, we expect that all functional forms of the variables loading on the same component will be similar.

Figure A.5 in Appendix A supports this conjecture. It shows the Shapley values based on the random forest when trained on the top two PCA components. The features in the top row have a high loading on the first component and low loadings on the second. The opposite holds for the features in the second row. In each row, the features show highly similar functional forms. The functional form of the features lagged unemployment and S&P 500—both included in the set of key features—differ from those shown in Figure 3. We do not observe a quadratic functional form for any of the 97 features when training the models on the PCA loadings, whereas two of the five most important features in our baseline experiment have such a functional form.

While we observe a small performance improvement when using a PCA instead of the hand-picked key features, this comes at the cost of a more complex model that arguably provides less economic insights. Our results partially support the idea of an "illusion of sparsity" (Giannone, Lenza, and Primiceri 2021). The authors used linear models to show that making a model sparse by picking a small set of predictors from the larger set comes with the cost of an inferior predictive performance. We observe the same for our ridge regression for which the absolute error falls by 0.06 and 0.1, respectively,

[22]It is important to note that the variables within the same group are not redundant. For example, the median absolute correlation (after transformation) of all variables within the labor market group and within the output and income group only is 0.29 and 0.43, respectively.

when using all features directly or training the model on three PCA components.

However, the performance gains from exploiting all variables are smaller for the machine learning models. Further, our set of key features was selected based on economic considerations rather than empirical selection and is thus probably not the best possible subset. This suggest that the trade-off between sparsity and accuracy might be less pronounced when using non-linear models because these are able to extract more information from sparse models.

5.3.2 A Richer Lag Structure

Finally, we extend the number of features by adding more lags of the key variables. The minimum lag of 12 months constitutes the prediction horizon. We add additional yearly lags from 24 to 72 months.[23] Table 7 shows the results of that experiment. While most models improve when adding more lags, the performance of the SVR and the neural network does not.

The best performance is achieved by the gradient boosting model when trained on annual lags of the last four years. We take a closer look at this model. Figure A.6 in Appendix A shows the functional forms for the different lags of the top features. The 12-month lags of the variables contribute most to the predictions. The other lags mostly make only small contributions to the predictions. It is interesting to observe that the functional forms differ not only in their size between the lags of the same variables but also in their shape. For example, comparing the lags of 12 months and 24 months, we observe contrary directions of the functional form for both industrial production and S&P 500. Larger lags thus provide a form of correction to the main effects (first lag), explaining the somewhat better model performance. Whether this improvement in performance warrants the more complex interpretation of the resulting models depends on the practical situation at hand.

5.4 Robustness of Shapley Values

We have shown in Figure 6 that the performance of the prediction models can be quite sensitive to random seeds. Here, we investigate

[23] We also experimented with adding monthly lags (e.g., $12, 13, \ldots, 23, 24$) but this richer set of features produced inferior results.

Table 7. Performance Comparison When Using Different Lag Structures

	Lags h (in months and steps of 12)					
	12	**12–24**	**12–36**	**12–48**	**12–60**	**12–72**
Gradient Boosting	0.56	0.59	0.54	**0.52**	0.53	0.54
SVR	**0.57**	0.58	0.60	0.61	0.63	0.64
Forest	0.58	0.58	0.56	0.57	**0.56**	**0.54**
Neural Network	**0.59**	0.62	0.61	0.60	**0.59**	**0.59**
Lasso	0.64	0.63	0.62	0.62	**0.61**	**0.61**
Ridge	0.64	0.64	0.63	0.60	0.59	0.60
OLS	**0.65**	**0.65**	0.68	0.70	0.75	0.78

Note: Lags are shown in months and are incremented in steps of 12. For example, the lag structure 12–48 contains lags 12, 24, 36, 48 of our key features. The lag of 12 corresponds to the baseline experiment. The error metric is mean absolute error. The lag structure that leads to the lowest error for each model (row) is highlighted in bold. Target audience: analysts.

whether random seeds also affect the global and local feature attributions and with that the economic interpretation.

Figure A.7 in Appendix A presents the Shapley shares of 10 different gradient boosting models, each based on a different random seed. For each variable, there is little variance in the Shapley share between the models. The functional forms learned by the models are also rather robust. Figure A.8 in Appendix A shows the Shapley values of the four most important predictors based on the 10 gradient boosting models with different random initializations. There are only minor differences between the fitted third-degree polynomials. However, the Shapley values of single data points can differ substantially between the different model realizations. This is indicated by the vertical lines which show, for each observation, the range of Shapley values across the 10 iterations. This shows the benefits of model averaging, which leads to more stable estimates.

Computing exact Shapley values is computationally expensive. It requires testing the predictions of all possible coalitions of features (see Appendix B). The number of coalitions grows exponentially with the number of features so that, in practice and as implemented by the *kernel* approach in the SHAP Python library (Lundberg and Lee 2017), coalitions are sub-sampled to approximate the true Shapley values. When a coalition does not include a feature k, it is integrated out by using its values within a background data set (see again Appendix B). Ideally, the background set is big and represents the data set well. However, the bigger the background sample, the more expensive the computation of the Shapley values becomes. In practice, the background sample is often summarized by using a random sub-sample of the training set, or by approximating the training set with a few representative centroids using k-means clustering.

In all experiments above, we have used the kernel method with 2,000 coalitions, and 25 centroids when estimating Shapley values for all machine learning models. Here, we investigate the robustness of the Shapley values when altering these two parameters. Figure A.9 in Appendix A shows the Shapley values of industrial production,[24] the most important predictor, for our most accurate models, gradient boosting and the SVR. We order all observations by increasing

[24]The results are qualitatively similar for the other features.

Shapley values. When varying the number of coalitions (top row), the gradient boosting model is insensitive to this parameter. Sampling only 50 coalitions suffices for an accurate estimation of Shapley values. In contrast, we see some variation in Shapley values for the SVR if only 50 coalitions are sampled. There is almost no variation anymore if 100 coalitions are used.

The middle row of Figure A.9 shows the effect of varying the size of the background sample. Here for both the gradient boosting model and the SVR, we only observe some errors in the estimates for a small background sample size of five.

An alternative to the kernel-based computation of Shapley values is Tree SHAP (Lundberg et al. 2020). It is not model agnostic and can only be used on tree-based models such as gradient boosting and random forests. It does not estimate Shapley values by enumerating all possible coalitions of features but by only considering those that actually are observed in the tree models, which makes this approach computationally much cheaper by construction.[25] The bottom panel of Figure A.9 compares the Shapley values of industrial production based on the kernel method with those based on Tree SHAP for both tree-based models. Both methods produce very similar estimates of Shapley values for gradient boosting and the random forest.

Table 8 shows the computation time required to obtain Shapley values for the 359 data points of the whole test period between 1990 and 2019.[26] With the baseline parameters of 2,000 coalitions, and a background sample of 25, the computation time is about one minute for the gradient boosting model and eight minutes for the SVR. But by reducing the coalitions to 100, the computation time for both methods drops substantially. Using Tree SHAP, Shapley values are estimated within two seconds.

This analysis suggests that, while the exact estimation of Shapley values can be computationally prohibitive, they can be approximated accurately and efficiently. For instance, in our case study,

[25] We set the parameter *feature_perturbation* to *interventional*. In this way, Tree SHAP, like the kernel method, ignores dependencies between features (see Appendix B). As another robustness check, we set this parameter to *tree_path_dependent* and thus consider correlations between features. Doing this, the Shapley values of our tree models do not change markedly.

[26] We train a single model every year and do not use bagging. The computation was conducted on a single kernel (not parallelized) of an AMD Ryzen 7 3700X.

Table 8. Computation Time (in seconds) When Estimating the Shapley Values of the SVR and Gradient Boosting for the Whole Test Period (1990–2019) Containing 359 Observations

Method	Background Sample	Coalitions	Computation Time (Seconds)	
			Gradient Boosting	SVR
Kernel	25	100	16.35	76.76
Kernel (Baseline Parameters)	25	2,000	69.74	451.73
Kernel	100	2,000	292.06	1,772.17
Tree SHAP	—	—	1.70	—

Note: Target audience: analysts.

the Shapley value computation for the full test period can be reduced to the order of one minute without noteworthy differences in attribution.

6. Discussion

This paper presents a workflow for using machine learning to inform decision making in situations where transparency and ease of communicating results are key. The three steps of the workflow are a model comparison, a decomposition of predictions into feature contributions, and statistical inference on those contributions.

We applied the workflow to an economic forecasting exercise, predicting the change of the U.S. unemployment rate one year ahead for the past 30 years.

In the first step of this case study, we found a significantly better performance of machine learning models compared with linear benchmarks. In the second step, we observed pronounced nonlinearities learned by the machine learning models that have clear economic interpretations. In the third step of the workflow, we use the Shapley regression framework to find that a larger number of variables is statistically significant when using machine learning

models than for the linear benchmark, which demonstrates that non-linear machine learning models can extract and uncover a richer set of robust predictive relationships in the data.

Machine learning methods are increasingly used in economic and social science research. However, most studies using machine learning focus on maximizing predictive accuracy and accept the black box nature of the models. Research that does attempt statistical inference on machine learning models often uses controlled data—for example, from randomized controls trials. Our study shows that the use of machine learning models and statistical inference can be combined to learn from real-world prediction problems. But our study also revealed challenges of machine learning modeling on rather small data. First, we showed that the performance of some of the machine learning models is rather sensitive to random seeds. Second, the machine learning models differ in how they are affected by experimental parameters such as the type of hyperparameter search (e.g., k-fold cross-validation versus blocked cross-validation) or winsorization. These challenges can be addressed by model averaging to increase robustness, and by rigorous robustness checks, respectively. At the same time, our diverse set of machine learning methods, which differ significantly in their predictive performance, all learned highly similar functional forms from the data.

When using a broad set of predictive variables instead of a small hand-picked selection, the predictive performance increased slightly. But as the non-linear machine learning models do not learn a sparse model but use most features for prediction, interpreting their predictions becomes considerably more challenging. Further, we showed that the functional forms learned were less consistent across model families. When we trained the models on PCA components, the interpretability was reduced as well, because the components do not have a clear interpretation. At the same time, being trained on just a few components limits the differential functional forms that the models can learn from each underlying feature. Finally, the Shapley regression will suffer from a reduction of statistical power when being calibrated on a large number of features. Collectively, these considerations suggest that, while our workflow is general and works in principle for problems with a large number of features, it will deliver more robust and easier to interpret results when based on a smaller set of features.

More generally, our case study is reminiscent of many real-world settings where one has a considerable number of features to learn from but a limited number of observations. Our results suggest that a combination of expert domain knowledge to select key indicators and robust model properties may lead to the best trade-off between model performance and the interpretability of results—if such a trade-off exists in the first place.

Machine learning methods may be unfamiliar to many decision-makers, but their predictive accuracy and ability to detect richer, more nuanced signals in data justify their use to inform decisions. While the Shapley regression framework cannot *fully* communicate complex and nonlinear models with a single statistic, like the Shapley share coefficient, it allows statistical inference on nonlinear outputs and can be communicated similarly to well-understood regression results.

Appendix A. Results

Figure A.1. Forecasting Performance across 10 Iterations as a Function of the Number of Models in the Bagging Ensenble

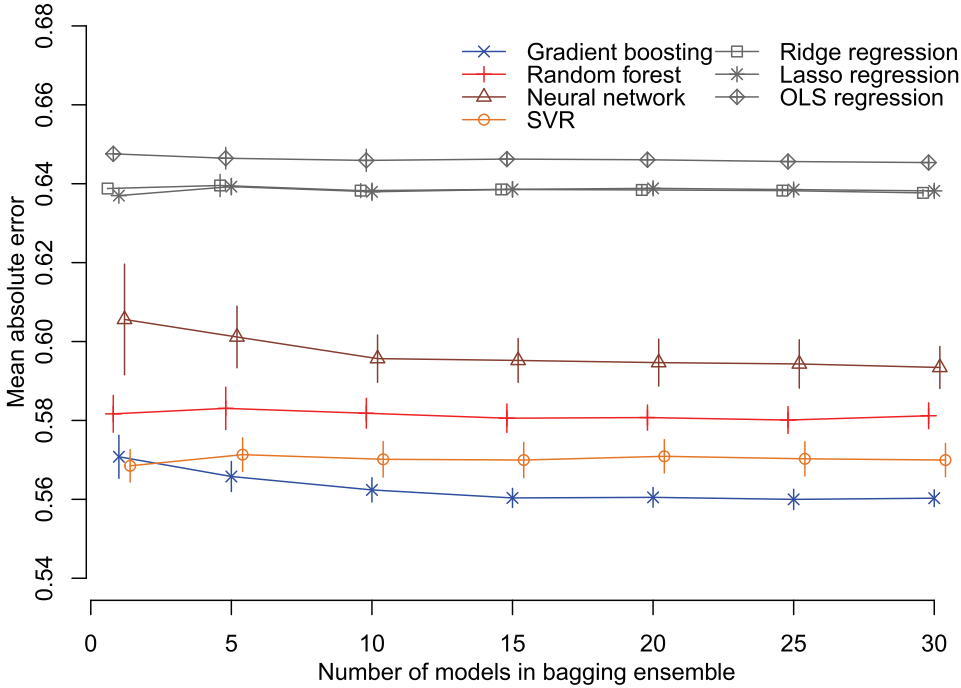

Note: The horizontal lines show the mean performance across all 10 iterations; the vertical bars show ± two standard errors around that mean. Target audience: analysts.

Figure A.2. Comparison of Real-Time and Revised Series after Transformations that Make Them Stationary (see Table 1)

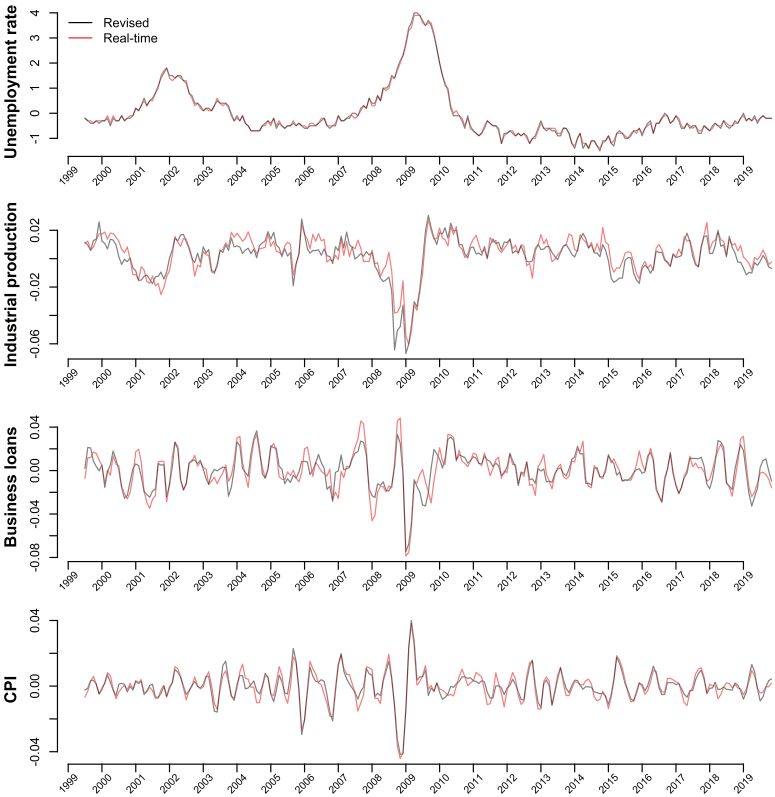

Note: Target audience: analysts.

Figure A.3. Shapley Share When Machine Learning Models Are Calibrated on All 97 Features of the Data Set

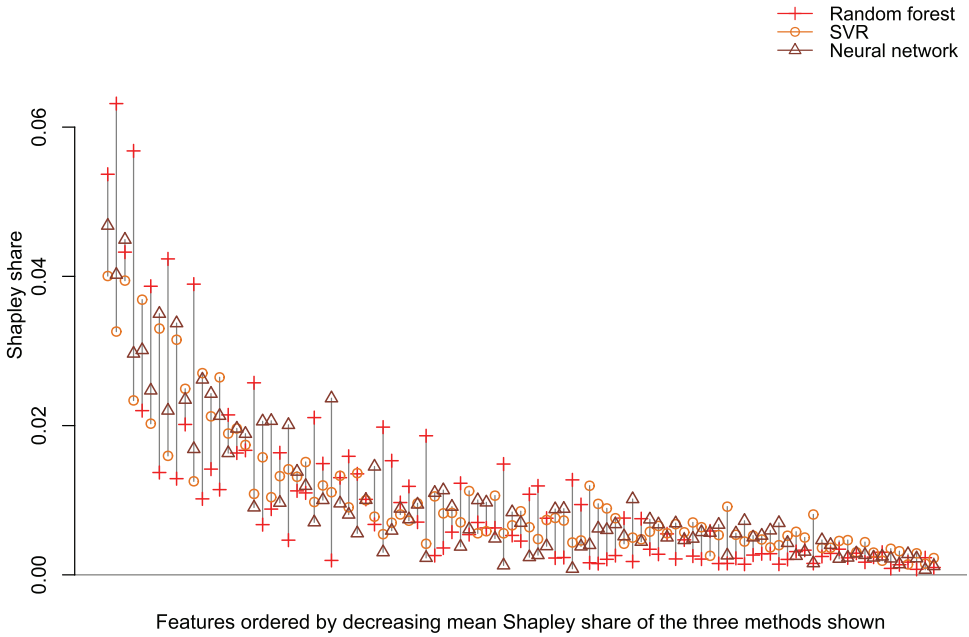

Note: The three shown models perform best in prediction when calibrated on all features. Target audience: analysts.

Figure A.4. Selected Learned Functional Forms for the Three Best-Performing Models When Using All 97 Features

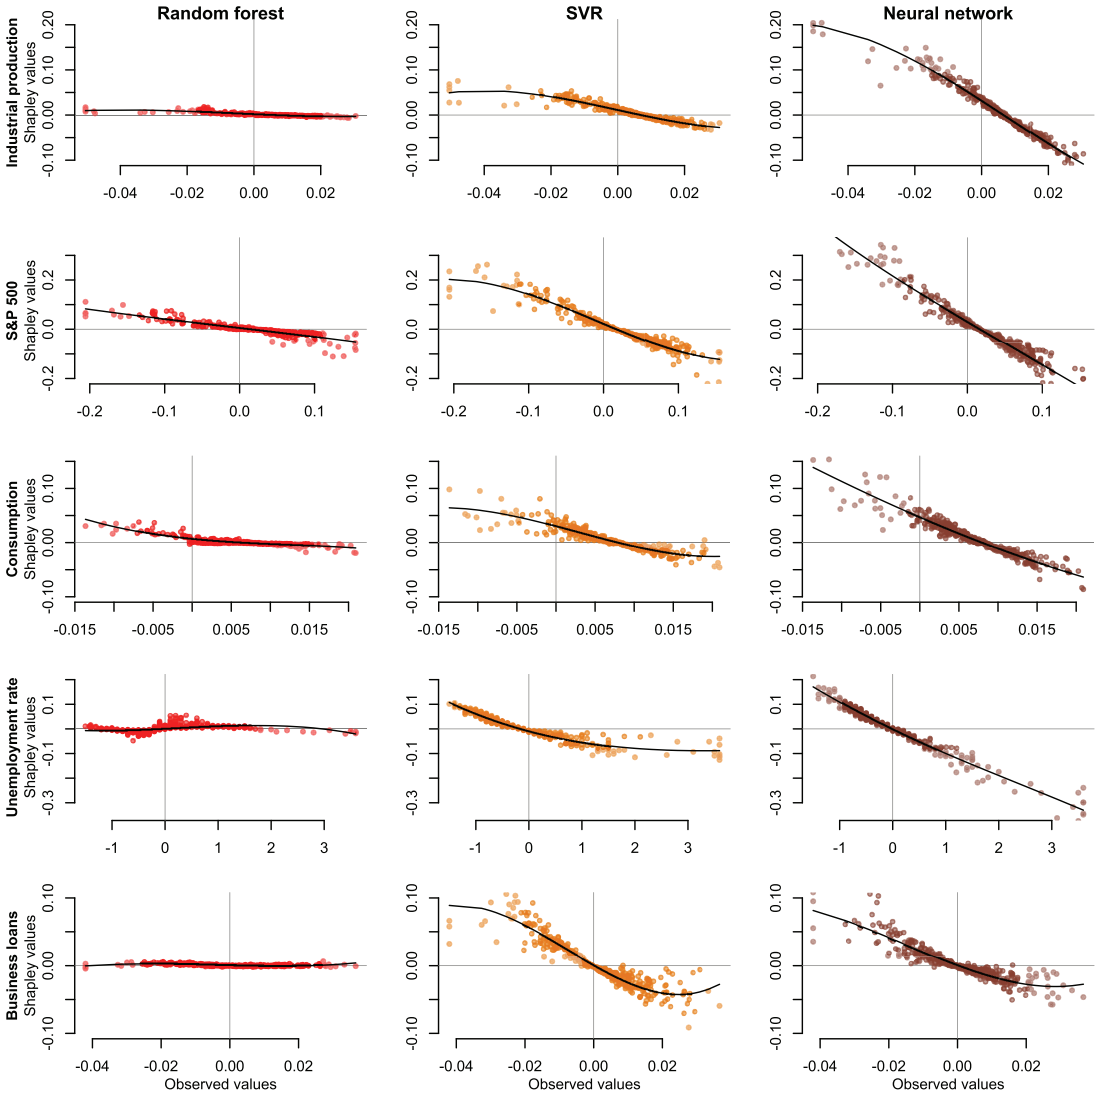

Note: The lines show best-fit third-degree polynomials. Note that the scale of the vertical axis differs between rows. Target audience: analysts.

Figure A.5. Learned Functional Forms of Selected Features Based on the Predictions of a Random Forest Trained on the Two First Components of a PCA

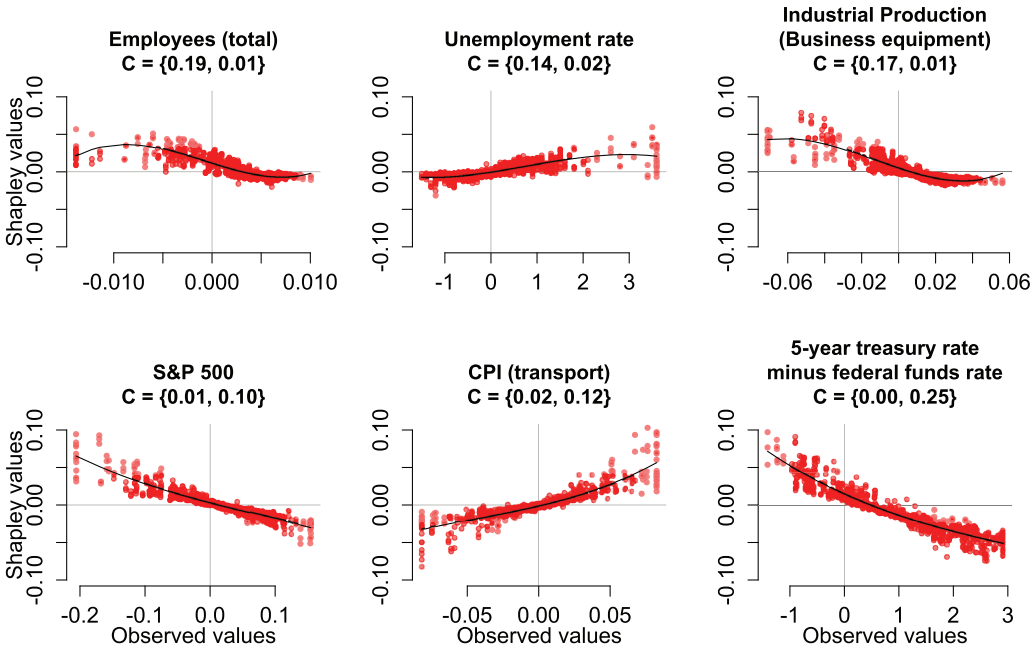

Note: The term C shows the loadings of the features on the first and second principal component. The first row shows features with a high loading on the first principal component and low loadings on the second component. The second row shows features with a high loading on the second principal component and low loadings on the first component. The lines show best-fit third-degree polynomials. Target audience: analysts.

Figure A.6. Learned Functional Forms of Key Predictors at Different Lags as Learned by the Gradient Boosting Model

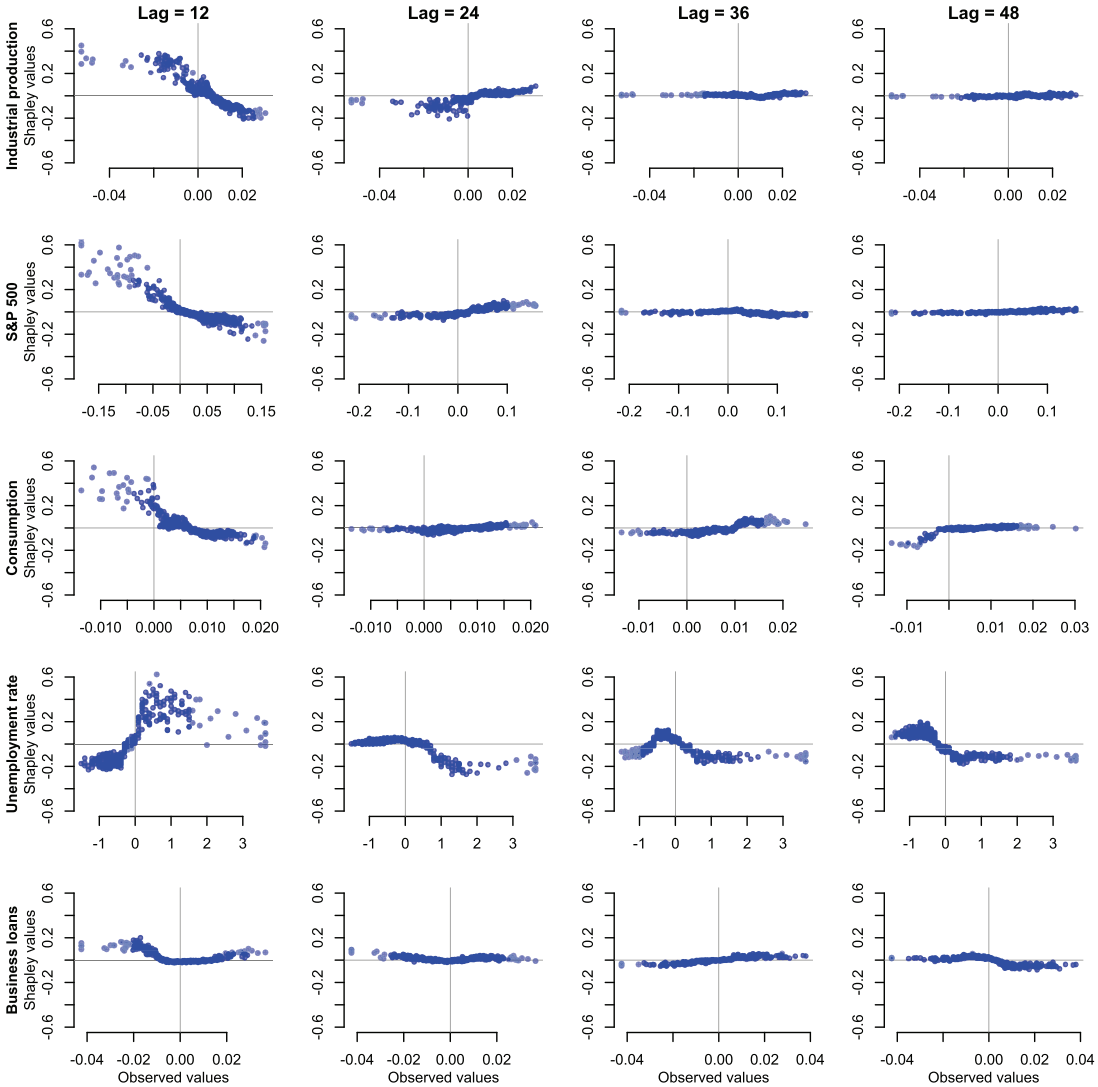

Note: The results shown are based on a model trained on 40 features: four lags (12, 24, 36, 48 months) for each of the 10 key features. Target audience: analysts.

Figure A.7. Robustness of the Shapley Share

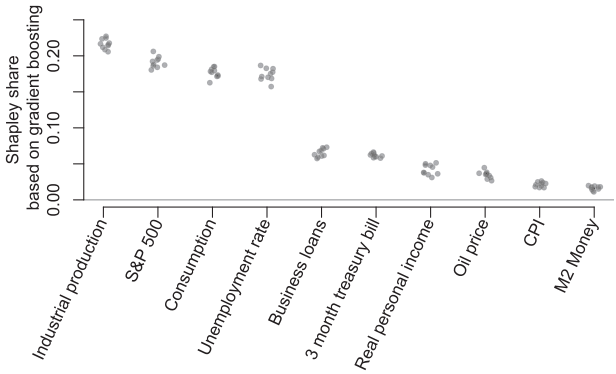

Note: The figure shows the Shapley shares according to 10 different gradient boosting models, each trained with a different random seed. Target audience: analysts.

Figure A.8. Robustness of Individual Shapley Values

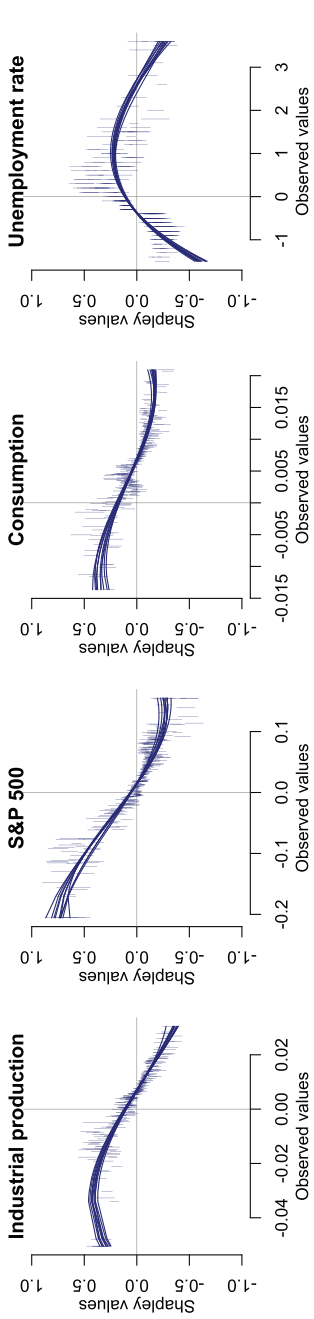

Note: Each line shows the functional form learned by one of 10 gradient boosting models, each trained with a different random seed. The vertical lines show the maximum range in Shapley values across the 10 iterations for each observation. Target audience: analysts.

Figure A.9. Accuracy of Shapley Value Computations

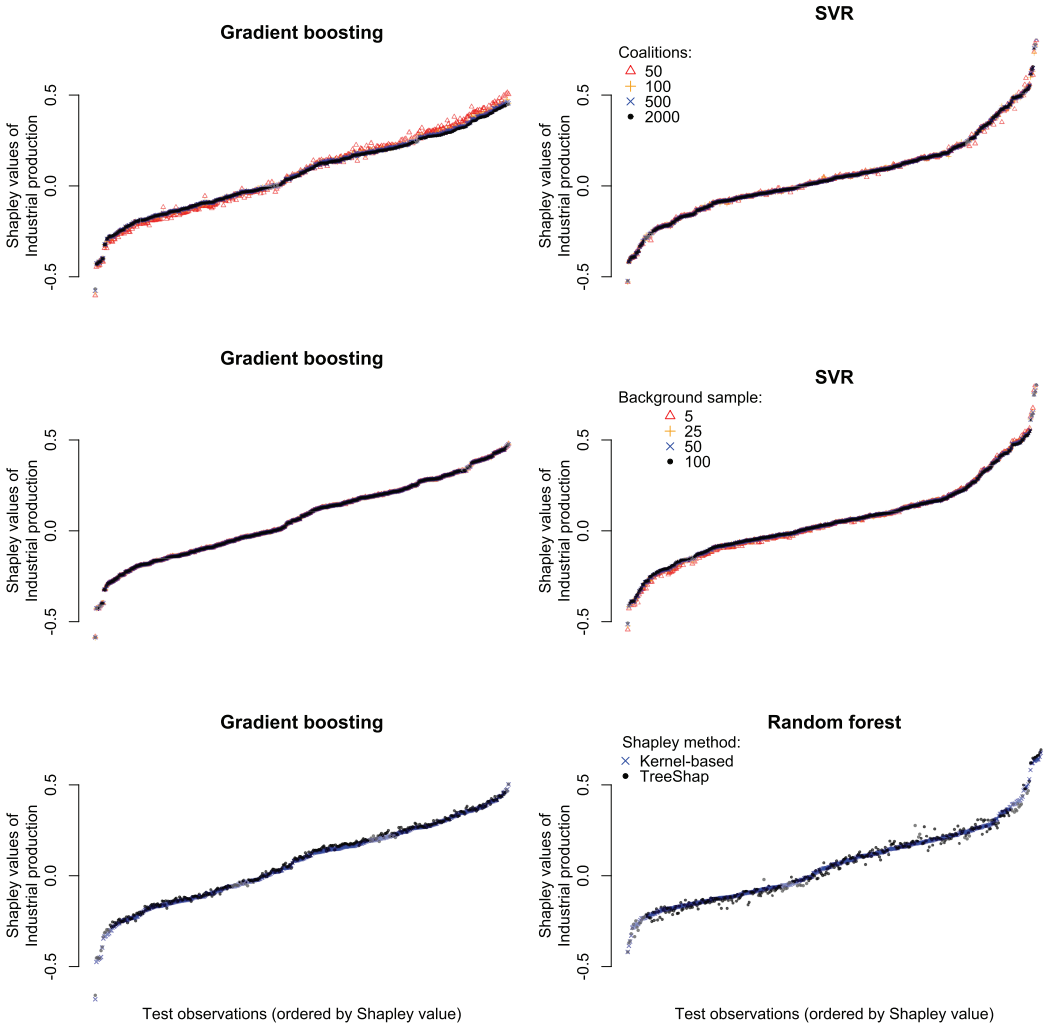

Note: The top row compares Shapley values estimated by the kernel method for different coalition sizes. The middle row shows the Shapley values for different background sample sizes. The bottom row compares Shapley values estimated by the kernel-based method and the Tree SHAP method for the two tree-based methods. From the top to bottom row, observations are ordered by increasing Shapley values of the largest number of coalitions, the largest background, sample, and the kernel-based method, respectively. Target audience: analysts.

Appendix B. Feature Importance Measures

We present a concise description of the computation of the two feature importance measure used in this paper, permutation importance and Shapley values. We also discuss computational and conceptual challenges.

B.1 Permutation Importance

The permutation importance of a variable measures the change of model performance when the values of that variables are randomly scrambled, i.e., permuted. If a model has learnt a strong dependency between the model outcome and a given variable, scrambling the value of the variable leads to very different model predictions and thus affects performance. A variable k is said to be important in a model if the error $\mathcal{P}$ after scrambling feature k is substantially higher than when using the original values for k, i.e., $\mathcal{P}_k^{perm} >> \mathcal{P}_k^{baseline}$. The value of the permutation error $\mathcal{P}_k^{perm}$ depends on the realization of the permutation, and variation in its value can be large, particularly in small data sets. Therefore, it is recommended to average $\mathcal{P}_k^{perm}$ over several random draws for more accurate estimation and to assess sampling variability. Note that it is intractable in most applications to evaluate all $M!$ permutations in a test set of size M. However, the average of multiple realizations of $\mathcal{P}_k^{perm}$ will mostly converge quickly with the number of permutation, making permutation importance a computationally cheap way to assess feature importance.[27]

The following procedure estimates the permutation importance.

(i) For each feature x_k:
   a. Generate a permutation sample x_k^{perm} with the values of x_k permuted across observations.

   b. Re-evaluate the test performance for x_k^{perm}, resulting in $\mathcal{P}_k^{perm}$.

[27] Given a large test set, bootstrap sub-samples may suffice.

c. The permutation importance of x_k is given by $I(x_k) = \mathcal{P}_k^{perm}/\mathcal{P}_{baseline}^k$. Alternatively, the difference $\mathcal{P}_k^{perm} - \mathcal{P}_k^{baseline}$ can be considered.

d. Repeat and average over Q iterations and average $\bar{I}_k = 1/Q \sum_q I(x_k)$.

(ii) If I_k is based on the ratio of errors $\mathcal{P}_k^{perm}/\mathcal{P}_k^{baseline}$, consider the normalized quantity $\bar{I}_k = (I_k - 1)/\sum_k (I_k - 1) \in (0, 1)$.[28]

This ease of use comes at some cost. For example, if two features contain similar information, permuting either of them will not reflect the actual importance of this feature relative to all other features. Only permuting both or excluding one would do so. This motivates our comparison with Shapley values because they identify the individual marginal effect of a feature, accounting for its interaction with all other features. More generally, permutation importance does not come with the same analytical guarantees as Shapley values. Additionally, the computation of permutation importance requires access to true outcome target values to evaluate performance. In many situations, e.g., when working with models trained on sensitive or confidential data, these may not be available.

B.2 Shapley Values

Shapley values originate from game theory as a general solution to the problem of attributing a joint pay-off obtained in a cooperative game to the individual players of a coalition based on their contribution to the game (Shapley 1953). Štrumbelj and Kononenko (2010) introduced the analogy between players in a cooperative game and variables in a general supervised model. In the latter, variables jointly generate a prediction, the pay-off.

The Shapley values of a model offer a local decomposition of each model prediction[29] x_i of the form given in Equation (1), $f(x_i) = \sum_{k=0}^{N} \phi_k(x_i)$. Here $\phi_k^S(x_i)$ is the Shapley value associated

[28] Note, $I_k \geq 1$ in general. If not, there may be problems with model optimization.

[29] We label observations by $i \in \{, \ldots, M\}$ here, which is more general than the time-series notation used in Section 3.

with predictor k and ϕ_0^S an intercept, usually the mean prediction of the model. Shapley values come with a host of appealing analytical properties which are inherited from their game-theoretic origins. Moreover, the decomposition in Equation (1) does not need to refer to single variables but can also include interactions or even higher-order terms of interest as introduced by Agarwal, Dhamdhere, and Sundararajan (2019).

The Shapley attribution $\phi_k^S(x_i; f)$ for variable k in observation x_i and model f in Equation (1) is given by

$$\phi_k^S(x_i; f) = \sum_{x' \subseteq \mathcal{C}(x) \setminus \{k\}} \frac{|x'|!(N - |x'| - 1)!}{N!} \big[f(x_i | x' \cup \{k\}) - f(x_i | x') \big], \tag{B.1}$$

where $\mathcal{C}(x) \setminus \{k\}$ is the set of all possible variable combinations (coalitions) when excluding variable k and $|x'|$ denotes the number of variables included in that coalition. Equation (B.1) is the weighted sum of marginal contributions of variable k to all possible variable coalitions not including k itself. For example, assuming we have three variables (players) $\{A, B, C\}$, the Shapley value of player C would be $\phi_C^S(f) = 1/3[f(\{A, B, C\}) - f(\{A, B\})] + 1/6[f(\{A, C\}) - f(\{A\})] + 1/6[f(\{B, C\}) - f(\{B\})] + 1/3[f(\{C\}) - f(\{\emptyset\})]$.

There are challenges in evaluating Equation (B.1), which may be called the no-free-lunch theorem of Shapley values. First, the number of coalitions x' to evaluate grows exponentially with the number of variables. This means that an exhaustive enumeration is not feasible for already moderate variable sets. Second, we need to evaluate conditional model predictions of the form $f(x_i|x')$, i.e., models where only variables in x' are "active," and information from the other variables is excluded.[30] Under the assumption of feature independence, $f(x_i|x')$ can be evaluated by conditional expectations over an informative background data set b. That is, non-active variables are integrating out numerically using b. Let $\omega_{x'} \equiv |x'|!(N - |x'| - 1)!/N!$

[30]This does not mean a different model which only uses variables in x'. Such a model would need to be retrained and would generate different predictions, i.e., not be the model we want to evaluate.

be the combinatorial weighting factor and $\bar{x'}$ the set of variables among all not included in x'; then Equation (B.1) can be written as

$$\phi_k^S(x_i; f) = \sum_{x' \subseteq \mathcal{C}(x) \setminus \{k\}} \omega_{x'} \left[\mathbb{E}_b[f(x_i)|x' \cup \{k\}] - \mathbb{E}_b[f(x_i)|x'] \right], \tag{B.2}$$

$$\text{with} \qquad \mathbb{E}_b[f(x_i)|x'] \equiv \int f(x_i)\, \mathrm{d}b(\bar{x'}) = \frac{1}{|b|} \sum_b f(x_i|\bar{x'}). \tag{B.3}$$

The intercept ϕ_0^S is determined by the background data set b motivating its name,

$$\phi_0^S = \mathbb{E}_b[f(\emptyset)] = \mathbb{E}_b[f(\bar{x'} = b)]. \tag{B.4}$$

This means that the components ϕ_k^S measure variable contributions relative to the mean model prediction in b and that ϕ_0^S is a reference point. The choice of b will also affect the interpretation of Shapley components. That is why b should be informative, e.g., as being representative of the whole data-generating process, or to represent a sub-group of a population, like the group of untreated in an experimental setting.

We have not yet discussed the problem of computational complexity and the case when features are not independent. These are briefly discussed with further references to the literature.

- *Computational Complexity:* The sum over variable coalitions becomes impractical or even infeasible for already relatively small sets of variables of about 8–10 depending on the application. For larger sets of variables some form of coalition sampling or variable selection is needed. The Kernel SHAP algorithm Lundberg and Lee (2017) provides an efficient sampling and evaluation of Shapley contributions in the form of a weighted least square calculation. We have shown in Section 5.4 that only 100 coalitions suffice to accurately estimate Shapley values.

  A variable selection method which provides an exact computation for a subset of features is presented in Joseph (2019).

Often one is not interested in the contributions of all variables of a model, but only a subset. In this case variables *not* of interest can be grouped into a single *other* component, or sub-groups may serve as *super-variables* until an exhaustive evaluation of Equation (B.1) is possible. We have used this approach in Section 4.2, when computing the Shapley values of an interaction of features.

Additionally, the computation of Shapley components can be reduced by limiting the size of the background b. A default option is the whole training data set representing all information the model has learned from. However, this can be impractical for large data sets. Here either sub-samples or summary points, e.g., cluster centroids, can be used. We have shown in Section 5.4 that even small background samples of 25 observations suffice to accurately estimate Shapley values.

- *Feature Dependence:* The evaluation of conditional expectations (Equation (B.3)) does not consider dependencies between features, which can lead to feature value combinations that are non-sensical and would not occur in the real world. We discuss three ways to address this. First, one can estimate Shapley values of tree-based models for which there exists an efficient algorithm that accounts for feature dependence (Lundberg et al. 2020). By comparing Shapley values when respecting or ignoring feature dependence, one can gauge the importance of feature dependencies. However, caution is warranted when transferring the findings to other model families, e.g., artificial neural networks. These may have learned different feature attributions for which the comparison of Shapley components evaluated under the independence assumption is indicative.

  Second, one can net out the effect of higher-order feature interactions using the Shapley-Taylor index (Agarwal, Dhamdhere, and Sundararajan 2019) to understand dependencies between features. This is an expansion of a function over sets of active features including interactions of any order of interest analogous to the Taylor expansion of differentiable functions. This means that the importance of a feature is either its main effect or the main effect in addition to different interaction terms. Interaction terms also provide additional

information, as shown in the analysis presented in the main text.

Third, the dependence structure within a variable set can be estimated (Aas, Jullerm, and Lølund 2021). While adding an extra potentially computationally costly step, this has the advantage of providing a single attribution for each feature which accounts for the dependence of the data at least approximately. It may, however, be argued that it is not necessarily clear what a single component in a non-linear model with feature dependence captures. This is more intuitive for Shapley contributions of feature interactions, which capture the effect of co-movements of two or more variables.

- *Expectation Consistency:* As shown by Sundararajan and Najmi (2020), attribution consistency which, casually put, avoids logical contradictions in feature attribution, can be violated when using conditional expectations for the computation of Equation (B.3),[31] and a single reference value is advocated for. However, this discards much of the potentially rich information a model has learned, such as non-linearities. A solution to this is provided in Joseph (2019) by an additional condition when comparing different models against a common background. The models' intercepts ϕ_0^S over the background data b need to coincide. This is fulfilled in many practical situations where models optimize the same objective functions, like the mean squared error. Variations in ϕ_0^S are of concern as soon as they reach the order of magnitude of the Shapley components ϕ_k^S.

None of the above challenges is fatal for the application of Shapley values for model interpretability, as we have shown in detail. However, one has to be aware of the possible pitfalls and consequences of approximations for model interpretations and any decisions based on them.

[31] See also Janzing, Minorics, and Blöbaum (2020).

Table B.1. Implementation Details on the Prediction Models

Method	Implementation	Fixed Parameters	Hyperparameter Space
Random Forest	`scikit-learn` RandomForestRegressor	n_estimators: 500 criterion: mae	max_depth: [2, 3, 4, 5, 6, 8, 10, 20, 30]* max_features: [1, 3, 5, 7, 9, 11, 15, 30, 50]*
Gradient Boosting	`lightgbm` LGBMRegressor		subsample: [0.05, .1, .2, .3, .4, .6, .6, .7, .8, .9, 1] reg_lambda: [0, 0.1, 1, 10, 20, 50, 100] reg_alpha: [0, 0.1, 1, 2, 7, 10, 50, 100] num_leaves: [2, 3, 4, 5, 10, 20, 40, 70, 100] n_estimators: [1, 3, 5, 10, 20, 30, 40, 50, 75, 100]* max_depth: [1, 2, 3, 5, 8, 15]* colsample_bytree: [0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 1]
SVR	`scikit-learn` SVR		C: $[2^1, 2^{1+1\frac{4}{9}}, 2^{1+2\frac{4}{9}}, \ldots, 2^{1+9\frac{4}{9}}]$ gamma: $[2^{-7}, 2^{-7+1\frac{6}{9}}, 2^{-7+2\frac{6}{9}}, \ldots, 2^{-7+9\frac{6}{9}}]$ epsilon: [0.0001, 0.001, 0.005, 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5]*
Neural Network	`scikit-learn` MLPRegressor	solver: lbfgs max_iter: 2000	hidden_layer_sizes: [5, (5, 5), (5, 5, 5), 10, (10, 10), ..., (10, 10, 10), 15, (15, 15), ..., 25, (25, 25), (25, 25, 25)] activation: [relu, tanh] alpha: $[10^{-5}, 10^{-4}, \ldots, 10^3]$
Lasso Regression	`scikit-learn` Lasso		alpha: $[10^{-5}, 10^{-5+1\frac{9}{99}}, 10^{-5+2\frac{9}{99}}, \ldots, 10^{-5+99\frac{9}{99}}]$
Ridge Regression	`scikit-learn` Ridge		alpha: $[10^{-5}, 10^{-5+1\frac{9}{99}}, 10^{-5+2\frac{9}{99}}, \ldots, 10^{-5+99\frac{9}{99}}]$
OLS Regression	`scikit-learn` LinearRegression		

*We initially used a large parameter space but have refined it to these values without sacrificing performance.
Note: The second column shows the Python package and the respective name of the machine learning method. The third column shows parameters that we set to a different value than the default. The fourth column shows the hyperparameter space.

References

Aas, K., M. Jullum, and A. Løland. 2021. "Explaining Individual Predictions when Features are Dependent: More Accurate Approximations to Shapley Values." *Artificial Intelligence* 298 (September): Article 103502.

Agarwal, A., K. Dhamdhere, and M. Sundararajan. 2019. "A New Interaction Index Inspired by the Taylor Series." arXiv: 1902.05622.

Armstrong, S., K. Green, and A. Graefe. 2015. "Golden Rule of Forecasting: Be Conservative." *Journal of Business Research* 68 (8): 1717–31.

Bergmeir, C., and J. Benítez. 2012. "On the Use of Cross-validation for Time Series Predictor Evaluation." *Information Sciences* 191 (May 15): 192–213.

Bianchi, D., M. Büchner, and A. Tamoni. 2019. "Bond Risk Premia with Machine Learning." USC-INET Research Paper No. 19-11.

Bluwstein, K., M. Buckmann, A. Joseph, M. Kang, S. Kapadia, and O. Simsek. 2020. "Credit Growth, the Yield Curve and Financial Crisis Prediction: Evidence from a Machine Learning Approach." Staff Working Paper No. 848, Bank of England.

Breiman, L. 1996. "Bagging Predictors." *Machine Learning* 24 (2): 123–40.

———. 2001a. "Random Forests." *Machine Learning* 45 (1): 5–32.

———. 2001b. "Statistical Modeling: The Two Cultures (with Comments and a Rejoinder by the Author)." *Statistical Science* 16 (3): 199–231.

Burgess, S., E. Fernandez-Corugedo, C. Groth, R. Harrison, F. Monti, K. Theodoridis, and M. Waldron. 2013. "The Bank of England's Forecasting Platform: COMPASS, MAPS, EASE and the Suite of Models." Staff Working Paper No. 471, Bank of England.

Chakraborty, C., and A. Joseph. 2017. "Machine Learning at Central Banks." Staff Working Paper No. 674, Bank of England.

Chen, J., A. Dunn, K. Hood, A. Driessen, and A. Batch. 2022. "Off to the Races: A Comparison of Machine Learning and Alternative Data for Predicting Economic Indicators." In *Big Data for 21st Century Economic Statistics,* ed. K. G. Abraham, R. S. Jarmin, B. Moyer, and M. D. Shapiro. University of Chicago Press.

Chernozhukov, V., M. Demirer, E. Duo, and I. Fernandez-Val. 2018. "Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments." NBER Working Paper No. 24678.

Chetverikov, D., Z. Liao, and V. Chernozhukov. 2020. "On Cross-validated Lasso in High Dimensions." arXiv: 1605.02214.

Coulombe, P. G., M. Leroux, D. Stevanovic, and S. Surprenant. 2020. "How is Machine Learning Useful for Macroeconomic Forecasting?" arXiv: 2008.12477.

Cybenko, G. 1989. "Approximation by Superpositions of a Sigmoidal Function." *Mathematics of Control, Signals, and Systems* 2: 303–14.

D'Amour, A., K. Heller, D. Moldovan, B. Adlam, B. Alipanahi, A. Beutel, C. Chen et al. 2020. "Underspecification Presents Challenges for Credibility in Modern Machine Learning." arXiv: 2011.03395.

Doerr, S., L. Gambacorta, and J. M. S. G. Garralda. 2021. "Big Data and Machine Learning in Central Banking." BIS Working Paper No. 930.

Döpke, J., U. Fritsche, and C. Pierdzioch. 2017. "Predicting Recessions with Boosted Regression Trees." *International Journal of Forecasting* 33 (4): 745–59.

Doshi-Velez, F., and B. Kim. 2017. "Towards a Rigorous Science of Interpretable Machine Learning." arXiv: 1702.08608.

Drucker, H., C. Burges, L. Kaufman, A. Smola, and V. Vapnik. 1997. "Support Vector Regression Machines." In *Advances in Neural Information Processing Systems*, Vol. 9, ed. M. C. Mozer, M. Jordan, and T. Petsche, 155–61. Cambridge, MA: MIT Press.

Einhorn, H. J., and R. M. Hogarth. 1985. "Prediction, Diagnosis, and Causal Thinking in Forecasting." In *Behavioral Decision Making*, ed. G. Wright, 311–28. Springer.

Elliott, G., and A. Timmermann. 2008. "Economic Forecasting." *Journal of Economic Literature* 46 (1): 3–56.

Federal Reserve Bank of St. Louis. 2020. "NBER Based Recession Indicators for the United States from the Period Following the Peak through the Trough [USREC]."

Fernández-Delgado, M., E. Cernadas, S. Barro, and D. Amorim. 2014. "Do We Need Hundreds of Classifiers to Solve Real World

Classification Problems?" *Journal of Machine Learning Research* 15 (1): 3133–81.

Fisher, A., C. Rudin, and F. Dominici. 2019. "All Models are Wrong, but Many are Useful: Learning a Variable's Importance by Studying an Entire Class of Prediction Models Simultaneously." *Journal of Machine Learning Research* 20 (177): 1–81.

Friedman, J. 2001. "Greedy Function Approximation: A Gradient Boosting Machine." *Annals of Statistics* 29 (5): 1189–1232.

Friedman, J., T. Hastie, and R. Tibshirani. 2009. *The Elements of Statistical Learning.* Springer Series in Statistics. Berlin: Springer.

George, E. 1999. "Economic Models at the Bank of England." Technical Report, Bank of England.

Giannone, D., M. Lenza, and G. E. Primiceri. 2021. "Economic Predictions with Big Data: The Illusion of Sparsity." *Econometrica* 89 (5): 2409–37.

Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep Learning.* MIT Press.

Haldane, A. 2018. "Will Big Data Keep Its Promise?" Speech at the Bank of England Data Analytics for Finance and Macro Research Centre, King's Business School, London, April 19.

Independent Evaluation Office. 2015. "Evaluating Forecast Performance Report." Technical Report, Bank of England.

Janzing, D., L. Minorics, and P. Blöbaum. 2020. "Feature Relevance Quantification in Explainable AI: A Causal Problem." In *Proceedings of Machine Learning Research*, Vol. 108, ed. S. Chiappa and R. Calandra, 2907–16. Proceedings of the International Conference on Artificial Intelligence and Statistics, August 26–28.

Joseph, A. 2019. "Parametric Inference with Universal Function Approximators." arXiv: 1903.04209.

Kazemitabar, J., A. Amini, A. Bloniarz, and A. S. Talwalkar. 2017. "Variable Importance Using Decision Trees." In *Advances in Neural Information Processing Systems,* Vol. 30, ed. I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 426–35. Red Hook, NY: Curran Associates, Inc.

Kim, H. H., and N. R. Swanson. 2018. "Mining Big Data Using Parsimonious Factor, Machine Learning, Variable Selection and

Shrinkage Methods." *International Journal of Forecasting* 34 (2): 339–54.

Kock, A. B., and T. Teräsvirta. 2014. "Forecasting Performances of Three Automated Modelling Techniques during the Economic Crisis 2007–2009." *International Journal of Forecasting* 30 (3): 616–31.

Lemaire, V., R. Féraud, and N. Voisine. 2008. "Contact Personalization Using a Score Understanding Method." In *2008 IEEE International Joint Conference on Neural Networks*, 649–54.

Lipovetsky, S., and M. Conklin. 2001. "Analysis of Regression in Game Theory Approach." *Applied Stochastic Models in Business and Industry* 17 (4): 319–30.

Lipton, Z. C. 2016. "The Mythos of Model Interpretability." arXiv e-prints, Vol. 1606.03490.

Ludvigson, S., and S. Ng. 2009. "A Factor Analysis of Bond Risk Premia." NBER Working Paper No. 15188.

Lundberg, S., G. Erion, H. Chen, A. DeGrave, J. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. 2020. "From Local Explanations to Global Understanding with Explainable AI for Trees." *Nature Machine Intelligence* 2 (1): 56–67.

Lundberg, S., and S.-I. Lee. 2017. "A Unified Approach to Interpreting Model Predictions." In *Advances in Neural Information Processing Systems,* Vol. 30, ed. I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 4765–74. Red Hook, NY: Curran Associates, Inc.

Makridakis, S., E. Spiliotis, and V. Assimakopoulos. 2018a. "The m4 Competition: Results, Findings, Conclusion and Way Forward." *International Journal of Forecasting* 34 (4): 802–8.

———. 2018b. "Statistical and Machine Learning Forecasting Methods: Concerns and Ways Forward." *PLOS ONE* 13 (3).

McCracken, M. W., and S. Ng. 2016. "FRED-MD: A Monthly Database for Macroeconomic Research." *Journal of Business and Economic Statistics* 34 (4): 574–89.

Medeiros, M., G. Vasconcelos, A. Veiga, and E. Zilberman. 2021. "Forecasting Inflation in a Data-Rich Environment: The Benefits of Machine Learning Methods." *Journal of Business and Economic Statistics* 39 (1): 98–119.

Miller, T. 2019. "Explanation in Artificial Intelligence: Insights from the Social Sciences." *Artificial Intelligence* 267 (February): 1–38.

Ng, S., and J. Wright. 2013. "Facts and Challenges from the Great Recession for Forecasting and Macroeconomic Modeling." *Journal of Economic Literature* 51 (4): 1120–54.

Pagan, A. 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors." *International Economic Review* 25 (1): 221–47.

Parmezan, A. R. Sabino, V. Souza, and G. Batista. 2019. "Evaluation of Statistical and Machine Learning Models for Time Series Prediction: Identifying the State-of-the-Art and the Best Conditions for the Use of Each Model." *Information Sciences* 484 (May): 302–37.

Racine, J. 2000. "Consistent Cross-validatory Model-selection for Dependent Data: *hv*-block Cross-validation." *Journal of Econometrics* 99 (1): 39-61.

Ribeiro, M., S. Singh, and C. Guestrin. 2016. "Why Should I Trust You?: Explaining the Predictions of Any Classifier." In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Vol. 22, 1135–44. New York, NY: Association for Computing Machinery.

Robnik-Šikonja, M., and I. Kononenko. 2008. "Explaining Classifications for Individual Instances." *IEEE Transactions on Knowledge and Data Engineering* 20 (5): 589–600.

Runkle, D. 1998. "Revisionist History: How Data Revisions Distort Economic Policy Research." *Quarterly Review* (Federal Reserve Bank of Minneapolis) 22 (4): 3–12.

Sermpinis, G., C. Stasinakis, K. Theofilatos, and A. Karathanasopoulos. 2014. "Inflation and Unemployment Forecasting with Genetic Support Vector Regression." *Journal of Forecasting* 33 (6): 471–87.

Shapley, L. 1953. "A Value for n-person Games." *Contributions to the Theory of Games*, Vol. 2, ed. H. W. Kuhn and A. W. Tucker, 307–17. Princeton University Press.

Shrikumar, A., P. Greenside, and K. Anshul. 2017. "Learning Important Features through Propagating Activation Differences." arXiv: 1704.02685.

Snijders, T. 1988. "On Cross-validation for Predictor Evaluation in Time Series." In *On Model Uncertainty and Its Statistical Implications*, ed. T. Dijkstra, 56–69. Springer.

Stock, J., and M. Watson. 2002. "Forecasting Using Principal Components from a Large Number of Predictors." *Journal of the American Statistical Association* 97 (460): 1167–79.

Stone, C. 1982. "Optimal Global Rates of Convergence for Nonparametric Regression." *Annals of Statistics* 10 (4): 1040–53.

Štrumbelj, E., and I. Kononenko. 2010. "An Efficient Explanation of Individual Classifications Using Game Theory." *Journal of Machine Learning Research* 11 (March): 1–18.

Sundararajan, M., and A. Najmi. 2020. "The Many Shapley Values for Model Explanation." In *Proceedings of Machine Learning Research*, Vol. 119, ed. H. Daumé and A. Singh, 9269–68. Proceedings of the International Conference on Artificial Intelligence and Statistics, July 13–18.

Wang, M., and C. Manning. 2013. "Effect of Non-linear Deep Architecture in Sequence Labeling." In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, 1285–91. Asian Federation of Natural Language Processing.

Young, P. 1985. "Monotonic Solutions of Cooperative Games." *International Journal of Game Theory* 14 (2): 65–72.