Volume 21, Number 1 January 2025

IJCB
INTERNATIONAL JOURNAL OF CENTRAL BANKING

IJCB INTERNATIONAL JOURNAL OF CENTRAL BANKING Volume 21, Number 1 January 2025

IJCB International Journal of Central Banking

Volume 21, Number 1 January 2025

*The contents of this journal, together with additional materials provided by article authors, are available without charge at **www.ijcb.org**.*

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

Requests for permission to reprint material from this journal should be addressed to:

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

Monetary Policy, Firms' Extensive Margin, and Productivity*

Benny Hartwig and Philipp Lieberknecht
Deutsche Bundesbank

This paper explores the effect of conventional monetary policy on aggregate productivity through firms' decisions to enter into or exit production. In a general equilibrium model with heterogeneous firms, we show that a monetary easing lowers productivity if it raises corporate profits: a rise in profitability allows low-productivity incumbents to remain active and unproductive new firms to enter production. Empirically, we find that expansionary monetary policy indeed raises profits, reduces firm exit, and increases entry. However, we do not find compelling evidence of an associated fall in aggregate productivity. Productivity decreases for small firms only. Entry and exit of unproductive firms induced by monetary policy hence appear of less quantitative importance for aggregate productivity than the theory suggests.

JEL Codes: E24, E32, E52, E58, L11.

1. Introduction

Since the Global Financial Crisis, the U.S. economy has featured two prominent characteristics: a protracted slowdown of productivity growth and an unprecedented degree of monetary stimulus. This

*Corresponding author: philipp.lieberknecht@bundesbank.de. Contact address: Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main, Germany. We would like to thank the editor Keith Kuester, Céline Poilly, and two anonymous referees for insightful reports, Refet Gürkaynak for kindly sharing his data set on monetary policy announcement surprises and helpful comments, Aydan Dogan and Federico Ravenna for excellent discussions, as well as Michael Binder, Ambrogio Cesa-Bianchi, Jannic Alexander Cutura, Ulrich Grosch, Sören Karau, Joachim Keller, Vivien Lewis, and participants of the Bundesbank's research center seminar, the 2020 VfS Jahrestagung, the Norges Bank 23rd Central Bank Macroeconomic Modelling Workshop and the 2020 ESCB Research Cluster on Monetary Economics Workshop for helpful comments and discussions. The views expressed in this paper are those of the authors and do not necessarily coincide with the views of the Deutsche Bundesbank, the European Central Bank, or the Eurosystem.

observation has sparked considerable interest in the role of central banks in driving productivity. The existing literature finds several channels favoring a *positive* effect in the sense that expansionary monetary policy increases productivity (Evans 1992; Moran and Queralto 2018; Garga and Singh 2021). These explanations focus on the extent to which monetary policy affects the productivity of *incumbent* firms. In contrast, the influence of monetary policy on the *composition* of actively producing firms and their individual productivity has received limited attention. This is fairly surprising given that the long-known notion of "zombification" (Hoshi and Kashyap 2004; Caballero, Hoshi, and Kashyap 2008)—which recently regained prominence amid the COVID-19 pandemic—is potentially associated with a *negative* productivity effect: expansionary monetary conditions might facilitate the survival of unproductive "zombie firms."

In this paper, we explore the influence of conventional monetary policy on the decisions of firms of heterogeneous productivity to enter into or exit production. The key idea is that monetary policy inherently affects corporate profitability whenever it alters aggregate demand and production costs. Changes in corporate profitability, in turn, have repercussions on firms' decisions about whether to produce or to become idle. This *firm-extensive margin* determines the *composition* of active production units and alters aggregate productivity—in addition to effects on incumbents—if entering and exiting firms differ in their productivity.[1]

As a first contribution, we showcase this notion in a dynamic stochastic general equilibrium (DSGE) model with heterogeneous firms and endogenous entry and exit. Upon entry into the market, firms draw an individual productivity level. This remains constant throughout their entire life cycle. We thereby isolate the composition channel and eliminate effects on the productivity of incumbents. Firms that are unable to generate profits in a given period remain

[1] We use the term firm-extensive margin to refer to firms' decisions to produce actively or to be inactive, similar to Bergin and Corsetti (2008) and analogous to the canonical extensive margins of labor supply (Heckman 1993) and exports (Hummels and Klenow 2005). We label these decisions as (production) entry and exit. Our theoretical model features a second form of entry and exit (into and out of the market) in the sense that some firms are newly created or closed down permanently.

inactive. As profits are strictly increasing in individual productivity, production entry and exit depend on a productivity threshold: only firms that feature a higher productivity than the threshold are profitable and produce actively.

We show that this framework is able to generate opposing views on how entry and exit shape the role of conventional monetary policy for productivity. If an expansionary interest rate shock raises profitability, it lowers aggregate productivity: higher corporate profitability allows unproductive incumbents to remain active and new low-productivity firms to start production. As a result, exit rates fall while entry increases. This scenario occurs with sticky wages, which implies that the effect of higher sales on profits outweighs the rise in production costs.[2] The converse case of expansionary monetary policy increasing aggregate productivity requires that the rise in production costs dominates such that corporate profitability decreases, which happens in the absence of wage rigidities. In this scenario of a "survival of the fittest," exit increases and entry declines.

As a second contribution, we test these theoretical predictions empirically by analyzing the dynamic effects of monetary policy on corporate profits, firm dynamics, and productivity. We employ a structural macrofinancial vector autoregression (VAR) based on sign restrictions and high-frequency asset price movements to identify monetary policy shocks, similar to Jarociński and Karadi (2020). We find that expansionary monetary policy increases corporate profits and affects both sides of the firm-extensive margin: Following an exogenous decrease in nominal interest rates, entry rises whereas exit initially declines and then overshoots in the medium run. These empirical findings regarding entry and exit are in line with the first model variant featuring wage stickiness.

However, we do not find compelling evidence of a systematic effect of monetary policy on aggregate productivity: Our empirical results feature insignificant responses of aggregate measures of productivity (derived from growth accounting) and of average firm-level productivity (estimated from firm-level balance sheet data) to monetary policy shocks. As such, the empirical results regarding overall productivity are not in line with either variant of the model.

[2]This case is also studied by Colciago and Silvestrini (2022) and Hamano and Zanetti (2022); see below for a discussion of similarities and differences.

If anything, the productivity effect of monetary policy seems to be empirically relevant and negative for small firms (classified according to sales), and especially for the two largest sectors, manufacturing and services.

What can be learned from our findings? First, expansionary monetary policy reduces average productivity across small firms in the data. The simple view of firm heterogeneity in the theoretical framework thus appears to be an appropriate representation for smaller firms. Second, however, aggregate productivity does not fall in the data following a monetary easing, contrary to the model predictions. This suggests a quantitatively limited importance of entry and exit of unproductive firms for the overall productivity effect of monetary policy. It also indicates that the theoretical framework misses a key dimension of reality. One obvious element that our model is deliberately silent on is the notion that monetary policy affects firms' productivity at different stages of their life cycle.

A third insight is hence that our findings also feature implications for the complementary literature on monetary policy's effect on the productivity of incumbent firms. This research strand argues that a monetary easing raises productivity through various mechanisms such as increased variable capital utilization (Christiano, Eichenbaum, and Evans 2005), enhanced incentives for research and development (Moran and Queralto 2018; Garga and Singh 2021), lower financial frictions (Midrigan and Xu 2014; Moll 2014),[3] or heterogeneous price pass-throughs (Meier and Reinelt 2022). A possible interpretation of the insignificant response of aggregate productivity in the data is thus that productivity-raising channels for incumbents (from which our theoretical framework abstracts) and the firm-extensive margin are simultaneously active in the real world. For small firms, the latter may dominate, implying that a monetary easing reduces productivity. Across all firms, the impact on incumbents and firm dynamics counteract each other, potentially explaining an overall productivity effect near zero.

Besides the literature on monetary policy and productivity, this paper connects to existing research on firm dynamics in general

[3]These papers posit a negative effect of financial frictions on productivity, and expansionary monetary policy typically alleviates financial frictions (Bernanke, Gertler, and Gilchrist 1999; Gertler and Karadi 2011).

equilibrium models along the lines of Hopenhayn (1992), Melitz (2003), and Ghironi and Melitz (2005).[4] Assuming exogenous exit, such models have been used to study the influence of endogenous entry on business cycles (Jaimovich and Floetotto 2008; Bilbiie, Ghironi, and Melitz 2012), the transmission of monetary policy (Lewis 2009; Lewis and Poilly 2012), and optimal monetary policy (Bergin and Corsetti 2008; Lewis 2013; Bilbiie, Fujiwara, and Ghironi 2014; Cacciatore, Fiori, and Ghironi 2016). In the context of endogenous exit, several papers analyze aggregate productivity shocks (Clementi and Palazzo 2016; Hamano and Zanetti 2017, 2018; Rossi 2019). In contrast to these papers, we investigate the transmission of monetary policy in a framework featuring endogenous entry and exit to analyze the effect on productivity. A similar approach is taken in two recent studies by Colciago and Silvestrini (2022) and Hamano and Zanetti (2022), who focus on market concentration and optimal monetary policy, respectively. In comparison, we discuss the flexibility of this framework to generate opposing views on aggregate productivity and provide empirical evidence on the role of entry and exit for productivity.

Lastly, this paper relates to the banking literature on "zombie lending." The key notion within this literature is that banks may grant new credit or prolong existing loans to financially distressed corporate borrowers (Hoshi and Kashyap 2004; Peek and Rosengren 2005). Existing studies explore whether policy choices such as non-standard monetary policy measures (Acharya et al. 2019; Antoni and Sondershaus 2021; Bittner, Fecht, and Georg 2021) and bank regulation (Andrews and Petroulakis 2019; Acharya, Lenzu, and Wang 2021) encourage lending to "zombie firms." Our analysis complements this microeconometric research by providing a different and macroeconomic perspective: we investigate whether expansionary conventional interest rate policy allows unproductive firms to remain active through its effect on firm profits and the associated implications for productivity.

The rest of the paper is structured as follows. Section 2 outlines the theoretical framework. In Section 3, we discuss the interplay

[4]Hopenhayn (1992) considers perfect competition, whereas the latter two papers introduce monopolistic competition and focus on international trade, i.e., entry and exit to export markets.

between monetary policy, profits, firm dynamics, and productivity within the model. Section 4 presents the empirical analysis and results. Section 5 discusses policy implications and concludes.

2. Theoretical Framework

We first present our theoretical framework, a DSGE model with endogenous firm entry and exit à la Hopenhayn (1992), Melitz (2003), and Ghironi and Melitz (2005). The economy is populated by a continuum of households that consume a variety of differentiated goods. These goods are produced by heterogeneous firms, who enter and exit production according to their (expected) profits. Exit from production takes two forms: firms can decide to suspend production temporarily (be idle) or may be forced to close down permanently due to an exogenous exit shock. In line with this, entry also takes two forms: entering the market for the first time and returning from idleness. Upon entry into the market for the first time, firms have to pay fixed entry costs and draw an individual productivity level, which remains constant throughout their life cycle. A firm that is alive (i.e., has entered the market and has not been forced to exit permanently) decides to produce or to be idle in each period. Production is subject to additional per-period fixed operational costs. Firms need to cover these costs by obtaining loans from financial intermediaries. Prices and wages are subject to nominal rigidities.

2.1 Firms

There is a continuum of firms, each producing a different good $\omega \in \Omega$ using labor as the only production factor.[5] Overall firm productivity is given by aggregate productivity A_t and idiosyncratic productivity z, with the latter remaining constant over the entire life cycle of the firm. The production function of a given firm can hence be written as

$$y_t^C(z) = A_t z l_t^C(z), \tag{1}$$

[5] We abstract from physical capital to keep the model as simple as possible. An extended model features qualitatively and quantitatively similar business cycle properties (Bilbiie, Ghironi, and Melitz 2012). An interesting avenue for future research is the interaction of individual productivity and firm-specific physical capital.

where $y_t^C(z)$ denotes consumption output produced by a firm with individual productivity z, and $l_t^C(z)$ is the corresponding amount of labor demand. Aggregate productivity evolves according to an autoregressive AR(1) process in logs.

Labor Demand and Price Setting. Production is subject to fixed operational costs f of effective labor units at the beginning of each period. At this stage, firms do not have funds available and hence obtain loans from financial intermediaries at the nominal gross interest rate R_t to prepare production. This reflects a working capital channel in the spirit of Ravenna and Walsh (2006).[6] Total real costs of production TC_t are given by

$$TC_t(z) = w_t \left(l_t^C(z) + f \frac{R_t}{A_t} \right), \tag{2}$$

where w_t is the real wage. Cost minimization then yields

$$mc_t(z) = \frac{w_t}{A_t z}, \tag{3}$$

which shows that marginal costs differ across firms depending on idiosyncratic productivity. As outlined below, household demand for a specific good is given by

$$y_t^C(z) = \left(\frac{p_t(z)}{P_t} \right)^{-\theta} Y_t^C, \tag{4}$$

where $p_t(z)$ is the nominal individual price, P_t is the aggregate price index, Y_t^C is overall consumption demand, and θ is the constant elasticity of substitution between goods.

The goods market is monopolistically competitive. Each firm chooses its price to maximize the sum of current profits and the firm value (the expected discounted value of the profit stream from $t+1$ onward)[7] subject to its production function, taking the household demand schedule and aggregate variables as given. Firms face

[6]The results are qualitatively and quantitatively similar if all wages have to be paid before production.

[7]The firm value is given by $v_t(z) = E_t \sum_{s=t+1}^{\infty} \Lambda_s d_s(z)$, where Λ_s denotes the household's stochastic discount factor. The pricing problem hence implicitly accounts for the probability to exit production via the expectations operator.

quadratic price adjustment costs following Rotemberg (1982). The costs of adjusting prices in real terms, pac_t, are

$$pac_t(z) = \frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right)^2 \rho_t(z) y_t^C(z), \tag{5}$$

where

$$\rho_t(z) = \frac{p_t(z)}{P_t} \tag{6}$$

denotes the real price of firm z. Real profits of a given firm can hence be written as

$$\begin{aligned} d_t(z) &= \rho_t(z) y_t^C(z) - w_t l_t^C(z) \\ &\quad - \frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right)^2 \rho_t(z) y_t^C(z) - f\frac{w_t R_t}{A_t}, \end{aligned} \tag{7}$$

where the first term captures revenues, and the remaining terms are different costs (discussed further in the next section). As shown in Appendix A.2, the optimal real price satisfies

$$\rho_t(z) = \mu_t(z) mc_t(z), \tag{8}$$

where the markup over marginal costs is given by

$$\mu_t(z) = \frac{\theta}{(\theta - 1)\left[1 - \frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right)^2\right] + \tau \Upsilon_t(z)}, \tag{9}$$

where the auxiliary term $\Upsilon_t(z)$ is defined as

$$\begin{aligned} \Upsilon_t(z) &= \frac{p_t(z)}{p_{t-1}(z)}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right) \\ &\quad - E_t\left[\Lambda_{t+1}\frac{y_{t+1}^C(z)}{y_t^C(z)}\frac{P_t}{P_{t+1}}\left(\frac{p_{t+1}(z)}{p_t(z)} - 1\right)\left(\frac{p_{t+1}(z)}{p_t(z)}\right)^2\right], \end{aligned} \tag{10}$$

and where Λ_{t+1} denotes the household stochastic discount factor (defined further below). Optimal prices are thus heterogeneous across firms of differing productivity, as both marginal costs and optimal markups differ.[8]

[8]The markup would be identical across firms absent nominal rigidities ($\tau = 0$) and equivalent to the textbook expression $\theta/(\theta - 1)$.

Entry and Exit. Each period, firms enter and exit production depending on their current and expected profitability. There is an unbounded mass of ex ante homogeneous prospective entrants. When entering the market for the first time, each firm draws an idiosyncratic productivity level z from a distribution $G(z)$ with support on $[z_m, \infty)$ and starts to produce in the next period after some time to build. Market entry is subject to entry costs f_E of effective labor units. Following Lewis and Poilly (2012), we assume that some market entries fail.[9] Denoting the total number of new firms entering the market by H_t, the success probability is given by

$$\Psi_t(H_t, H_{t-1}) = 1 - F_{H,t}\left(\frac{H_t}{H_{t-1}}\right), \tag{11}$$

which has the properties $F_H(1) = F'_H(1) = 0, F''_H(1) = \psi > 0$. Potential entrants are forward-looking and decide to enter the market based on the firm value, i.e., the expected value of operation v_t (which is determined via a household asset pricing equation; see below). In equilibrium, firm dynamics yield the following free entry condition:

$$f_E \frac{w_t}{A_t} = v_t(\Psi_t + \Psi'_t H_t) + \beta E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-1} v_{t+1}\Psi''_{t+1}H_{t+1}\right]. \tag{12}$$

In Equation (12), the left-hand side represents the costs associated with market entry. These are equated to the expected value of operation on the right-hand side, accounting for changes in the entry success probability induced by the number of entrants.

Turning to firm exit, an incumbent firm produces actively in a given period if its profits are positive, i.e., if $d_t(z) > 0$, and decides to be idle if profits are zero or negative. As such, only a subset of firms $\Omega_t \in \Omega$ are actively producing in any given period. The decision to exit production hence depends on firms' idiosyncratic productivity. The cutoff level of productivity $\bar{z}_t$ is defined by a zero profit condition given by

$$\bar{d}_t \equiv d_t(\bar{z}_t) = 0. \tag{13}$$

[9]This assumption guarantees a gradual response of entry in response to exogenous disturbances. As shown below, this is in line with our empirical findings.

Firms with $z > \bar{z}_t$ make positive profits and thus produce actively, whereas low-productivity firms with $z \leq \bar{z}_t$ decide to be idle. As a result, only relatively more productive firms are active in any period, and some exit from production takes place endogenously. In addition, each firm faces an exogenous exit shock at the end of each period, which occurs with probability δ, and forces firms to close down permanently. The total number of existing firms, N_t, thus evolves according to

$$N_t = (1-\delta)(N_{t-1} + \Psi_{t-1}(H_t, H_{t-1})H_{t-1}), \tag{14}$$

whereas the number of actively producing firms is given by[10]

$$S_t = (1 - G(\bar{z}_t))N_t. \tag{15}$$

2.2 Households

There is a continuum of infinitely lived identical and atomistic households. The representative household maximizes expected utility U_t, given by

$$U_t = E_t \left[\sum_{s=t}^{\infty} \beta^{s-t} \left(log(C_s) - \chi \frac{L_s^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}} \right) \right], \tag{16}$$

where C_s is consumption and L_s denotes labor supply.[11] The discount factor is given by β, and η is the elasticity of labor supply to wages. Consumption is defined as a basket of individual varieties ω over a continuum of goods Ω. Consumption preferences follow Dixit and Stiglitz (1977), such that the elasticity of substitution between individual goods (θ) is constant. This yields the demand schedule for a given variety as shown in Equation (4).

Households can invest in (government) bonds and equity shares in a mutual fund of firms. Bonds yield a safe gross nominal interest rate R_t in the next period. The mutual fund pays out dividends equal

[10]The notation of S_t follows Hamano and Zanetti (2017), who embrace the notion that these firms are "surviving" product destruction. However, firms that decide to remain idle in a given period are also "alive," but are not included in S_t.

[11]The utility function follows King, Plosser, and Rebelo (1988) and ensures the existence of a balanced growth path.

to total firm profits in each period. In period t, the representative household obtains equity shares x_{t+1} at the real share price v_t (the firm value).[12] In addition to interest income and dividend income, the household receives income by selling its existing shareholdings and by supplying labor at the real wage w_t. The budget constraint in real terms hence reads

$$C_t + x_{t+1}v_t(N_t + H_t) + B_{t+1} = w_t L_t + x_t N_t v_t + x_t S_t \widetilde{d}_t + \frac{R_{t-1}}{\pi_t^C} B_t, \tag{17}$$

where B_t are real holdings of bonds, $\widetilde{d}_t$ is the average dividend across active firms, and π_t^C denotes the gross consumption-based inflation rate:

$$\pi_t^C = \frac{P_t}{P_{t-1}}. \tag{18}$$

The household maximizes expected utility by choosing consumption, labor supply, and its portfolio allocation subject to the budget constraint in Equation (17). The first-order condition with respect to bond holdings is a standard Euler equation given by

$$1 = \beta E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-1} \frac{R_t}{\pi_{t+1}^C} \right]. \tag{19}$$

The optimality condition with respect to shareholdings is given by

$$v_t = E_t \left[\Lambda_{t+1} \left(v_{t+1} + \frac{S_t}{N_t} \widetilde{d}_{t+1} \right) \right], \tag{20}$$

where the stochastic discount factor, Λ_{t+1}, is defined by

$$\Lambda_{t+1} = \beta(1-\delta) E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-1} \right]. \tag{21}$$

The labor market is monopolistically competitive: households have some market power and set their wages. The differentiated labor supplied by each household is aggregated by a union and hired by firms on a competitive market. The real wage is then given by

[12] By assumption, the household does not know which firms operate in the next period. As a result, it finances all incumbent and new firms during a given period.

$$
\begin{aligned}
w_t &= \left(\int_0^1 (w_t(j))^{1-\theta_W} dj \right)^{\frac{1}{1-\theta_W}} \\
&= \left(\lambda_W \left(\frac{w_{t-1}}{\pi_t^C} \right)^{1-\theta_W} + (1-\lambda_W)\,(w_t^*)^{1-\theta_W} \right)^{\frac{1}{1-\theta_W}} . \qquad (22)
\end{aligned}
$$

2.3 Aggregation

Following Melitz (2003) and Ghironi and Melitz (2005), we specify that individual firm productivity is drawn from a Pareto distribution

$$
G(z) = 1 - \left(\frac{z_m}{z} \right)^{\kappa}, \qquad (23)
$$

where z_m is the minimum possible productivity level and κ governs the shape and dispersion of the distribution. Since the cutoff level of productivity $\bar{z}_t$ varies over the business cycle, the average productivity across active firms is time-varying as well:

$$
\widetilde{z}_t \equiv \left[\frac{1}{1-G(\bar{z}_t)} \int_{\bar{z}_t}^{\infty} z^{\theta-1} dG(z) \right]^{\frac{1}{\theta-1}} = \bar{z}_t \left[\frac{\kappa}{\kappa-(\theta-1)} \right]^{\frac{1}{\theta-1}} . \qquad (24)
$$

Variables referring to firms with average productivity are denoted similarly in the following, i.e., $\widetilde{a}_t \equiv a(\widetilde{z}_t)$ for a generic variable a. The average markup is given by

$$
\widetilde{\mu}_t = \frac{\theta}{\begin{array}{c}(\theta-1)\left(1-\frac{\tau}{2}(\pi_t-1)^2\right) \\ +\,\tau\left(\pi_t(\pi_t-1) - E_t\left[\Lambda_{t+1}\frac{Y_{t+1}^C}{Y_t^C}\frac{S_t}{S_{t+1}}(\pi_{t+1}-1)\,\pi_{t+1}\right]\right)\end{array}} . \qquad (25)
$$

Equation (25) is the nonlinear Phillips curve in our model,[13] relating average markups to producer price inflation π, which is linked to consumer price inflation by

[13]One can show that a log-linear version of Equation (25) reduces to an augmented New Keynesian Phillips curve. In contrast to the model by Bilbiie, Ghironi, and Melitz (2008) with exogenous exit, the number of *active* firms (S) determines inflation dynamics, instead of the total number of firms (N). We briefly discuss the implications of endogenous exit for inflation dynamics in Appendix A.4.

$$\pi_t = \frac{\widetilde{\rho}_t}{\widetilde{\rho}_{t-1}} \pi_t^C. \tag{26}$$

The number of active firms can be written as

$$S_t = (1 - \zeta_t) N_t, \tag{27}$$

where the endogenous fraction ζ_t of exits from production due to low productivity is

$$\zeta_t \equiv 1 - G(\bar{z}_t) = 1 - \left(\frac{z_m}{\bar{z}_t} \right)^{\kappa}. \tag{28}$$

Finally, the price index captures a variety effect stemming from consumer preferences:

$$\widetilde{\rho}_t = S_t^{\frac{1}{\theta - 1}}. \tag{29}$$

Market Clearing. Equilibrium on the goods market requires that aggregate consumption output equals the sum of private consumption and price adjustment costs:

$$Y_t^C = C_t + S_t \widetilde{pac}_t = \left(1 - \frac{\tau}{2} (\pi_t - 1)^2 \right)^{-1} C_t. \tag{30}$$

The aggregate accounting identity equates aggregate output to the sum of labor and dividend income:

$$C_t + v_t H_t = w_t L_t + S_t \widetilde{d}_t. \tag{31}$$

Aggregate output is consumption plus investment:

$$Y_t = C_t + I_t, \tag{32}$$

and investment is the creation of new firms:

$$I_t = v_t H_t. \tag{33}$$

The equilibrium on the labor market requires that

$$L_t = S_t \left(\widetilde{l}_t^C + \frac{f}{A_t} \right) + H_t \frac{v_t}{w_t}. \tag{34}$$

To close the model,[14] we assume a central bank interest rate rule given by

$$log\left(\frac{R_t}{R}\right) = \phi_R \, log\left(\frac{R_{t-1}}{R}\right) + (1-\phi_R)\left[\phi_\pi \, log\left(\frac{\pi_t}{\pi}\right) + \phi_y \, log\left(\frac{Y_t}{Y_{t-1}}\right)\right] + \varepsilon_t^M. \tag{35}$$

The central bank thus responds to deviations of producer price inflation from steady state and output growth.[15] ε_t^M is a monetary policy shock, which we analyze in the following.

3. The Theoretical Effect of Monetary Policy

In this section, we analyze how conventional monetary policy affects firm entry and exit in our theoretical model. We show that the resulting productivity effect depends crucially on the reaction of corporate profitability to the monetary shock. We contrast two model variants yielding opposing predictions, which we test empirically in the next section.

3.1 Calibration

The following numerical analysis is based on standard parameter values and estimates for the U.S. economy. We interpret periods as quarters and set $\beta = 0.99$, equivalent to an annualized steady-state real interest rate of 4 percent, and consider a steady-state gross inflation rate $\pi = 1$. Regarding household preferences, we set the elasticity of labor supply $\eta = 2$ and calibrate $\chi = 0.90$ to normalize steady-state labor supply $L = 1$.

With respect to the firm parameters, the entry cost f_E and the minimum productivity level z_m are set to unity, without loss of

[14] Appendix A.1 shows the equilibrium equations and the steady-state computation.

[15] As discussed by Bilbiie, Ghironi, and Melitz (2008), a response to welfare-based consumer price inflation π^C is infeasible in reality due to infrequent updating of the baskets used to measure inflation. Actual consumer price inflation is closer to p_t than P_t. Aghion et al. (2019) discuss how firm dynamics raise difficulties for measuring inflation and growth.

generality. We follow Ghironi and Melitz (2005) and calibrate the elasticity of substitution between goods $\theta = 3.8$ and the shape parameter of the productivity Pareto distribution $\kappa = 3.4$. As in Hamano and Zanetti (2018), we choose the fixed costs f and the exogenous exit rate δ to match annual U.S. entry and exit rates, which were 12.3 percent and 10.6 percent over 1977–2016 according to the Business Dynamics Statistics (BDS). Using this entry rate in Equation (14) implies $\delta = 0.03$. Together with the exit rate, this yields $f = 0.009$ and a steady-state ratio between average and cutoff productivity of $\widetilde{z}/\bar{z} = 1.86$ (see Appendix A.2). Following Lewis and Poilly (2012), we calibrate the firm entry costs parameter $\psi = 8.31$.

Turning to the parameters for nominal rigidities, the elasticity of substitution between differentiated labor is set to $\theta_W = 21$, implying a steady-state wage markup of 1.05 as in Christiano, Eichenbaum, and Evans (2005). We calibrate the fraction of non-adjusting firms as $\lambda_W = 0.75$. The Rotemberg price adjustment parameter τ is set to $\tau = 77$, in line with Bilbiie, Fujiwara, and Ghironi (2014). The monetary policy parameters are calibrated as $\phi_R = 0.8$, $\phi_\phi = 1.5$, $\phi_y = 0.5/4$.

3.2 Monetary Policy Shocks and the Role of Profits

We now aim to understand how conventional monetary policy shocks affect firm dynamics in the model. The framework implies that entry and exit decisions are tightly linked to corporate profits and expectations thereof. On the one hand, the entry condition in Equation (12) stipulates that more firms decide to enter the market if the expected firm value rises. In turn, Equation (20) specifies that the firm value is given by the sum of discounted average profits. On the other hand, entry into and exit from production depends on the zero profit condition in Equation (13). It is thus instructive to consider the different components of profits in closer detail. As outlined in Equation (7), profits of a firm with productivity z are given by

$$d_t(z) = \underbrace{\rho_t(z) y_t^C(z)}_{(1)} - \underbrace{w_t l_t^C(z)}_{(2)} \\ \underbrace{- \frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right)^2 \rho_t(z) y_t^C(z)}_{(3)} - \underbrace{f\frac{w_t R_t}{A_t}}_{(4)}. \tag{36}$$

One can thus decompose profits into four components: (1) sales revenues, (2) labor costs, (3) price adjustment costs, and (4) fixed costs. The overall response of profits to changes in monetary conditions thus depends on the relative effects on these components. Inspecting the latter two components directly reveals that their quantitative importance is limited. Price adjustment costs are a squared function of price changes, which makes this term generally small (as also argued by Bilbiie, Fujiwara, and Ghironi 2014). In the fourth term, wages and the interest rate move in opposite directions after monetary policy shocks.

The key question of interest is thus how strongly revenues and labor costs, labeled (1) and (2) above, react to changes in monetary conditions. On the one hand, expansionary monetary policy stimulates aggregate demand from households by decreasing the real interest rate. The additional demand is—ceteris paribus, i.e., before price changes and other general equilibrium effects—distributed proportionally across all firms. As a result, demand for the individual variety increases. This raises sales, profits, and the firm value directly. On the other hand, labor costs rise as well after a monetary expansion. Firms need to hire more workers to expand their production to satisfy the additional demand. As a compensation for the higher labor supply, workers demand a higher wage. To showcase this notion, we contrast two model variants in the following. Variant A features wage rigidities as outlined in the previous section. Variant B assumes perfectly flexible wages ($\lambda_W = 0$).[16]

Figure 1 shows the impulse responses to an expansionary monetary policy shock that decreases the nominal interest rate in both model variants.[17] The monetary policy shock increases output and inflation in both cases.[18] As the real interest rate falls, households

[16]Variant A is similar to the frameworks used by Colciago and Silvestrini (2022) and Hamano and Zanetti (2022). Both contributions abstract entirely from price stickiness, such that they are inherently silent about real effects of monetary policy in the absence of wage stickiness (Variant B).

[17]The model's monetary policy variable is the interest rate on one-period (a quarter) bonds. In our empirical analysis, we use the one-year constant-maturity Treasury yield. While not identical, these measures are closely related, such that their reactions to monetary policy shocks are expected to be qualitatively and quantitatively similar.

[18]Note that the graph shows producer price inflation. Within our closed-economy model, this corresponds to GDP deflator inflation, which is the variable employed in our empirical exercise.

Figure 1. Expansionary Monetary Policy Shock in Model Variants

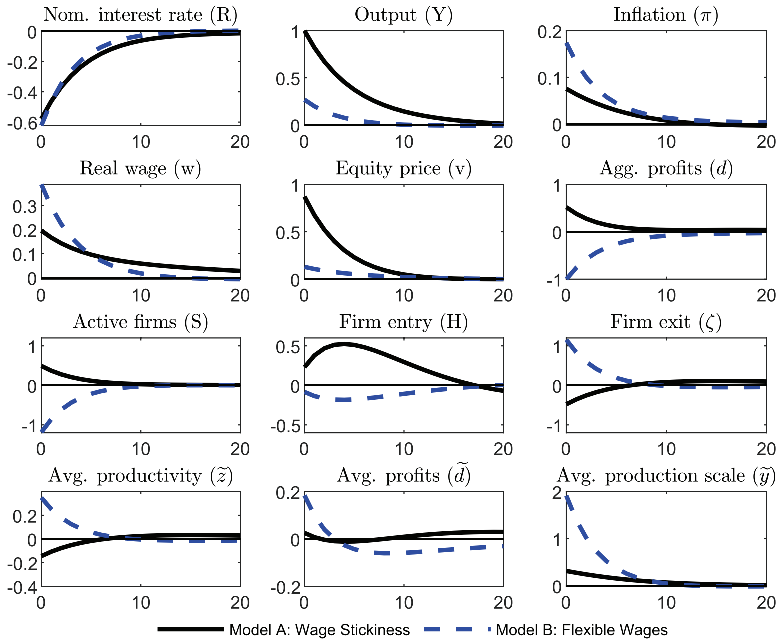

Note: Impulse response functions to an expansionary monetary policy shock. Solid lines refer to Variant A (wage stickiness), dashed lines to Variant B (flexible wages). The shock size is calibrated to yield a 1 percent increase of output in Variant A. All variables are shown in percent deviations from steady state, except for inflation, interest rate, and exit (percentage point deviations).

increase consumption and reduce bond holdings. Firms demand more labor to accommodate the higher demand for consumption goods. The tighter labor market implies higher real wages to compensate for higher labor supply.

While the macroeconomic picture is qualitatively similar across the model variants, they feature completely different views on profits, firm dynamics, and productivity. In Variant A, the economic expansion is accompanied by a procyclical response of the number of active firms (as in Colciago and Silvestrini 2022 and Hamano and Zanetti 2022). This reflects an increase of entry into production (ΔS_t), a rise in the number of firms entering the market (H_t), and a decrease of the exit rate from production (ζ_t). At the same time, aggregate profits increase, indicating that the revenue channel dominates the labor cost channel. As a result, firms become more profitable on impact. This lowers the productivity threshold level that guarantees non-negative profits. Incumbent low-productivity firms thus produce actively, such that average productivity *decreases*. At the same time, the higher corporate profitability increases firms' expected value and thereby renders equity investment more attractive to households. This induces more firms to enter the market. Among these new entrants, low-productivity firms also make positive profits, albeit featuring a lower scale of production. In general equilibrium, average firm profits rise only marginally since actively producing firms charge higher prices, which *lowers* average profits (as the elasticity of substitution is larger than unity), and because they face higher marginal costs (due to a lower average productivity).

The expansionary monetary policy shock thus allows low-productivity incumbents to produce actively and facilitates the entry into production of relatively unproductive firms. As a result, the average productivity of active firms declines. While the favorable monetary conditions prevail, unproductive firms remain profitable and thus continue to produce actively. However, as the monetary stimulus and the associated economic boom fade, the cutoff level of productivity for profitability increases again. As a result, low-productivity firms become unprofitable and decide to become idle, leading to an overshooting of firm exit in the medium run.[19]

Variant B features diametrically opposite predictions on profits, firm dynamics, and productivity: because labor costs rise sharply, firm exit is procyclical and firm entry is countercyclical in this variant. To understand this observation, note that the labor cost channel is tightly linked to the markup decisions (see Equation 36). In the presence of price adjustment costs, optimal markups are inversely related to inflation: as firms raise prices following expansionary monetary policy shocks, markups decrease. The resulting downward pressure on profits makes low-productivity firms unprofitable, such that they decide to become idle and exit increases. At the same time,

[19] Other expansionary shocks such as aggregate productivity shocks yield similar firm dynamics; see Hamano and Zanetti (2017), Rossi (2019), and Appendix A.5.

the rise in real wages implies that entry costs are higher, reducing firm entry (in line with Bilbiie, Ghironi, and Melitz 2008; Lewis 2009).[20] The fewer active firms are on average more productive, more profitable, and larger in terms of their production scale.

The theoretical framework is thus able to generate two opposing views regarding the effect of monetary policy on productivity through firm entry and exit. If expansionary monetary policy raises profitability by stimulating aggregate demand, it allows unproductive firms to be active, thus reducing overall productivity. The converse case that a monetary easing increases productivity requires that the rise in production costs dominates such that profitability decreases. As a consequence, only productive firms are profitable, such that exit increases while entry declines.

4. Empirical Analysis

In this section, we test the theoretical predictions empirically: we analyze the effects of monetary policy on (1) firm dynamics and (2) various measures of productivity.

4.1 Data

Our sample covers U.S. data from 1993:Q2 through 2017:Q4.[21] More recent observations are excluded due to the lack of availability of high-frequency financial surprises (used to identify monetary policy shocks; see below), earlier observations due to the data on firm dynamics. To capture firm entry and exit, we use quarterly data on the number of establishment births and deaths from the Bureau of Labor Statistics (BLS).[22]

We consider a variety of productivity measures. On the one hand, we use aggregate series by Fernald (2014): total factor productivity

[20] The decrease of interest rates implies a fall in the return to bonds. To restore no-arbitrage across different investments, the return to shareholdings also decreases slightly. This happens via a slight increase in the equity prices today relative to tomorrow.

[21] In Appendix B.1, we provide descriptive statistics and data charts.

[22] An establishment is a single physical location; a firm is an establishment or a combination of establishments. Rossi (2019) similarly uses the establishment series to proxy firm entry and exit decisions.

(TFP)—the Solow residual, utilization-adjusted TFP (TFPu)—a cleaner measure of pure technological change, and labor productivity in the business sector (LP).[23] On the other hand, as a closer counterpart to the theoretical productivity variable, we compute average firm-level productivity based on microdata from Compustat. To this end, we consider balance sheet data of all firms in the nonfinancial sectors except utilities, construct a firm-specific measure of the real capital stock using the perpetual inventory method, and estimate firm-level productivity using a fixed-effects regression of sales on production inputs (see Clementi and Palazzo 2019; Ottonello and Winberry 2020; and Appendix B.2).

Our monetary policy variable is the one-year constant-maturity Treasury yield, following Gertler and Karadi (2015) and Jarociński and Karadi (2020). This measure captures the effects of forward guidance and moved sufficiently even during the zero lower bound (ZLB) period.[24,25] The block of macroeconomic variables consists of real GDP and the GDP implicit price deflator. As financial variables, we include the S&P 500 stock price index (deflated by the GDP price deflator) as well as the excess bond premium of Gilchrist and Zakrajšek (2012). Including a measure of financial frictions is crucial to identify the transmission channel of monetary policy (Gertler and Karadi 2015) and the monetary policy rule (Caldara and Herbst 2019). In robustness checks, we consider alternative macrofinancial variables (see Section 4.4 and Appendix B.5).

[23]The aggregate productivity series are based on growth accounting techniques proposed by Basu, Fernald, and Kimball (2006). TFP growth is output growth not explained by (observed) input growth $\Delta TFP = \Delta Y - \alpha\Delta K - (1-\alpha)\Delta L$, where ΔY is real output growth, ΔK is capital growth, ΔL is labor growth, and α is the capital share on output. Utilization-adjusted TFP growth is TFP not explained by capital and labor utilization growth $\Delta TFPu = \Delta TFP - \Delta Util$. Labor productivity growth is defined as growth in output per hour $\Delta LP = \Delta Y - \Delta H$, where ΔH is hours worked in business sector.

[24]Forward guidance became important for U.S. monetary policy after the FOMC started issuing press releases in February 1994 (Gürkaynak, Sack, and Swanson 2005), which almost coincides with the start of our sample.

[25]Ikeda et al. (2024) show that the ZLB is empirically relevant when identifying the transmission of a monetary policy shock and find that the shadow short rate appropriately captures unconventional monetary policy. We consider the shadow short rate in a robustness exercise.

The theoretical model suggests that wages and profits constitute further variables of interest. We construct per capita wages by dividing aggregate wages and salaries by the total number of employees. Aggregate profits are measured by corporate profits after taxes with inventory valuation and capital consumption adjustment from the Bureau of Economic Analysis (BEA). We deflate both series using the GDP implicit price deflator.

4.2 Methodology

Our baseline empirical model is a VAR with high-frequency surprises along the lines of Jarociński and Karadi (2020). The high-frequency surprises are yield and stock price changes around monetary policy announcements by the Federal Open Market Committee (FOMC), which we use to identify monetary policy shocks (see further below). Let m_t be a vector of surprises in quarter t[26] and y_t be a vector of macroeconomic and financial variables. We add further variables of interest (e.g., firm exit) one by one to adopt a parsimonious estimation approach. The baseline VAR model is given by

$$\begin{pmatrix} m_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0 \\ c_Y \end{pmatrix} + \sum_{p=1}^{4} \begin{pmatrix} 0 & 0 \\ A_{p,YM} & A_{p,YY} \end{pmatrix} \begin{pmatrix} m_{t-p} \\ y_{t-p} \end{pmatrix} + \begin{pmatrix} u_t^m \\ u_t^y \end{pmatrix}, \begin{pmatrix} u_t^m \\ u_t^y \end{pmatrix} \sim \mathcal{N}(0, \Sigma), \tag{37}$$

where $\mathcal{N}$ denotes the normal distribution. The zero restrictions for m_t imply a zero mean and independence from lags of m_t and y_t; these restrictions are plausible if high-frequency financial surprises are unpredictable. We estimate this VAR using Bayesian methods in log-levels for all variables except interest rates, spreads, and high-frequency surprises, set the maximum lag length to four, and use a flat prior for our benchmark results.

To identify monetary policy shocks, we adopt a sign-restriction approach in the spirit of Jarociński and Karadi (2020). We use changes in the three-month federal funds future and the S&P

[26] m_t is the sum of intraday surprises on the days with FOMC announcements occurring in quarter t.

500 index within a tight window around FOMC announcements.[27] Changes in the three-month federal funds future reflect both surprises about actual rate-setting and near-term forward guidance and therefore constitute a broad measure of conventional monetary policy. Our identification procedure imposes opposite sign restrictions on both high-frequency surprises and their low-frequency counterparts (the interest rate and the stock price). While not the focus of our analysis, we also identify central bank information shocks to avoid confounding effects.[28,29] Table 1 shows our set of identification restrictions.

Our sign restrictions are more stringent than the approach of Jarociński and Karadi (2020), which remains agnostic about low-frequency variables. In our application, these additional restrictions are necessary to ensure a proper identification. As outlined above, the lack of data availability restricts our sample start to 1993:Q2. For this sample, sign restrictions on high-frequency variables only yield implausible interest rate dynamics after monetary policy shocks (see Figure B.6 in Appendix B.4).[30] We hence add low-frequency sign restrictions in line with our theoretical framework, standard DSGE models (e.g., Smets and Wouters 2007), and the empirical literature (e.g., Liu et al. 2019).

[27]We use an updated version of the data set by Gürkaynak, Sack, and Swanson (2005). The window starts 10 minutes before the announcement and ends 20 minutes after. Gürkaynak, Sack, and Swanson (2005) show that these changes are not driven by confounding factors like macroeconomic releases on that day.

[28]High-frequency interest rate surprises may not only reflect monetary policy shocks, but also contain information about the state of the economy (Miranda-Agrippino and Ricco 2021). It is hence essential to control for this information channel (Romer and Romer 2000; Melosi 2017; Nakamura and Steinsson 2018) when identifying monetary policy shocks, as it may bias impulse responses. Figure B.4 in Appendix B shows that an identification procedure that does not account for central bank information shocks yields a decline of real GDP and stock prices following a monetary easing.

[29]We enlarge the rotation space of orthonormal matrices to include the interest rate and the stock price to increase the set of structural models that potentially exhibit a strong link between high-frequency surprises and their low-frequency counterparts.

[30]The data from 1990:M2 through 1993:M3 (as used in Jarociński and Karadi 2020) are particularly informative about monetary policy shocks. This period coincides with the U.S. savings and loan crisis, which featured large and surprising interest rate cuts by the FOMC and associated positive stock surprises.

Table 1. Sign Restrictions

Variable	Shock		
	Monetary Policy (Negative Comovement)	**CB Information (Positive Comovement)**	**Other**
m_t, High Frequency			
Interest Rate Surprise	+	+	0
Stock Price Surprise	−	+	0
y_t, Low Frequency			
Interest Rate	+	+	0
Stock Price Index	−	+	0

Note: Sign restrictions imposed on the respective variable's impact response to shocks. Empty fields denote an unrestricted response.

Aside from the VAR, we use panel local projections (PLPs) à la Jordà (2005) and Jordà, Schularick, and Taylor (2015) to exploit the cross-sectional information of the firm-level data. The PLP model is given by

$$y_{i,t+h} - y_{i,t-1} = \alpha_{i,h} + D_t\eta_h + x_t\beta_h + \sum_{j=1}^{2} \Delta y_{i,t-j}\theta_{j,h} + \sum_{j=1}^{2} w_{t-j}\gamma_{j,h} + u_{i,t+h}, \tag{38}$$

where $y_{i,t}$ is productivity of firm i, D_t is a dummy vector to control for seasonal patterns, x_t is a monetary policy shock, and w_t is a vector of additional controls. To match the information set of the VAR, we consider the same macrofinancial controls (interest rate, output, price level, stock price, and the excess bond premium). β_h is the coefficient of interest and measures the response of firm productivity at time $t+h$ to a shock at time t. The lags of the dependent variables in first differences control for autoregressive dynamics. Standard errors are clustered two-way at the firm level and the time level.

4.3 Results

Figure 2 shows our empirical results regarding the effect of an expansionary monetary policy shock on macrofinancial variables and firm

Figure 2. Expansionary Monetary Policy and Firm Dynamics

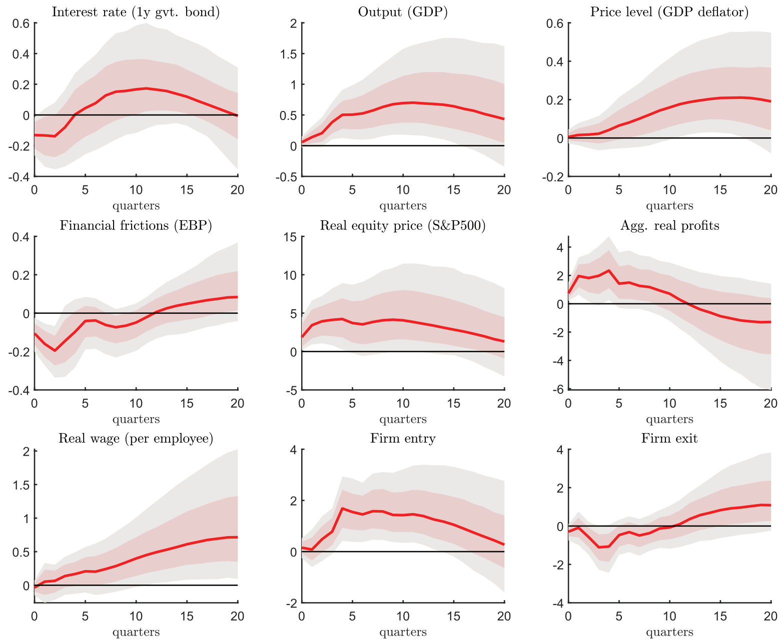

Note: Impulse response functions to a monetary policy shock identified in a VAR with FOMC announcement surprises using sign restrictions on the comovement between high- and low-frequency variables. The thick lines are the median estimates; the shaded areas depict the 68 percent and 90 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

dynamics based on the VAR. On impact, the one-year Treasury yield decreases by roughly 10 basis points. This interest rate response to the shock is very short-lived. Output and the aggregate price level increase in response to the monetary stimulus with a delay of a couple of quarters, consistent with standard theory. Stock prices increase over a prolonged period, while financial frictions decline on impact, in line with the credit channel of monetary policy. Albeit somewhat smoother, these macrofinancial impulse responses are very similar to the monthly estimates reported in recent contributions (Gertler and Karadi 2015; Caldara and Herbst 2019; Jarociński and Karadi 2020; Miranda-Agrippino and Ricco 2021).

With respect to firm dynamics, our results show that a monetary policy shock significantly affects the firm-extensive margin. Firm entry is procyclical and increases following a monetary easing. The peak effect occurs after one year at close to 2 percent. The rise in firm entry is persistent and lasts around three years. At the same time, firm exit is countercyclical and decreases following more favorable monetary conditions.[31] The number of active production units (sometimes labeled net business formation, i.e., entry minus exit) peaks one year after the monetary shock. After around two years, firm exit overshoots its long-run level and gradually reverts afterward, while the expansion of economic activity fades. These empirical results are qualitatively in line with the theoretical predictions of Variant A (see Figure 1). This empirical support for Variant A regarding firm entry and exit is corroborated by the impulse responses of wages and profits: corporate profits increase persistently after the monetary easing (in line with Lewis and Poilly 2012), while wages rise sluggishly over the medium run, likely reflecting nominal rigidities. The first set of empirical results regarding firm entry and exit thus supports the notion invoked by Variant A (with sticky wages), while contradicting Variant B (with flexible wages).

Our result regarding firm entry is consistent with Lewis (2009), Lewis and Poilly (2012), and Bergin, Feng, and Lin (2018), while Hamano and Zanetti (2022) document similar results for entry and exit. All of these studies are based on pre-2000 data to proxy entry and exit. As such, our analysis provides new empirical evidence that an expansionary monetary policy shock induces a rise in firm entry in more recent times, shows that both sides of the firm-extensive margin are affected, and documents an overshooting of firm exit.[32] Our

[31] Firm exit is also unconditionally countercyclical in our sample; see Appendix B.1, in line with earlier findings by Campbell (1998) and Jaimovich and Floetotto (2008).

[32] These studies use short-run restrictions on output and prices for identification. In our more recent sample characterized by forward guidance, such restrictions are insufficient for identification, as they do not fully capture the central bank information set (Gürkaynak, Sack, and Swanson 2005). Short-run restrictions are also problematic when including financial variables, as these may

results share similarities with earlier findings for productivity shocks by Rossi (2019), who documents that firm exit falls initially after a positive aggregate productivity shock, but overshoots its long-run level after approximately two years. She rationalizes these findings in a model similar to our Variant A.

Turning to the empirical response of productivity, Figure 3 shows the effects of an expansionary monetary policy shock on various productivity measures. The upper panels (A) are based on the VAR, whereas the lower panels (B) originate from PLPs using the sign-identified monetary policy shocks from the VAR.

Overall, the empirical results regarding productivity differ across measures and methods. Among the aggregate measures (shown in the first row), TFP rises persistently and significantly following more favorable monetary conditions. In contrast, the responses of utilization-adjusted TFP and labor productivity are largely insignificant, which suggests that the rise of TFP is driven by (nontechnological) variable input utilization.[33] The response of average firm-level productivity is also insignificant at conventional levels in the VAR (second row, left panel). However, a closer inspection (second row, middle and right panels) reveals that the productivity response depends on firm size: productivity declines somewhat for comparably small firms (classified as having sales below $10 million).

The PLP impulse responses are qualitatively and quantitatively highly similar to the VAR responses. However, exploiting the rich cross-sectional variation by using PLPs sharpens the identification of the impulse response functions considerably. In particular, the decline of average productivity across small firms is both substantial—about 1 percent after two years—and statistically

respond simultaneously with policy (Gertler and Karadi 2015). Figure B.5 in Appendix B illustrates this issue.

[33]These results stand in contrast to those of Christiano, Eichenbaum, and Evans (2005), Moran and Queralto (2018), and Meier and Reinelt (2022), who document that aggregate productivity rises after a monetary easing. Christiano, Eichenbaum, and Evans (2005) and Moran and Queralto (2018) identify monetary policy shocks using short-run restrictions in a sample in which forward guidance became important, which is problematic (see Footnote 32). Meier and Reinelt (2022) employ local projections without macrofinancial controls. Figure B.17 in Appendix B highlights the importance of including such control variables.

Figure 3. Expansionary Monetary Policy and Productivity

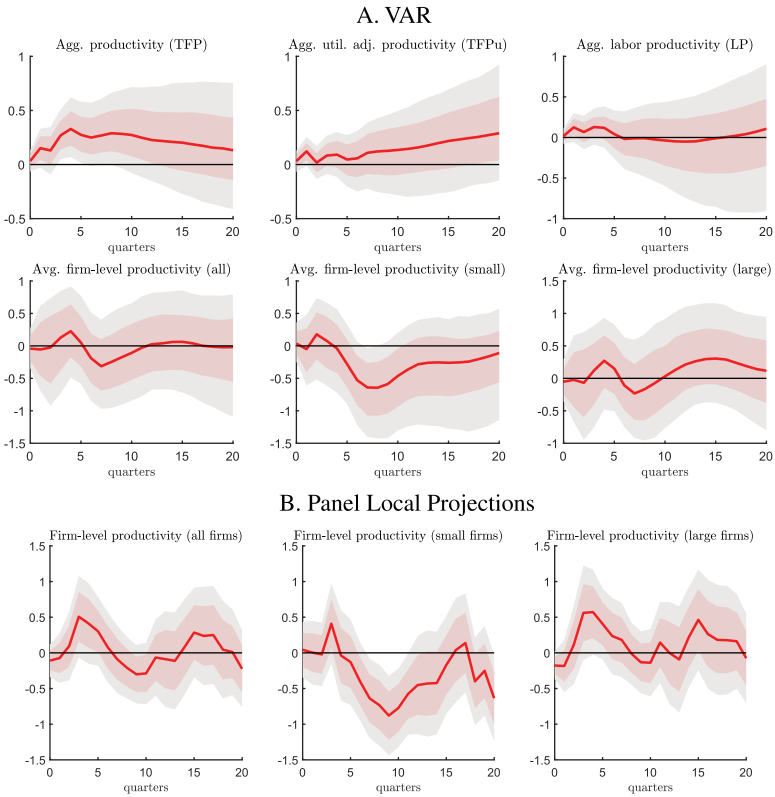

Note: Panel A shows impulse response functions to the sign-identified monetary policy shock in a VAR with FOMC announcement surprises. Panel B shows impulse response functions to the sign-identified monetary policy shock in PLPs using a full set of macrofinancial controls. The thick lines are the median estimates; the shaded areas depict the 68 percent and 90 percent credible (confidence) intervals for the VAR (PLPs). Responses are shown in percent deviations.

highly significant (third row, middle panel). In contrast, the response of average productivity across all firms is nonsystematic and not significant at the 10 percent level (third row, left panel). This largely resembles the behavior of average productivity of large firms (third row, right panel).

Our empirical findings hence suggest that expansionary monetary policy facilitates the entry and profitability of small, unproductive firms. The simple view of firm heterogeneity in the theoretical framework, in particular in Variant A, thus appears to be an appropriate representation for smaller firms. Intuitively, small firms are particularly likely to be the "marginal" firms, i.e., the ones entering or exiting production. This notion is also embedded in the theoretical framework: the firm-level production scale is proportional to productivity, such that low-productivity firms are smaller than productive firms (see Appendix A.1). In contrast, new firms are naturally rarely large in reality, and large firms tend to be more successful and thus less likely to exit production.

However, the empirical results regarding *aggregate* productivity are clearly not in line with either variant of the model. While the responses of entry and exit in the data are consistent with Variant A, one would accordingly expect a significant decline of aggregate productivity after a monetary easing. Contrary to this theoretical prediction, we do not find compelling evidence of a systematic effect of monetary policy on aggregate productivity. This finding suggests that entry and exit of unproductive firms are of limited quantitative importance for the overall productivity effect of monetary policy: since smaller firms are primarily affected, they hardly have an impact on the aggregate.

The empirical exercise also indicates that a key dimension of reality is not included in the theoretical framework, i.e., that the model is incomplete and misses an important feature. A model in line with the empirical results could, for example, combine an endogenous firm-extensive margin with firm-specific productivity evolving endogenously over the life cycle. Such a unifying framework hence constitutes an interesting avenue for future research with a view to examining the relative importance of the different productivity channels.

4.4 Robustness

We verify the robustness of our empirical results along several dimensions, reported in Appendix B.5. Figure B.8 shows that the low-frequency sign restrictions guarantee a proper identification of monetary policy shocks in our setup. Figure B.9 highlights that an alternative identification, the so-called poor man's proxy by Jarociński and Karadi (2020), yields results similar to our baseline. The same is true when using surprises from scheduled FOMC meetings only or assuming that surprises are predictable by macrofinancial factors (Figures B.10 and B.11). Our VAR results are also robust when considering a monthly frequency (B.12), to different monetary policy indicators (B.13), to alternative measures for output and financial frictions (B.14), and in a sample up to or excluding the Great Recession (B.15 and B.16). Figure B.17 confirms our VAR results by using local projections and highlights the importance of macrofinancial controls. We hence explore the PLP results' sensitivity to alternative sets of controls in Figure B.18, which reveals considerable robustness of the significantly negative response of average productivity in the subset of small firms. In Figure B.19, we show that controlling for sectoral heterogeneity yields results broadly in line with our baseline findings. Finally, we confirm that the baseline VAR results are robust to different procedures to construct the firm-level productivity series, in particular alternative sector compositions (B.20) and estimation approaches (B.21).

5. Conclusion

There exists a notion that accommodative monetary conditions may allow unproductive firms to remain active or start production. We explore this effect of monetary policy on productivity through firm entry and exit in the context of conventional policy measures. In a general equilibrium model of heterogeneous firms, we show that an exogenous decrease of nominal interest rates allows low-productivity incumbents to remain active and unproductive firms to enter production if the looser monetary conditions stimulate corporate profitability. Empirically, we find that a monetary easing indeed raises profits, reduces firm exit, and increases entry. However, we find compelling evidence for a negative productivity effect only for small firms, whereas the response of aggregate productivity is insignificant and nonsystematic. These results imply that a negative impact of expansionary monetary policy on productivity through entry and exit is primarily a concern for small firms, and hence quantitatively less important for aggregate productivity than suggested by the theory.

Appendix A. Theoretical Analysis

A.1 Equilibrium Equations

The equilibrium is characterized by 33 endogenous and 3 exogenous variables, $(A_t, \varepsilon_t^C, \varepsilon_t^M)$. See Table A.1.

A.2 Steady State

We first normalize technology, labor, and inflation in the steady state to 1:

$$A = 1, \tag{A.1}$$

$$L = 1, \tag{A.2}$$

$$\pi = \pi^C = 1. \tag{A.3}$$

From the household bond Euler equation (19), we get

$$R = \beta^{-1}\pi, \tag{A.4}$$

and from the definition of the stochastic discount factor (21),

$$\Lambda = \beta(1-\delta). \tag{A.5}$$

Average markup and markups at the cutoff then follow from (9) and (25) as

$$\widetilde{\mu} = \bar{\mu} = \frac{\theta}{(\theta-1)\left(1-\frac{\tau}{2}(\pi-1)^2\right)+\tau(1-\Lambda)\pi(\pi-1)}. \tag{A.6}$$

We now want to obtain an expression for the total number of products N. Starting from the average profit in (28), inserting (31) and (36) and using (A.1) yields

$$\widetilde{d} = \frac{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}{1-\frac{\tau}{2}(\pi-1)^2}\frac{C}{S} - fwR. \tag{A.7}$$

The aggregate resource constraint, obtained by combining (32) and (A.2), is given by

$$C + vH = w + \widetilde{d}S. \tag{A.8}$$

Table A.1. Equilibrium Equations

Firms		
Average Pricing	(E1)	$\widetilde{\rho}_t = \widetilde{\mu}_t \widetilde{mc}_t$
Average Markup	(E2)	$\widetilde{\mu}_t = \frac{\theta}{(\theta-1)\left(1-\frac{\tau}{2}(\pi_t-1)^2\right)+\tau\left(\pi_t(\pi_t-1)-E_t\left[\Lambda_{t+1}\frac{Y_{t+1}^C}{Y_t^C}\frac{S_t}{S_{t+1}}(\pi_{t+1}-1)\pi_{t+1}\right]\right)}$
Average Marginal Costs	(E3)	$\widetilde{mc}_t = \frac{w_t}{A_t \widetilde{z}_t}$
Real Price	(E4)	$\widetilde{\rho}_t = S_t^{\frac{1}{\theta-1}}$
Average Profit	(E5)	$\widetilde{d}_t = \left(1-\widetilde{\mu}_t^{-1}-\frac{\tau}{2}(\pi_t-1)^2\right)\frac{Y_t^C}{S_t} - f\frac{w_t R_t}{A_t}$
Entry Condition	(E6)	$f_e \frac{w_t}{A_t} = v_t(\Psi_t + \Psi_t' H_t) + \beta E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-1} v_{t+1}\Psi_{t+1}'' H_{t+1}\right]$
Entry Success Probability	(E7)	$\Psi_t = 1 - g_3\left(exp\left(g_1\left(\frac{H_t}{H_{t-1}}-1\right)\right)+\frac{g_1}{g_2}exp\left(-g_2\left(\frac{H_t}{H_{t-1}}-1\right)\right)-2\right)$
1st Derivative Entry Success Prob.	(E8)	$\Psi_t' = -g_1 g_3\left(exp\left(g_1\left(\frac{H_t}{H_{t-1}}-1\right)\right)-exp\left(-g_2\left(\frac{H_t}{H_{t-1}}-1\right)\right)\right)\frac{1}{H_{t-1}}$
2nd Derivative Entry Success Prob.	(E9)	$\Psi_t'' = g_1 g_3\left(exp\left(g_1\left(\frac{H_t}{H_{t-1}}-1\right)\right)-exp\left(-g_2\left(\frac{H_t}{H_{t-1}}-1\right)\right)\right)\frac{H_t}{(H_{t-1})^2}$
Profit at Cutoff	(E10)	$\bar{d}_t = \left(1-\bar{\mu}_t^{-1}-\frac{\tau}{2}\left(\frac{\bar{\rho}_t}{\bar{\rho}_{t-1}}\pi_t^C-1\right)^2\right)\bar{\rho}_t\bar{y}_t - f\frac{w_t R_t}{A_t}$
Exit Condition	(E11)	$\bar{d}_t = 0$
Price at Cutoff	(E12)	$\bar{\rho}_t = \bar{\mu}_t \bar{mc}_t$
Markup at Cutoff	(E13)	$\bar{\mu}_t = \frac{\theta}{(\theta-1)\left(1-\frac{\tau}{2}\left(\frac{\bar{\rho}_t}{\bar{\rho}_{t-1}}\pi_t^C-1\right)^2\right)+\tau\left(\frac{\bar{\rho}_t}{\bar{\rho}_{t-1}}\pi_t^C\left(\frac{\bar{\rho}_t}{\bar{\rho}_{t-1}}\pi_t^C-1\right) -E_t\left[\Lambda_{t+1}\frac{\bar{y}_{t+1}}{\bar{y}_t}\left(\frac{\bar{\rho}_{t+1}}{\bar{\rho}_t}\pi_{t+1}^C-1\right)\left(\frac{\bar{\rho}_{t+1}}{\bar{\rho}_t}\right)^2\pi_{t+1}^C\right]\right)}$
Marginal Costs at Cutoff	(E14)	$\bar{mc}_t = \frac{w_t}{A_t \bar{z}_t}$
Exit Rate	(E15)	$\zeta_t = 1 - \left(\frac{z_m}{\bar{z}_t}\right)^{\kappa}$
Average Productivity	(E16)	$\widetilde{z}_t = \bar{z}_t\left(\frac{\kappa}{\kappa-(\theta-1)}\right)^{\frac{1}{\theta-1}}$
Active Firms	(E17)	$S_t = (1-\zeta_t)N_t$
Evolution of Firms	(E18)	$N_t = (1-\delta)(N_{t-1}+\Psi_{t-1}H_{t-1})$
Average Output	(E19)	$\widetilde{y}_t = \frac{Y_t^C}{\widetilde{\rho}_t S_t}$
Output at the Cutoff	(E20)	$\bar{y}_t = \widetilde{y}_t\left(\frac{\bar{z}_t}{\widetilde{z}_t}\right)^{\theta}$

(continued)

Table A.1. (Continued)

Households		
Wage Setting 1st FOC	(E21)	$g_t = \frac{\theta_W - 1}{\theta_W} (w_t^*)^{1-\theta_W} w_t^{\theta_W} \varepsilon_t^C C_t^{-1} L_t + \beta\lambda_W E_t \left[\left(\pi_{t+1}^C \frac{w_{t+1}^*}{w_t^*} \right)^{\theta_W - 1} g_{t+1} \right]$
Wage Setting 2nd FOC	(E22)	$g_t = \chi \left(\frac{w_t}{w_t^*} \right)^{\theta_w \left(1+\frac{1}{\eta}\right)} L_t^{1+\frac{1}{\eta}} + \beta\lambda_W E_t \left[\left(\pi_{t+1}^C \frac{w_{t+1}^*}{w_t^*} \right)^{\theta_W \left(1+\frac{1}{\eta}\right)} g_{t+1} \right]$
Real Wage	(E23)	$w_t = \left(\lambda_W \left(\frac{w_{t-1}}{\pi_t^C} \right)^{1-\theta_W} + (1-\lambda_W)(w_t^*)^{1-\theta_W} \right)^{\frac{1}{1-\theta_W}}$
Euler Equation Shares	(E24)	$v_t = E_t \left[\Lambda_{t+1} (v_{t+1} + \frac{S_t}{N_t} \widetilde{d}_{t+1}) \right]$
Euler Equation for Bonds	(E25)	$1 = \beta E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-1} \frac{R_t}{\pi_{t+1}^C} \right]$
Stochastic Discount Factor	(E26)	$\Lambda_{t+1} = \beta(1-\delta) E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-1} \right]$
CPI Inflation	(E27)	$\pi_t^C = \frac{\widetilde{\rho}_{t-1}}{\widetilde{\rho}_t} \pi_t$
Aggregation and Monetary Policy		
Market Clearing	(E28)	$Y_t = C_t + v_t H_t$
Accounting	(E29)	$Y_t = w_t L_t + \widetilde{d}_t S_t$
Aggregate Consumption Output	(E30)	$Y_t^C = \left(1 - \frac{\tau}{2}(\pi_t - 1)^2\right)^{-1} C_t$
Investment	(E31)	$I_t = v_t H_t$
Aggregate Profits	(E32)	$d_t = \widetilde{d}_t S_t$
Taylor Rule	(E33)	$log\left(\frac{R_t}{R}\right) = \phi_R \, log\left(\frac{R_{t-1}}{R}\right) + (1-\phi_R)\left[\phi_\pi \, log\left(\frac{\pi_t}{\pi}\right) + \phi_y \, log\left(\frac{Y_t}{Y_{t-1}}\right)\right] + \varepsilon_t^M$

Rearranging (A.7) for C and inserting this expression in (A.8) gives

$$\frac{1-\frac{\tau}{2}(\pi-1)^2}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}\left(\widetilde{d}+fwR\right)S+vH=w+\widetilde{d}S. \tag{A.9}$$

In steady state, the entry condition (12) implies under the normalization $f_E=1$, (A.1) and the steady-state properties of the success probability (11):

$$v=w. \tag{A.10}$$

Inserting this in (A.9) and rearranging yields

$$1=\left(\frac{\widetilde{\mu}^{-1}}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}\frac{\widetilde{d}}{w}+\frac{1-\frac{\tau}{2}(\pi-1)^2}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}fR\right)S+H. \tag{A.11}$$

Now, we want to replace the term $\frac{\widetilde{d}}{w}$. Combining (7) at the cutoff and (13) gives

$$\left(1-\bar{\mu}^{-1}-\frac{\tau}{2}\left(\pi-1\right)^2\right)\bar{\rho}^{1-\theta}Y^C=fwR. \tag{A.12}$$

Using (3) and (8) at the cutoff while inserting (24) and (A.6) gives

$$\left(1-\widetilde{\mu}^{-1}-\frac{\tau}{2}\left(\pi-1\right)^2\right)\widetilde{\rho}^{1-\theta}Y^C=f\frac{\kappa}{\kappa-(\theta-1)}wR. \tag{A.13}$$

Note that the left-hand side is the first term in the average profit in (27). We can use this to rewrite (A.13) as

$$\frac{\widetilde{d}}{w}=f\frac{\theta-1}{\kappa-(\theta-1)}R. \tag{A.14}$$

This is the term we wanted to replace in (A.11), which we can now write as

$$\begin{aligned}1=f\Bigg(&\frac{\widetilde{\mu}^{-1}}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}\frac{\theta-1}{\kappa-(\theta-1)}\\&+\frac{1-\frac{\tau}{2}(\pi-1)^2}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}\Bigg)RS+H.\end{aligned} \tag{A.15}$$

Inserting (14) and rearranging yields

$$N^{-1} = f\left(\frac{\widetilde{\mu}^{-1}}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}\frac{\theta-1}{\kappa-(\theta-1)} + \frac{1-\frac{\tau}{2}(\pi-1)^2}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}\right)R\frac{S}{N}+\frac{\delta}{1-\delta}. \tag{A.16}$$

This provides the steady state of the number of firms N, given the endogenous destruction rate S/N. From the Euler equation in (20), we have

$$1 = \Lambda\left(1+\frac{S}{N}\frac{\widetilde{d}}{v}\right). \tag{A.17}$$

Again using $v = w$ from (A.10) and inserting (A.16) yields

$$1 = \Lambda\left(1+f\frac{\theta-1}{\kappa-(\theta-1)}R\frac{S}{N}\right). \tag{A.18}$$

Rearranging yields

$$\frac{S}{N} = \frac{1}{fR}\frac{\kappa-(\theta-1)}{\theta-1}\frac{1-\Lambda}{\Lambda}. \tag{A.19}$$

Inserting this into (A.16) yields the steady state for the total number of firms:

$$N = \left(\frac{\widetilde{\mu}^{-1}+\left(1-\frac{\tau}{2}(\pi-1)^2\right)\frac{\kappa-(\theta-1)}{\theta-1}}{1-\widetilde{\mu}^{-1}-\frac{\tau}{2}(\pi-1)^2}\frac{1-\Lambda}{\Lambda}+\frac{\delta}{1-\delta}\right)^{-1}. \tag{A.20}$$

The number of active firms follows directly from (A.19). The steady-state values of all other variables can be solved recursively.

A.3 Firms' Pricing Decision

This section derives the expression for the firm markup in Equation (9) of the main text. Firms choose prices, $p_t(z)$, and labor, $l_t^C(z)$, to maximize the sum of current profits, $d_t(z)$, and the firm value, $v_t(z)$ (the expected discounted value of the profit stream from $t+1$ onward) in period t subject to its production function, taking the household demand schedule and aggregate variables as given. The Lagrangian of this problem is given by

$$\mathcal{L}_t(z) = d_t(z) + v_t(z) + \Xi_t(z)[A_t z l_t^C(z) - y_t(z)], \tag{A.21}$$

where $\Xi_t(z)$ denotes the Lagrange multiplier on the production constraint (the term in square brackets). Firm profits in real terms are given by

$$d_t(z) = \frac{p_t(z)}{P_t} y_t(z) - w_t l_t^C(z) - pac_t(z) - f\frac{w_t R_t}{A_t}, \tag{A.22}$$

with price adjustment costs, $pac_t(z)$, being defined as

$$pac_t(z) = \frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right)^2 \frac{p_t(z)}{P_t} y_t^C(z). \tag{A.23}$$

The first-order condition of the Lagrangian with respect to labor is

$$-w_t + \Xi_t(z) A_t z = 0, \tag{A.24}$$

which implies that

$$\Xi_t(z) = \frac{w_t}{A_t z}. \tag{A.25}$$

The Lagrange multiplier is hence equivalent to real marginal costs, $mc_t(z)$, at the optimum.

The first-order condition with respect to the product price is

$$\frac{\partial d_t(z)}{\partial p_t(z)} + \frac{\partial v_t(z)}{\partial p_t(z)} - mc_t(z)\frac{\partial y_t(z)}{\partial p_t(z)} = 0. \tag{A.26}$$

We now derive these three expressions one by one. First, the derivative of firm profits with respect to the product price is

$$\begin{aligned}\frac{\partial d_t(z)}{\partial p_t(z)} &= (1-\theta)\,\frac{y_t(z)}{P_t}\left[1 - \frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right)^2\right] \\ &\quad - \tau\frac{p_t(z)}{p_{t-1}(z)}\left(\frac{p_t(z)}{p_{t-1}(z)} - 1\right)\frac{y_t(z)}{P_t},\end{aligned} \tag{A.27}$$

which uses the insight that the elasticity of firm output with respect to the firm price is equal to $-\theta$, the negative constant elasticity of substitution.

Going to the second expression, the firm value at time t equals the expected discounted value of the profit stream from $t+1$ to infinity,

$$v_t(z) = E_t \sum_{s=t+1}^{\infty} \Lambda_s d_s(z), \tag{A.28}$$

where the representative household's discount factor is given by

$$\Lambda_s = [\beta(1-\delta)]^{s-t} \left(\frac{C_s}{C_t}\right)^{-1}. \tag{A.29}$$

The firm value at time t implicitly accounts for the probability of exiting production in any given period via the expectations operator. The only term in the infinite sum that depends on $p_t(z)$ is $d_{t+1}(z)$ through $pac_{t+1}(z)$. The second expression we are searching for is hence given by

$$\frac{\partial v_t(z)}{\partial p_t(z)} = \frac{\partial E_t[\Lambda_{t+1} d_{t+1}(z)]}{\partial p_t(z)} \tag{A.30}$$

$$= -\frac{\partial E_t[\Lambda_{t+1} pac_{t+1}(z)]}{\partial p_t(z)} \tag{A.31}$$

$$= \tau E_t \left[\Lambda_{t+1} \left(\frac{p_{t+1}(z)}{p_t(z)}\right)^2 \left(\frac{p_{t+1}(z)}{p_t(z)} - 1\right) \frac{y_{t+1}(z)}{P_{t+1}}\right]. \tag{A.32}$$

The third expression can be rewritten as follows:

$$-mc_t(z)\frac{\partial y_t(z)}{\partial p_t(z)} = -mc_t(z)\frac{\partial y_t(z)}{\partial p_t(z)}\frac{p_t(z)}{y_t(z)}\frac{y_t(z)}{p_t(z)} = \theta mc_t(z)\frac{y_t(z)}{p_t(z)}. \tag{A.33}$$

The first-order condition hence becomes

$$(1-\theta)\frac{y_t(z)}{P_t}\left[1-\frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)}-1\right)^2\right]$$

$$-\tau\frac{p_t(z)}{p_{t-1}(z)}\left(\frac{p_t(z)}{p_{t-1}(z)}-1\right)\frac{y_t(z)}{P_t} \tag{A.34}$$

$$+\tau E_t\left[\Lambda_{t+1}\left(\frac{p_{t+1}(z)}{p_t(z)}\right)^2\left(\frac{p_{t+1}(z)}{p_t(z)}-1\right)\frac{y_{t+1}(z)}{P_{t+1}}\right] \tag{A.35}$$

$$+\theta mc_t\frac{y_t(z)}{p_t(z)}=0. \tag{A.36}$$

Dividing both sides by $\dfrac{y_t(z)}{P_t}$ yields

$$(1-\theta)\left[1-\frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)}-1\right)^2\right]-\tau\Upsilon_t(z)+\theta mc_t\frac{P_t}{p_t(z)}=0, \tag{A.37}$$

where

$$\Upsilon_t(z)=\frac{p_t(z)}{p_{t-1}(z)}\left(\frac{p_t(z)}{p_{t-1}(z)}-1\right)$$
$$-E_t\left[\Lambda_{t+1}\frac{y^C_{t+1}(z)}{y^C_t(z)}\frac{P_t}{P_{t+1}}\left(\frac{p_{t+1}(z)}{p_t(z)}-1\right)\left(\frac{p_{t+1}(z)}{p_t(z)}\right)^2\right]. \tag{A.38}$$

Rearranging yields

$$p_t(z)=mc_tP_t\frac{\theta}{(\theta-1)\left[1-\frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)}-1\right)^2\right]+\tau\Upsilon_t(z)}, \tag{A.39}$$

which can be used to define the markup $\mu_t(z)$ over real marginal costs as in Equation (9) of the main text as

$$\mu_t(z)=\frac{\theta}{(\theta-1)\left[1-\frac{\tau}{2}\left(\frac{p_t(z)}{p_{t-1}(z)}-1\right)^2\right]+\tau\Upsilon_t(z)}. \tag{A.40}$$

A.4 Endogenous Exit and Inflation

In this section, we briefly analyze the implications of endogenous firm exit for inflation dynamics. Log-linearizing Equation (25) yields

$$\widehat{\pi}_t = -\frac{\theta - 1}{\tau}\widehat{\widetilde{\mu}}_t + \beta(1-\delta)E_t[\widehat{\pi}_{t+1}], \tag{A.41}$$

where variables with a hat denote log-deviations from steady state. This is the familiar linearized New Keynesian Phillips curve, relating inflation to variations in the (average) firm markup. Using the optimal pricing condition (8) with the definitions of marginal costs (3) and the variety effect (36), we can substitute the markup and write

$$\widehat{\pi}_t = \frac{\theta - 1}{\tau}(\widehat{w}_t - \widehat{A}_t - \widehat{\widetilde{z}}_t) - \frac{1}{\tau}\widehat{S}_t + \beta(1-\delta)E_t[\widehat{\pi}_{t+1}]. \tag{A.42}$$

Equation (A.42) is a New Keynesian Phillips curve relating producer price inflation to marginal costs and the number of active firms in the economy. Intuitively, firms' price setting crucially depends on their marginal costs. As such, changes in aggregate (A) or firm-specific ($\widetilde{z}$) productivity affect effective marginal costs and thus inflation. Furthermore, the number of active firms influences relative prices (the price of each good relative to the consumption basket) and thus markups, which translates into an effect on inflation. This may be interpreted as representing the effect of heightened competition.

As an illustration, we compare Variant A to an economy where all firms are homogeneous and exit occurs only exogenously. We set idiosyncratic firm productivity $z = 1$ for all firms and abstract from fixed costs of production by setting $f = 0$. This model variant is essentially the one considered by Bilbiie, Ghironi, and Melitz (2008) and Bilbiie, Fujiwara, and Ghironi (2014), but additionally includes entry frictions and wage rigidities. Figure A.1 compares the transmission of an expansionary monetary policy shock across the two models.

In Variant A, expansionary monetary policy shocks increase firm profits and allow unproductive firms to be active. As a consequence, average productivity declines, and average marginal costs increase sharply on impact of the shock. The decline in average productivity yields an initially stronger inflation response in the first year

Figure A.1. Monetary Policy Shock and the Role of Endogenous Exit

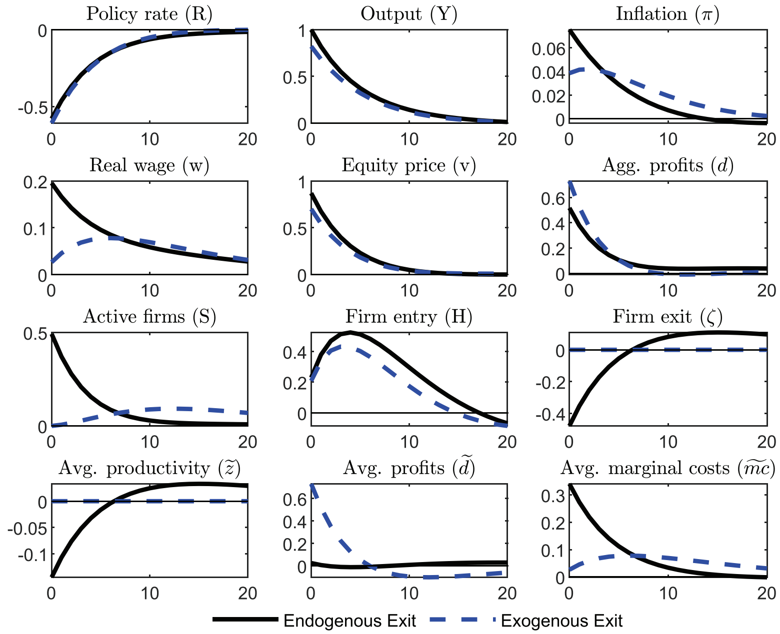

Note: Impulse response functions to an expansionary monetary policy shock in Variant A (with endogenous exit, solid lines) and a variant with exogenous exit (dashed lines). The shock size is calibrated to yield a 1 percent increase of output in Variant A. Inflation, interest rate, and exit are shown in percentage point deviations from steady state, all other variables in percentage deviations.

(compared with the model with exogenous exit). At the same time, the overall number of firms increases and exit rates decline. Via the competition effect, this translates into lower markups and thus lower inflation. After the first year, the competition effect dominates the productivity effect such that the overall inflation response is lower in the case of endogenous exit. Interestingly, this shares similarities with the microeconometric findings by Acharya et al. (2020), who document that a rise in lending to "zombie firms" is associated with disinflation. In this respect, the demand-side and preference-based variety effect in our framework may be interpreted as operating similarly to a supply-side competition effect, whereby excess capacity creates downward pressure on prices.

Figure A.2. Technology Shock

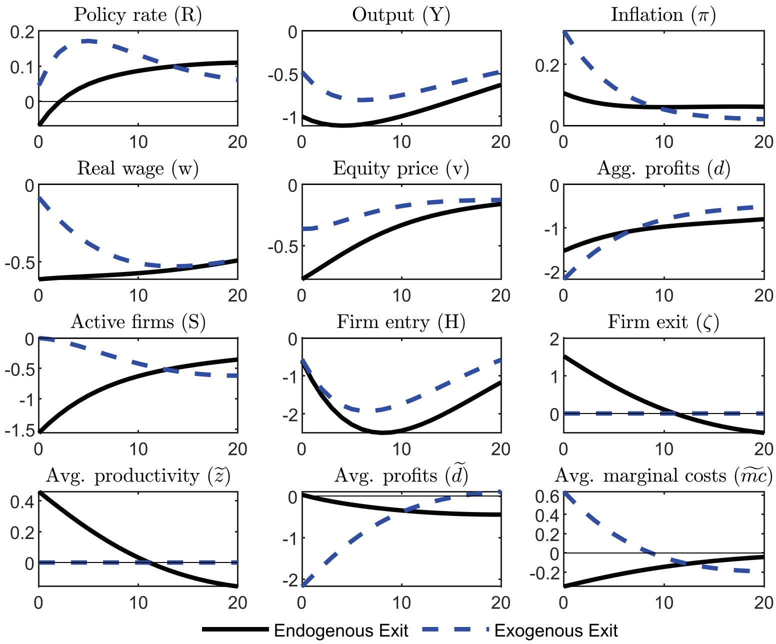

Note: Impulse response functions to a contractionary technology shock with an autoregressive coefficient of 0.9 in Variant A (with wage stickiness and endogenous exit, solid lines) and a variant with exogenous exit (dashed lines). The shock size is calibrated to yield a 1 percent increase of output in Variant A. Inflation, interest rate, and exit are shown in percentage point deviations from steady state, all other variables in percentage deviations.

The amplification of the output response via endogenous exit is largely due to higher investment in new firms. Intuitively, entering production becomes profitable for firms with relatively low productivity. As a result, investment in new firms and firm entry respond more strongly to monetary policy shocks. Over the medium term, lower inflation and real interest rates also contribute to slightly higher consumption relative to the model with exogenous exit.

A.5 Technology Shock

Figure A.2 shows the transmission of a contractionary technology shock, comparing Variant A (with wage stickiness) to a model where firm exit is entirely exogenous and constant. As also described by Hamano and Zanetti (2017) and Rossi (2019), negative technology shocks increase real marginal costs and thus lower expectations of future profits in Variant A, thereby disincentivizing the entry of new firms. The firm-specific productivity cutoff required for profitability increases, such that more firms exit production. As a result, the contraction is more pronounced relative to a model with exogenous exit. Only relatively more productive firms are able to remain active, causing average productivity to increase initially. As the economy reverts to the initial equilibrium, firm exit drops below baseline, reflecting a decreasing cutoff level of productivity.

Appendix B. Empirical Analysis

B.1 Data

Figure B.1 shows time-series plots of the data. Firm entry is procyclical and exit countercyclical around recessions. Profits show some procyclical patterns. Aggregate TFP displays some mild signs of procylicality, while utilization-adjusted TFP and labor productivity evolve rather independently of the cycle. The comovement of average firm-level productivity depends on the firm size (see definition in Section B.2): it is procyclical for all firms and the subset of large firms, but hardly reacts to the cycle for small firms.

Table B.1 presents descriptive statistics on business cycle fluctuations as measured by the cyclical component of all variables in log-levels using the regression-based filter of Hamilton (2018). Panel A reports volatility, relative volatility to the cyclical component of real GDP, persistence, and contemporaneous comovement. Firm entry and exit dynamics are almost three times more volatile than fluctuations in output and profits are about seven times more volatile. Aggregate productivity measures are less volatile than output, whereas average firm-level productivity is more volatile—at a similar level to firm dynamics. Firm entry and exit dynamics are the least persistent series. Profits and productivity measures are slightly less persistent than output. Firm entry is procyclical, while exit is countercyclical, in line with the evidence presented by Campbell (1998) and Jaimovich and Floetotto (2008), who consider a different data set and study an earlier period. Interestingly, aggregate TFP

Figure B.1. Data

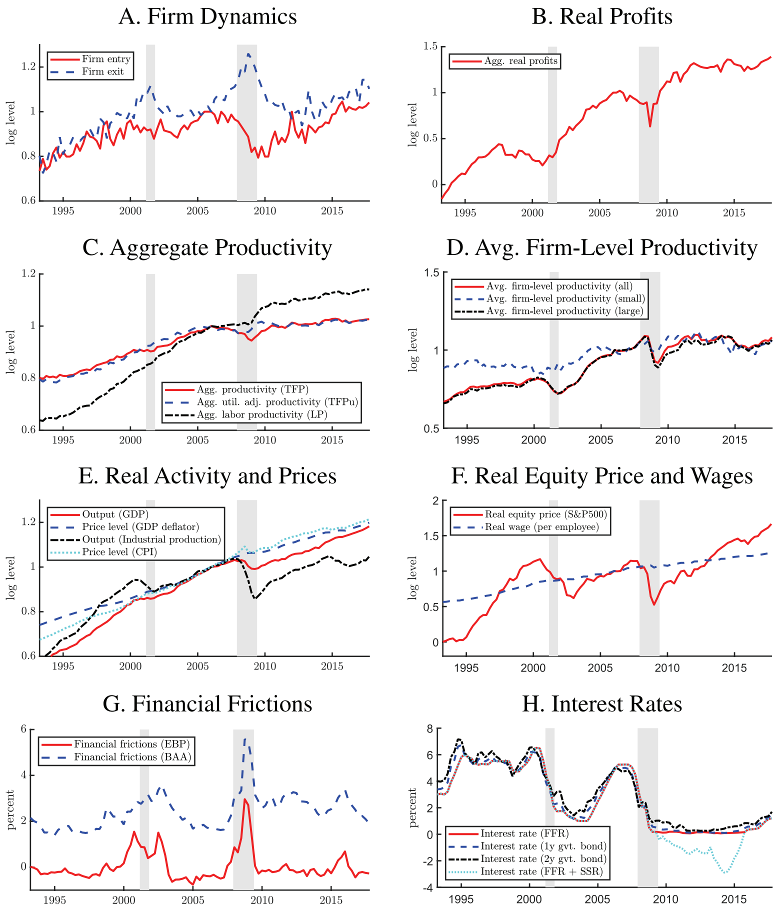

Note: Time-series plot of the data. All variables are in log-levels and normalized to 1 in 2006:Q1, except for the proxies of financial frictions and the interest rates. Measures of financial frictions and interest rates are in percent. Shaded gray areas indicate National Bureau of Economic Research (NBER) recession dates.

is strongly procyclical while profits, utilization-adjusted TFP, and labor productivity hardly comove with the cycle. In fact, the estimated correlations are insignificant and thus, these series may be considered as acyclical. Average firm-level productivity is procyclical for the overall aggregate and even more so for large firms, while it is countercyclical for small firms.

Table B.1. Business Cycle Statistics for the U.S. Economy

	$\sigma(Y_{i,t})$	$\frac{\sigma(Y_{i,t})}{\sigma(X_t)}$	$\rho(Y_{i,t}, Y_{i,t-1})$	$\rho(Y_{i,t}, X_t)$
A. Business Cycle Moments				
(1) Output (GDP)	2.29	1.00	0.89	1.00
(2) Price Level (GDP Deflator)	1.02	0.45	0.86	0.30
(3) Firm Entry	5.54	2.42	0.75	0.70
(4) Firm Exit	6.71	2.93	0.79	−0.24
(5) Agg. Real Profits	15.32	6.69	0.86	0.00
(6) Agg. Productivity (TFP)	1.82	0.79	0.87	0.64
(7) Agg. Util. Adj. Productivity (TFPu)	1.60	0.70	0.83	−0.13
(8) Agg. Labor Productivity (LP)	1.81	0.79	0.82	−0.09
(9) Avg. Firm-Level Productivity (all)	6.40	2.80	0.90	0.45
(10) Avg. Firm-Level Productivity (small)	4.58	2.00	0.77	−0.18
(11) Avg. Firm-Level Productivity (large)	6.52	2.85	0.90	0.54
(12) Real Wage (per Employee)	1.77	0.77	0.79	0.62
(13) Real Equity Price (S&P 500)	21.57	9.43	0.88	0.81
(14) Interest Rate (1y Gvt. Bond)	2.18	0.95	0.97	0.28
(15) Financial Frictions (EBP)	0.67	0.29	0.83	−0.48

(*continued*)

Table B.1. (Continued)

B. Contemporaneous Correlations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
(1)	1.00														
(2)	0.30	1.00													
(3)	0.70	0.35	1.00												
(4)	−0.24	0.10	−0.01	1.00											
(5)	0.00	0.08	0.03	−0.60	1.00										
(6)	0.64	0.22	0.36	−0.67	0.58	1.00									
(7)	−0.13	−0.17	−0.11	−0.20	0.26	0.36	1.00								
(8)	−0.09	−0.21	−0.29	−0.47	0.51	0.59	0.78	1.00							
(9)	0.45	0.61	0.32	−0.18	0.35	0.48	−0.25	−0.01	1.00						
(10)	−0.18	0.20	−0.07	−0.23	0.51	0.22	0.09	0.27	0.57	1.00					
(11)	0.54	0.68	0.40	−0.17	0.28	0.49	−0.26	−0.08	0.97	0.43	1.00				
(12)	0.62	0.34	0.46	0.25	−0.57	0.14	−0.29	−0.29	0.26	−0.29	0.34	1.00			
(13)	0.81	0.10	0.55	−0.14	−0.05	0.45	−0.26	−0.15	0.36	−0.16	0.43	0.57	1.00		
(14)	0.28	0.20	0.09	0.04	−0.40	0.01	−0.14	−0.11	−0.03	−0.40	0.02	0.49	0.06	1.00	
(15)	−0.48	−0.14	−0.28	0.55	−0.49	−0.63	0.01	−0.18	−0.44	−0.20	−0.44	−0.03	−0.54	−0.03	1.00

Note: All variables are in log-levels. The cyclical component is estimated using the regression-based filter of Hamilton (2018), except for the interest rate and the excess bond premium. Panel A reports business cycle moments for each variable: (1) standard deviation, (2) relative standard deviation to output, (3) first-order autocorrelation, and (4) contemporaneous correlation with output. Panel B depicts the contemporaneous correlation matrix of all variables.

Panel B provides more details on contemporaneous correlations between real GDP, firm dynamics, profits, and various productivity measures. While there is no comovement between firm entry and exit, firm dynamics are strongly correlated with aggregate TFP. A cyclical upswing of firm entry is associated with higher aggregate productivity, while firm exit and productivity tend to move in opposite directions, in line with VAR estimates of Rossi (2019). Moreover, profits show negative comovement with firm exit and positive comovement with productivity measures, while they are unrelated to entry.

Firm dynamics are contemporaneously hardly related to pure technological progress. Because of the growth accounting definitions in Basu, Fernald, and Kimball (2006), this implies that the covariation of firm dynamics with aggregate TFP is driven by variable capital and labor utilization. Moreover, both firm entry and exit are countercyclical to labor productivity. Roughly speaking, labor productivity rises if workers have more capital or better skills, or if aggregate TFP rises (Fernald 2015). Thus, the negative correlation of firm entry suggests that variations in capital and labor quality dominate the positive effects of aggregate TFP. For firm exit, these effects only slightly affect the negative correlation.

The comovement between average firm-level productivity and firm dynamics depends on the size of the firm. An increase in the rate of firm entry is associated with an increase of average productivity for large firms, while for small firms it is associated with a decline. Firm exit, on the other hand, is mildly negatively correlated with all average firm-level productivity measures.

Turning from unconditional to conditional comovements, Figure B.2 shows Burns-Mitchell diagrams; these diagrams depict the average behavior of selected time series around the start of U.S. recessions. Chart A shows the average behavior of firm dynamics. Firm entry remains high during the expansion, but drops substantially after the turning point of the cycle. Firm exit, on the other hand, starts to increase prior to the start of the recession and peaks after four quarters. During the recovery, firm exit starts to diminish while firm entry remains subdued for a prolonged period. Chart B shows that profits are acyclical to real activity but start to decline prior to a recession. After the economy reaches its trough after four to six quarters and the economy starts to recover, profits increase strongly.

Figure B.2. Burns-Mitchell Diagrams

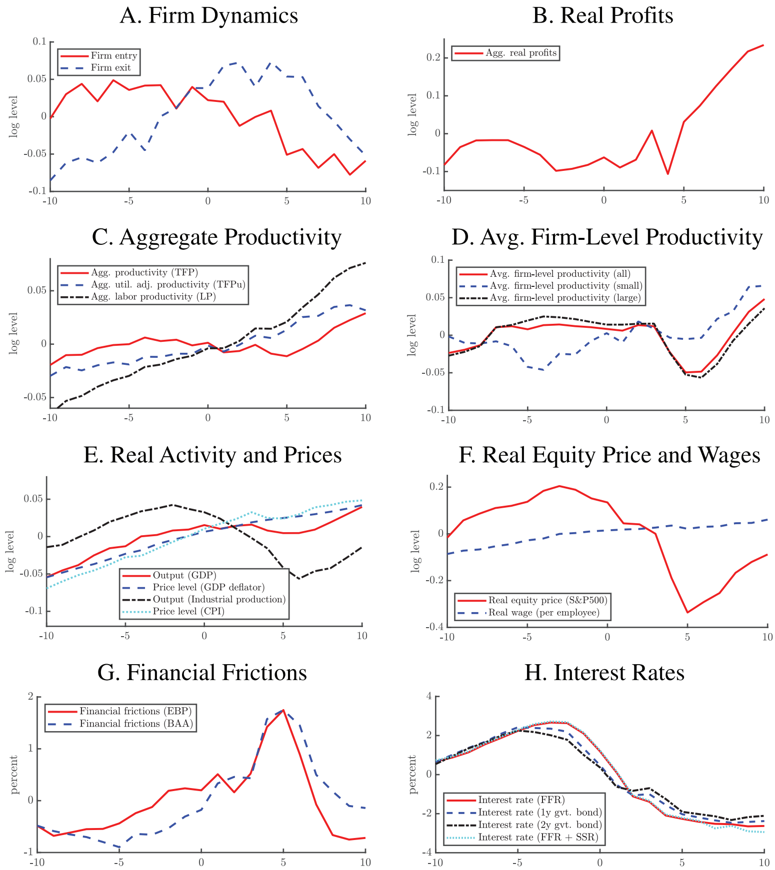

Note: Average behavior of variables around cyclical peaks, as measured by the start of a U.S. recession. $x_t = \frac{1}{M}\sum_{i=1}^{M}\left(y_{i,t} - \frac{1}{21}\sum_{t=-10}^{10} y_{i,t}\right)$, where $y_{i,-10}, y_{i,-9}, \ldots, y_{i,0}, y_{i,1}, \ldots, y_{i,10}$, $i = 1, 2, \ldots, M$, and $y_{i,0}$ is quarter of business cycle peak. All variables enter in log-levels, except interest rates, the excess bond premium, and corporate BAA spread, which are in percent.

Charts C and D show the average conditional behavior of productivity measures. Aggregate TFP is procyclical and leading, peaking several quarters before the turning point, while utilization-adjusted TFP and labor productivity show no strong cyclical patterns. Their average growth is uninterrupted during recessions. Average firm-level productivity of all and large firms behave similarly during a recession. They hardly react prior to the recession but eventually decline after the turning point. Average productivity of small firms, in contrast, declines somewhat prior to the recession but increases continuously throughout this period.

The remaining charts present the average behavior of key macroeconomic and financial variables. Real activity and prices, Chart E, move as expected. Real activity contracts while prices react rather sluggishly. The sluggish behavior of prices is a common feature of more recent recessions; see Figure B.1. Equity peaks prior to the turning point in real GDP and substantially declines in the downturn, while real wages show hardly any reaction; see Chart F. Financial frictions, Chart G, are low prior to a recession but increase substantially when the economy dips further into the recession. Chart H shows that the policy rate and longer-term interest rates decline in response to subdued economic activity.

B.2 Construction of Firm-Level Productivity

We combine annual and quarterly Compustat data on U.S. public firms—incorporated in the U.S. and doing business in U.S. dollars—from 1990:Q1 to 2019:Q4 to estimate TFP at the firm level. We exclude financial firms (due to their special balance sheets) and utilities (due to their dependence on commodity prices) from our data set.

For each firm, we construct the capital stock using the perpetual inventory method. First, we initialize the capital stock with the first available entry of PPEGT (total gross property, plant, and equipment). Second, we iterate using the initial value of the firms' capital stock using the accumulation equation

$$k_{i,t} = k_{i,t}(1 - \delta) + i_{i,t},$$

where we use PPENT (total net property, plant, and equipment) as our measure of net investment $(i_{i,t} - \delta k_{i,t})$. In case of missing values for PPENT, we replace them using a log-linear interpolation. Moreover, we deflate our constructed measure of firm-level capital stock by the investment goods deflator from the Bureau of Economic Analysis (BEA).

We construct a quarterly measure of employment by merging annual and quarterly Compustat series using the firm identifier

(GVKEY) and time (DATADATE). We use a log-linear interpolation for the missing observations. Further, we construct a real measure of sales by deflating the nominal series by the GDP deflator from the BEA. Last, we exclude observations with negative sales, capital stock or employment (measurement error) and winsorize sales, capital, and employment at the 1st and 99th percentile (outliers).

We then estimate a standard growth accounting equation using panel ordinary least squares (OLS):

$$\log(y_{i,t}) = \mu_i + \mu_t + \alpha \log(k_{i,t-1}) + \beta \log(n_{i,t}) + \epsilon_{i,t}, \tag{B.1}$$

where $y_{i,t}$ is real sales, $k_{i,t-1}$ is the constructed capital stock (Compustat capital is recorded at the end of the period), $n_{i,t}$ is employment, μ_i is a firm fixed effect, and μ_t is a time fixed effect. Then, TFP at the firm level is given by

$$\log(\hat{\nu}_{i,t}) = \log(y_{i,t}) - \hat{\alpha} \log(k_{i,t-1}) - \hat{\beta} \log(n_{i,t}) = \hat{\mu}_i + \hat{\mu}_t + \hat{\epsilon}_{i,t},$$

which we use to compute average firm-level productivity as

$$\hat{\nu}_t = N_t^{-1} \sum_i^{N_t} \hat{\nu}_{i,t},$$

with N_t being the number of firms in time t. We winsorize estimated firm-level productivity at the 1st and 99th percentile to control for the effect of potentially very large outliers.[34] We adjust the average firm-level productivity series for seasonality using the x13 program of the U.S. Census Bureau.

Table B.2 reports the number of firms for each sector and by firm size. We classify a firm to be small (large) if it has sales below (above) \$10 million. The manufacturing and services sectors make up 70 percent of the firms in the Compustat universe. Moreover, panel A in Figure B.3 shows that average productivity dynamics of nonfinancial ex utilities firms (Chart (I)) are primarily driven by manufacturing and services sector firms (Chart (II)).

[34]Without winsorization, average firm-level productivity exhibits somewhat distinct dynamics and occasional breaks relative to the median and other percentiles of firm-level productivity.

Table B.2. Number of Firms

Sector	All Firms	Small Firms	Large Firms
Agriculture*	69	43	26
Construction	205	61	144
Manufacturing	6,866	3,334	3,532
Mining	831	491	340
Retail Trade	1,128	274	854
Services	3,918	1,979	1,939
Transportation*	1,054	362	692
Wholesale Trade	676	230	446
Total	14,747	6,774	7,973

Note: Number of firms per sector and firm size. Agriculture* denotes the sector agriculture, fishing, and forestry, and Transportation* denotes the sector transportation, communications, electricity, and sanitary services except utilities.

Table B.3 reports cross-sectional moments of firm-level productivity based on the period 1990:Q2–2019:Q4. Our overall sample consists of 559,796 observations. Small firms make up about one-third and large firms two-thirds of the observations.[35] Average firm-level productivity depends positively on the firm size, i.e., larger firms tend to be more productive. The productivity distribution has a very long right tail due to some extremely large companies (measured by sales). The distribution is less dispersed for small firms. On average, the most productive firms are in the construction and wholesale trade sector.

As a robustness check, we consider two alternative estimation methods: (1) imposing constant returns to scale in production, i.e., $\alpha = 1 - \beta$, estimated via restricted panel OLS, and (2) using the Olley and Pakes (1996) method, which controls for input factor endogeneity. Table B.4 shows that these alternative estimation methods affect the estimated share in production factors. Specifically, they lead to an increase in the share of capital and labor as compared with the baseline panel regression with fixed effects. Nevertheless, the resulting firm-level productivity averages (and other moments of

[35] We lose 14,919 observations as $k_{i,t}$ enters with a lag in (B.1).

Figure B.3. Alternative Estimates of Average Firm-Level Productivity

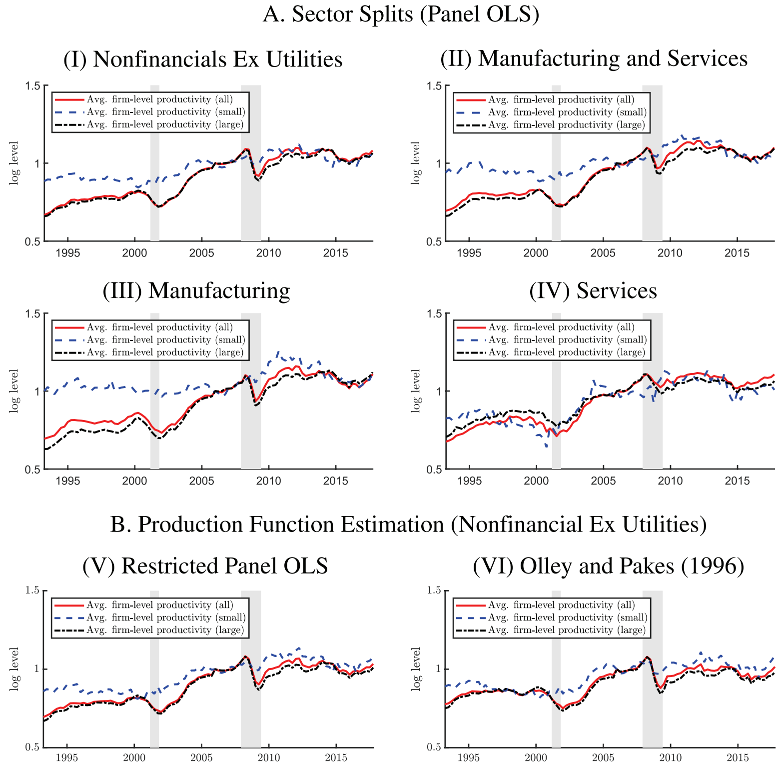

Note: Time-series plot of average firm-level productivity for different sector splits (panel A) and alternative production function estimation approaches (panel B). All variables are in log-levels and normalized to 1 in 2006:Q1. Shaded gray areas indicate NBER recession dates.

the distribution) exhibit similar cyclical variation as compared with the baseline estimates; see panel B in Figure B.3.

B.3 Importance of the Central Bank Information Effect

Figure B.4 shows impulse response functions to a monetary policy shock identified by short-run zero restrictions on the interest rate surprises ordered first in the VAR (blue dashed) and those of a central bank information shock (black dashed-dotted).

Table B.3. Cross-Sectional Moments of Firm-Level Productivtiy

	Total				Small Firms				Large Firms			
Sector	**N**	**Mean**	**Median**	**Std.**	**N**	**Mean**	**Median**	**Std.**	**N**	**Mean**	**Median**	**Std.**
Agriculture*	2,649	20.2	12.1	33.8	1,161	9.5	5.7	21.4	1,488	28.6	18.6	38.9
Construction	8,116	76.0	40.6	73.9	1,603	27.0	14.4	37.8	6,513	88.1	52.0	75.6
Manufacturing	283,060	28.7	22.7	29.2	101,069	15.5	11.0	21.2	181,991	36.1	28.4	30.5
Mining	31,235	32.9	21.0	40.1	14,477	13.8	8.2	20.6	16,758	49.5	34.9	45.2
Retail Trade	43,955	28.5	19.9	33.0	5,154	18.3	8.1	34.2	38,801	29.9	21.2	32.6
Services	128,133	25.7	19.1	28.9	45,954	16.6	11.6	25.0	82,179	30.7	23.9	29.8
Transportation*	37,438	31.6	22.5	33.9	6,782	19.1	10.6	31.1	30,655	34.3	24.6	33.9
Wholesale Trade	25,210	73.5	51.7	65.9	5,831	28.9	18.5	38.2	19,378	86.9	63.3	66.6
Total	559,796	31.1	22.1	35.7	182,031	16.3	11.0	24.1	377,763	38.2	27.4	38.2

Note: The table reports, for each sector, the total number of observations, and the mean, median, and standard deviation of firm-level productivity over the sample period 1990:Q1–2019:Q4. Agriculture* denotes the sectors agriculture, fishing, and forestry, and Transportation* denotes the sectors transportation, communications, electricity, and sanitary services except utilities.

Table B.4. Firm-Level Productivity Regressions

Sale	**Panel OLS (1)**	**Rest. Panel OLS (2)**	**Olley and Pakes (3)**
Capital	0.23***	0.27***	0.34***
	(0.01)	(0.01)	(0.00)
Employment	0.68***	0.73***	0.73***
	(0.02)	(0.01)	(0.01)
Fixed Effects	Firm, Quarter	Firm, Quarter	Firm, Quarter
Observations	559,796	559,796	558,500

Note: The table reports the estimated share of capital and labor from (1) our baseline panel regression in (B.1), (2) a restricted panel regression with constant returns to scale, and (3) the Olley and Pakes (1996) production function approach. Standard errors are clustered two-way at the firm level and time level for (1) and (2) and bootstrapped in (3). *, **, *** indicate that the coefficient is significant at the 10 percent, 5 percent, and 1 percent level.

Figure B.5 shows impulse response functions to a monetary policy shock identified by short-run zero restrictions on real GDP and prices (blue dashed) and those of a central bank information shock (black dashed-dotted).

B.4 Additional Information on the Identification

In this section, we investigate how the results are affected by the different data frequency, the different sample size, and the additional sign restrictions on the low-frequency variables (coupled with the enlargement of the rotation space) as compared with Jarociński and Karadi (2020). Figure B.6 shows estimates based on monthly data in panel A and based on quarterly data in panel B. We exclude our main variables of interest due to the slightly longer sample. Each panel shows the estimates for our considered sample starting in 1993:M4 and the slightly longer sample starting in 1990:M2, as well as for the different identification schemes. The red solid lines and the black dashed-dotted lines correspond to sign restriction on both high-frequency and low-frequency variables, while the blue dashed lines and the cyan dotted lines correspond to sign restrictions on high-frequency variables only, as in Jarociński and Karadi (2020).

Figure B.4. VAR Short-Run Restrictions on Interest Rate Surprises

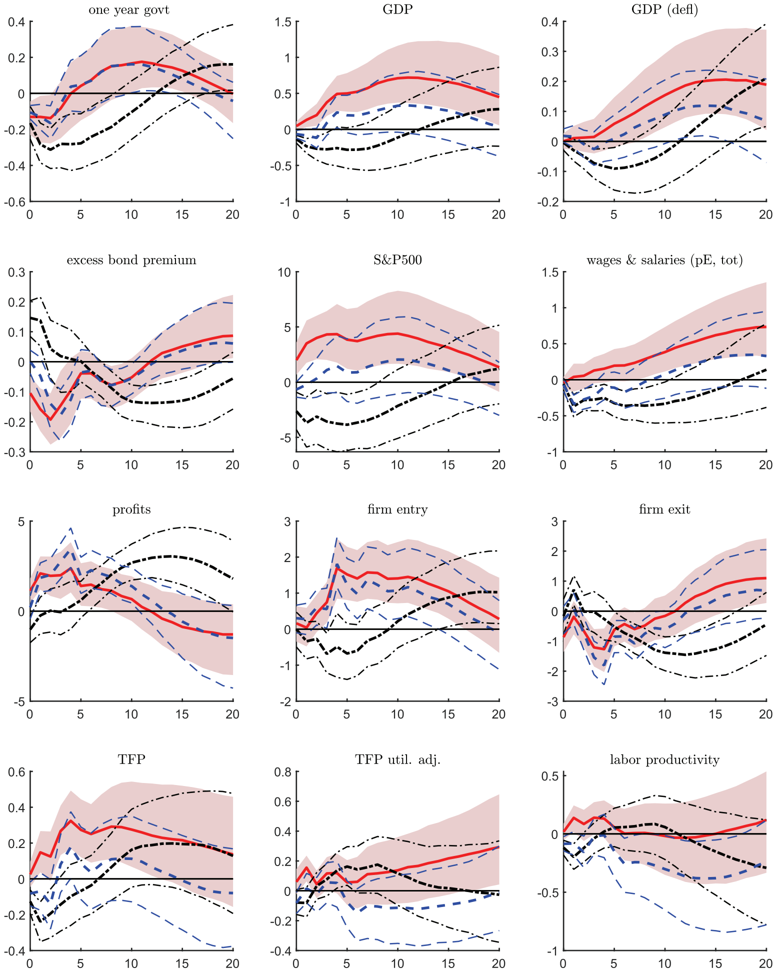

Note: Impulse response functions to a monetary policy shock according to the baseline identification (solid lines), a shock identified by zero restrictions on the interest rate surprises ordered first (dashed), and a central bank information shock identified by sign restrictions (dashed-dotted). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.5. VAR Short-Run Restrictions on Interest Rate

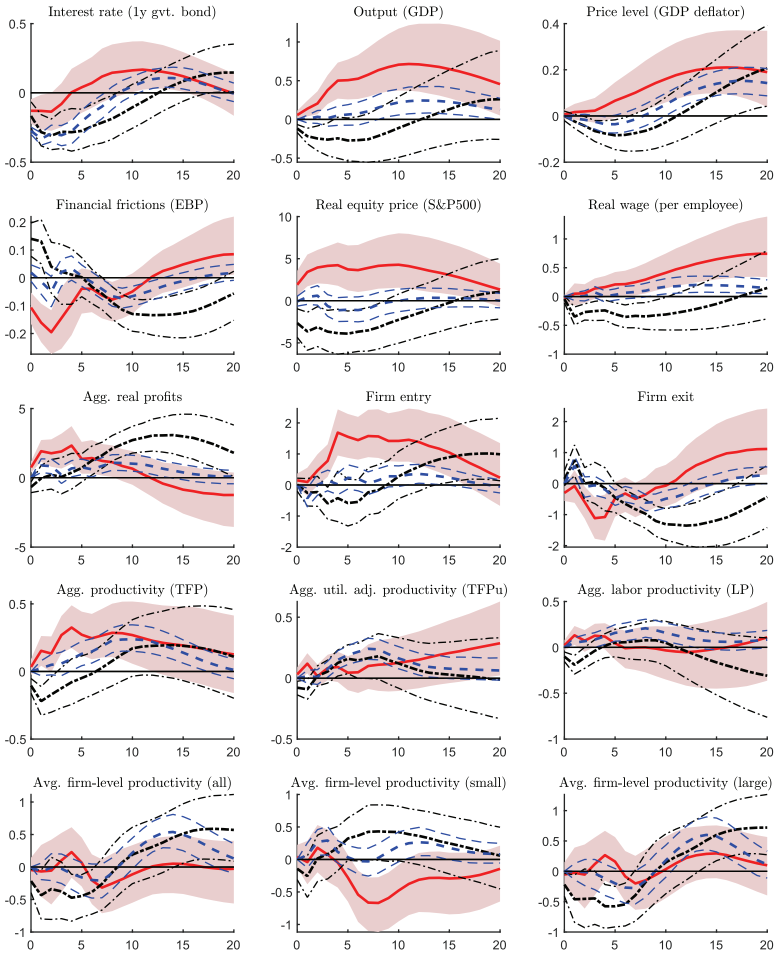

Note: Impulse response functions to a monetary policy shock according to the baseline identification (solid lines), a shock identified by zero restrictions on the contemporaneous comovement of GDP and prices for the interest rate (dashed), and a central bank information shock identified by sign restrictions (dashed-dotted). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.6. Monetary Transmission at Different Frequency and Sample Size

A. Monthly

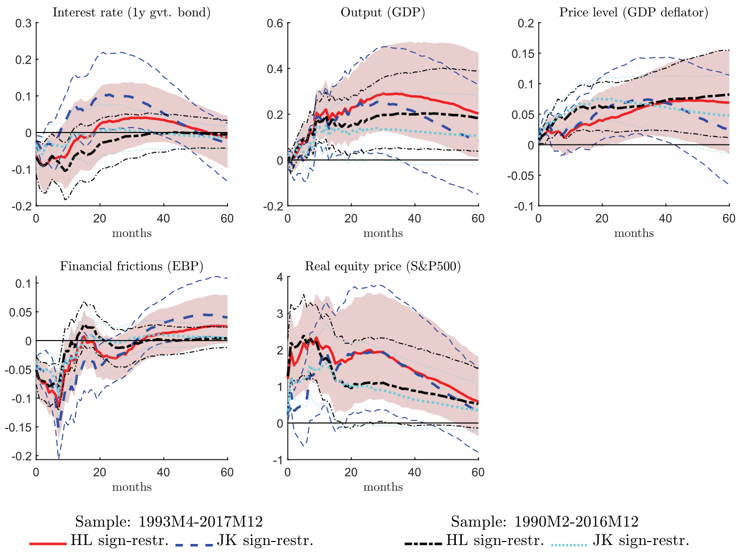

B. Quarterly

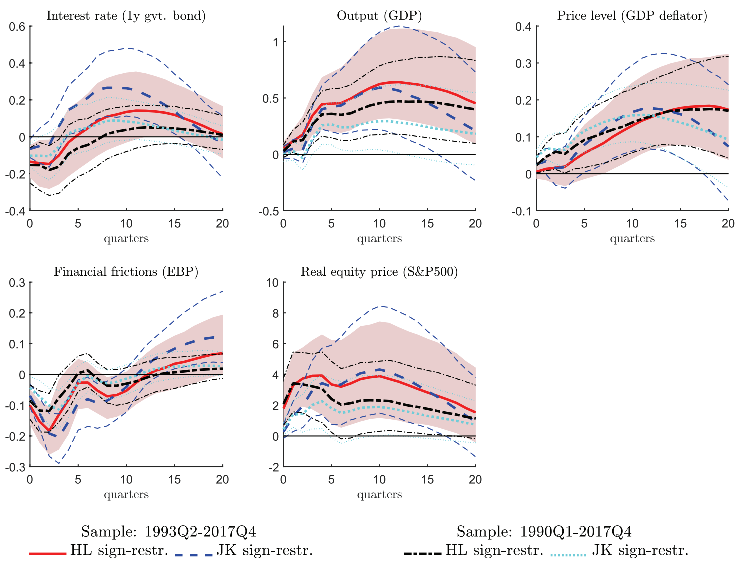

Note: Impulse response functions to a monetary policy shock identified by alternative sign restrictions, for different sample sizes and different data frequencies. HL corresponds to the baseline sign restrictions and JK to the sign restrictions used by Jarociński and Karadi (2020). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Panel A shows that the estimated impulse response functions at the monthly frequency are affected by the different sample size and by the alternative identification restrictions. The monetary impulse identified by sign restrictions on high-frequency variables only and for the sample starting in 1993:M4 (blue dashed) differs substantially from the other median estimates. Particularly, the initial impulse to the interest rate is rather small and disproportionally compensated in the medium term. Notable, furthermore, is the insignificant response of the stock market under this specification.

In contrast, the median estimate of the interest rate response based on the longer sample starting in 1990:M2 (cyan dotted) is more comparable to the median estimates obtained under our identification restrictions, which are qualitatively similar in both samples. Note that for the longer sample the response of the interest rate and the stock price index are significant and last over several months when using sign restrictions only on high-frequency variables. Apart from that, it should be noted that both identification schemes yield qualitatively similar estimates of the responses of macroeconomic and financial variables to a monetary policy shock.

Turning to panel B, the chart shows that median estimates of the interest rate response also differ across different sample size and identification restrictions at the quarterly frequency. However, the estimated impulse response functions for a specific sample size and identification scheme are very similar at different data frequencies. In particular, the quarterly estimates can be interpreted as a smoothed version of the monthly estimates. Nevertheless, it should be noted that the stock market response is only very marginally significant when sign restrictions are imposed on high-frequency variables only.

Based on these considerations, we conclude that excluding the sample from 1990:M2 to 1993:M3 from the estimation may obscure the relationship between high-frequency and low-frequency variables. The lack of these data in our sample makes it more difficult to identify a plausible monetary transmission channel when structural parameters are identified using sign restriction on high-frequency variables only. By imposing additional restrictions on the low-frequency variables (coupled with the enlarged rotation space), we are able to identify a plausible monetary transmission channel.

Figure B.7. Interest Rate and Stock Price Surprises

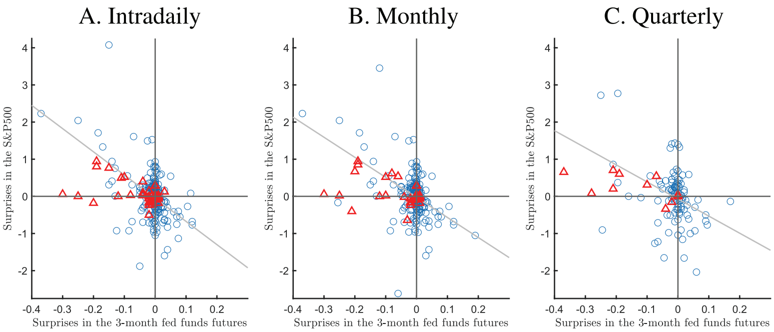

Note: Changes in the three-month federal funds futures and the S&P 500 stock index around FOMC announcements, in percent. For plot A, each dot represents one FOMC announcement. For plots B and C, each dot represents the sum of intradaily surprises of FOMC announcements in the current month and quarter, respectively. The gray line is the fitted least-squares prediction. Red triangles correspond to the period 1990:M2–1993:M3 and blue circles correspond to 1993:M4 through 2017:M12.

To further investigate the effects of a different sample size on the relationship between high-frequency surprises and their low-frequency counterparts, Figure B.7 depicts scatter plots of interest rate and stock price surprises across (A) intradaily frequency, (B) monthly frequency, and (C) quarterly frequency for the sample 1993:M4–2017:M12 in blue dots and the pre-sample 1990:M2–1993:M3 in red triangles.

Two notable features stand out. First, the pre-sample period from 1990:M2 through 1993:M3 features relatively large negative interest rate surprises as well as positive stock market surprises. Thus, the pre-sample is dominated by surprises that classify as a monetary policy shock according to the comovement restrictions. The Federal Reserve lowered the interest rate during several intermeeting moves to cushion the effects of the savings and loan crisis on the U.S. economy during this time. Therefore, the absence of this relatively important episode may be the reason why the sign-restriction approach on high-frequency variables only lacks the power to identify a reasonably sized monetary impulse for the sample starting in 1993:M4.

Second, there are fewer large interest rate and stock price surprises at the quarterly frequency as compared with the monthly and the intradaily frequency. The similarity between monthly and intradaily frequency can be rationalized by the fact that there is rarely more than one FOMC announcement per month.[36] However, there are several surprises within a given quarter that might potentially offset each other, thus leading to smaller surprises in the aggregate. This loss in variability might make it more difficult to identify a relationship between high-frequency and low-frequency variables.

B.5 Robustness

This section presents a series of robustness checks regarding our baseline empirical strategy.

B.5.1 Specifications Exploiting the Aggregate Time-Series Dimension

Sign Restrictions on Low-Frequency Variables. To explore whether our VAR results are driven by the low-frequency sign restrictions, we analyze the sensitivity of our results by using the identification strategy of Jarociński and Karadi (2020), which imposes sign restrictions on the comovement of high-frequency surprises only. Figure B.8 shows that this identification yields a rather implausible interest rate impulse response in our sample (as also discussed above in Section 4.2). In particular, the initial impulse is small and disproportionally compensated in the medium term. However, our main results on the firms' extensive margin are robust. While the responses of aggregate productivity measures are closer to zero, average productivity now declines significantly in the short run.

Poor Man's Proxy. Jarociński and Karadi (2020) also propose a simpler identification of monetary policy shocks based on sign restrictions on the comovement of surprises in a given month. We follow their approach and construct the so-called poor man's proxy at quarterly frequency: we impose sign restrictions on the sum of daily surprises in a quarter. The implicit assumption is that each

[36]Since 1994, most FOMC announcements are regularly scheduled meetings and take place monthly or every six weeks. The remaining FOMC announcements are unscheduled meetings and conference calls, which are, however, rare in the sample we consider.

Figure B.8. VAR with Sign Restrictions on High-Frequency Variables Only

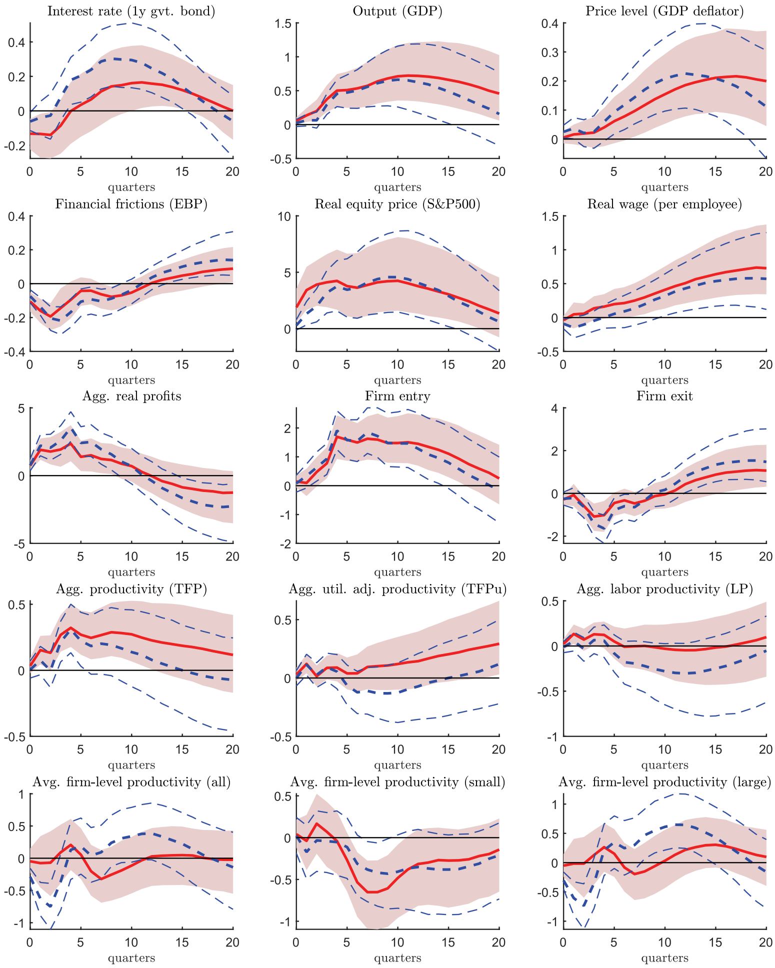

Note: Impulse response functions to a monetary policy shock according to the baseline identification (solid lines) and a shock identified by using sign restrictions on high-frequency variables only (dashed). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

quarter features a monetary policy shock or a central bank information shock.[37] We incorporate the poor man's proxy of monetary policy shock into a VAR with zero restrictions, ordering the shock first while imposing short-run restrictions. Figure B.9 shows that the qualitative impulse response patterns for all variables remain roughly unchanged.

Surprises from Scheduled FOMC Announcements Only. FOMC decisions at unscheduled meetings and conference calls often occur when economic conditions deteriorate abruptly, i.e., may constitute an endogenous response of monetary policy to contemporaneous shocks (Nakamura and Steinsson 2018). This raises concerns with respect to the proper identification of monetary policy shocks. Figure B.10 shows that excluding the unscheduled decisions particularly affects the interest rate response. The qualitative response pattern of all other variables is roughly unchanged.

Are Surprises Unpredictable? Miranda-Agrippino and Ricco (2021) show that the three-month federal funds future surprises are serially correlated and predictable by macrofinancial factors. We thus explore the sensitivity of our results by abandoning the zero restrictions of the VAR in Equation (37) and estimate a fully parameterized VAR. Figure B.11 shows that our main results are broadly unchanged when relaxing the restrictions.

Monthly Frequency. Our empirical specification uses quarterly data, while several other contributions use monthly data. We hence build a monthly data set by interpolating our baseline quarterly variables using monthly proxies (if available) or by cubic splines.[38] Figure B.12 shows that our results are robust to the

---

[37]In practice, monetary policy and information shocks occur simultaneously in a month (Jarociński and Karadi 2020). We obtain similar results when we impose a weaker version of the poor man's sign restriction that allows monetary policy shocks and information to occur simultaneously in a quarter. This procedure involves sign restrictions on (1) daily and (2) monthly surprises to identify pure monetary policy and information shocks. These shocks are converted to quarterly frequency by summation.

[38]Specifically, we include the core variables from the Jarociński and Karadi (2020) data set, i.e., surprises, interest rate, activity, prices, excess bond premium, and the stock price, in our data set and interpolate profits, wages, measures of firm dynamics, and productivity by cubic splines. Real GDP and GDP deflator are each interpolated by industrial production and consumer prices.

Figure B.9. VAR with Poor Man's Proxy

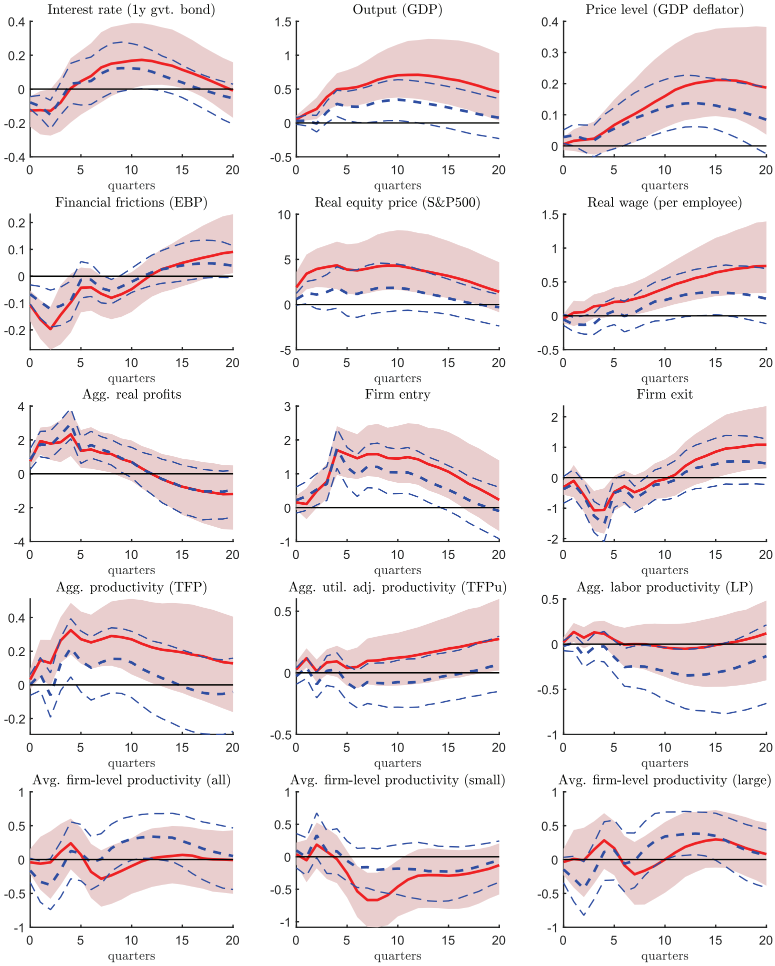

Note: Impulse response functions to a monetary policy shock according to the baseline identification (solid lines) and a shock identified using the poor man's proxy of a monetary policy shock in a VAR with zero restrictions (dashed). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.10. Surprises from Scheduled FOMC Meetings Only

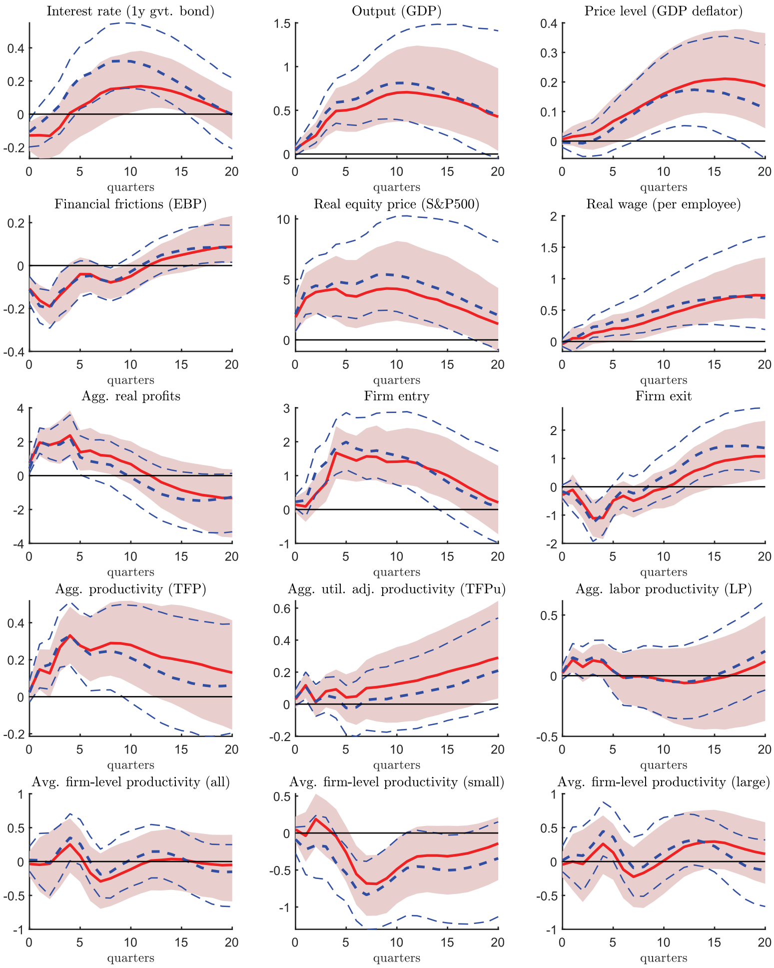

Note: Impulse response functions to a monetary policy shock according to the baseline using all FOMC announcements (solid lines) and when using only surprises from scheduled FOMC announcements (dashed). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.11. Unrestricted VAR

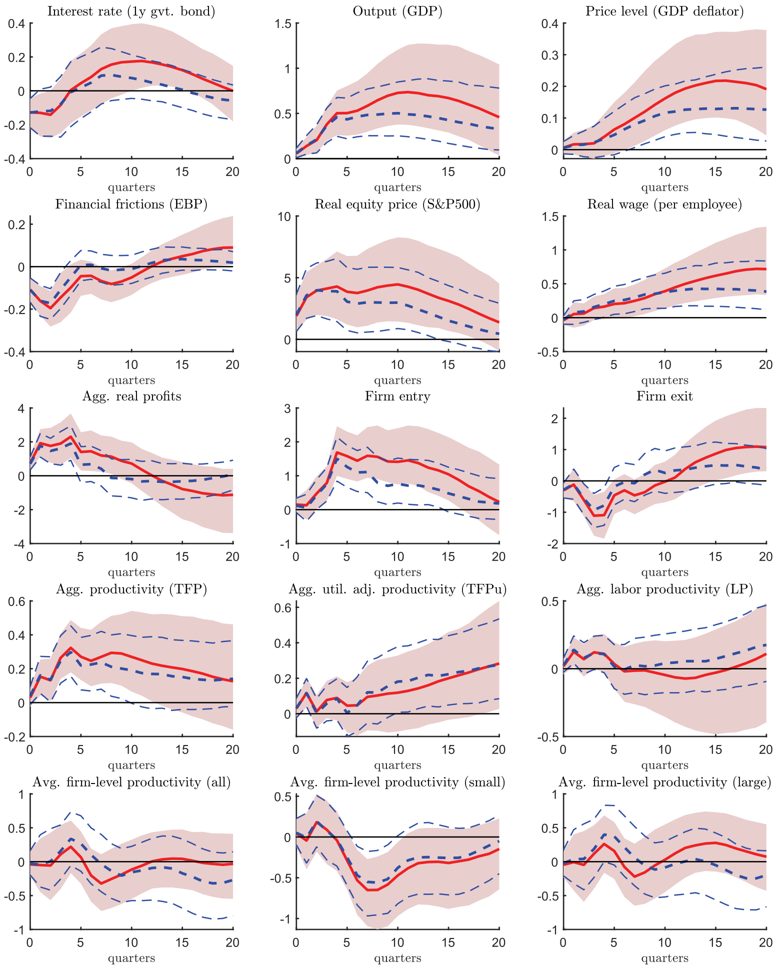

Note: Impulse response functions to a monetary policy shock according to the baseline using a restricted VAR (solid lines) and an unrestricted VAR (dashed). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.12. VAR with Monthly Interpolated Time Series

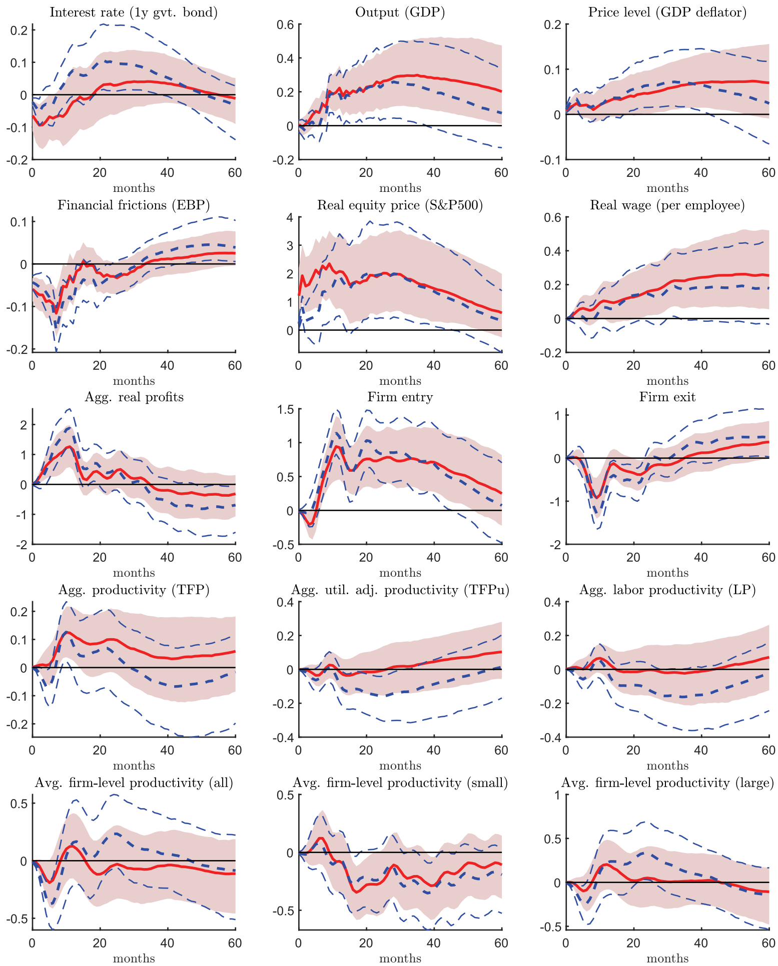

Note: Impulse response functions to a monetary policy shock according to the baseline identification (solid lines) and a shock identified by using sign restriction on high-frequency variables only (dashed) for the data set at monthly frequency. The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

alternative data frequency for both our baseline identification procedure and that of Jarociński and Karadi (2020).

Alternative Measures and Sample Splits. Identifying monetary policy shocks at the ZLB is associated with potential issues (Ikeda et al. 2024). This suggests using measures of the interest rate that specifically account for the ZLB and nonstandard measures. Figure B.13 shows that our baseline results are robust when using the Wu and Xia (2016) shadow rate, the simple federal funds rate or the two-year government bond yield. Our results are also unaffected by alternative measures for output (industrial production), the price level (consumer price index), and financial frictions (spread of BAA-rated corporate bonds relative to 10-year Treasury yield); see Figure B.14. They are furthermore robust to different sample splits, i.e., when considering a sample up to the Great Recession only or excluding the Great Recession, as shown in Figures B.15 and B.16.

Local Projections. Estimated VAR impulse responses may be biased for more distant lags if the selected lag order is too small. We hence use the local projection method by Jordà (2005), which is more flexible and imposes weaker dynamic restrictions.[39] The local projection (LP) model is given by

$$y_{t+h} = \alpha_h + x_t\beta_h + \sum_{j=1}^{2} y_{t-j}\theta_{j,h} + \sum_{j=1}^{2} w_{t-j}\gamma_{j,h} + u_{t+h}, \qquad \text{(B.2)}$$

where y_t is the dependent variable, x_t is a monetary policy shock and w_t is a vector of controls. Figure B.17 shows the estimates of the VAR with the poor man's proxy, the LP estimates with macrofinancial controls, and LP estimates without additional controls. Overall, the estimated LP impulse responses with controls are qualitatively similar, though somewhat more erratic.[40] This confirms our main

[39]LPs and VARs estimate the same impulse responses in a recursive VAR with unrestricted lag structure (Plagborg-Møller and Wolf 2021). As we use a flat prior in the VAR, the estimates are directly comparable for $h \leq 2$.

[40]The erratic pattern of LP impulse response functions as compared with a VAR is due to a loss in efficiency in the estimation and fewer dynamic restrictions (Barnichon and Brownlees 2019).

Figure B.13. VAR with Various Interest Rate Measures

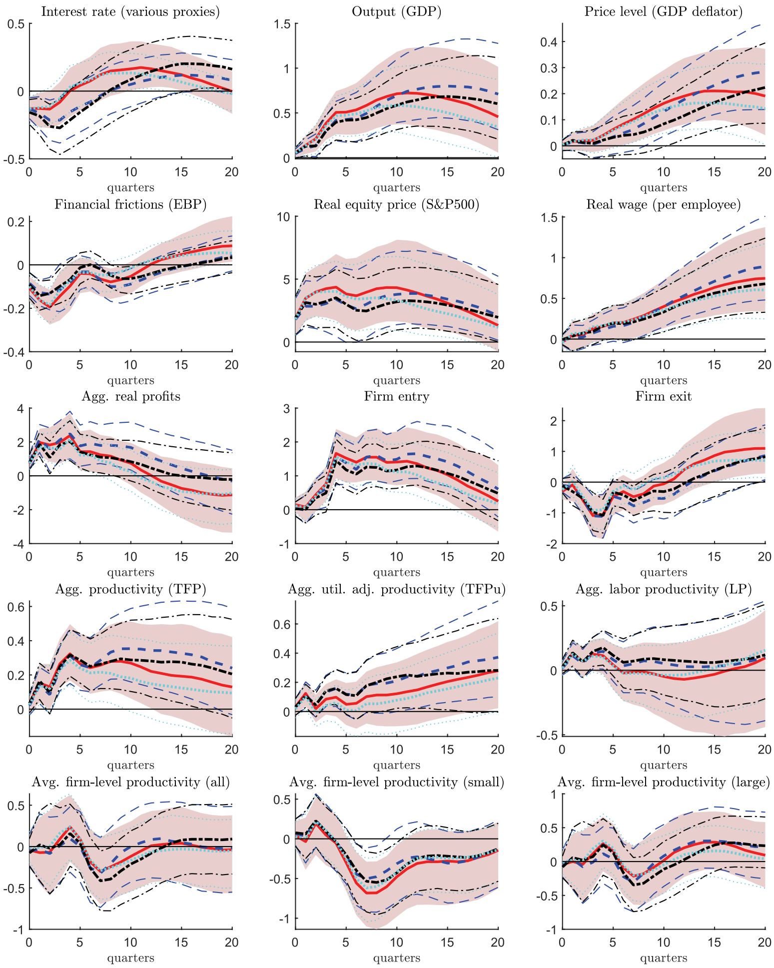

Note: Impulse response functions to a monetary policy shock according to the baseline using the one-year government bond yield (solid lines) and using alternative measures of monetary policy: the federal funds rate (dashed), the federal funds rate extended by the shadow short rate of Wu and Xia (2016) (dashed-dotted), and the two-year government bond yield (dotted). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.14. VAR with IP, CPI, BAA

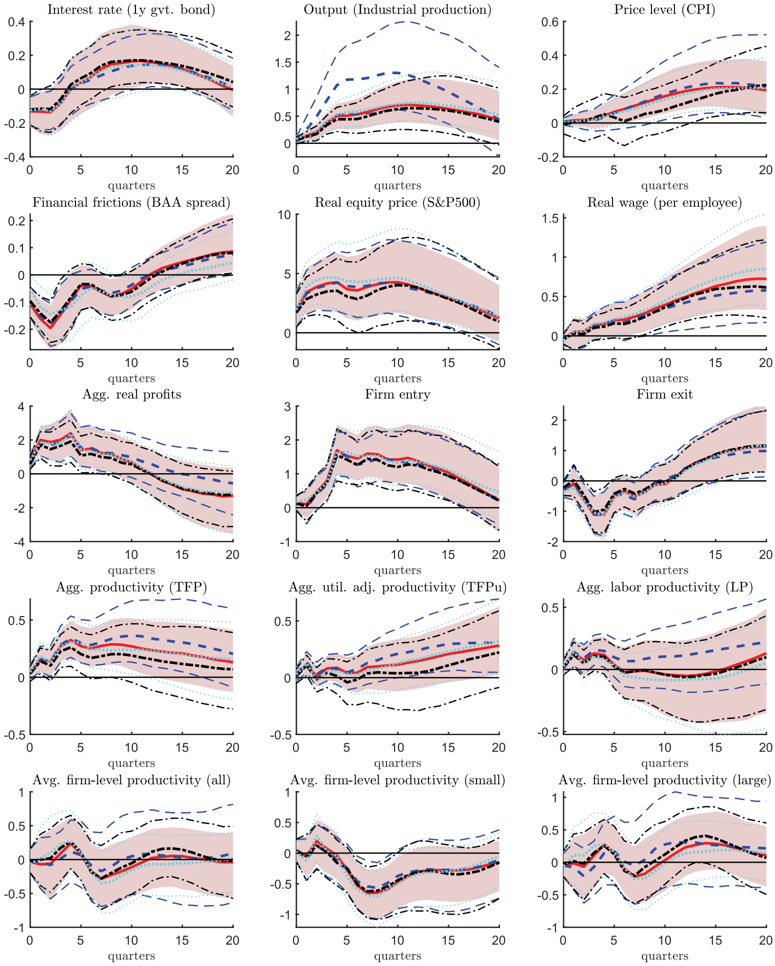

Note: Impulse response functions to a monetary policy shock using the baseline measures (solid lines) and when using alternative measures of activity (dashed), prices (dashed-dotted), and financial frictions (dotted). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

results. In contrast, the estimates without controls exhibit substantial output and price puzzles and may hence be regarded as implausible. This highlights the importance of including macrofinancial

Figure B.15. VAR Pre–Great Recession

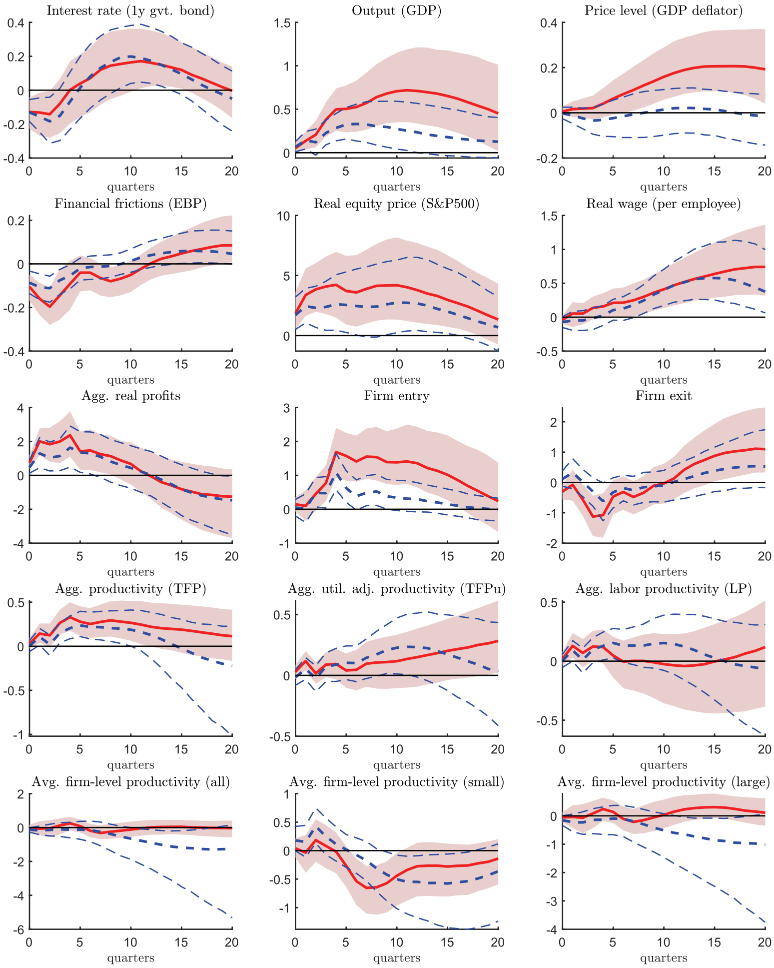

Note: Impulse response functions to a monetary policy shock using the baseline sample (solid lines) and when using a sample until 2008:Q2 (dashed). For the pre–Great Recession sample, a moderately loose Minnesota prior is used with overall tightness of $\lambda = 0.7$ and the federal funds rate is used as the policy indicator. The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.16. VAR Excluding Surprises from Apex of Great Recession

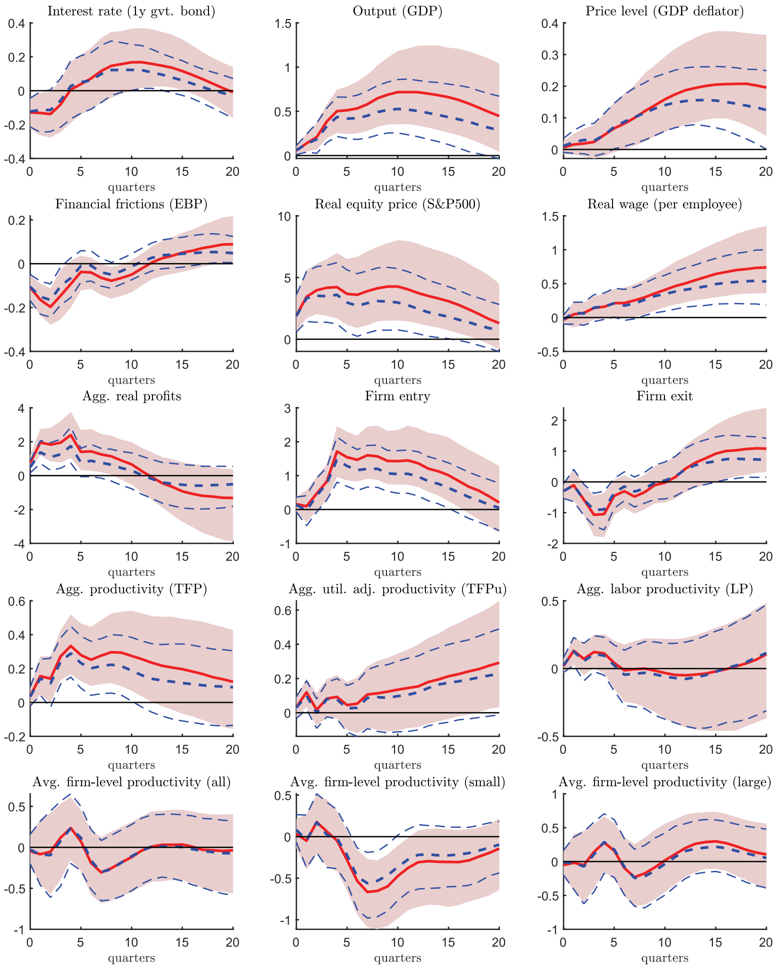

Note: Impulse response functions to a monetary policy shock according to the baseline using the full sample (solid lines) and a sample excluding the apex of the Great Recession, i.e., ex 2008:Q3–2009:Q2 (dashed). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

Figure B.17. Poor Man's Proxy: VAR and Local Projections

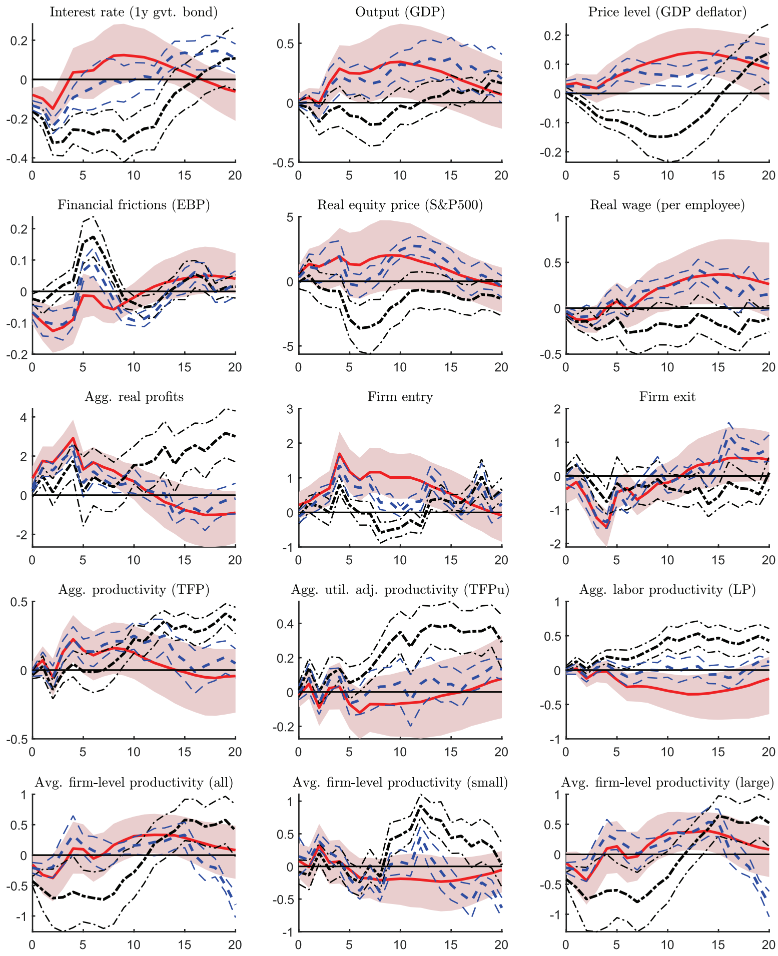

Note: Impulse response functions to a monetary policy shock identified using the poor man's proxy in a VAR with zero restrictions (solid lines), in a local projection with a set of macroeconomic and financial controls (dashed), and in a local projection without additional controls (dashed-dotted). The thick lines are the median (point) estimates; the shaded areas (thin lines) are the 68 percent credible intervals (confidence intervals) of the VAR (local projection). Responses are shown in percent deviations, except for the interest rate and the measure of financial frictions (percentage point deviations).

controls when estimating the effects of a monetary policy shock via local projections. Further below, we hence also discuss the impact of alternative controls on our PLP results.

B.5.2 Specifications Exploiting the Panel Dimension

Alternative Sets of Controls. As discussed above, the set of controls may have a decisive impact on the estimated impulse responses in local projections. We hence explore how our PLP results change when including only macro controls (i.e., excluding stock price and financial frictions from the set of controls) and when including no macrofinancial controls. Figure B.18 shows that the qualitative patterns and the confidence intervals indeed hinge on the set of control variables. Most importantly, the response of firm-level productivity across small firms is robust to the set of controls and remains significant throughout. In contrast, the responses for all firms and large firms change considerably. We view this as further evidence that expansionary monetary policy decreases average productivity of small firms. At the same time, this exercise highlights the importance of accounting for the macrofinancial state when estimating the effects of monetary policy using PLPs to avoid misguided inference.

Heterogeneous Effects by Firm Size. The sample split by firm size in the baseline PLP does not account for potentially heterogeneous responses of different sectors to monetary policy, and that such sectoral heterogeneity could be correlated with firm size. To investigate whether small and large firms respond differently to monetary policy, we interact our sign-identified monetary policy shock with a firm-size dummy in a PLP while including various time-by-sector fixed effects to control for sectoral heterogeneity following Anderson and Cesa-Bianchi (2020) and Ottonello and Winberry (2020). We estimate the following model:

$$\begin{aligned} y_{i,t+h} - y_{i,t-1} &= \alpha_{i,h} + D_t\eta_h + \Psi + (x_t \cdot size_{i,t})\beta_h \\ &\quad + \sum_{j=1}^{2} \Delta y_{i,t-j}\theta_{h,j} + \sum_{j=1}^{2} size_{i,t-j}\delta_{h,j} + \sum_{j=1}^{2} w_{t-j}\gamma_{h,j} + u_{i,t+h}, \end{aligned} \tag{B.3}$$

Figure B.18. Alternative Set of Controls in PLPs

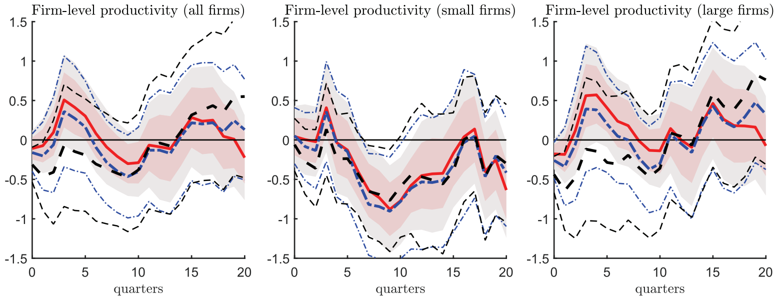

Note: Impulse response functions to the sign-identified monetary policy shock for different specifications of the PLP, using a full set of macrofinancial controls as a baseline (solid lines), macro controls only (dashed), and no additional controls (dashed-dotted). Shaded areas depict the 68 percent and 90 percent confidence intervals for the baseline; the thin lines show 68 percent confidence intervals for the alternative specifications. Responses are shown in percent deviations.

where $size_{i,t}$ is a dummy variable that is 1 for firms with sales lower than \$10 million and $\Psi \in \{\Xi_s, \Xi_{y,s}, \Xi_{t,s}\}$ is a time-by-sector fixed effect for all t (Ξ_s), by year ($\Xi_{y,s}$), and by quarter ($\Xi_{t,s}$). Note that Ξ_s is simply a sector fixed effect and thus identical to $\alpha_{i,h}$ since firms do not switch sectors. Following Ottonello and Winberry (2020) and Anderson and Cesa-Bianchi (2020), we include the lagged firm size dummy as an additional control.

Figure B.19 shows the corresponding relative impulse response functions of firm-level productivity for three different time-by-sector fixed-effects specifications (sector only, year-by-sector, quarter-by-sector). In all cases, the relative responses are first positive and turn negative after around six quarters. This suggests that the productivity of small firms is initially less responsive and becomes subsequently more responsive to monetary policy compared with that of large firms. This result is broadly in line with our baseline findings using a sample split by firm size (Figure 3, third row, middle and right panel). However, the relative responses are largely not significant at conventional levels. In interpreting these results, one needs to keep in mind that the year-by-sector fixed effects absorb a substantial amount of information contained in the macro and financial controls (middle panel); the quarter-by-sector fixed effects even

Figure B.19. Relative Impulse Response Functions of Small Firms

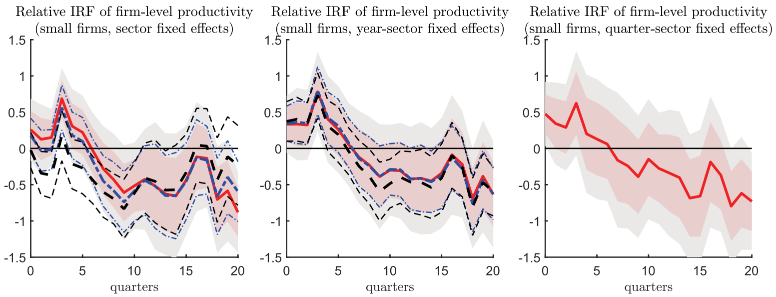

Note: Relative impulse response functions to the sign-identified monetary policy shock for different specifications of the PLP, using a full set of macrofinancial controls as a baseline (solid lines), macro controls only (dashed), and no additional controls (dashed-dotted), and various degrees of time-by-sector fixed effects. Shaded areas depict the 68 percent and 90 percent confidence intervals for the baseline; the thin dashed and dashed-dotted lines show 68 percent confidence intervals for the alternative specifications. Responses are shown in percent deviations.

absorb it completely (right panel). In addition, small firms make up only one-third of our sample. Therefore, we view the evidence from this exercise to be somewhat mixed and inconclusive, which also guides our interpretation of the results in Section 4.3.

B.5.3 Construction of Firm-Level Productivity Series

Sector Composition. Firms in the manufacturing and services sectors are the closest counterparts to the theoretical goods-producing firms and constitute more than 70 percent of the firms in the Compustat sample (see Table B.2).[41] We hence investigate how average firm-level productivity in these sectors responds to monetary policy. Figure B.20 shows the effects of a monetary policy shock using the baseline VAR with all nonfinancial firms ex utilities and alternative sectoral splits: (A) manufacturing and services, (B) manufacturing

[41]The productivity dynamics for nonfinancial firms excluding utilities are mainly driven by these sectors (see Panel A in Figure B.3 and the related discussion in Section B.2).

Figure B.20. Sector Splits and Firm-Level Productivity

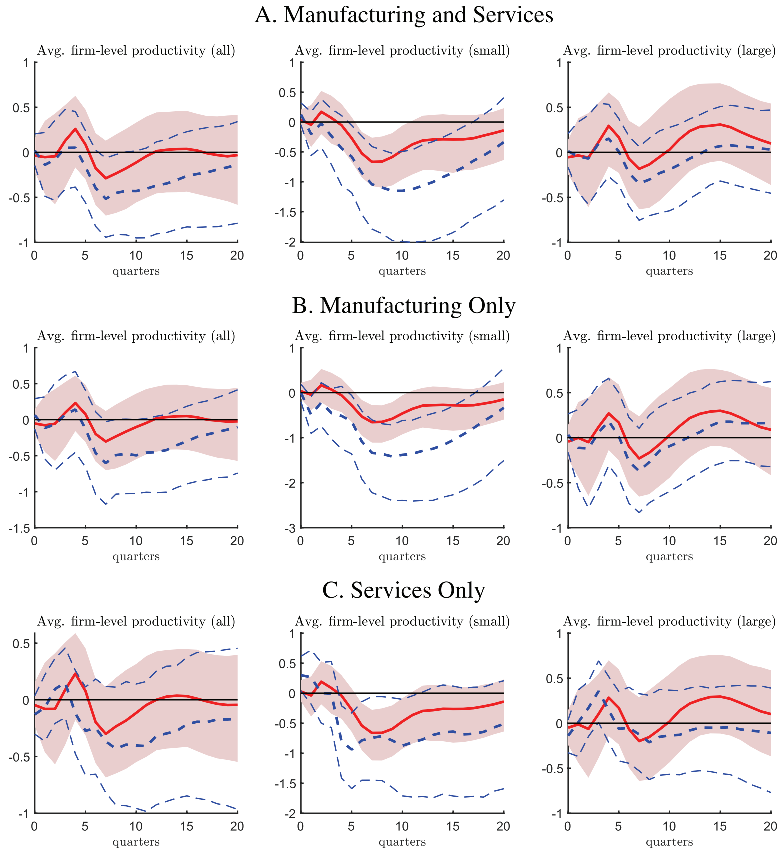

Note: Impulse response functions to a monetary policy shock according to the baseline using all nonfinancial sectors ex utilities (solid lines) and alternative sector splits (dashed). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations.

only, and (C) services only. The responses of average productivity for these three sector splits are somewhat shifted downward compared with the benchmark, but are still barely significant. However, the previously documented decline for small firms is even more pronounced for the manufacturing and service sectors.

Production Function Estimation. Our micro productivity measures are based on estimations of firm-level production functions.

Figure B.21. Production Function and Firm-Level Productivity

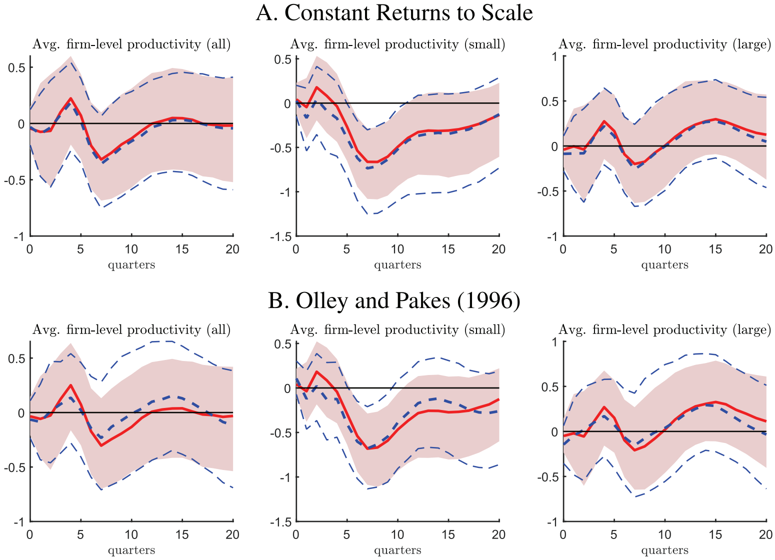

Note: Impulse response functions to a monetary policy shock according to the baseline using a panel OLS regression to estimate firm-level productivity (solid lines) and alternative production function estimations (dashed). The thick lines are the median estimates; the shaded areas and thin lines are the 68 percent credible intervals. Responses are shown in percent deviations.

We explore the robustness of our baseline—a panel OLS regression with fixed effects—by considering two alternative methods. First, we impose constant returns to scale for the share of capital α and labor β, i.e., $\alpha = 1 - \beta$, using a restricted panel OLS regression with fixed effects. Second, we use the semiparametric estimation approach of Olley and Pakes (1996) to control for input factor endogeneity. Table B.4 shows that these alternative methods lead to an increase in the estimated shares of capital and labor as compared with the baseline. Nevertheless, the resulting productivity time series are highly similar to the baseline estimate; see panel B in Figure B.3. Figure B.21 confirms that these alternative productivity measures do not respond to monetary policy shocks differently than the baseline series.

References

Acharya, V. V., M. Crosignani, T. Eisert, and C. Eufinger. 2020. "Zombie Credit and (Dis-)Inflation: Evidence from Europe." NBER Working Paper No. 27158, May.

Acharya, V. V., T. Eisert, C. Eufinger, and C. Hirsch. 2019. "Whatever It Takes: The Real Effects of Unconventional Monetary Policy." *Review of Financial Studies* 32 (9): 3366–411.

Acharya, V. V., S. Lenzu, and O. Wang. 2021. "Zombie Lending and Policy Traps." NBER Working Paper No. 29606.

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li. 2019. "Missing Growth from Creative Destruction." *American Economic Review* 109 (8): 2795–822.

Anderson, G., and A. Cesa-Bianchi. 2020. "Crossing the Credit Channel: Credit Spreads and Firm Heterogeneity." IMF Working Paper No. 20/267.

Andrews, D., and F. Petroulakis. 2019. "Breaking the Shackles: Zombie Firms, Weak Banks and Depressed Restructuring in Europe." ECB Working Paper No. 2240.

Antoni, M., and T. Sondershaus. 2021. "Do Asset Purchase Programmes Shape Industry Dynamics? Evidence from the ECB's SMP on Plant Entries and Exits." Discussion Paper No. 12/2019, Halle Institute for Economic Research.

Barnichon, R., and C. Brownlees. 2019. "Impulse Response Estimation by Smooth Local Projections." *Review of Economics and Statistics* 101 (3): 522–30.

Basu, S., J. G. Fernald, and M. S. Kimball. 2006. "Are Technology Improvements Contractionary?" *American Economic Review* 96 (5): 1418–48.

Bergin, P. R., and G. Corsetti. 2008. "The Extensive Margin and Monetary Policy." *Journal of Monetary Economics* 55 (7): 1222–37.

Bergin, P. R., L. Feng, and C. Lin. 2018. "Firm Entry and Financial Shocks." *Economic Journal* 128 (609): 510–40.

Bernanke, B. S., M. Gertler, and S. Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics*, Vol. 1 (C), ed. J. B. Taylor and M. Woodford, 1341–93.

Bilbiie, F. O., I. Fujiwara, and F. Ghironi. 2014. "Optimal Monetary Policy with Endogenous Entry and Product Variety." *Journal of Monetary Economics* 64: 1–20.

Bilbiie, F. O., F. Ghironi, and M. J. Melitz. 2008. "Monetary Policy and Business Cycles with Endogenous Entry and Product Variety." In *NBER Macroeconomics Annual 2007*, Vol. 22, ed. D. Acemoglu, K. Rogoff, and M. Woodford, 299–379. Chicago: University of Chicago Press.

———. 2012. "Endogenous Entry, Product Variety, and Business Cycles." *Journal of Political Economy* 120 (2): 304–45.

Bittner, C., F. Fecht, and C.-P. Georg. 2021. "Contagious Zombies." Working Paper No. 15/2021, Deutsche Bundesbank.

Caballero, R. J., T. Hoshi, and A. K. Kashyap. 2008. "Zombie Lending and Depressed Restructuring in Japan." *American Economic Review* 98 (5): 1943–77.

Cacciatore, M., G. Fiori, and F. Ghironi. 2016. "Market Deregulation and Optimal Monetary Policy in a Monetary Union." *Journal of International Economics* 99: 120–37.

Caldara, D., and E. Herbst. 2019. "Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs." *American Economic Journal: Macroeconomics* 11 (1): 157–92.

Campbell, J. R. 1998. "Entry, Exit, Embodied Technology, and Business Cycles." *Review of Economic Dynamics* 1 (2): 371–408.

Christiano, L. J., M. Eichenbaum, and C. L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

Clementi, G. L., and B. Palazzo. 2016. "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations." *American Economic Journal: Macroeconomics* 8 (3): 1–41.

———. 2019. "Investment and the Cross-Section of Equity Returns." *Journal of Finance* 74 (1): 281–321.

Colciago, A., and R. Silvestrini. 2022. "Monetary Policy, Productivity, and Market Concentration." *European Economic Review* 142: Article 103999.

Dixit, A. K., and J. E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67 (3): 297–308.

Evans, C. L. 1992. "Productivity Shocks and Real Business Cycles." *Journal of Monetary Economics* 29 (2): 191–208.

Fernald, J. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Working Paper No. 2012-19, Federal Reserve Bank of San Francisco.

———. 2015. "Productivity and Potential Output Before, During, and After the Great Recession." In *NBER Macroeconomics Annual 2014,* Vol. 29, ed. J. A Parker and M. Woodford, 1–51. Chicago: University of Chicago Press.

Garga, V., and S. R. Singh. 2021. "Output Hysteresis and Optimal Monetary Policy." *Journal of Monetary Economics* 117: 871–86.

Gertler, M., and P. Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58 (1): 17–34.

———. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76.

Ghironi, F., and M. J. Melitz. 2005. "International Trade and Macroeconomic Dynamics with Heterogeneous Firms." *Quarterly Journal of Economics* 120 (3): 865–915.

Gilchrist, S., and E. Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102 (4): 1692–720.

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.

Hamano, M., and F. Zanetti. 2017. "Endogenous Product Turnover and Macroeconomic Dynamics." *Review of Economic Dynamics* 26: 263–79.

———. 2018. "On Quality and Variety Bias in Aggregate Prices." *Journal of Money, Credit and Banking* 50 (6): 1343–63.

———. 2022. "Monetary Policy, Firm Heterogeneity, and Product Variety." *European Economic Review* 144: Article 104089.

Hamilton, J. D. 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *Review of Economics and Statistics* 100 (5): 831–43.

Heckman, J. J. 1993. "What Has Been Learned About Labor Supply in the Past Twenty Years?" *American Economic Review* 83 (2): 116–21.

Hopenhayn, H. A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–50.

Hoshi, T., and A. K. Kashyap. 2004. "Japan's Financial Crisis and Economic Stagnation." *Journal of Economic Perspectives* 18 (1): 3–26.

Hummels, D., and P. J. Klenow. 2005. "The Variety and Quality of a Nation's Exports." *American Economic Review* 95 (3): 704–23.

Ikeda, D., S. Li, S. Mavroeidis, and F. Zanetti. 2024. "Testing the Effectiveness of Unconventional Monetary Policy in Japan and the United States." *American Economic Journal: Macroeconomics* 16 (2): 250–86.

Jaimovich, N., and M. Floetotto. 2008. "Firm Dynamics, Markup Variations, and the Business Cycle." *Journal of Monetary Economics* 55 (7): 1238–52.

Jarociński, M., and P. Karadi. 2020. "Deconstructing Monetary Policy Surprises — The Role of Information Shocks." *American Economic Journal: Macroeconomics* 12 (2): 1–43.

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

Jordà, O., M. Schularick, and A. M. Taylor. 2015. "Betting the House." *Journal of International Economics* 96: S2–S18.

King, R. G., C. I. Plosser, and S. T. Rebelo. 1988. "Production, Growth and Business Cycles: II. New Directions." *Journal of Monetary Economics* 21 (2): 309–41.

Lewis, V. 2009. "Business Cycle Evidence on Firm Entry." *Macroeconomic Dynamics* 13 (5): 605–24.

———. 2013. "Optimal Monetary Policy and Firm Entry." *Macroeconomic Dynamics* 17 (8): 1687–710.

Lewis, V., and C. Poilly. 2012. "Firm Entry, Markups and the Monetary Transmission Mechanism." *Journal of Monetary Economics* 59 (7): 670–85.

Liu, P., K. Theodoridis, H. Mumtaz, and F. Zanetti. 2019. "Changing Macroeconomic Dynamics at the Zero Lower Bound." *Journal of Business and Economic Statistics* 37 (3): 391–404.

Meier, M., and T. Reinelt. 2022. "Monetary Policy, Markup Dispersion, and Aggregate TFP." *Review of Economics and Statistics* 106 (4): 1–45.

Melitz, M. J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–725.

Melosi, L. 2017. "Signalling Effects of Monetary Policy." *Review of Economic Studies* 84 (2): 853–84.

Midrigan, V., and D. Y. Xu. 2014. "Finance and Misallocation: Evidence from Plant-Level Data." *American Economic Review* 104 (2): 422–58.

Miranda-Agrippino, S., and G. Ricco. 2021. "The Transmission of Monetary Policy Shocks." *American Economic Journal: Macroeconomics* 13 (3): 74–107.

Moll, B. 2014. "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?" *American Economic Review* 104 (10): 3186–221.

Moran, P., and A. Queralto. 2018. "Innovation, Productivity, and Monetary Policy." *Journal of Monetary Economics* 93: 24–41.

Nakamura, E., and J. Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–330.

Olley, G. S., and A. Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–97.

Ottonello, P., and T. Winberry. 2020. "Financial Heterogeneity and the Investment Channel of Monetary Policy." *Econometrica* 88 (6): 2473–502.

Peek, J., and E. S. Rosengren. 2005. "Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan." *American Economic Review* 95 (4): 1144–66.

Plagborg-Møller, M., and C. K. Wolf. 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica* 89 (2): 955–80.

Ravenna, F., and C. E. Walsh. 2006. "Optimal Monetary Policy with the Cost Channel." *Journal of Monetary Economics* 53 (2): 199–216.

Romer, C. D., and D. H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90 (3): 429–57.

Rossi, L. 2019. "The Overshooting of Firms' Destruction, Banks and Productivity Shocks." *European Economic Review* 113: 136–55.

Rotemberg, J. J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *Review of Economic Studies* 49 (4): 517–31.

Smets, F., and R. Wouters. 2007. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Wu, J. C., and F. D. Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (2–3): 253–91.

ECB Communication as a Stabilization and Coordination Device: Evidence from Ex Ante Inflation Uncertainty*

Cecilia Melo Fernandes
International Monetary Fund

This paper investigates the impact of ECB communication of its assessment of the economic outlook on ex ante inflation uncertainty and sheds light on how central bank information shocks operate. The results suggest that central bank information acts as a "coordination device" able to influence opinions and actions. Most importantly, it generates a "stabilizer effect" by substantially decreasing the dispersion among the inflation point forecasts, which converge toward their aggregate mean. The paper not only helps to explain the impact of central bank information but is also useful for policymakers to define a communication strategy that attenuates ex ante inflation uncertainty.

JEL Codes: D83, E52, E58, E65, G14.

1. Introduction

In the past decades, central bank communication has gained increasing importance. It has evolved from a reluctance of central banks to provide precise information on the policy process to a facilitator of conventional monetary policy, eventually becoming a new instrument of monetary policy itself (Blinder 2018; Weidmann 2018; Issing 2019). Central bank communication steers expectations, and the better expectations are aligned with the monetary policy

*The opinions expressed in this paper do not necessarily reflect the views of the International Monetary Fund or of the European Central Bank. Author contact: Monetary and Capital Markets Department, International Monetary Fund, 1900 Pennsylvania Avenue NW, Washington, DC, 20431, United States. E-mail: CMeloFernandes@imf.org.

objective, the more likely it is that the central bank will stabilize aggregate demand and therefore inflation (Clarida, Gali, and Gertler 1999).

Recently, a nascent literature has been shifting the attention from quantifying and estimating the implications of several aspects of communication, such as transparency, clarity, and tone,[1] toward the informative nature of central bank communication. By using high-frequency surprises around central bank announcements, recent research seeks to isolate the communication of assessments of the economy from information about monetary policy, which are conveyed simultaneously in policy announcements (see Andrade and Ferroni 2016; Cieslak and Schrimpf 2019; Kerssenfischer 2019; Jarociński and Karadi 2020).

In this context, there are at least two important gaps in the literature. First, existing studies focus mainly on assessing the impact of central bank information on aggregate measures of expectations and on the economy. While this is consistent with the consensus that disentangling communication about the economic outlook from monetary policy information in central bank communication is important to prevent bias in the estimated effects of monetary policy, there has so far been no attempt to understand the effects of news communicated by the central bank on measures of ex ante uncertainty about the economy, particularly ex ante uncertainty about inflation.

Ex ante uncertainty refers to measurements of uncertainty which does not include the realization of events, in contrast to ex post (or realized) uncertainty, which does. Investigating the relationship between central bank communication and ex ante inflation uncertainty is important because if the latter is exacerbated by communication, it may harm economic activity and the effectiveness of monetary policy in maintaining price and/or financial stability. Inflation uncertainty can increase the costs related to a contractionary monetary policy or counteract an expansionary stimulus by,

[1]These elements are typically proxied by indices or dictionary approaches (see, for example, Eijffinger and Geraats 2006; Minegishi and Cournède 2009; Jegadeesh and Wu 2017; Picault and Renault 2017; Dincer, Eichengreen, and Geraats 2019).

for example, slowing investments and affecting wealth allocation.[2] In addition, increasing inflation uncertainty can be a sign of a central bank's weakening credibility. Therefore, assessing whether central bank communication mitigates or exacerbates inflation uncertainty is very important for monetary policy strategy.

Second, the channels through which central bank information shocks operate and how they affect the ex ante inflation uncertainty are unknown. The closest related discussion in the literature is about how central bank information affects the economy and expectations, focusing on the levels and first moment of inflation. In particular, the discussion revolves around whether central banks convey new information that directly affects forecasts or whether their announcements help market participants and forecasters focus on one particular equilibrium, thereby serving as an impactful coordination device. This debate still remains unresolved.

By making use of the European Central Bank (ECB) Survey of Professional Forecasters (SPF) and the central bank information shocks provided by Jarociński and Karadi (2020), this paper provides a twofold contribution. First, for the first time in the context of the central bank communication literature, the paper disentangles the effects of ECB communication on three different types of ex ante inflation uncertainty: disagreement, average individual uncertainty, and aggregate uncertainty.

In particular, by using local projection methods (Jordà 2005), I find evidence that the ECB's outlook information shocks not only reduce the dispersion across agents' average point forecasts (disagreement) but also make agents less uncertain about their own beliefs (ex ante average individual uncertainty). Both effects result in a lower aggregate ex ante inflation uncertainty. This decomposition across different types of ex ante uncertainties is possible because, in contrast with other surveys used in the literature, the ECB SPF

[2]There is substantial evidence in the literature on the negative impact of inflation uncertainty on financial and macroeconomic variables. Inflation uncertainty may induce agents to postpone investment or savings decisions and reduce market efficiency due to an increase in the volatility of both relative prices and risks regarding income streams from nominal financial and wage contracts (Friedman 1977; Bloom 2009). Furthermore, inflation uncertainty can lead to shifts in wealth allocation between creditors and debtors (see Fama 1976; Barnea, Dotan, and Lakonishak 1979; Grauer and Litzenberger 1979).

provides both point (mean) forecasts and their distributions for each individual forecaster.

Second, given that there is evidence that ECB communication affects ex ante inflation uncertainty, the next question is: how does it happen? In answering this question, this paper also sheds light on the channels through which central bank communication operates. The particularities and the complementarities of each ex ante uncertainty measure provide unique insights when interpreting the results of the reactions of these measures to central bank information shocks. Most importantly, disagreement reflects the dispersion of projections across forecasters but does not provide information about each forecaster's uncertainty regarding their own forecast. In contrast, average individual uncertainty assesses the uncertainty of each individual regarding their own projections, so it is often considered a better proxy for uncertainty (see Abel et al. 2016; Glas and Hartmann 2016; Glas 2020). Some studies even show that disagreement in survey forecasts could be more reflective of differences in opinion than of uncertainty (see Diether, Malloy, and Scherbina 2002; Mankiw, Reis, and Wolfers 2004).

Given that central bank information shocks lead agents to disagree less among each other about their inflation projections and also to become less uncertain about their own projections, I find evidence that they act as a public signal, which is effective in coordinating opinions and actions. Furthermore, forecasters are comfortable with incorporating the public signal emitted by the central bank in the assessment of their analysis. This also implies that this signal is valuable and on average contributes to strengthen their confidence in their predications.

In addition, after a central bank information shock, the point forecasts converge toward their mean. This convergence implies that the central bank communication generates a "stabilizer effect" in which the dispersion among the point forecasts decreases and, most importantly, this convergence moves toward the mean. This convergence is very important, as it induces a steady consensus among the forecasters more in line with the ECB's objectives, in contrast to the alternative, which would imply a convergence of the point forecasts toward one of the tails.

This paper is organized as follows: Section 2 provides a review of the related literature. Section 3 provides a detailed description of the databases and how uncertainty measures and the central bank communication shocks used in this study are estimated. Section 4 summarizes the estimation methodology using local projections. Section 5 explains the identification strategy for the econometric analysis. Sections 6 and 7, respectively, show the results and the robustness checks. Section 8 concludes.

2. Related Literature

Typically, empirical studies exploiting the relationship between central bank communication and uncertainty focus on the transparency aspect of central bank communication as the object of study. In most cases, these studies use survey-based data to measure uncertainty as the dispersion of individual forecasts around the average forecast (disagreement) or around the forecast outcome (mean forecast error). Likewise, most of the studies employ panel data for different economies. Within this framework, the literature provides evidence that greater central bank transparency reduces inflation uncertainty (Ehrmann, Eijffinger, and Fratzscher 2012;[3] Siklos 2013; Naszodi et al. 2016).

This paper is the first to investigate the relationship between the ECB communication and ex ante inflation uncertainty in the euro area using survey-based measures of inflation uncertainty. As explained in Section 3, in order to measure ECB communication, I use the new data set on central bank information shocks from Jarociński and Karadi (2020), which are estimated using high-frequency data. These shocks ultimately consist of ECB communication about the economy. Furthermore, by following Engelberg, Manski, and Williams (2009) and Melo Fernandes and Kenny (2024), I estimate three ex ante uncertainty measures using the ECB SPF: disagreement, average individual uncertainty, and aggregate uncertainty.

Another common approach for estimating inflation uncertainty in the literature is from an ex post perspective, either by

[3]In addition to transparency, Ehrmann, Eijffinger, and Fratzscher (2012) also construct a measure of central bank communication based on dummy variables, which specify whether or not a central bank has announced a quantified inflation objective.

estimating conditional variance using generalized autoregressive conditional heteroskedasticity (GARCH) models (Grier and Perry 2000; Fountas, Ioannidis, and Karanasos 2004; Kontonikas 2004; Conrad and Karanasos 2005) or stochastic volatility (see Berument, Yalcin, and Yildirim 2009; Chan 2017). To the best of my knowledge, the paper by Kliesen and Schmid (2004) is the first to investigate how ex post inflation uncertainty reacts to central bank communication. They define inflation uncertainty as the conditional volatility of inflation compensation, i.e., the additional yield that investors require to hold nominal assets that are exposed to inflation risk, and following a common event analysis approach based on Kohn and Sack (2003), they find that Federal Reserve communication reduces ex post inflation uncertainty.

In contrast to market-based measures, expectations and uncertainty measures derived from survey-based sources do not incorporate any additional compensation for risk and liquidity premia that may cause distortions in the signals and drivers of the measures.[4] On the other hand, the information content of survey data on inflation expectations is sometimes questioned because these expectations might not correspond to those on which economic decisions are based or to those that economic agents truly think. In addition, these measures are more subject to mistakes. These arguments are, however, unlikely to apply in the case of professionals who make macroeconomic forecasts as part of their regular duties (see Garcia 2003). Furthermore, survey-based measures have a clear advantage in that regard, as they contain direct estimates of future inflation outcomes. Therefore, ex ante survey-based inflation uncertainty measures are arguably the most appropriate for the purpose of this paper.

The closest study related to this paper is by Jitmaneeroj, Lamla, and Wood (2019), who analyze the impact of central bank transparency on three types of uncertainty: disagreement, aggregate uncertainty, and common uncertainty. In contrast to this paper, which focuses on the euro area, they use panel data for 25 economies and provide evidence that greater transparency reduces uncertainty

[4]Grothe and Meyler (2015) show that both market-based and survey-based measures have a non-negligible predictive power for inflation developments, as compared with statistical benchmark models.

of interest rates and inflation, primarily by reducing common uncertainty rather than disagreement. Rather than estimating a measure for common uncertainty, in this paper I estimate the ex ante average individual uncertainty. I find that, of the three measures, the reduction in disagreement is the most prominent response in terms of magnitude.

More recently, a new strand of literature has emerged focusing on the relevance of central bank communication in non-conventional times and its implications for uncertainty. Coenen et al. (2017) find evidence that announcements of asset purchase programs have lowered market uncertainty (measured by the VSTOXX index), particularly when accompanied by a contextual release of implementation details of the program. Ehrmann et al. (2019) find that while forward guidance directly decreases forecast disagreement, the way that it is implemented matters for uncertainty. In particular, the implementation of weak types of forward guidance makes market prices less informative and may increase uncertainty.

Other related studies investigate the effect of central bank communication on other types of uncertainty. Swanson (2006) finds that increased transparency by the U.S. Federal Reserve reduces ex ante uncertainty about the future course of short-term interest rates. Hüning (2017) shows that Swiss National Bank communications indicating a future rate cut reduce stock market uncertainty, measured as the abnormal stock market variance derived from the Swiss Market Index. In contrast, communication indicating future policy tightening does not affect it.

The main novelty of this paper is that, in addition to gaining new insights into the implications of the ECB communication on ex ante inflation uncertainty, it sheds some light on understanding the channels through which central bank information shocks operate. So far, to the best of my knowledge, the mechanism through which central bank communication affects ex ante inflation uncertainty has not yet been explored.

There is, however, a similar debate in the literature about how central bank information shocks affect market expectations and the economy. There are two hypotheses when it comes to addressing this point, but no concrete answer has so far been provided on which of them is more plausible. The first hypothesis is based on a Bayesian approach, in which central bank information shocks could contain new information about how the central bank interprets the state of the economy and/or predicts future economic developments. Once this new information is communicated, financial market participants and forecasters would use this information to update their expectations as long as the central bank analysis is credible.

There are several explanations for the central bank's information advantage in the literature. Romer and Romer (2000) argue that the Federal Reserve has an advantage compared with the market in terms of resources and chooses to use more of these inputs than any commercial forecasters find profitable. Therefore, the private sector considers the information provided by the central bank to be valuable, since the forecasts and analyses are conducted by well-trained staff with a high degree of specialization.

Another explanation is that because most central banks function both as supervisors and as liquidity providers, central banks have tighter links with the financial sector in particular after the crisis. This provides a comparative advantage in collecting detailed information about current and recent developments in the economy. Furthermore, the central bank has the knowledge advantage of its own probable policy actions, so it plays some role in determining the variables it is forecasting (see Jung and Uhlig 2019). Nakamura and Steinsson (2018) and Jarociński and Karadi (2020) suggest that the central bank also simply announces information earlier than other sources. This interpretation implies that if the central bank would not have communicated some specific information, this content would have become known to the market via other sources anyway at a later stage. Nevertheless, their interpretation ultimately suggests that central bank information shocks convey new information and the market learns from it.

The second hypothesis is that central bank information might also contain little or no new information about the current or future state of the economy in terms of hard data. But in a world of possible multiple equilibria, the released information could help market participants and forecasters to focus on one particular equilibrium, supported by the central bank, and therefore serve as an impactful coordination device. This hypothesis thus implies that the public nature of certain signals (in the case of this paper, the communication itself) acts as a signal that can guide expectations and individual decisions even if they contain minimal information, as in Morris and Shin (2002). From this perspective, public signals serve as a coordination device.

Interestingly, Born, Ehrmann, and Fratzscher (2011), when investigating how central bank communication about financial stability influences financial markets, find that it works primarily as a coordination device, highlighting that markets also perceive it to contain relevant information.

While the assessment of whether central bank information shocks convey new information about the economy is beyond the scope of this paper, I provide evidence that they do act as a public signal, able to coordinate and influence opinions and actions. I thereby explore how central bank information operates on the second moments, focusing on the role of central bank communication as a coordination device. In addition, I also document that central bank information shocks do not significantly affect inflation expectations, but they do decrease all three measures of ex ante inflation uncertainty. More precisely, these shocks help to align opinions across forecasters, generating a “stabilizer effect,” as the convergence of these measures is toward their mean.

3. Data Description

The research question of this paper centers on four main variables of interest: the three types of ex ante inflation uncertainty and the central bank communication shocks. Subsections 3.1 and 3.2, respectively, provide detailed explanations of how these measures and shocks are estimated.

To estimate the three measures of ex ante inflation uncertainty, both the aggregate and the individual histograms of the ECB SPF are exploited. The ECB SPF gathers information on the expected rates of inflation, real gross domestic product (GDP) growth, and unemployment in the euro area at different horizons. These expectations are reported both as point forecasts and as probability distributions. The ECB SPF provides both the aggregate histogram containing the median of the responses of the forecasters and the individual histograms containing the anonymized distribution of projections provided by each forecaster. In order to measure central bank communication, I use the central bank information shocks from Jarociński and Karadi (2020) as a proxy.

As the SPF is conducted on a calendar quarter basis, the central bank information shocks—which are on a daily basis—are added together to make a quarterly frequency (see, for example, Kerssenfischer 2019; Jarociński and Karadi 2020). Adding the information shocks is preferable to other methods of aggregation (such as the average) because information accumulates over time and the sum makes sure that there are no losses in terms of content. Given the nature of a shock, which is exogenous and does not anticipate the dependent variable, I assume that ex ante inflation uncertainty in t is affected by all shocks that occurred since the previous survey in t–1. Therefore, these shocks are aggregated on a quarterly basis, always respecting the deadlines to reply to the SPF. As shown in detail in Section 5, this approach assures that all publicly available central bank information is known by the forecasters by the deadline to reply to the survey, that is, when the uncertainty measures are estimated.

The analysis covers the period between 2002:Q1 and 2019:Q1.[5] The structure of the SPF database allows a clear distinction between the specific horizons over which uncertainty is measured, since the participants are asked to provide their inflation forecasts for one-, two-, and five-year horizons. This paper focuses on forecasts for two years ahead, which is the relevant horizon for monetary policy. In other words, the benchmark analysis evaluates how central bank information shocks affect the current uncertainty of the forecasters about inflation on a two-year horizon.

The remaining variables employed in the analysis reflect the control variables identified in the literature as potential influencing factors on forecast uncertainty and disagreement. They are the quarterly change in crude oil prices, inflation (year-over-year Harmonised Index of Consumer Prices, HICP), real GDP, the unemployment rate, the output gap, and the term spread defined as the difference between the euro-area 10-year government benchmark bond yield

[5]The earlier part of the sample dating back to 1999:Q3 is characterized by a relatively low market liquidity, which affects the reliability of the surprises. This is reflected in the very small and negative correlation between the series of daily shocks aggregated to a quarterly frequency using the SPF deadlines and the monthly shocks aggregated to quarterly frequency not using the SPF deadlines. The correlation becomes high and positive only from 2002:Q1 onwards.

and the euro interbank offered rate (EURIBOR) three-month money market rate. Table 1 shows the data used in the analysis, including definitions and sources.

3.1 Estimating Ex Ante Inflation Uncertainties

This section shows how I estimate the three ex ante uncertainty inflation measures. These measures closely relate to each other, as the ex ante aggregate inflation uncertainty (EAU in the equations, and from now on referred to as "aggregate" in the text) incorporates both individual uncertainty and disagreement (see, for example, Wallis 2005). Nonetheless, they carry different meanings and are all estimated separately. Table 2 presents the key statistics for each of the three measures.

The forecasts are reported in the SPF not only as point forecasts but also as probability distributions. In other words, for each horizon, the forecasters should provide the estimation of the HICP inflation as a single number and assign probabilities for different predefined ranges of possible outcomes for the HICP inflation. I exploit both features to construct the ex ante inflation uncertainty measures.

Aggregate is the proxy for the overall ex ante inflation uncertainty. It is the resulting variance after fitting a generalized beta distribution to the aggregate SPF histogram, as in Engelberg, Manski, and Williams (2009) and Melo Fernandes and Kenny (2024). The other two measures are more specific proxies for ex ante inflation uncertainty. Disagreement d_{t+h} is defined as the variance of the point forecasts of a variable y performed in t for a specific horizon h. In other words, disagreement is the dispersion of the point forecasts, indicating how much the individuals diverge among each other regarding the future values of inflation, as shown in Equation (1):

$$d_{t+h} = \mathrm{N}^{-1} \sum_{\mathrm{i}=1}^{\mathrm{N}} \left[\mathrm{E}_{\mathrm{i},t}\left[\mathrm{y}_{t+\mathrm{h}}\right] - \overline{\mathrm{y}}_{t+\mathrm{h}}\right]^2, \tag{1}$$

where $\mathrm{E}_{\mathrm{i},t}$ is the expectation of the forecaster i in time t with respect to the variable y for a specific horizon h and $\overline{\mathrm{y}}_{t+\mathrm{h}}$ is the average forecast of variable y in time t for a specific horizon h.

Table 1. Data Information

Variable	Units	Definitions	Data Sources
Aggregate Uncertainty	Index	Own calculations	ECB SPF[1]
Average Individual Uncertainty	Index	Own calculations	ECB SPF
Disagreement	Index	Variance of forecasts	ECB SPF
Central Bank Information Shocks (Benchmark)	Index	Positive co-movement between EuroStoxx 50 and the first principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years	Jarociński and Karadi (2020)
Inflation Expectations	Percent per annum	Average of point forecasts	ECB SPF
GDP Expectations	Percent per annum	Average of point forecasts	ECB SPF
Unemployment Expectations	Percent per annum	Average of point forecasts	ECB SPF
Real GDP	Percentage change	Gross domestic product at market prices—annual rate of change	Eurostat
Output Gap	Percent	Deviations of actual output from potential output	Estimated based on Hamilton (2018)
Unemployment Rate	Percent	Standardized unemployment, total, percentage of labor force	Eurostat

(*continued*)

Table 1. (Continued)

Variable	Units	Definitions	Data Sources
Crude Oil Prices	Percent per annum	Bloomberg European Dated Brent Forties Oseberg Ekofisk (BFOE) crude oil spot price	ECB SDW[2]
Consumer Prices	Percentage change	Harmonised Index of Consumer Prices—annual rate of change	Eurostat
Term Spread	Percent per annum	Own calculations—spread between the euro-area 10-year government benchmark bond yield and the three-month EURIBOR rate	ECB SDW
Three-Month EURIBOR Rate	Percent per annum	Euro interbank offered rate—historical close, average of observations through period	ECB SDW
10-year Government Benchmark Bond Yield	Percent per annum	Benchmark bond—yield	ECB SDW

[1]Survey of Professional Forecasters: http://www.ecb.europa.eu/stats/prices/indic/forecast/html/index.en.html.
[2]ECB Statistical Data Warehouse: https://sdw.ecb.europa.eu/.

Table 2. Summary Statistics: Ex Ante Inflation Uncertainty Measures

Horizon	Measure	Mean	St. Deviation	Skewness	Kurtosis
One-Year Horizon	Aggregate	0.30	0.11	0.19	2.01
	AIU	0.27	0.07	–0.18	1.64
	Disagreement	0.09	0.06	1.79	6.55
Two-Year Horizon	Aggregate	0.34	0.10	–0.18	1.61
	AIU	0.32	0.08	–0.29	1.61
	Disagreement	0.06	0.04	1.87	7.64
Five-Year Horizon	Aggregate	0.36	0.09	–0.37	1.7
	AIU	0.38	0.08	–0.19	1.82
	Disagreement	0.05	0.04	5.05	35.7

The average individual uncertainty (AIU) is the average of the individual variances, which can be interpreted as how assured individuals are with respect to their own forecasts:

$$\overline{\sigma}_{t+h} = \mathrm{N}^{-1} \sum_{\mathrm{i}=1}^{\mathrm{N}} \mathrm{E}_{\mathrm{i,t}} \left[\left(\mathrm{y}_{\mathrm{t+h}} - \mathrm{E}_{\mathrm{i,t}} \left[\mathrm{y}_{\mathrm{t+h}} \right] \right)^2 \right]. \tag{2}$$

Finally, EAU incorporates both individual uncertainty and disagreement as shown below:

$$\mathrm{EAU}_{t+h} = \overline{\sigma}_{t+h} + d_{t+h}. \tag{3}$$

Looking at Equation (3), one could calculate AIU as simply the residual between EAU and d, as in Abel et al. (2016). However, conditioning the estimation of AIU to disagreement is not ideal. First, the literature documents that disagreement may on its own be a relatively poor proxy for uncertainty as compared with AIU (see further discussion in Section 6.1). Therefore, estimating AIU as the residual of Equation (3) might lead to a less accurate measure of AIU compared with using the individual data independently of disagreement. Indeed, as shown in Figure 1, when AIU is calculated as the residual after plugging in aggregate and disagreement in Equation (3), it reflects, for example, a point forecast outlier in

Figure 1. Ex Ante Inflation Uncertainties—AIU as Residual (Two-Year Horizons)

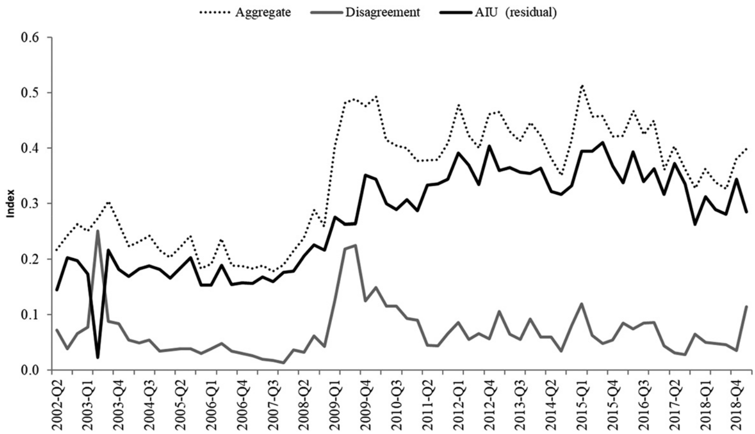

Note: Ex ante average individual uncertainty is estimated as the residual between aggregate and disagreement.

2003:Q2.[6] When subtracting disagreement from the aggregate, this outlier is reflected in both a temporary fall in AIU and a peak in disagreement, which does not make economic sense. Likewise, in situations where disagreement increases more than EAU, the residual AIU falls, which also leads to a misleading measurement of average individual uncertainty.

Therefore, instead of employing Equation (3), I compute AIU by first estimating the respective variances using a similar approach to the estimation of aggregate uncertainty. I follow Engelberg, Manski, and Williams (2009) in estimating the measure in three steps. First, I fit distributions in each individual histogram provided by each forecaster. These distributions are determined according to the intervals at which the forecasters place their probabilities. In the second step, I extract the variance of each histogram after fitting these

[6] In that quarter, the average of the forecast for the year-over-year change in inflation for a two-year horizon was 1.7 percent, while one specific forecaster reported a projection of −1 percent. Note that this outlier in disagreement was removed before performing the regressions.

distributions. In the third stage, I take the average of these resulting variances.

When estimating the variances, two different distributions are fitted. When the probabilities are placed in three or more histogram intervals, the assumption is that each subjective distribution has the generalized beta form. Just as in the case of the aggregate histogram, I estimate the variance by using the interval probability data to fit the parameters.

In contrast, when a forecaster places probabilities in only one or two intervals, the assumption is that the distribution has the shape of an isosceles triangle. The placement of probabilities in fewer bins can be interpreted as if these forecasters have relatively more conviction about the outcome of the future inflation than those that place their probabilities in more bins. This happens in approximately only 3 percent of the total cases in the database.[7] Furthermore, 88 percent of these cases occur before the Great Financial Crisis.

Finally, in cases where the forecaster is 100 percent convinced that the outcome of inflation will be within a particular range, the base of the triangle includes the interval correspondent to this range and part of the adjacent interval. In cases where the forecaster places the probabilities in two intervals, they are always adjacent to one another and the base of the triangle includes the entire interval with the greater probability mass and part of the neighboring interval. This assumption gives one parameter to be fit, which fixes the center and height of the triangle.

Despite providing similar outcomes to the residual estimation method, the direct AIU estimation method results in a slightly higher level of AIU and does not reflect potentially noisy observations coming from other estimation sources. Therefore, extracting AIU directly from the histograms leads to a more accurate and cleaner measure of AIU (see Figure 2).

Table 3 shows that the different nature of individual uncertainty and disagreement are reflected in the low correlation between these measures (0.28, 0.37, and 0.09 for the one-, two-, and five-year horizons, respectively). In contrast, aggregate uncertainty has a very high correlation with AIU (0.93 for the two-year horizon) and a

[7]Estimates are calculated based on the sample composed by forecasts for one- and two-year horizons.

Figure 2. Ex Ante Inflation Uncertainties—AIU Estimated (Two-Year Horizons)

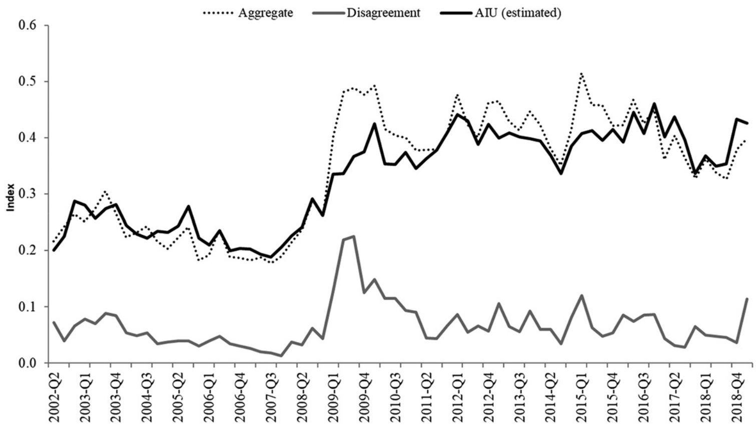

lower correlation with disagreement (0.61 for the two-year horizon). Indeed, Figure 2 shows that unlike disagreement, both AIU and aggregate uncertainty show a clearer level shift and much higher persistence in the period since the Great Financial Crisis. Such differences highlight the importance of variation in uncertainty at the individual level as a key driver of aggregate ex ante uncertainty. In addition, for all ex ante uncertainty variables, one can observe that the longer the time horizon, the lower the correlation between all measures. That might reflect the fact that given the relatively high degree of persistence in inflation, shorter horizons are more influenced by data realizations on which forecasters agree, while the impact of present developments fades away over longer-term projections.

3.2 Central Bank Information Shocks

Central bank announcements simultaneously convey information about monetary policy and the central bank's assessment of the economic outlook. Jarociński and Karadi (2020) distinguish between these two types of information quantitatively and provide a measure of ECB communication by identifying high-frequency co-movement

Table 3. Correlation: Ex Ante Inflation Uncertainty Measures

		Disagreement			**AIU**			**Aggregate**		
	Horizons	**One Year**	**Two Years**	**Five Years**	**One Year**	**Two Years**	**Five Years**	**One Year**	**Two Years**	**Five Years**
Disagreement	One Year	1								
	Two Years	0.72	1							
	Five Years	0.36	0.56	1						
AIU	One Year	0.28	0.36	0.18	1					
	Two Years	0.30	0.37	0.18	0.98	1				
	Five Years	0.07	0.21	0.09	0.90	0.93	1			
Aggregate	One Year	0.72	0.65	0.31	0.85	0.84	0.65	1		
	Two Years	0.50	0.61	0.33	0.92	0.93	0.81	0.92	1	
	Five Years	0.14	0.31	0.25	0.91	0.93	0.96	0.70	0.87	1

of interest rates and stock prices in a narrow window around ECB policy announcements.

The reasoning behind it is that when interest rates go up, stock prices are expected to go down for two reasons: first, after a policy tightening, investors foresee a relative slowdown in the economy, which discourages the appetite for investments, and second, the discount rate increases with higher real interest rates and rising risk premia (the denominator effect). However, if instead stock prices increase following an increase in interest rates, the authors attribute this unexpected move to the impact of information shocks containing positive economic news. Therefore, central bank information shocks are identified when interest rates and stock prices co-move positively. As the scope of the shocks is limited to communication about economic outlook assessments only, one can exclude any type of direct effect involving forward guidance.[8]

In order to capture these co-movements, Jarociński and Karadi (2020) construct a data set of euro-area high-frequency financial market surprises,[9] which are defined as financial asset price changes around the ECB announcements. These announcements are delimited within windows of 30 minutes around press statements and 90 minutes around press conferences, both starting 10 minutes before and ending 10 minutes after the event. The assumption is that within this narrow window only two structural shocks can materialize and systematically influence the financial market surprises: a monetary policy shock, which is defined as the negative co-movement between interest rate and stock price changes, and a central bank information shock, defined as the positive co-movement of interest rates and stock prices. In the euro area, this is the case for approximately 46 percent of the data points. The data set contains more than 300 ECB policy announcements from 1999 to 2019.

In this paper, I use the shocks estimated by Jarociński and Karadi (2020) using the "poor man's" sign restrictions method. In

[8]The information shocks by Jarociński and Karadi (2020) carry information about the economy, not about future monetary policy.

[9]This novel data set for the euro area is based on Gürkaynak, Sack, and Swanson (2005), who constructed a similar data set for the United States.

a nutshell, the poor man's sign restrictions use the interest rate surprises in the days in which announcements resulted in stock price surprises with the same sign as the interest rate change as the proxy for central bank information shocks. Otherwise, the proxy is zero.

The measure used to compute changes in stock valuation is the EuroStoxx 50 index. The proxy for interest rates is a combination of different maturities of euro overnight index average (EONIA) swaps. In particular, the measure used as a benchmark in this paper is the first principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years. The reason to choose this proxy as a benchmark rather than one single and shorter maturity is that by including longer maturities one can capture higher volatilities that might occur in the zero lower bound period. Typically, in this period the value of assets with longer maturities changes more than those with shorter maturities.

Kerssenfischer (2019) follows the same standard sign restrictions approach of Jarociński and Karadi (2020) and builds central bank information shocks using two-year German bond yields as a proxy for interest rates and the EuroStoxx 50 index as a proxy for stock valuations.[10] Furthermore, he replaces the narrow window with a wider window around the ECB's press release that also includes the market reaction to the press conference. Table 4 shows all the communication shocks that were employed as robustness checks in Section 7. As explained in Section 5, all shocks were aggregated to quarterly frequency using the dates of the ECB survey deadlines in order to obtain an accurate identification. Figure 3 shows the final aggregation.

4. The Empirical Model

The primary objective of the analysis is to estimate the impact of central bank information shocks on ex ante inflation uncertainty. I use local projections (see Jordà 2005) to estimate the

[10]The study encompasses 186 scheduled ECB Governing Council meetings between March 2002 and December 2018.

Table 4. Central Bank Information Shocks Information

Shock	Methodology	Definition of Interest Rate	Definition of Stock Prices	Data Sources
Benchmark Shock	Poor Man's Sign Restrictions	First principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years	EuroStoxx 50 Index	Jarociński and Karadi (2020)
Robustness 1	Sign Restrictions	First principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years	EuroStoxx 50 Index	Jarociński and Karadi (2020)
Robustness 2	Poor Man's Sign Restrictions	First principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years	EuroStoxx 50 Index	Jarociński and Karadi (2020)
Robustness 3	Sign Restrictions	First principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years	EuroStoxx 50 Index	Jarociński and Karadi (2020)
Robustness 4	Poor Man's Sign Restrictions	Three-month EONIA swaps	EuroStoxx 50 Index	Jarociński and Karadi (2020)
Robustness 5	Sign Restrictions	Three-month EONIA swaps	EuroStoxx 50 Index	Jarociński and Karadi (2020)
Robustness 6	Sign Restrictions	Two-year German bond yields	EuroStoxx 50 Index	Kerssenfischer (2019)

Figure 3. Central Bank Information Shocks (Baseline)

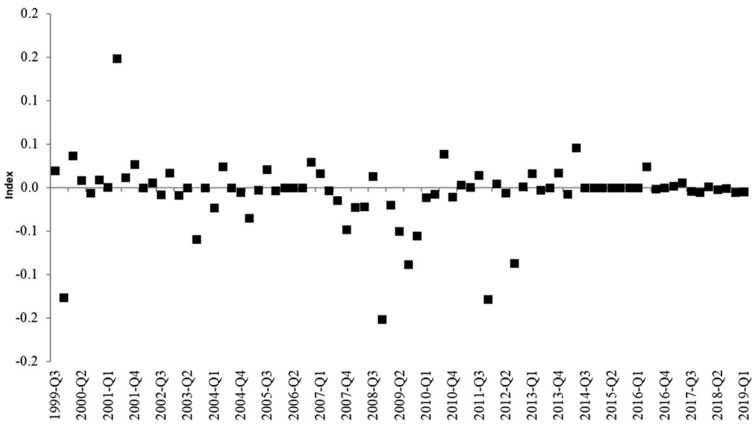

Note: Central bank information shocks estimated by Jarociński and Karadi (2020) using the "poor man's method." The measure used to compute changes in stock valuation is the EuroStoxx 50 index and the proxy for interest rates is the first principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years. The daily shocks were aggregated to a quarterly frequency by summing the shocks in between the deadlines to reply to the SPF.

impulse responses. Local projections consist of the estimation of a series of regressions for each variable in each horizon h. Therefore, the linear regression of the benchmark model is designated as follows:

$$\triangle x_{t+h} = \beta_{0,h} + \beta_{1,h} shock_t + \beta_{n,h}(L) y_{n,t-1} + \varepsilon_{t+h},$$
$$\text{for} \quad h = 0, 1, 2, \ldots \tag{4}$$

where $\triangle x_{t+h}$ is defined as $x_{t+h} - x_{t-1}$, where x_{t+h} and x_{t-1} are in logs. The changes for each ex ante uncertainty type are shown in Figures 4, 5, and 6. $\beta_{0,h}$ is a constant, $\beta_h(L)$ is a polynomial in the lag operator, *shock* is the identified shock, and y is the vector of control variables. The coefficient $\beta_{1,h}$ gives the response of the changes in x at time $t + h$ with respect to $t - 1$ to the shock that happens at time t. This calculation ensures that the direct impact of the shock is isolated, enabling the analysis to focus on the net change that has occurred over that time span.

Figure 4. Ex Ante Average Individual Uncertainty: Changes by Impulse Response Function Horizons

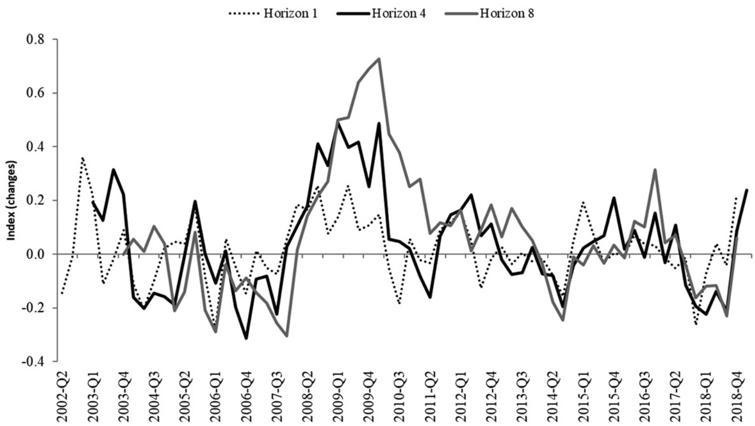

Figure 5. Ex Ante Aggregate Uncertainty: Changes by Impulse Response Function Horizons

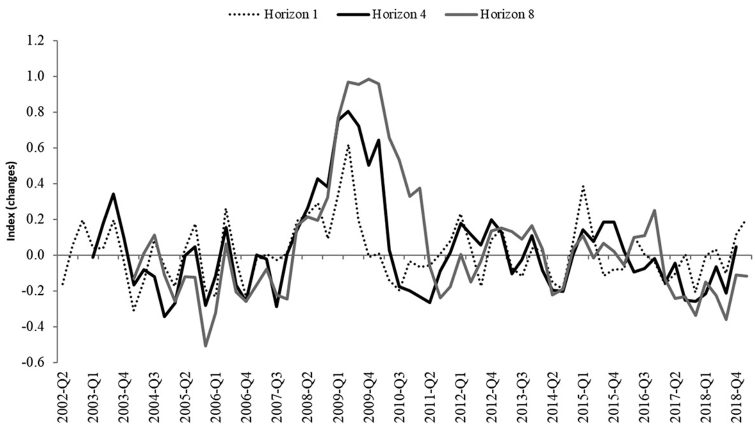

The baseline shock is estimated using the poor man's sign restrictions method, which ultimately calculates the co-movement between the EuroStoxx 50 index and the first principal component of the EONIA swaps with maturities of one month, three months, six months, one year, and two years (see Section 3.3). In essence, they

Figure 6. Disagreement: Changes by Impulse Response Function Horizons

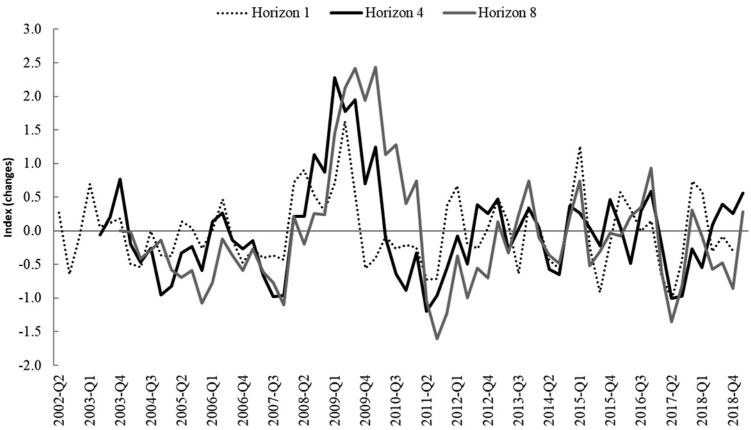

consist of market reactions to unanticipated communications about the state of the economy and are unrelated to other factors likely to influence ex ante inflation uncertainty in the near term.

The first specification relies on the exogenous nature of these shocks, which leads to a simple regression in which each ex ante inflation uncertainty measure is regressed on a constant, on the shock, and on the lagged ex ante inflation uncertainty.

From this starting point, the model is progressively augmented to include different sets of controls in vector y as well as a variety of lags for robustness check purposes. The control variables and the other specifications are further detailed in Section 6.

In all cases, the coefficients of interest are the sequence $\beta_{1,h}$, which gives the response of x at time $t+h$ to the shock that happened at time t. Hence, the results are presented as impulse responses built on this sequence of $\beta_{1,h}$ estimated by single regressions for each horizon. As central bank communication on economic outlooks is often focused on a short-term period, the horizon of the estimated effects is limited to eight quarters. Furthermore, given the limited number of observations in the sample due to the relatively short time series (70 quarters in total), I opt for a more parsimonious approach, as the higher the number of horizons, the shorter the sample of observations available for estimations in the later horizons.

Table 5. Representation of the Timing for the Aggregation of Shocks

Quarters	Months	Deadline to Reply to SPF (Day)	Date in Which Shocks Were Recorded (Day of the Month)
Q2	April		4
	April	19	—
	May		**2**
	June		**6**
Q3	July		**4**
	July	19	—

Note: This table illustrates a case in which the deadline to reply to the SPF in 2013:Q3 was on July 19, 2013. Therefore, only shocks that happened between the deadline to reply to the SPF in Q2 (i.e., April 19) and July 19 were summed. The dates of the days in which shocks were aggregated to build the shocks for Q3 are highlighted in bold.

5. Identification Strategy

An important aspect of the identification is that surveyed probabilities used in the estimation of ex ante inflation uncertainty are on average collected in the middle of the first month of quarter t. Therefore, it is important to make sure that all the information available is known by the forecasters by the deadline to reply to the survey.

The alignment between the timing of the survey deadlines and the timing of the information shocks is made possible by combining the daily data set of the shocks and the quarterly deadlines to reply to the SPF. This alignment is achieved by summing the shocks that occurred between the deadline to reply to the SPF in the quarter $t-1$ and the deadline to reply to the next survey round in quarter t, thereby ensuring that all shocks that happened within this period have been observed by the forecasters and potentially included in their projections, and are consequently reflected in their replies to the survey in quarter t. In summary, I regress this aggregated sum on the ex ante inflation uncertainty estimated from the survey of quarter t.

Table 5 shows an example of the timing framework used to aggregate the shocks in 2013:Q3. In this case, the deadline to reply to the

SPF in Q3 was on July 19. Therefore, only shocks that happened between the deadline to reply in Q2 (i.e., April 19) and July 19 were summed. The corresponding days are in bold.

If instead one opted to add the shocks by calendar quarter, ignoring the SPF deadlines, two issues would arise: first, one would miss some information that was released in the following quarter just before the SPF deadline (in this example, the shock on July 4), and second, one's models would incorporate information that had already been absorbed in the former survey (in this case, the shock on April 4).

Another relevant point to consider in the identification strategy is the timing of the control variables. Following the same logic described above, I also define the timing of the real variables in the regressions using the SPF deadlines as a reference. I use the Eurostat calendar to extract the latest information of real variables that was available for the forecasters. For example, for inflation I use the latest value released before each survey. The same applies for the change in crude oil prices and unemployment. These variables, which are available at a monthly frequency, are therefore included in $t-1$ when the survey deadline was in t. The most recent release of real GDP information prior to the SPF deadline always contains the real GDP value for the two previous quarters. Therefore, real GDP is included in the timing $t-2$.

Finally, the approach used to calculate the changes in the dependent variable as shown in Equation (4)—in which the changes on the measures of uncertainty in time to $t+h$ are always with respect to $t-1$ in response to a shock that happens at time t—ensures consistency between the timing of the survey data collection and the aggregation of the shocks. This rationale is elucidated through the illustrative example provided in Table 5. The survey deadline for 2013:Q3 (time t) was on July 19, meaning that the estimation for uncertainty in 2013:Q3 was produced in the beginning of that quarter. Meanwhile, the shocks for 2013:Q3 were aggregated using central bank information that started to be collected right after the survey deadline for 2013:Q2 (time $t-1$), with most of the information being collected still in this quarter. In light of this dynamic, in order to measure the effect of the shock in 2013:Q3 (time t) on the subsequent horizons, it becomes crucial to utilize the uncertainty data preceding the shock as the foundational reference for computing changes. This entails employing 2013:Q2 (time $t-1$) as the benchmark for such calculations.

6. Results

Figure 7 summarizes the results of estimating the benchmark specification of Equation (4), which includes a constant and the central bank information shock on the right side of the equation. It is also important to control for the normal dynamics of ex ante inflation uncertainty and for several other factors that are likely to be serially correlated and may affect the dependent variable. Hence, the benchmark model also includes the lagged ex ante inflation uncertainty as a control variable. I adopt the results from this specification as the baseline. Other specifications including different lags and controls are explored in Section 7. In Figure 7, each column shows the cumulative responses for each ex ante inflation uncertainty measure to a central bank information shock. The estimations rely on 90 percent confidence bands and are based on Newey-West standard errors to account for serial correlation.

After a central bank information shock, all three types of ex ante inflation uncertainty fall significantly after one quarter. Two interesting observations can be made based on this result: first, these findings suggest that central bank communication decreases both the average individual uncertainty and the divergence of opinions among the forecasters. Second, the reaction of ex ante inflation uncertainties systematically happens with a delay. This delay is in line with the findings of Coibion and Gorodnichenko (2012), who document evidence of a delayed response of mean forecasts to macroeconomic shocks for professional forecasters in the United States, reflecting information rigidities.

The impact of the central bank information shocks is most prominent on disagreement, which decreases 5.5 percentage points in the first quarter—approximately more than five times the drop of average individual uncertainty. While average individual and aggregate uncertainty retract from their peak in the third quarter, disagreement falls nearly half a percentage point further. Aggregate ex ante inflation uncertainty decreases 1.7 percentage points in the first quarter, with some persistence in the last horizons. Clearly, the

Figure 7. Impulse Response Functions: Baseline Specification

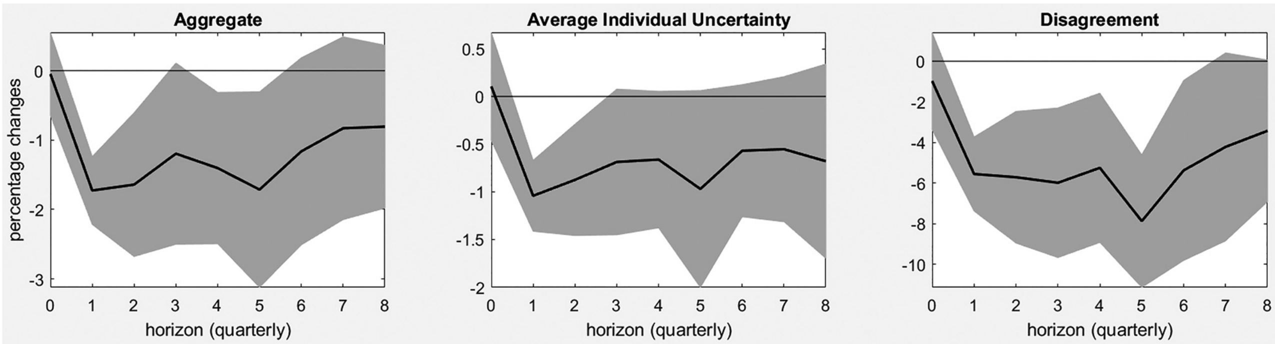

Note: This specification includes constant and the lag of ex ante inflation uncertainty.

results for the aggregate uncertainty are driven by the stronger magnitude and persistence of the reaction of disagreement.

When interpreting these results, the first conclusion is that after analyzing the same new public information provided by a credible central bank, agents become more aligned in their views even as they also become more certain about their own predictions.[11] However, in addition to that, the nature of each ex ante uncertainty measure can provide interesting insights into the mechanism behind the impact of the central bank information shocks on ex ante uncertainty.

6.1 The Role of Disagreement in Understanding How Central Bank Information Shocks Operate

As shown in Section 3.1, disagreement reflects the dispersion of projections across forecasters but does not provide information about each forecaster's uncertainty regarding their own forecast. For example, it could be that each forecaster is extremely uncertain about future events; however, they could still have very similar point estimates. In this case, disagreement fails to accurately capture the actual level of inflation uncertainty.

In fact, although used as a common approach to estimate ex ante uncertainty in the literature, disagreement survey-based measures have been criticized as a relatively poor proxy for uncertainty.[12] In particular, some studies show that disagreement in survey forecasts could be more reflective of differences in opinion than of uncertainty (see Diether, Malloy, and Scherbina 2002; Mankiw, Reis, and Wolfers

[11] This is in contrast to the findings of Johnstone (2016), who shows that the best available information can often leave decisionmakers less certain about future events.

[12] For discussion, see Zarnowitz and Lambros (1987), Grier and Perry (1998, 2000), Giordani and Söderlind (2006), Lahiri and Sheng (2010), Abel et al. (2016), Glas and Hartmann (2016), and Clements, Rich, and Tracy (2023). These studies highlight the absence of a theoretical foundation to link disagreement with uncertainty and document empirical deviations between disagreement and ex ante average individual uncertainty. Lahiri and Sheng (2010) establish a simple relationship connecting forecast uncertainty to disagreement and show that disagreement is found to be a reliable measure for uncertainty in a stable period, but not in periods with a large volatility of aggregate shocks.

2004). Despite being often seen as a criticism, this feature is particularly useful for understanding how central bank information shocks operate in reducing ex ante uncertainty.

Specifically, the substantial fall in disagreement in response to central bank information shocks implies that these shocks are able to influence forecasters' opinions. In particular, the shocks help opinions to converge. However, it is also important to understand whether these opinions converge in a direction that contributes to market stabilization—i.e., whether these opinions converge to the mean, leading to a "stabilizer effect"—or whether this convergence goes toward a point that may cause instability. For example, if after a central bank communication shock the opinions converged toward one of the tails of the distribution of inflation expectations rather than toward the mean—which is aligned with the ECB objective of 2 percent inflation in the medium term—that could be a detrimental outcome given the risk of de-anchorage of inflation expectations. Since central bank communication undoubtedly plays a fundamental role in steering expectations (see Blinder et al. 2008), it is important also to understand the response of forecasters' expectations to these shocks in order to answer this question. Interestingly, the literature addressing the effects of central bank information shocks first shows that central bank information shocks generate an increase in inflation expectations; however, this effect is not significant, as shown by Jarociński and Karadi (2020) for the United States and Kerssenfischer (2019) for the euro area.

Therefore, in order to have a precise interpretation of what the results for ex ante inflation uncertainty mean, it is useful to understand how central bank information shocks affect the changes in the level of inflation expectations. Thus, I repeat the exercise using the baseline specification with inflation expectations being the dependent variable to verify how it reacts to central bank information shocks.[13] Figure 8 shows inflation expectations in levels. Figures 9 and 10 depict the resulting impulse response functions, illustrating the effect of central bank information shocks on inflation expectations. It is noteworthy that this effect lacks statistical significance, aligning with previous findings in the literature.

[13]In contrast to the baseline specification for ex ante inflation uncertainties, no dummies were included for inflation expectations.

Figure 8. Inflation Expectations—Two-Year Horizon

Figure 9. Response of Inflation Expectations to Central Bank Information Shocks

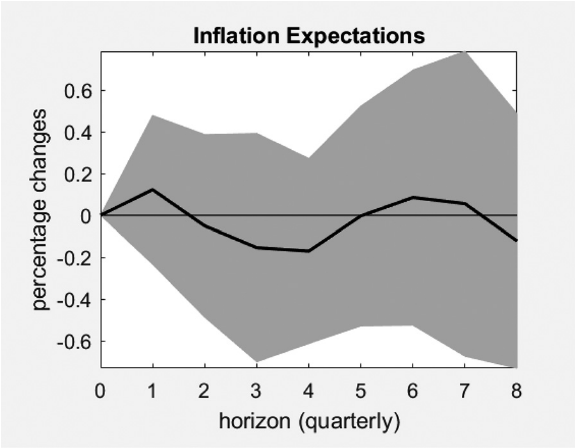

Note: This figure shows the response of inflation expectations for the two-year horizon measured as the average of point forecasts to central bank information shocks. As the figure shows, the effect of central bank information shocks on inflation expectations is not significant.

These findings lead to some interesting reflections. First, the muted responses from inflation expectations and the strong decline of disagreement suggest that after being affected by a central bank

Figure 10. Response of Inflation Expectations to Central Bank Information Shocks (Robustness)

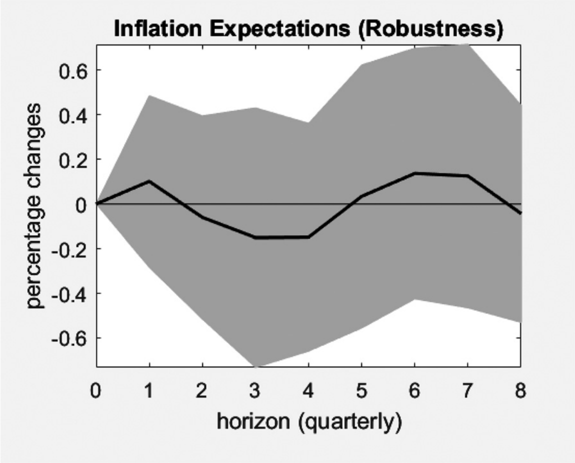

Note: This figure shows the response of inflation expectations for the two-year horizon measured as the average of point forecasts to central bank information shocks. The specification also includes inflation as a control variable for a robustness check. The result also shows that the effect of central bank information shocks on inflation expectations is not significant.

information shock, agents do not necessarily update their expectations, but they converge toward the mean of the point forecasts, which remains close to the ECB objective of 2 percent inflation over the medium term. This convergence implies that the central bank communication has a "stabilizer effect" in which the dispersion among the point forecasts decreases and, most importantly, this convergence moves toward the mean. This convergence is very important, as it induces a steady consensus among the forecasters more in line with the ECB's objective of price stability. In contrast, if the point forecasts responded significantly with a steep increase or decrease to central bank communication shocks, that would indicate that inflation expectations could converge toward one of the tails, which could ultimately lead to the de-anchorage of inflation expectations, undermining the ECB's price stability goals. This result is also consistent with the high credibility of the ECB.

It has been shown by some studies that one important reason why professional forecasters disagree is that they may interpret public information in different ways (see Lahiri and Sheng 2008; Manzan 2011). The decrease in disagreement after a central bank information shock implies that these shocks help to better align how forecasters interpret public information, providing evidence that the content of the shocks in this case is more related to clarifications or reinforcements of previous messages. Another well-known reason why forecasters disagree is that forecasters are presumed to have asymmetric loss functions (see Capistrán and Timmermann 2009).

Therefore, the response of disagreement to central bank information shocks indicates that central bank information shocks operate as some sort of public signal able to influence and coordinate forecasters' opinions. Public signals can often serve as a focal point for the beliefs of market players (Morris and Shin 2002).

6.2 The Role of Average Individual Uncertainty in Understanding How Central Bank Information Shocks Operate

As demonstrated in Section 3.1, average individual uncertainty is the uncertainty of an individual forecaster averaged across all forecasters. In contrast to disagreement, it disregards how forecasters' projections are positioned in comparison with their peers. This measure is often considered a better proxy for uncertainty than disagreement in the literature (Abel et al. 2016; Glas and Hartmann 2016; Glas 2020). The responses of both measures are complementary for understanding how central bank information shocks operate.

The decrease of average individual uncertainty after central bank information shocks means that forecasters became more confident about their own projections. This suggests that forecasters are comfortable with incorporating the public signal emitted by the central bank in the assessment of their analysis, which also implies that this signal is valuable and on average contributes to strengthen the confidence in their predications. This is in line with Morris and Shin's (2002, p. 1521) statement that "when prevailing conventional wisdom or consensus impinge on people's decision-making process, public information may serve to reinforce their impact on individual decisions to the detriment of private information."

Concerning what we can learn from average individual uncertainty with respect to the content of central bank information, there are the following possibilities: it might consist either of clarifications or reinforcements of previous messages and/or of new information that is incorporated by the forecasters, which helps to improve their confidence about their own assessments. As central bank information shocks induce forecasters to sharpen their own beliefs about possible outcomes, one cannot exclude the possibility that these emitted signals also contain relevant information that ultimately increases the forecasters' confidence in their own forecasts. However, the assessment of whether central bank information shocks indeed convey new information about the economy to the agents requires further empirical exercises and is beyond the scope of this paper.

7. Robustness

It is important to account for potential remaining information in the estimated residuals that might influence ex ante inflation uncertainty. Therefore, this section explores the potential sensitivity of the results to other specification choices and to the addition of other controls.

First, I estimate the baseline equation adding different lags of the correspondent dependent variable in levels. Figure 11 shows that the findings for the three types of ex ante inflation uncertainty are robust to different lag specifications and therefore aligned with the reasoning of the baseline results.

Next, by closely following Jitmaneeroj, Lamla, and Wood (2019), I augment the baseline specification with control variables that have been identified in the literature as potential real, nominal, and financial impact factors on forecast uncertainty and disagreement.[14] These variables are the lagged inflation levels year-over-year (HICP), lagged unemployment rate, lagged output gap, and lagged

[14] As listed by Jitmaneeroj, Lamla, and Wood (2019), see Döpke and Fritsche (2006), van der Cruijsen and Demertzis (2007), Patton and Timmermann (2010), Dovern, Fritsche, and Slacalek (2012), Ehrmann, Eijffinger, and Fratzscher (2012), Lamla and Maag (2012), Hartmann and Roestel (2013), Posso and Tawadros (2013), and Siklos (2013).

Figure 11. Robustness Check—Different Lags

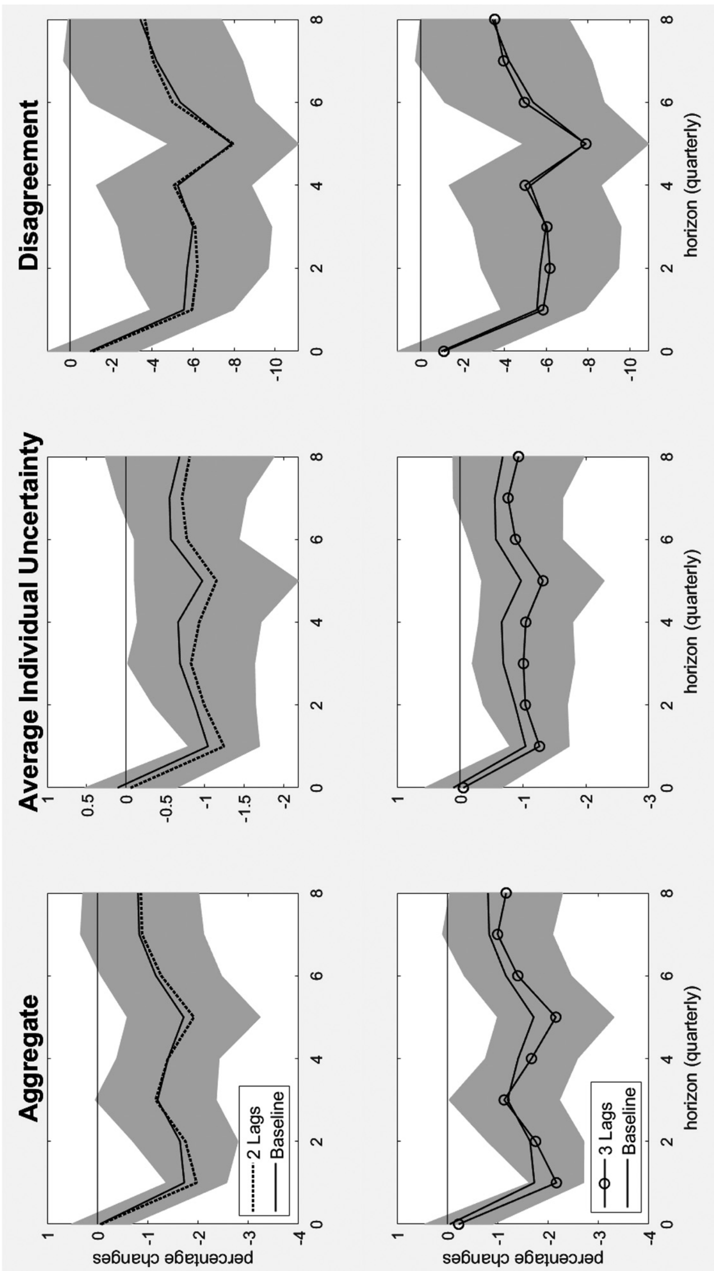

Note: The confidence intervals correspond to the impulse responses of the series employed as a robustness check.

term spread, which is defined as the difference between the euro-area 10-year government benchmark bond yield and the EURIBOR three-month money market rate.

The inclusion of these control variables results in slightly milder responses for aggregate uncertainty and disagreement, while it is marginally more pronounced for average individual uncertainty, though it remains closely aligned with the baseline. As shown in Figure 12, interestingly, disagreement has the same drop as the baseline specification in the first quarter (−5.5 percentage points). The same specification is only slightly modified by replacing inflation with changes in crude oil prices, and the responses remain robust (Figure 13).

In addition, I estimate the baseline specification using other central bank information shocks. Specifically, I compare different versions of the poor man's shocks from Jarociński and Karadi (2020), and central bank information shocks as estimated by Kerssenfischer (2019). As explained in Section 3.3 and shown in Table 4, different versions of the poor man's shocks are estimated by employing EONIA swaps with different maturities. Kerssenfischer (2019) follows the same sign restriction methodology but sticks to the immediate change in two-year German bond yields. The measure used to compute changes in stock valuation is the EuroStoxx 50 index for all cases.

Figure 14 shows the comparisons for the different shocks and maturities. The first row shows the comparison between the responses to the short-maturity version of the benchmark poor man's sign restriction shock to the baseline shock, both estimated using the three-month EONIA swap. The second row shows the responses to another version of these shocks, using the first principal component of the EONIA swaps with maturities of one month, three months, six months, and one year. The third row shows the responses to the shocks estimated by Kerssenfischer (2019). The responses of the three ex ante inflation uncertainty measures are fairly robust to all versions of the shocks.

As a further robustness exercise, it is interesting to see whether these findings hold for different ex ante inflation uncertainty horizons. As shown in Figure 15, the responses of one-year horizon ex ante inflation uncertainties are in general milder but still similar to the baseline, while the response of five-year aggregate is notably

Figure 12. Robustness Check—Different Controls (1)

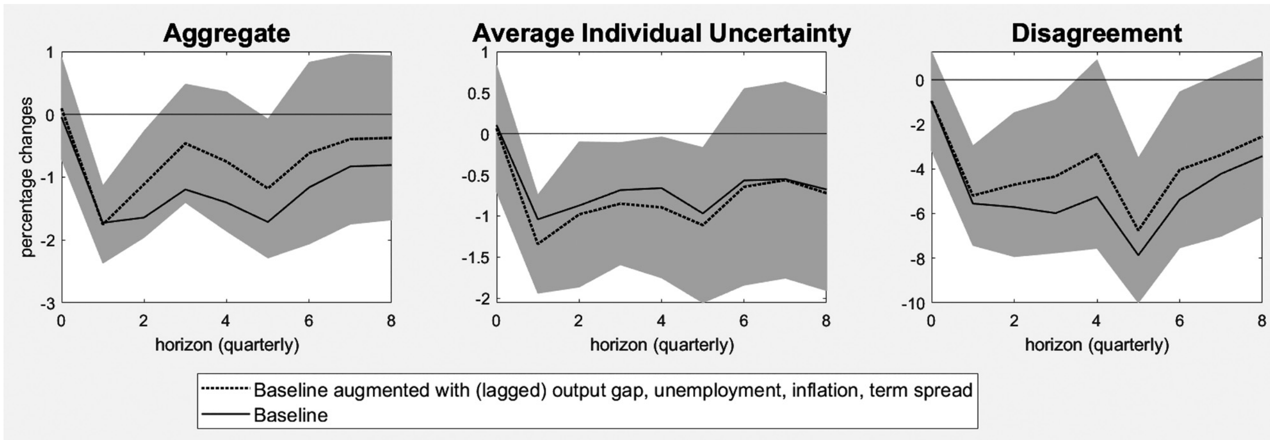

Note: The confidence intervals correspond to the impulse responses of the series employed as a robustness check.

Figure 13. Robustness Check—Different Controls (2)

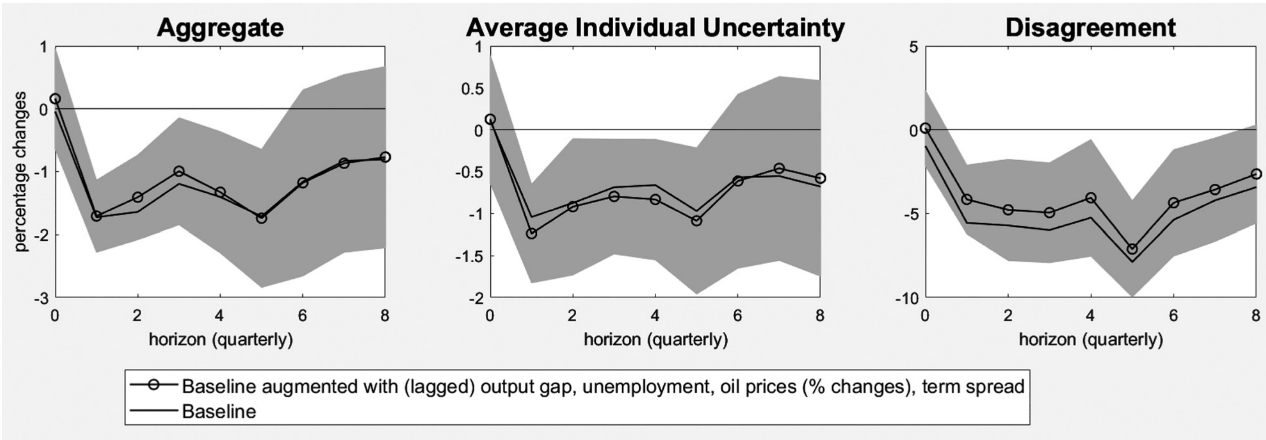

Note: The confidence intervals correspond to the impulse responses of the series employed as a robustness check.

Figure 14. Robustness Check—Different Shocks

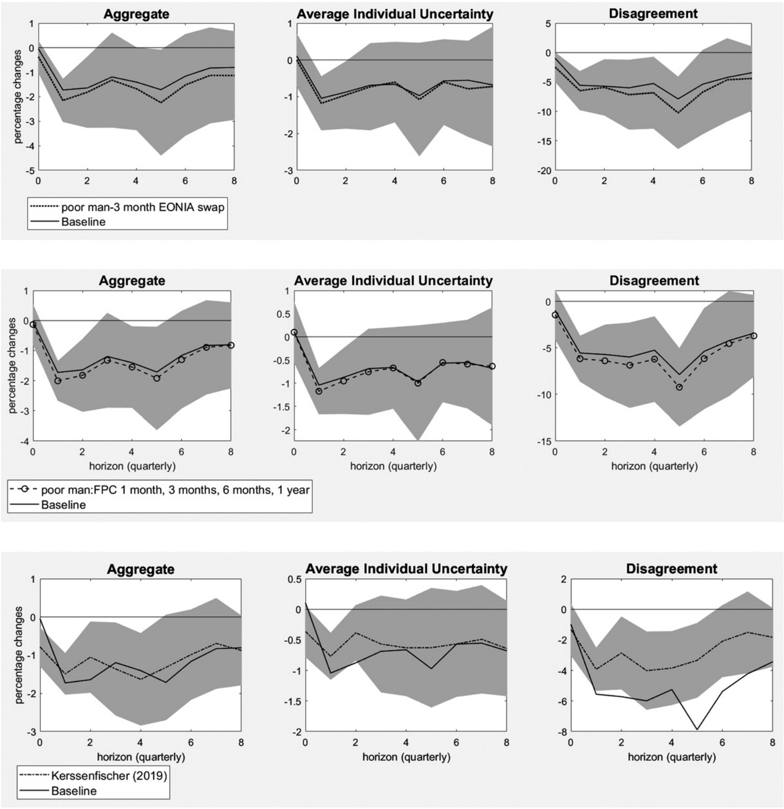

Note: The confidence intervals correspond to the impulse responses of the series employed as a robustness check.

less prominent in the first quarter, becoming larger than the benchmark in the following quarters, a trend also present in the response of average individual uncertainty. In addition, disagreement reacts with a larger delay than the benchmark measure: the first significant reactions appear only after two quarters. These results also provide reassurance regarding the robustness of the benchmark estimation.

In the appendix I show the impact of central bank information shocks on ex ante uncertainty about the other two variables that are also included in the Survey of Professional Forecasters, that is, GDP

Figure 15. Robustness Check—Different Horizons

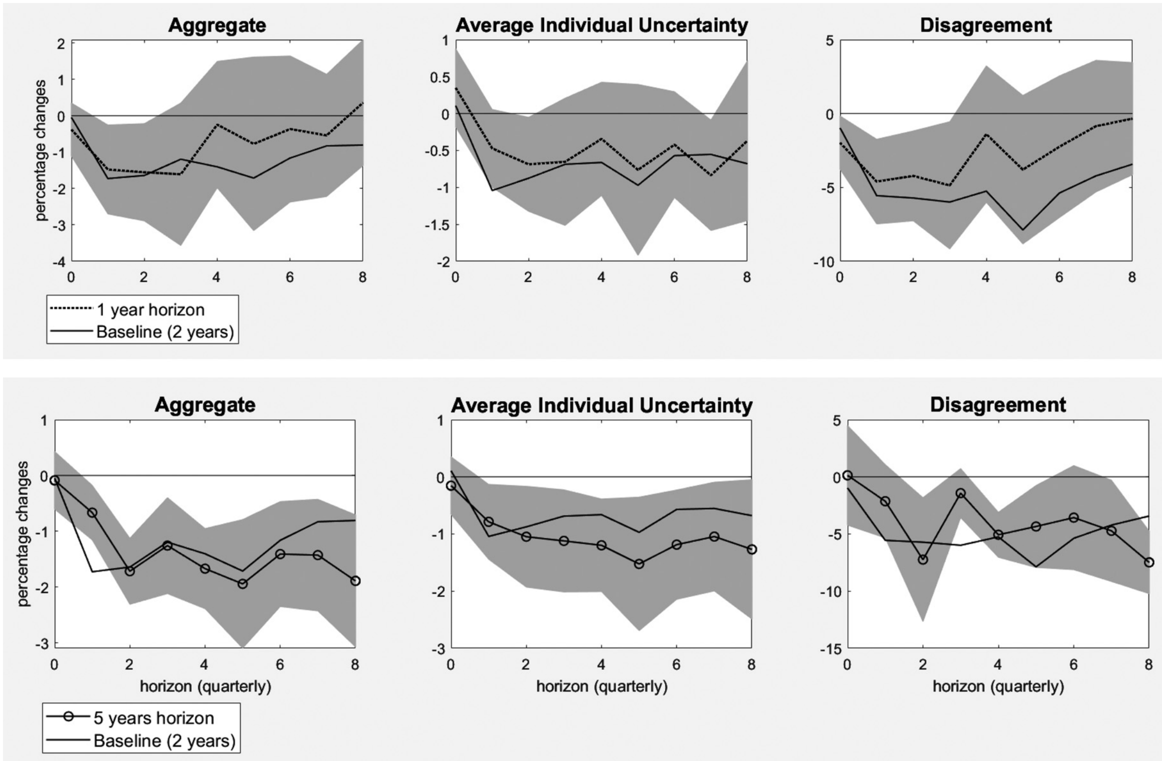

Note: The confidence intervals correspond to the impulse responses of the series employed as a robustness check.

Figure 16. Robustness Check—Dummies

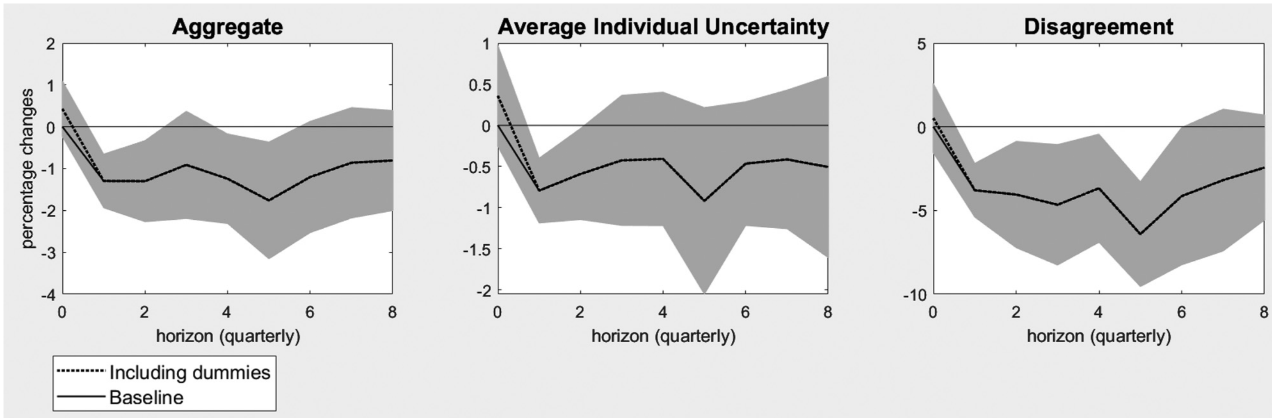

Note: The confidence intervals correspond to the impulse responses of the series employed as a robustness check.

and unemployment. The results are aligned with the uncertainties about inflation and discussed in detail in the appendix.

Finally, I include a set of dummies from 2009:Q1 to 2009:Q4 to account for the effect of the Great Financial Crisis, as the changes computed during that period gain more prominence in each horizon (see Figures 4, 5, and 6). Following the Great Financial Crisis, there was a steep fall in inflation, which contributed to an upward shift in aggregate and average individual uncertainty after 2008:Q4 and resulted in an unprecedent level of disagreement in 2009:Q3. In fact, annual HICP change reached –0.6 percent in July 2009, the lowest level since the beginning of the series in 1999. The inclusion of dummies does not have any relevant impact either on the shape or on the magnitude of the impulse responses (see Figure 16).

8. Conclusions

This paper investigates how the ECB communication of its assessment of the economic outlook affects three types of ex ante inflation uncertainty in the euro area by making use of the ECB SPF and the central bank information shocks provided by Jarociński and Karadi (2020). In addition, the paper also sheds some light on understanding the channels through which central bank information shocks operate.

The results can be summarized as follows. First, I find evidence that ECB communication of its assessment of the economic outlook reduces the dispersion across agents' average point forecasts (disagreement) and at the same time makes agents less uncertain about their own beliefs (ex ante average individual uncertainty). The decrease of disagreement following an ECB information shock suggests that these shocks help opinions to converge, while the reduction of the average individual uncertainty indicates that this signal is valuable and on average contributes to strengthen the confidence in their predications.

Second, a remarkable aspect of this finding is the direction in which inflation forecasts converge. As the point forecasts move toward the mean instead of toward the tails, one can conclude that ECB communication has a "stabilizer effect" on inflation forecasts. Therefore, this result reinforces the idea that central bank information shocks operate as some sort of public signal that is able to influence and coordinate forecasters' opinions and might contribute to market stabilization.

Finally, the muted reaction of inflation expectations to central bank information shocks provides evidence that medium-term inflation expectations remain anchored, reinforcing the institutional credibility aspect of the ECB.

Deciphering how each type of ex ante inflation uncertainty responds to ECB announcements can help policymakers define a communication strategy that attenuates inflation uncertainty in the most effective way possible. One well-known reason for why forecasters disagree is that forecasters may interpret public information in a different way. Therefore, the ECB could tailor its communication to mitigate potential increases in forecast disagreement in volatile times as well as to minimize the possibility of different interpretations among the group of forecasters. Likewise, it is important to sharpen communication when further clarifications or reinforcements of previous messages are necessary, as it helps to improve the forecasters' confidence about their own assessments.

Appendix. Further Robustness Checks

In this appendix, I report in more detail the results related to other variables available in the Survey of Professional Forecasters. Hence, I build the equivalent uncertainty measures for GDP and unemployment[15] for the two-year horizon using the same method described in Section 3.1 (see Figures A.1 and A.2).

Then, I do the same exercise using local projections as shown in Equation (4) to investigate whether the central bank information shocks yield similar results for GDP and unemployment ex ante uncertainties. Figures A.3 and A.4 show that they do: following a central bank information shock, all types of uncertainties decrease for both variables.

In addition, as is also the case in the analysis for ex ante inflation uncertainty, both average individual uncertainty and disagreement

[15]For unemployment average individual uncertainty, in cases where the forecaster placed probabilities in one or two bins, the simple variance was calculated instead of fitting the triangle distribution.

Figure A.1. Ex Ante Unemployment Uncertainties—Two-Year Horizon

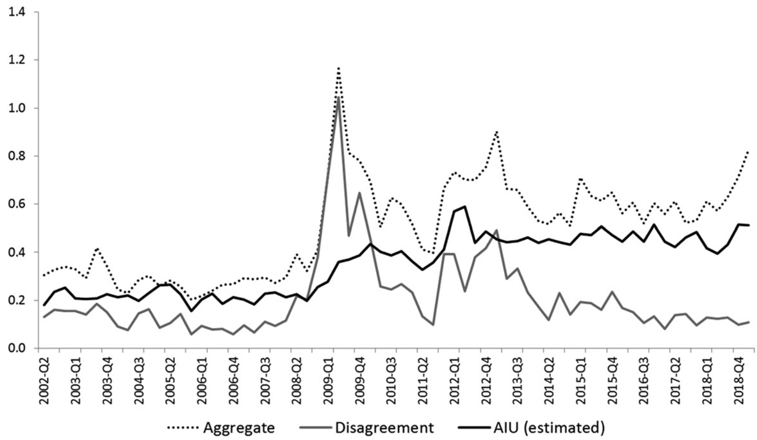

Figure A.2. Ex Ante GDP Uncertainties—Two-Year Horizons

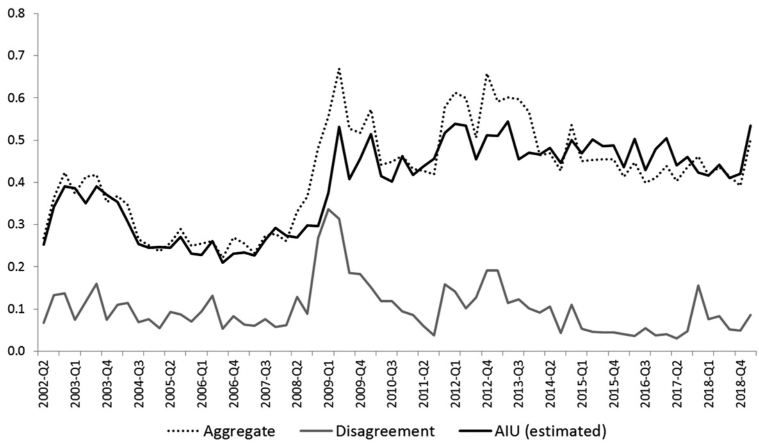

are reduced, with the effect on disagreement being the most prominent, persistent, and immediate. The persistent effect of central bank information shocks on both GDP and unemployment disagreement confirms the influence of central bank communication on aligning opinions across forecasters.

Figure A.3. Response of Ex Ante GDP Uncertainties to Central Bank Information Shocks

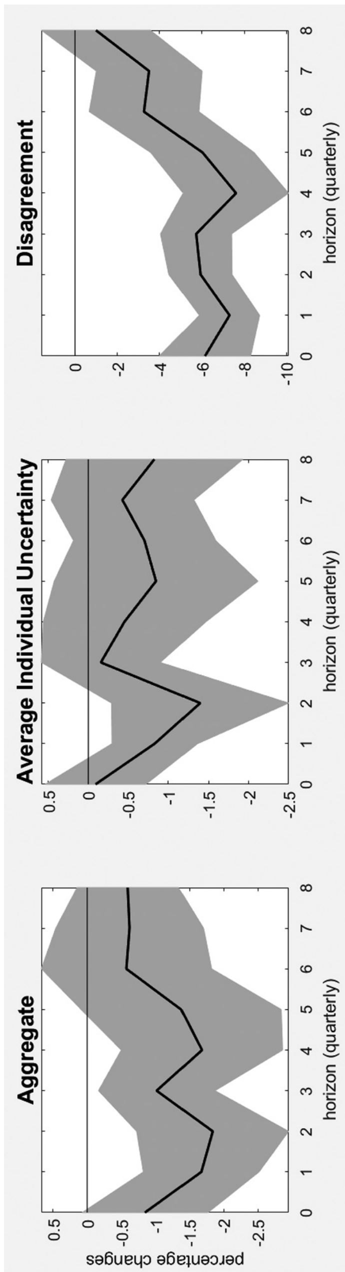

Figure A.4. Response of Ex Ante Unemployment Uncertainties to Central Bank Information Shocks

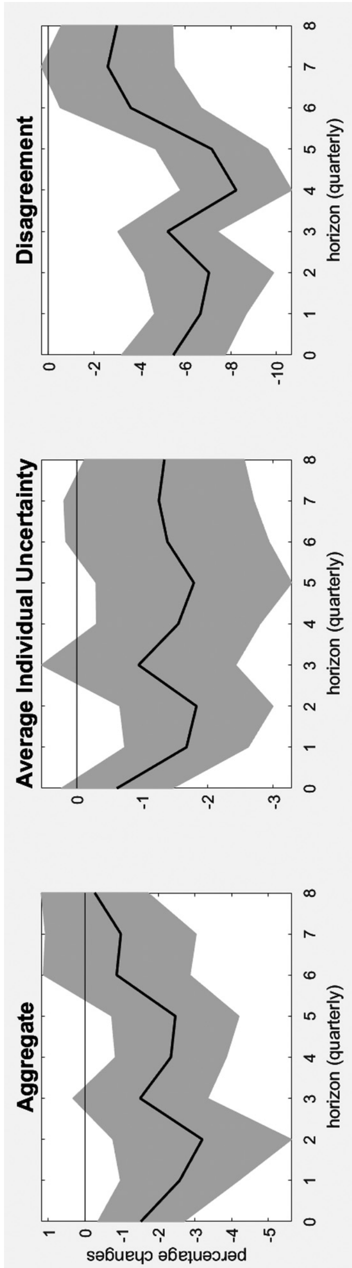

References

Abel, J., R. Rich, J. Song, and J. Tracy. 2016. "The Measurement and Behavior of Uncertainty: Evidence from the ECB Survey of Professional Forecasters." *Journal of Applied Econometrics* 31 (3): 533–50. https://doi.org/10.1002/jae.2430.

Andrade, P., and F. Ferroni. 2021. "Delphic and Odyssean Monetary Policy Shocks: Evidence from the Euro Area." *Journal of Monetary Economics* 117 (January): 816–32. https://doi.org/10.1016/j.jmoneco.2020.06.002.

Barnea, A., A. Dotan, and J. Lakonishok. 1979. "The Effect of Price Level Uncertainty on the Determination of Nominal Interest Rates: Some Empirical Evidence." *Southern Economic Journal* 46 (2): 609–14. https://doi.org/10.2307/1057432.

Berument, H., Y. Yalcin, and J. Yildirim. 2009. "The Effect of Inflation Uncertainty on Inflation: Stochastic Volatility in Mean Model within a Dynamic Framework." *Economic Modelling* 26 (6): 1201–7. https://doi.org/10.1016/j.econmod.2009.05.007.

Blinder, A. S. 2018. "Through a Crystal Ball Darkly: The Future of Monetary Policy Communication." *AEA Papers and Proceedings* 108 (May): 567–71. https://doi.org/10.1257/pandp.20181080.

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen. 2008. "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence." *Journal of Economic Literature* 46 (4): 910–45.

Bloom, N. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–85. https://doi.org/10.3982/ECTA6248.

Born, B., M. Ehrmann, and M. Fratzscher. 2011. "Central Bank Communication on Financial Stability." Working Paper No. 93, National Bank of Poland. https://doi.org/10.2139/ssrn.1927149.

Capistrán, C., and A. Timmermann. 2009. "Disagreement and Biases in Inflation Expectations." *Journal of Money, Credit and Banking* 41 (2–3): 365–96. https://doi.org/10.1111/j.1538-4616.2009.00209.x.

Chan, J. C. 2017. "The Stochastic Volatility in Mean Model with Time-Varying Parameters: An Application to Inflation Modeling." *Journal of Business and Economic Statistics* 35 (1): 17–28. https://doi.org/10.1080/07350015.2015.1052459.

Cieslak, A., and A. Schrimpf. 2019. "Non-Monetary News in Central Bank Communication." *Journal of International Economics* 118 (May): 293–315. https://doi.org/10.1016/j.jinteco.2019.01.012.

Clarida, R., J. Gali, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707. https://doi.org/10.3386/w7147.

Clements, M. P., R. W. Rich, and J. S. Tracy. 2023. "Surveys of Professionals." In *Handbook of Economic Expectations,* ed. R. Bachmann, G. Topa, and W. van der Klaauw, 71–106. Elsevier.

Coenen, G., M. Ehrmann, G. Gaballo, P. Hoffmann, A. Nakov, S. Nardelli, E. Persson, and G. Strasser. 2017. "Communication of Monetary Policy in Unconventional Times." CFS Working Paper No. 578. https://doi.org/10.2139/ssrn.3043098.

Coibion, O., and Y. Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1): 116–59. https://doi.org/10.1086/665662.

Conrad, C., and M. Karanasos. 2005. "On the Inflation-Uncertainty Hypothesis in the USA, Japan and the UK: A Dual Long Memory Approach." *Japan and the World Economy* 17 (3): 327–43. https://doi.org/10.1016/j.japwor.2004.03.002.

Diether, K. B., C. J. Malloy, and A. Scherbina. 2002. "Differences of Opinion and the Cross Section of Stock Returns." *Journal of Finance* 57 (5): 2113–41. https://doi.org/10.1111/0022-1082.00490.

Dincer, N., B. Eichengreen, and P. Geraats. 2019. "Transparency of Monetary Policy in the Post-Crisis World: Revised Central Bank Transparency Data Set for 1998-2015." In *The Oxford Handbook of the Economics of Central Banking*, ed. D. G. Mayes, P. L. Siklos, and J.-E. Sturm, 287–335. New York, NY: Oxford University Press.

Döpke, J., and U. Fritsche. 2006. "When Do Forecasters Disagree? An Assessment of German Growth and Inflation Forecast Dispersion." *International Journal of Forecasting* 22 (1): 125–35. https://doi.org/10.1016/j.ijforecast.2005.05.006.

Dovern, J., U. Fritsche, and J. Slacalek. 2012. "Disagreement among Forecasters in G7 Countries." *Review of Economics and Statistics* 94 (4): 1081–96. https://doi.org/10.1162/rest_a_00207.

Ehrmann, M., S. Eijffinger, and M. Fratzscher. 2012. "The Role of Central Bank Transparency for Guiding Private Sector Forecasts." *Scandinavian Journal of Economics* 114 (3): 1018–52. https://doi.org/10.1111/j.1467-9442.2012.01706.x.

Ehrmann, M., G. Gallo, P. Hoffmann, and G. Strasser. 2019. "Can More Public Information Raise Uncertainty? The International Evidence on Forward Guidance." *Journal of Monetary Economics* 108 (December): 93–112.

Eijffinger, S. C. W., and P. M. Geraats. 2006. "How Transparent Are Central Banks?" *European Journal of Political Economy* 22 (1): 1–21. https://doi.org/10.1016/j.ejpoleco.2005.09.013.

Engelberg, J., C. F. Manski, and J. Williams. 2009. "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters." *Journal of Business and Economic Statistics* 27 (1): 30–41. https://doi.org/10.1198/jbes.2009.0003.

Fama, E. F. 1976. "Inflation Uncertainty and Expected Returns on Treasury Bills." *Journal of Political Economy* 84 (3): 427–48. https://doi.org/10.1086/260453.

Fountas, S., A. Ioannidis, and M. Karanasos. 2004. "Inflation, Inflation Uncertainty and a Common European Monetary Policy." *The Manchester School* 72 (2): 221–42. https://doi.org/10.1111/j.1467-9957.2004.00390.x.

Friedman, M. 1977. "Nobel Lecture: Inflation and Unemployment." *Journal of Political Economy* 85 (3): 451–72. https://doi.org/10.1086/260579.

Garcia, J. A. 2003. "An Introduction to the ECB's Survey of Professional Forecasters." ECB Occasional Paper No. 8. https://doi.org/10.2139/ssrn.748971.

Giordani, P., and P. Söderlind. 2006. "Is There Evidence of Pessimism and Doubt in Subjective Distributions? Implications for the Equity Premium Puzzle." *Journal of Economic Dynamics and Control* 30 (6): 1027–43. https://doi.org/10.1016/j.jedc.2005.05.001.

Glas, A. 2020. "Five Dimensions of the Uncertainty–Disagreement Linkage." *International Journal of Forecasting* 36 (2): 607–27. https://doi.org/10.1016/j.ijforecast.2019.07.010.

Glas, A., and M. Hartmann. 2016. "Inflation Uncertainty, Disagreement and Monetary Policy: Evidence from the ECB Survey of Professional Forecasters." *Journal of Empirical Finance* 39 (Part B): 215–28. https://doi.org/10.1016/j.jempfin.2016.05.001.

Grauer, F. L., and R. H. Litzenberger. 1979. "The Pricing of Commodity Futures Contracts, Nominal Bonds and Other Risky Assets under Commodity Price Uncertainty." *Journal of Finance* 34 (1): 69–83. https://doi.org/10.1111/j.1540-6261.1979.tb02071.x.

Grier, K. B., and M. J. Perry. 1998. "On Inflation and Inflation Uncertainty in the G7 Countries." *Journal of International Money and Finance* 17 (4): 671–89. https://doi.org/10.1016/s0261-5606(98)00023-0.

———. 2000. "The Effects of Real and Nominal Uncertainty on Inflation and Output Growth: Some GARCH-M Evidence." *Journal of Applied Econometrics* 15 (1): 45–58.

Grothe, M., and A. Meyler. 2015. "Inflation Forecasts: Are Market-Based and Survey-Based Measures Informative?" ECB Working Paper No. 1865.

Gürkaynak, R. S., B. Sack, and E. Swanson. 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review* 95 (1): 425–36. https://doi.org/10.1257/0002828053828446.

Hamilton, J. D. 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *Review of Economics and Statistics* 100 (5): 831–43. https://doi.org/10.1162/rest_a_00706.

Hartmann, M., and J. Roestel. 2013. "Inflation, Output and Uncertainty in the Era of Inflation Targeting — A Multi-Economy View on Causal Linkages." *Journal of International Money and Finance* 37 (October): 98–112. https://doi.org/10.1016/j.jimonfin.2013.05.011.

Hüning, H. 2017. "Asset Market Response to Monetary Policy News from SNB Press Releases." *North American Journal of Economics and Finance* 40 (April): 160–77. https://doi.org/10.1016/j.najef.2017.02.003.

Issing, O. 2019. *The Long Journey of Central Bank Communication.* Cambridge, MA: The MIT Press.

Jarociński, M., and P. Karadi. 2020. "Deconstructing Monetary Policy Surprises—The Role of Information Shocks."

American Economic Journal: Macroeconomics 12 (2): 1–43. https://doi.org/10.1257/mac.20180090.

Jegadeesh, N., and D. (Andrew) Wu. 2017. "Deciphering Fedspeak: The Information Content of FOMC Meetings." https://api.semanticscholar.org/CorpusID:30415687.

Jitmaneeroj, B., M. J. Lamla, and A. Wood. 2019. "The Implications of Central Bank Transparency for Uncertainty and Disagreement." *Journal of International Money and Finance* 90 (February): 222–40. https://doi.org/10.1016/j.jimonfin.2018.10.002.

Johnstone, D. 2016. "The Effect of Information on Uncertainty and the Cost of Capital." *Contemporary Accounting Research* 33 (2): 752–74. https://doi.org/10.1111/1911-3846.12165.

Jordà, Ò. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82. https://doi.org/10.1257/0002828053828518.

Jung, A., and H. Uhlig. 2019. "Monetary Policy Shocks and the Health of Banks." ECB Working Paper No. 2303. https://doi.org/10.2139/ssrn.3429629.

Kerssenfischer, M. 2019. "Information Effects of Euro Area Monetary Policy: New Evidence from High-Frequency Futures Data." Deutsche Bundesbank Discussion Paper No. 07/2019. https://doi.org/10.2139/ssrn.3368356.

Kliesen, K. L., and F. A. Schmid. 2004. "Monetary Policy Actions, Macroeconomic Data Releases, and Inflation Expectations." *Review* (Federal Reserve Bank of St. Louis) 86 (3): 9–21. https://doi.org/10.3886/icpsr01301.

Kohn, D. L., and B. P. Sack. 2003. "Central Bank Talk: Does It Matter and Why?" Finance and Economics Discussion Series No. 2003-55, Board of Governors of the Federal Reserve System. https://doi.org/10.17016/feds.2003.55.

Kontonikas, A. 2004. "Inflation and Inflation Uncertainty in the United Kingdom, Evidence from Garch Modelling." *Economic Modelling* 21 (3): 525–43. https://doi.org/10.1016/j.econmod.2003.08.001.

Lahiri, K., and X. Sheng. 2008. "Evolution of Forecast Disagreement in a Bayesian Learning Model." *Journal of Econometrics* 144 (2): 325–40. https://doi.org/10.1016/j.jeconom.2008.02.002.

———. 2010. "Measuring Forecast Uncertainty by Disagreement: The Missing Link." *Journal of Applied Econometrics* 25 (4): 514–38. https://doi.org/10.1002/jae.1167.

Lamla, M. J., and T. Maag. 2012. "The Role of Media for Inflation Forecast Disagreement of Households and Professional Forecasters." *Journal of Money, Credit and Banking* 44 (7): 1325–50. https://doi.org/10.1111/j.1538-4616.2012.00534.x.

Mankiw, N. G., R. Reis, and J. Wolfers. 2004. "Disagreement about Inflation Expectations." In *NBER Macroeconomics Annual 2003,* Vol. 18, ed. M. Gertler and K. Rogoff, 209–70. Cambridge, MA: The MIT Press. https://doi.org/10.1086/ma.18.3585256.

Manzan, S. 2011. "Differential Interpretation in the Survey of Professional Forecasters." *Journal of Money, Credit and Banking* 43 (5): 993–1017. https://doi.org/10.1111/j.1538-4616.2011.00404.x.

Melo Fernandes, C. M., and G. Kenny. 2024. "Ask the Experts: Ex-Ante Uncertainty and the Euro Area Business Cycle." Forthcoming in ECB Working Paper Series.

Minegishi, M., and B. Cournède. 2009. "The Role of Transparency in the Conduct of Monetary Policy." OECD Economics Department Working Paper No. 274. https://doi.org/10.1787/221128317344.

Morris, S., and H. S. Shin. 2002. "Social Value of Public Information." *American Economic Association* 92 (5): 1521–34. https://doi.org/10.1257/000282802762024610.

Nakamura, E., and J. Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–1330. https://doi.org/10.1093/qje/qjy004.

Naszodi, A., C. Csavas, S. Erhart, and D. Felcser. 2016. "Which Aspects of Central Bank Transparency Matter? A Comprehensive Analysis of the Effect of Transparency of Survey Forecasts." *International Journal of Central Banking* 12 (4): 147–92.

Patton, A. J., and A. Timmermann. 2010. "Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion." *Journal of Monetary Economics* 57 (7): 803–20. https://doi.org/10.1016/j.jmoneco.2010.07.001.

Picault, M., and T. Renault. 2017. "Words Are Not All Created Equal: A New Measure of ECB Communication." *Journal of International Money and Finance* 79 (December): 136–56. https://doi.org/10.1016/j.jimonfin.2017.09.005.

Posso, A., and G. B. Tawadros. 2013. "Does Greater Central Bank Independence Really Lead to Lower Inflation? Evidence from Panel Data." *Economic Modelling* 33 (July): 244–47. https://doi.org/10.1016/j.econmod.2013.04.005.

Romer, C. D, and D. H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90 (3): 429–57. https://doi.org/10.1257/aer.90.3.429.

Siklos, P. L. 2013. "The Global Financial Crisis and Central Bank Speak: Central Bank Guidance in Good Times and in Bad." https://doi.org/10.2139/ssrn.2315009.

Swanson, E. T. 2006. "Have Increases in Federal Reserve Transparency Improved Private Sector Interest Rate Forecasts?" *Journal of Money, Credit, and Banking* 38 (3): 791–819. https://doi.org/10.1353/mcb.2006.0046.

van der Cruijsen, C., and M. Demertzis. 2007. "The Impact of Central Bank Transparency on Inflation Expectations." *European Journal of Political Economy* 23 (1): 51–66. https://doi.org/10.1016/j.ejpoleco.2006.09.009.

Wallis, K. F. 2005. "Combining Density and Interval Forecasts: A Modest Proposal." *Oxford Bulletin of Economics and Statistics* 67 (s1): 983–94. https://doi.org/10.1111/j.1468-0084.2005.00148.x.

Weidmann, J. 2018. "Central Bank Communication as an Instrument of Monetary Policy." Lecture at the Centre for European Economic Research, Mannheim, May 2. https://www.bis.org/review/r180511a.htm.

Zarnowitz, V., and L. A. Lambros. 1987. "Consensus and Uncertainty in Economic Prediction." *Journal of Political Economy* 95 (3): 591–621. https://doi.org/10.1086/261473.

Estimates of the Natural Rate of Interest Consistent with a Supply-Side Structure and a Monetary Policy Rule for the U.S. Economy*

Manuel González-Astudillo[a,b] and Jean-Philippe Laforte[a]
[a]Federal Reserve Board
[b]Escuela Superior Politécnica del Litoral (ESPOL)

We estimate the natural rate of interest (r^*) using a semi-structural model of the U.S. economy that jointly characterizes the trend and cyclical factors of key macroeconomic variables such as output, the unemployment rate, inflation, and short- and long-term interest rates. We specify a monetary policy rule and a 10-year Treasury yield equation to exploit the information provided by both interest rates to infer r^*. However, the use of a monetary policy rule with a sample that spans the Great Recession and its aftermath poses a challenge because of the effective lower bound. We devise a Bayesian estimation technique that incorporates a Tobit-like specification to deal with the censoring problem. We compare and validate our model specifications using pseudo-out-of-sample forecasting exercises. Our results show that the smoothed value of r^* declined sharply around the Great Recession, eventually falling below zero, and remained negative through early 2020. Our results also indicate that obviating the censoring

*We thank Robert Tetlow, Benjamin Johannsen, Matthias Paustian, Falk Mazelis, Claus Brand, Gabriel Pérez-Quirós, seminar participants at the European Central Bank and the Bank of Mexico, and an anonymous referee for insightful comments that led to significant improvements in earlier versions of our paper. Mark Wilkinson, Evan Corden, and Avi Lipton provided speedy and accurate research assistance. Any remaining errors are our own. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System. Author e-mails: manuel.p.gonzalez-astudillo@frb.gov (corresponding author), jean-philippe.laforte@frb.gov.

would imply higher estimates of r^* than otherwise. We also extend our results to the COVID-19 pandemic period, introducing stochastic volatility in the model and dealing with the massive swings in the data, to find that our estimate of r^* is slightly below 1 percent in early 2023.

JEL Codes: C32, C34, E32.

1. Introduction

The natural rate of interest has become a key concept to understand and characterize monetary policy in both theory and practice. As pointed out by Summers and Rachel (2019), monetary policymakers across the globe have highlighted it as a fundamental policy variable to assess the stance of monetary policy. For instance, Chair Jerome Powell cites this factor as one of the benchmarks of the Federal Reserve's monetary policy decisions:

> [W]e set our policy interest rate to achieve our goals of maximum employment and stable prices. In doing so, we often refer to certain benchmarks. One of these is the interest rate that would be neutral—neither restraining the economy nor pushing it upward. We call that rate "r^*" (pronounced "r-star"). A policy rate above r^* would tend to restrain economic activity, while a setting below r^* would tend to speed up the economy. A second benchmark is the natural rate of unemployment, which is the lowest rate of unemployment that would not create upward pressure on inflation. We call that rate "u^*" (pronounced "u-star"). You can think of r^* and u^* as two of the main stars by which we navigate. In an ideal world, policymakers could rely on these stars like mariners before the advent of GPS. But, unlike celestial stars on a clear night, we cannot directly observe these stars, and their values change in ways that are difficult to track in real time. (Powell 2019)

In this paper, we postulate and estimate a semi-structural model of the U.S. economy that allows us to jointly infer time-varying measures for r^* and u^* (denoted r_t^* and u_t^* from here onward) within a framework in which monetary policy is characterized by an inertial version of the Taylor (1993) rule. In particular, r_t^* is the time-varying intercept of the monetary policy rule. As such, it is the value of the real interest rate that would prevail in the long run, when the inflation rate is at its target and output is at its potential level (and the unemployment rate is at u_t^*). The specification of a policy rule requires that we account for the effective lower bound (ELB) on the federal funds rate, as otherwise the relationship between the rule and the observed short-term interest rate breaks down when the latter is at the ELB. It is particularly important to explicitly account for the censoring if one wants, like us, to analyze and estimate a sample that includes the Great Recession and its aftermath or the COVID-19 pandemic period.

A comprehensive literature on the estimation of a notion of the natural rate of interest for the U.S. already exists (see Lubik and Matthes 2015; Kiley 2020; Cúrdia et al. 2015; Del Negro et al. 2017; Christensen and Rudebusch 2019; Lewis and Vazquez-Grande 2019; Johannsen and Mertens 2021, for instance).[1] A seminal and original work on the estimation of the natural rate of interest for the U.S. economy is Laubach and Williams (2003) (LW hereafter), which has been subsequently updated and expanded to other advanced economies in Holston, Laubach, and Williams (2017) (HLW hereafter). They exploit the theoretical relationship between the real rate of interest and the growth rate of the economy to estimate r_t^* based on information from real gross domestic product (GDP), the inflation rate, and the short-term real interest rate. While their estimator is widely popular, several issues have been raised regarding their approach by subsequent work (see Beyer and Wieland 2019). First, a great deal of uncertainty pertains to the estimate of r_t^*. Second, there is a significant wedge between their output gap estimate and more conventional ones—such as that of the Congressional Budget Office (CBO)—starting in the early 2000s and widening from then on (strikingly, the output gap estimate casts the Great Recession as a rather shallow downturn, historically speaking). Lastly, the choice of relying on maximum likelihood methods exposes their estimates to the pile-up problem, as the estimated variances of some shocks may be biased toward zero. This problem remains even if the model is adequately identified given the data, which may not even be the case with the original LW model (see Fiorentini et al. 2018).

[1]We include a literature review in Appendix A.

We give consideration to these observations in our model. For instance, we use information from short- and long-run interest rates, as in Bauer and Rudebusch (2020), to better pin down the evolution of r_t^*. Similarly, we also introduce an inflation trend and tightly match its dynamics with that of a measure of survey-based inflation expectations, which helps discriminate between movements in yields due to r_t^* and those due to trend inflation. In addition to real GDP and the inflation rate, we also add information on the unemployment rate and use it to better identify the output gap through an Okun's law, as originally proposed by Clark (1989), as well as to estimate the natural rate of unemployment. Finally, we adopt a more robust estimation approach that relies on Bayesian techniques appropriate for state-space modeling.

As indicated earlier, we exploit the information provided by the federal funds rate by specifying its evolution as a Taylor (1993) rule with inertia. However, the binding of the ELB during the Great Recession and the recovery that followed complicates the use of a policy rule for any data set that extends beyond 2008. We tackle this issue by embedding the model with a Tobit-like specification for the Taylor rule and, hence, a shadow rate. The failure to account for the ELB can significantly distort the outcomes of the estimation in terms of both parameters and latent states—r_t^* among them. Our results indicate that this is the case.

Our approach is similar to but not the same as those in Wu and Xia (2016), Carriero et al. (2023), and Johannsen and Mertens (2021). Wu and Xia estimate the shadow rate implied by a discrete-time multifactor model of the term structure of interest rates, embedding an analytical approximation adjustment to account for the lower bound on the observed short-term interest rate. They use monthly frequency information from (and only from) a set of forward rates of different maturities. Carriero et al. also use rates at different maturities to estimate shadow interest rates, but without imposing no-arbitrage conditions, to improve the forecasting performance of vector autoregressive (VAR) models. Our approach differs, for instance, regarding both the breadth of the information set and the choice of the identifying assumptions and structures. While our data set includes only a short-term and a long-term maturity yield (the federal funds rate and the 10-year Treasury yield), it also includes information about macroeconomic variables as well as long-run inflation expectations. The combination of a more comprehensive information set and macroeconomic structures allows us to exploit comovements across key macroeconomic variables to possibly improve on the identification of r_t^*. Johannsen and Mertens also utilize macroeconomic variables to inform their estimate of the natural rate of interest in addition to interest rates of different maturities. Like us, they propose a flexible time-series approach that decomposes their data as trends and cycles, and that explicitly accounts for the presence of the ELB by simulating a shadow rate for the periods when the ELB is binding. However, and in contrast to our methodology, they do not infer the output gap based on the structure of their model and the data, but instead rely on the CBO estimates and treat it an as observed series. Neither do they estimate u_t^* consistent with their inference of r_t^*.

A last paper written around the same time as ours and worth mentioning is Zaman (2021). It is a comprehensive study of a semi-structural model of the U.S. economy that shares many features with our paper. The model is estimated with Bayesian methods, includes information from survey data (to a greater extent than ours), and specifies the cyclical component of the short-term interest rate using a Taylor-like policy rule. It also allows for time variation in some of the parameters, in addition to the variances of the innovations as done at the end of this paper. One notable difference between Zaman's and our approach is that the former does not sample a model-consistent distribution of the shadow rate at the ELB but instead relies on the estimate of Wu and Xia (2016) as an observable variable. Another distinction is that we use the 10-year Treasury yield to directly inform the estimate of r_t^*.

Based on data whose sample ends just before the pandemic (2020:Q1), our estimate of r_t^* gradually declines starting in the mid-1980s and enters negative territory in early 2008; r_t^* is estimated to have hovered around –1 percent since 2012, in line with simple estimates of the short-term real interest rate, before gradually edging down to –1.7 percent over the last year of our sample. We find that the shadow federal funds rate would have reached –5.9 percent at the trough of the Great Recession. Regarding the natural unemployment rate, we find that it has been steadily declining since 2010, when it reached 5.6 percent, to a level of 4.5 percent in 2020:Q1. Our measure of the (time-varying) potential output growth rate has declined over time and has been around 1.4 percent per year since 2012. Finally, our estimate of the output gap is somewhat similar to those from the CBO and the staff of the Federal Reserve Board (and significantly different from LW and HLW). It peaked at 1.7 percent in 2019 and declined to about 1 percent at the end of the sample.

Taking advantage of historical data decomposition techniques, we find that our negative estimates of the natural rate of interest since the Great Recession are based on the information from a small subset of observations. More specifically, the secular decline in the long-run interest rate and the persistently low realized inflation apply enough downward pressures on our estimate of r_t^* for it to turn negative around the Great Recession and remain below zero thereafter.

Following the estimation and analysis of our benchmark, we investigate the relevance and contribution of some of the assumptions behind the baseline model. First, we gauge the importance of allowing for correlated disturbances. In our baseline model, we specify that the shocks to the short- and long-run interest rates are correlated in order to introduce a conventional monetary policy channel by which shocks to the federal funds rate translate into changes to the long-run interest rate, affecting the output gap (through an IS-curve specification), inflation (through a Phillips-curve equation), and the unemployment rate (through an Okun's law relationship). We also allow correlation between the innovations of the r_t^* process and those of trend output growth to link these two variables in a way similar to LW. We find that accounting for the latter correlation is empirically significant but not so for the former, as the conventional marginal data densities strongly penalize the assumption of correlation between transitory shocks to the rates.

We also quantify the effects of ignoring the ELB and find that our estimate of r_t^* is about 35 basis points higher, on average, than in the model that takes into account the censoring of the policy rule. Moreover, we assess the impact of two material changes to our framework: (i) adding the CBO estimate of the output gap to our set of observable series (in line with Johannsen and Mertens 2021); and (ii) assuming that the federal funds rate follows a simple local-level model, similar to Fiorentini et al. (2018), rather than a Taylor rule. Then, we conduct pseudo-out-of-sample forecasting exercises to determine which specification performs best. The results indicate that the baseline specification—which incorporates a Taylor rule with a shadow rate and an IS curve, and assumes some correlations between key innovations—overall outperforms the other specifications.

Finally, we reestimate the baseline model with the variances of the innovations evolving according to a stochastic volatility specification. This version of the model has a forecasting performance similar to our baseline specification without stochastic volatility and implies an evolution of r_t^* above its homoskedastic counterpart, hovering slightly above zero since 2012 and dipping to about –1 percent at the very end of the sample. We then use this version of the model to assess the unfolding of r_t^* and other latent variables during the COVID-19 pandemic, using sample information through early 2023. The results indicate that the inferred value of r_t^* has rebounded from around –1 percent at the onset of the pandemic to a level close to 0.7 percent. This result implies that the neutral policy rate (obtained as the sum of our estimate of r_t^* and that of trend inflation) is close to 2.9 percent.

2. The Model

Our model of the U.S. economy includes equations for (the log of) real GDP, denoted as y_t, the unemployment rate, u_t, the core personal consumption expenditures (PCE) price inflation rate, π_t, the federal funds rate, i_t, the 10-year Treasury yield, i_t^{10}, and survey information about long-run inflation expectations, π_t^e.

2.1 Interest Rates

We begin our presentation of the model with a description of the policy rule and the model specification of the 10-year Treasury yield, as our main innovations mostly relate to and concentrate on this block of the model.

We assume that the dynamics of the (unconstrained) short-term interest rate are determined by a monetary policy rule specified as an inertial version of Taylor (1993):

$$R_t = \rho R_{t-1} + (1-\rho)\left(r_t^* + \pi_t^* + \alpha^{\pi}\left(\bar{\pi}_t - \pi_t^*\right) + \alpha^y c_t\right) + \eta_t^R, \quad (1)$$

where $\bar{\pi}_t$ is the four-quarter average of the inflation rate, π_t^* is its trend (assumed to be equal to the policymakers' inflation target), and c_t is the output gap.[2] Here, R_t is the nominal interest rate that would be set by the monetary authority in the absence of a lower bound on the target federal funds rate, also called the shadow rate. In this setup, $r_t^* + \pi_t^*$ is a measure of the trend policy rate, when inflation is at its target and the output gap is closed. This level of the short-term interest rate is called "neutral" or "equilibrium" because it is neither expansionary nor contractionary (see Yellen 2017). As a consequence, r_t^* may be viewed as the Wicksellian concept of the natural interest rate, compatible with stable prices and such that an increase of the real interest rate above r_t^* contracts economic activity (see Lubik and Matthes 2015 for an additional discussion). In addition, r_t^* can also be viewed as a measure of the trend real interest rate, also referred to as the natural real interest rate by Taylor (1993).

In Taylor (1993)'s proposal, r_t^* was modeled as a constant equal to 2 percent, close to the then-estimated steady-state growth rate of trend GDP. The choice of this value was supported at the time by the average historical value of the federal funds rate. However, the economic events that have taken place since the publication of the paper have led monetary policymakers and economists to reconsider the view and assumption of a constant level of r_t^*. For example, Yellen (2017) points out that a Taylor (1993) policy rule with r_t^* at 2 percent prescribes a path for the federal funds rate that is much higher than the median of Federal Open Market Committee (FOMC) participants' assessment of appropriate policy. Yellen mentions that, because overall growth has been quite moderate over the past few years, some recent estimates of the current value of r_t^* stand close to zero, citing HLW. Similarly, Bullard (2018) advocates for a modernized version of the Taylor (1999) rule in which the natural rate of interest varies over time. Lower labor productivity growth, a slow pace of labor force growth, and a stronger desire for safe assets than in the past would be factors that currently imply a lower equilibrium real interest rate.

[2]In addition, $\alpha^\pi > 1, \alpha^y > 0, \rho \in [0, 1)$, and $\eta_t^R \sim N(0, \sigma^2_{\eta^R})$.

We assume that r_t^* evolves as follows:

$$r_t^* = r_{t-1}^* + \eta_t^{r^*}, \tag{2}$$

with $\eta_t^{r^*} \sim N(0, \sigma^2_{\eta^{r^*}})$. Orphanides and Williams (2002) and Kiley (2020), among others, also use a random walk specification for r_t^*, as in (2). Additionally, and in contrast with LW and several subsequent papers by other authors that assume the trend output growth rate (denoted μ_t in our paper) loads with unit coefficient on r_t^*, we assume that their respective error terms are correlated, i.e., $\text{corr}\left(\eta_t^{r^*}, \eta_t^{\mu}\right) = \omega$, where η_t^{μ} is the shock to the trend output growth rate.[3]

The inclusion of a monetary policy rule to improve the identification of r_t^* has also been investigated in Brand and Mazelis (2019), which they append to a version of the LW model. However, they ignore the matter of the ELB binding and the consequences of this omission for their estimation results. In contrast, we explicitly account for the ELB and specify the observed federal funds rate, i_t, as the maximum between a lower bound, denoted as $\underline{i}$, and the shadow rate, as follows:

$$i_t = \max\{R_t, \underline{i}\}. \tag{3}$$

Several papers in the literature have built in a measure of the shadow rate in their estimation of the stance of monetary policy. Bauer and Rudebusch (2016) and Wu and Xia (2016), for instance,

[3]In HLW, r_t^* is given by the following specification:

$$r_t^* = \mu_t + z_t,$$
$$\mu_t = \mu_{t-1} + \eta_t^{\mu},$$
$$z_t = z_{t-1} + \eta_t^{z},$$

where z_t is meant to capture the net contribution of the other determinants of r_t^* beside μ_t, with $\eta_t^{\mu} \sim N(0, \sigma^2_{\eta^{\mu}})$ and $\eta_t^{z} \sim N(0, \sigma^2_{\eta^{z}})$. The correlation between changes in r_t^* and the trend output growth rate in HLW is given by

$$\text{corr}\left(\Delta r_t^*, \Delta\mu_t\right) \equiv \omega = \text{corr}\left(\eta_t^{z} + \eta_t^{\mu}, \eta_t^{\mu}\right) = \frac{\sigma_{\eta^{\mu}}}{\sqrt{\sigma^2_{\eta^{\mu}} + \sigma^2_{\eta_z}}}.$$

Hence, given the parameter estimates in their paper, $\omega = 0.63$.

use shadow rate term structure models (SRTSMs) to calculate the short-term interest rate during the zero lower bound episode of the U.S. economy. In the SRTSM, the short-term interest rate depends on latent factors extracted from yields at different maturities or from a combination of yields and macroeconomic variables. Our setup can be viewed as one in which the short-term interest rate depends on latent factors such as r_t^*, the inflation trend, and the output gap that are obtained from macroeconomic and financial variables.

To the best of our knowledge, Johannsen and Mertens (2021) is the only study that incorporates the concept of r_t^* within the framework of a shadow nominal interest rate. The authors impose a long-run Fisher equation in which the shadow rate trend is decomposed into an inflation trend and a real-rate trend that is modeled as in (2). Even though yields at different maturities are used to estimate the trends and cycles of the model, they do not impose any no-arbitrage condition.

In the spirit of Johannsen and Mertens (2021), we include in our set of variables the 10-year Treasury yield, denoted i_t^{10}, as in principle it provides information about the inflation trend and r_t^* beyond that given by the short-term interest rate. We specify its dynamics as follows:

$$i_t^{10} = r_t^* + \pi_t^* + p_t^{10} + c_t^{10}, \tag{4}$$

$$c_t^{10} = \psi_1 c_{t-1}^{10} + \psi_2 c_{t-2}^{10} + \varepsilon_t^{10}, \tag{5}$$

$$p_t^{10} = p_{t-1}^{10} + \eta_t^{p^{10}}, \tag{6}$$

where $\varepsilon_t^{10} \sim N(0, \sigma^2_{\varepsilon^{10}})$, $\eta_t^{p^{10}} \sim N(0, \sigma^2_{\eta^{p^{10}}})$, and c_t^{10} is a process representing any persistent but stationary deviations around the shifting endpoints $r_t^* + \pi_t^* + p_t^{10}$, which could be, for instance, the confluence of term premium and expected future short-run interest rate dynamics. As evidenced by Bauer and Rudebusch (2020), the term premium may display nonstationary dynamics even after accounting for a stochastic trend driving the term structure of interest rates. To allow and capture movements of that nature for the components of the 10-year Treasury yield beyond r_t^*, π_t^*, and c_t^{10}, we include a random walk component, p_t^{10}, in the specification of i_t^{10}. Working with reduced-form specifications rather than explicitly modeling expectations and no-arbitrage conditions is not without consequences with respect to the status and contribution of monetary policy in the model. For instance, the 10-year Treasury yield cycle process, c_t^{10}, conflates both the cyclical dynamics of the term premium and the expectations of the short-term interest rate. However, since the term premium and the identification of the expectational component of the 10-year Treasury yield are not our primary objects of interest, we are comfortable with that simplification.

Nonetheless, this specification of the short- and long-run interest rates does not allow for a direct effect from a conventional monetary policy shock to the long-term interest rate. In order to introduce such an effect, we assume a non-zero correlation between the innovation of the policy rate, η_t^R, and that of the cycle of the 10-year Treasury yield, ε_t^{10} (see Cochrane and Piazzesi 2002; Nakamura and Steinsson 2018, for instance). Under such a specification, a conventional contractionary monetary policy shock would result in a proportional change in the long-term interest rate via its cyclical component.[4]

The inclusion of a long-term interest rate beside the federal funds rate provides some signal not only about expected future variations in interest rates of shorter maturity but also about shifts in their common low-frequency component (r_t^* and π_t^*). The information provided by the interest rate of longer maturity can be particularly valuable when the short-term interest rate is at the ELB.

In the remainder of this section, we outline our setups of real activity (GDP and the unemployment rate) and inflation, and describe how r_t^* may influence and be influenced by these sectors of the economy through their effects on the short- and long-term interest rates.

2.2 Real GDP and the Unemployment Rate

We characterize real GDP and the unemployment rate using a trend-cycle decomposition approach, similar to that used by Clark (1989), as follows:

$$y_t = y_t^* + c_t, \tag{7}$$

$$u_t = u_t^* + \theta_1 c_t + \theta_2 c_{t-1} + v_t, \tag{8}$$

[4]We also explored an alternative specification in which the error term of the cyclical component of the 10-year Treasury yield, ε_t^{10}, is a linear function of the shadow rate shock, η_t^R, plus an i.i.d. disturbance. The results are very similar.

where (the log of) real GDP is decomposed as the sum of potential output, denoted as y_t^*, and the output gap, denoted as c_t. In turn, we assume that potential output is a local-linear trend, whereas the output gap is a stationary AR(2) process influenced by the cyclical component of the 10-year (real) Treasury yield, as shown below:

$$y_t^* = \mu_{t-1} + y_{t-1}^* + \eta_t^{y^*}, \tag{9}$$

$$\mu_t = \mu_{t-1} + \eta_t^{\mu}, \tag{10}$$

$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + \lambda_1 c_{t-1}^{10} + \lambda_2 c_{t-2}^{10} + \varepsilon_t, \tag{11}$$

where $\eta_t^{y^*} \sim N(0, \sigma_{\eta^{y^*}}^2), \eta_t^{\mu} \sim N(0, \sigma_{\eta^{\mu}}^2)$, $\varepsilon_t \sim N(0, \sigma_{\varepsilon}^2)$, and the shocks are independent of each other. Equation (10) allows potential output to exhibit a (time-varying) trend growth rate, denoted as μ_t. This feature is particularly important given the lower-than-average productivity growth rates observed, in particular, after the Great Recession. We ensure feedback from monetary policy to economic activity and inflation with the presence of the long-term real interest rate gap as in Roberts (2018). This assumption represents a departure from LW, who included a short-term real interest rate gap. The reasons behind this choice are rather straightforward: First, spending decisions more likely depend on the long-term than on the short-term interest rate gap; second, monetary policymakers used balance sheet policies as well as forward guidance to influence long rates during the Global Financial Crisis (GFC); and, finally, the relationship between short and long rates may have changed after the Great Recession.

The unemployment rate in (8) is determined by an Okun's law with coefficients θ_1 and θ_2. The natural rate of unemployment is given by u_t^*, which evolves according to the following random walk process:

$$u_t^* = u_{t-1}^* + \eta_t^{u^*}, \tag{12}$$

where $\eta_t^{u^*} \sim N(0, \sigma_{\eta^{u^*}}^2)$. The Okun's law error, denoted $\upsilon_t \sim N(0, \sigma_{\upsilon}^2)$, allows for deviations of the unemployment rate from its trend and cyclical components.[5]

[5] All the disturbances in this section are independent of each other. Some authors allow for correlated trend-cycle disturbances in a similar setting (see

2.3 Inflation

We specify the inflation process with the likes of a hybrid Phillips curve in which inflation expectations are treated as a latent variable specified as a weighted average of trend inflation, denoted as π_t^*, and actual lagged inflation (see Basistha and Nelson 2007). Inflation is also a function of the degree of slack (measured by c_t) in the economy. The specification appears below:

$$\pi_t = \beta\pi_t^* + (1-\beta)\,\pi_{t-1} + \kappa c_t + \eta_t^\pi, \tag{13}$$

with $\eta_t^\pi \sim N(0, \sigma_{\eta^\pi}^2)$ and where κ is the slope of the Phillips curve; we ensure long-run neutrality by assuming that $\beta \in (0,1]$. Notice that this approach allows the inflation rate to converge to trend inflation when the output gap is closed.[6]

Additionally, we assume that the inflation trend evolves as a random walk process, as follows (see Stock and Watson 2007; Aruoba and Schorfheide 2011; Cogley and Sargent 2015; Mertens 2016, for example):

$$\pi_t^* = \pi_{t-1}^* + \eta_t^{\pi^*}, \tag{14}$$

with $\eta_t^{\pi^*} \sim N(0, \sigma_{\eta^{\pi^*}}^2)$. We choose a random walk specification also because our sample includes the 1970s, which likely has associated a level of trend inflation much higher than what is implied by the readings of inflation in the last three decades. Furthermore, in a

Morley, Nelson, and Zivot 2003; Basistha and Nelson 2007, for example). González-Astudillo and Roberts (2021) allow for correlated disturbances in a similar setting and find that, even though the correlation coefficient is statistically significant, the results are broadly similar with respect to a model in which there is no correlation.

[6]Ascari and Sbordone (2014) show that when the inflation trend does not revert to zero in the long run, as is the case in this paper, the New Keynesian Phillips curve for inflation deviations from a nonzero steady state does not have the simple form $\hat{\pi}_t = \beta\mathbb{E}_t\hat{\pi}_{t+1} + \kappa\hat{mc}_t$, where $\hat{mc}_t$ are the firm's marginal costs, that we implicitly assume in this paper, but a more general form in which the coefficients vary over time as a function of trend inflation and an additional term that describes the discounted value of future marginal costs. We will nonetheless use our inflation setup, as it is a rather common one in the relevant literature, while including time-varying coefficients and a more sophisticated structure would significantly complicate the estimation of our model. We leave the time-varying coefficients approach to future research.

similar fashion to Del Negro, Giannoni, and Schorfheide (2015) and Bauer and Rudebusch (2020), we use information on 10-year-ahead inflation expectations, denoted as π_t^e, to pin down the inflation trend by assuming the following:

$$\pi_t^e = \pi_t^* + e_t, \tag{15}$$

with $e_t \sim N(0, \sigma_e^2)$. This specification explicitly assumes that survey long-run inflation expectations are an unbiased estimate of the inflation trend.

2.4 The Role of Monetary Policy in the Model

The identification of r_t^* in our model setup relies on the feedback from monetary policy to economic activity and vice versa. On the one hand, because of the correlated disturbances between the shocks to the federal funds rate and the 10-year Treasury yield's cyclical component—and because of the feedback from the latter to the output gap—a conventional monetary policy shock has an effect on output, the unemployment rate, and inflation. In particular, a positive correlation coefficient between these two aforementioned shocks implies that, ceteris paribus, an unexpected increase in the federal funds rate reduces output and inflation, and increases the unemployment rate, under the right configuration of parameter signs.

On the other hand, an unconventional monetary policy shock in our model—such as forward guidance or asset purchases by the Federal Reserve—would show, at least partially, through a change in the cyclical component of the 10-year Treasury yield which, in turn, will affect output, the unemployment rate, and inflation through its effect on the output gap. In addition to these explicit features of the model regarding the effects of monetary policy, by setting $\alpha^\pi > 1$, we impose the Taylor principle in our policy rule, which implies that the estimate of r_t^* is implicitly informed by changes in the federal funds rate that already have inflation and output stabilization features.

We would like to conclude the presentation of our model by noting that we do not see our setup as a simple extension of LW and HLW. Importantly, rather than identifying r_t^* by explicitly linking it to the trend growth rate of potential output, we instead rely on information from observed interest rates to identify their common real trend. As shown in Fiorentini et al. (2018), this alternative environment strengthens identification and prevents the possibility of a lack of observability, hence significantly reducing filtering uncertainty.

3. Data

We use data on real GDP, the civilian unemployment rate, the PCE price deflator inflation excluding food and energy, the effective federal funds rate, the 10-year Treasury constant maturity rate, and the 10-year-ahead PCE price deflator inflation expectations used in the FRB/US model (available as "PTR" in the public FRB/US package, a mnemonic that we will use henceforth).[7] All the variables come from the Federal Reserve Economic Data (FRED) database of the Federal Reserve Bank of St. Louis, except PTR, which comes from the publicly available FRB/US data set. In the estimation of the baseline specification described in the previous section, we use a sample that covers the period 1962:Q1 to 2020:Q1, except for the federal funds rate, for which we use a sample that starts in 1987:Q3.[8] When we assess the COVID-19 pandemic period with our model in Section 6.3, we extend the sample through 2023:Q1. Appendix C details the data used.

4. Estimation

We estimate the model with Bayesian methods. The Gibbs sampler alternates sampling between coefficients and latent states. The

[7]The FRB/US model is a large-scale estimated general equilibrium model of the U.S. economy that has been in use at the Federal Reserve Board since 1996. The model is designed for detailed analysis of monetary and fiscal policies. More details can be found at the following webpage: https://www.federalreserve.gov/econres/us-models-about.htm.

[8]Cúrdia et al. (2015) suggest using data from 1987:Q3 because this period coincides with the date on which Alan Greenspan became Chairman of the Federal Reserve, and monetary policy is generally viewed as having been relatively stable and consistent over time since then, and well-approximated by an interest rate rule. In our state-space model, we assume missing data on the federal funds rate prior to 1987:Q3.

explicit modeling of the ELB in the specification of the monetary policy rule implies that our state-space model is partially nonlinear. To deal with that situation, we embed the Bayesian estimation of Tobit models proposed by Chib (1992), called data augmentation, within the Gibbs sampler.

Broadly speaking, the procedure is as follows: First, given the (censored) data and initial latent states and parameters, we simulate the shadow rate, R_t, for the censored part of the sample from a truncated (from above) normal distribution with mean given by $\rho R_{t-1} + (1-\rho)\left(r_t^* + \pi_t^* + \alpha^\pi\left(\pi_t - \pi_t^*\right) + \alpha^y c_t\right)$ and variance $\sigma^2_{\eta^R}$. This is the data augmentation step suggested by Chib (1995). Second, we use the set of augmented data and obtain simulated states using the Durbin and Koopman (2002) simulation smoother from the state-space model. By construction, the sampled states deliver a shadow rate below the ELB. Third, with the sampled states, we obtain draws of the parameters of the model using the conventional independent normal-inverse-gamma posterior scheme, including for the equation of the shadow interest rate. Finally, with the newly sampled parameters and states, we simulate the shadow rate as indicated before and repeat the steps.[9] Appendix D describes the sampler in more detail. The choice of prior distributions appears in Appendix E.[10]

All told, following a burning-in set of 100,000 draws, we sample 200,000 observations, which, after thinning every 100th draw, gives us 2,000 draws to approximate the posterior distribution. The results have been checked for convergence and absence of autocorrelation of the posterior draws.

5. Estimation Results and Analysis

In this section, we present and discuss the estimation results of our benchmark model.

[9]Monte Carlo simulations confirm that this procedure produces unbiased trajectories of the latent variables.

[10]The inclusion of a Tobit step in our sampler is theoretically equivalent to but more efficient than the rejection sampling approach originally proposed by Johannsen and Mertens (2021). Carriero et al. (2023) also propose a similar sampler to generate the censored values.

5.1 Parameter Estimates

Key statistics of the model parameters' posterior distribution appear in columns 3 and 4 of Table E.1 in Appendix E. The posterior mean estimates of the cycle imply that it is highly persistent with hump-shaped dynamics. The Okun's law coefficients indicate a quantitative relationship between the output and unemployment gaps that is slightly less than the conventional 2-to-1 scaling. The Phillips curve coefficients imply a somewhat weak link between actual inflation and its trend, and a slope (with respect to the output gap) with a 68 percent credible interval between 0.05 and 0.09, which indicates a relatively weak response of inflation to the output gap compared with historical estimates, as documented by Blanchard (2016).

The posterior mean estimates of the monetary policy rule coefficients imply a relatively high degree of persistence in the rule—although, at 0.7, lower than the usual persistence coefficient of 0.85 (see Board of Governors 2018)—and sensitivities to inflation and the output gap that are consistent with the literature (and that obey the Taylor principle). For the r_t^* process, the estimate of the variance of its perturbation implies a standard deviation around 0.28 percent, in the vicinity of the estimate in Kiley (2020).[11]

The IS curve coefficients that link the cyclical components of the long-term interest rate and output have the expected overall negative sign. They imply a long-run sensitivity of the output gap to the interest rate gap around –7.8. According to the exercises and calculations presented in Roberts (2018), the magnitude of our model's response to changes in interest rate conditions lies between those of macroeconomic models that are usually considered as having lower interest rate elasticity (e.g., the FRB/US model) and those with higher interest rate elasticity like standard dynamic stochastic general equilibrium (DSGE) models (e.g., Smets and Wouters 2007).

The correlation coefficient between output growth and interest rate trend shocks is 0.50, with a 68 percent credibility interval between 0.21 and 0.75, which includes the implied estimate from

[11]Kiley (2020) points out that the data provide little information to estimate the variance of the r_t^* shock in his version of the LW model. We find that the posterior distribution of this parameter is significantly different than its prior, as can be seen in Appendix F.

Figure 1. Results of the Baseline Model

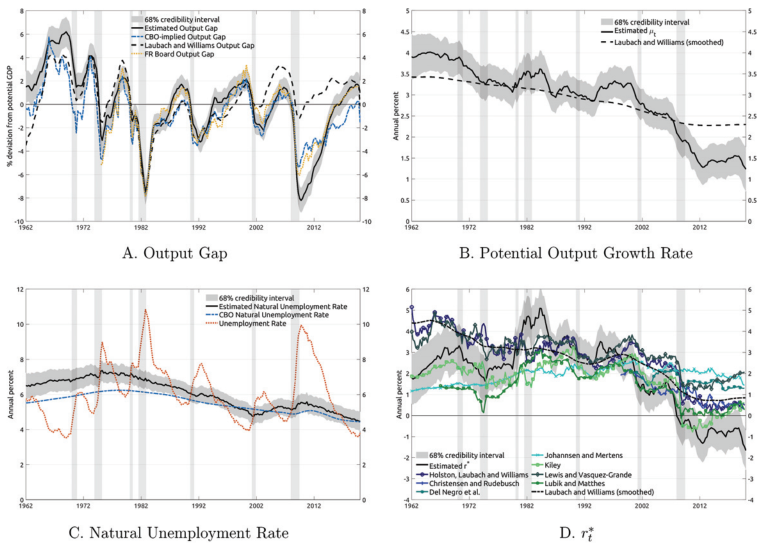

A. Output Gap

B. Potential Output Growth Rate

C. Natural Unemployment Rate

D. r_t^*

Note: Shaded vertical areas indicate NBER recession periods. Smoothed estimates are reported, except for the r_t^* estimates of other studies in the bottom right panel, which are the filtered estimates.

HLW. Finally, the correlation coefficient between the shocks to the shadow interest rate and the cyclical component of the 10-year Treasury yield has a posterior mean equal to 0.05 with a credible set that includes zero with 68 percent probability.[12]

5.2 Latent Factors Estimates

The results of the estimation with regard to the output gap, the growth rate of potential output, the natural unemployment rate, and r_t^* appear in Figure 1. Our estimate of the output gap in Figure 1A

[12]With this parameter configuration, an unexpected increase of 1 percentage point in the shadow interest rate—keeping all the other elements of the rule constant—causes a decline in the cyclical component of GDP of about 0.15 percentage point at the trough and a decline of 3 basis points in inflation. Impulse response functions after shocks to the cyclical component of the 10-year Treasury yield and to the output gap appear in Appendix H.

resembles those of the CBO—which is implied from their calculation of potential output—and the staff of the Board of Governors of the Federal Reserve System. Our estimate declines during NBER recession periods, but the magnitudes of the peaks and troughs can occasionally differ. For example, both the Board's staff and the CBO estimated an output gap around –6 percent during the Great Recession, whereas our estimate is close to –8 percent. Nonetheless, these three estimates imply sweeping output losses relative to its potential. In contrast, the output gap from LW casts the Great Recession as a relatively shallow one. At the end of the sample, the available estimates for the CBO and LW have turned negative whereas our posterior mean estimate has fallen by almost a full percentage point, but remains in positive territory.

Our estimate of the potential output growth rate, shown in Figure 1B, has declined over the sample period, just as that of LW. However, our estimate initiates a decline toward the end of the 1990s that is more pronounced than shown by their estimate. Our estimate stabilizes around 1.4 percent after 2012, about 0.9 percentage point below that of LW. The inclusion of data through 2020:Q1 also results in our smoothed estimate ticking down toward the end of our sample.

The natural unemployment rate estimate in Figure 1C shows some variation over time, fluctuating between 4.5 percent at the end of the sample and 7 percent during the 1970s; our estimate reached 5.6 percent during the Great Recession. We compare our measure with that from the CBO, which is lower in general throughout the sample. In 2020:Q1, the CBO estimate stands at 4.3 percent, within the 68 percent credible interval of our model, which covers the range 3.9–5.0 percent.

Finally, Figure 1D depicts our smoothed estimate of r_t^* along with filtered estimates of other models in the literature and the smoothed estimate from LW. From the plot, it is apparent that in the period 1962–82, the estimates that closely follow the approach of LW—in which r_t^* is explicitly linked to the growth rate of potential output (LW, HLW, and Lewis and Vazquez-Grande 2019)—are markedly above those that do not follow it (among those, our estimate). Higher-than-average economic growth during the 1960s and 1970s entails a similar pattern for the trend output growth rate, which, in turn, is more likely to hold for the equilibrium interest rate, unless the link is diminished through the contribution of the

Figure 2. Estimates Related to the Short-Term Nominal Interest Rate

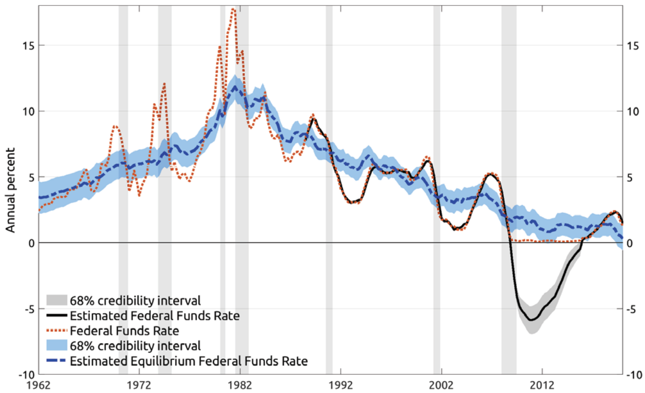

Note: In our estimation, data on the federal funds rate are treated as missing prior to 1987:Q3.

nongrowth component and at a price, statistically speaking. Our model, which only imposes a relationship between these two variables through correlated error terms, shows that the data prefer somewhat diverging patterns for the two trends over the first two decades of the sample. Our results suggest that the addition of a non-growth rate component, as in LW and Lewis and Vasquez-Grande, does not adequately account for the divergence implied by the data. Later in the sample, all the estimates in the existing literature trend down and have roughly stabilized in the last several years of our sample; they range between 0 percent and a bit above 2 percent in 2020:Q1. In contrast, our estimate shows a more pronounced downward trend that has put its 68 percent credibility interval in negative territory in recent years; our estimate of r_t^* is -1.7 percent at the end of the sample. To the best of our knowledge, only four estimates in the literature reach negative territory: that in Kiley (2020) does so after the Great Recession and (not shown in the figure) those of Brand and Mazelis (2019), Lopez-Salido et al. (2020), and Williams, Abdih, and Kopp (2020).

The estimated shadow-trend (or equilibrium) and federal funds rates (whenever the ELB binds) are shown in Figure 2 (recall that the shadow-trend or equilibrium federal funds rate is given by $r_t^* + \pi_t^*$). Starting in 1987:Q3, when data on the federal funds rate enter the model, the equilibrium federal funds rate is shown to be smoother than its observed counterpart, with the former above the latter in the later stage of expansions and below during or immediately after recessions. Also, the estimate of the shadow interest rate reaches –5.9 percent at the trough of the Great Recession. The decline in the equilibrium federal funds rate accelerates a bit at the end of our sample, and its estimate is slightly above zero with a 68 percent credibility interval between –0.6 percent and 1.3 percent at the beginning of 2020.

5.3 Why Has the Estimate of r_t^* Been Negative since the Great Recession?

Our estimate of the natural rate of interest is negative during the Great Recession and since then. This result is in contrast with most of the alternative estimates from the literature shown in Figure 1D. This difference seen in our results, as well as in the handful of studies mentioned in the previous section, warrants the question: What aspects of the data and models' structures drive the natural rate of interest negative around the 2008–09 recession and keep it below zero thereafter?

A feature common to all the studies that have estimated a negative r_t^* in recent decades is that, in contrast to LW and HLW, the Phillips curve used in the estimation assumes that current inflation is anchored to its trend and, more importantly, the latter is approximated with some measure of long-term inflation expectations. For instance, Lopez-Salido et al. (2020) use the Consensus Economics 10-year-ahead CPI inflation forecast extended back to 1961:Q2 by Blanchard, Cerutti, and Summers (2015). Kiley (2020) and Williams, Abdih, and Kopp (2020) use survey measures of long-run inflation expectations, as we do in this paper. Lastly, Brand and Mazelis (2019) use an inflation trend equal to 2 percent after the early 1990s; it is also the value of the inflation target in their Taylor rule.

Figure 3 shows the contributions of the observed variables, grouped in three categories, to the path of the estimated natural rate of interest.[13] Its examination suggests that the fluctuations in our

[13]This figure details the results of a historical data decomposition—i.e., a calculation of the contribution of each observed variable to the latent variables of the

Figure 3. Historical Data Decomposition of the Estimate of r_t^*

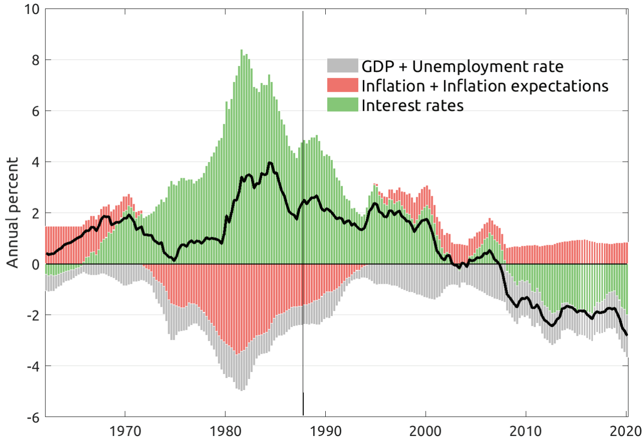

Note: The contributions of GDP and the unemployment rate (GDP + Unemployment rate) have been added together. The same is true for the inflation rate and PTR (Inflation + Inflation expectations), and for the federal funds rate and the 10-year Treasury yield (Interest rates). The gray vertical line indicates the period from which information on the federal funds rate was added to the system.

estimate of r_t^* are primarily coming from interest rate fluctuations (in particular, the 10-year Treasury yield), suggesting the importance of including these series in the information set and assuming roles for them in the economic model. We also observe that the substantial interest rate rise in the late 1970s and early 1980s only translated into a moderate rise in the natural real rate of interest, as these upward movements were in large part offset by similar increases in actual and expected inflation (as seen through their negative contributions). The decomposition shows that the gradual decline in the natural rate that began around the new millennium is primarily

transition equations of the model's state-space system. This kind of decomposition was proposed, building on the original work of Koopman and Harvey (2003), by Sander (2013) and Andrle (2013); these papers explain how to compute its elements by exploiting the linear structure of the model, as each observable variable has an independent effect on the smoothed estimates of a latent variable. We refer the readers interested in the more technical aspects of the decomposition to these papers as well as Chung et al. (2021). Notice that these results are obtained using the posterior mean of the parameters.

explained by the decline in interest rates (in particular, the 10-year Treasury yield), with a small offset from (low) inflation rates (as seen through their positive contribution).[14]

These results are consistent with the mechanism presented in Lopez-Salido et al. (2020) explaining why a negative inflation gap can contribute to a lower-than-otherwise-estimated r_t^*: All else equal, a lower inflation gap requires a lower output gap because of the link enforced by the Phillips curve.[15] In turn, our version of the IS curve equation compels a decline in the natural rate of interest to push up the interest rate gap for a given observed long-term real interest rate to account for the lower output gap on the left-hand side of the equation. Under these considerations, the key role played by the 10-year Treasury yield data in driving the dynamics of the natural interest rate is not really surprising, as our notion of the interest rate gap is defined for said long rate. It is worth noting that the policy rule in our model works as a counterweight to the aforementioned mechanism (which is absent from Lopez-Salido et al. 2020) as, everything else equal, negative inflation and output gaps compel an upward revision to the estimate of r_t^* in the same period.[16]

5.4 The Role of Particular Structures of the Model

The semi-structural model approach of this paper has both benefits and shortcomings. On the one hand, it provides flexibility in fitting the data and allows the modelers to choose selectively the economic relationships that will be used to impose structures on the data. On the other hand, the specific nature and validity of and motivations underlying choices are not always easy to establish and agree upon (e.g., one's preferred choice may be called "ad hoc" by another). Moreover, the imposition of economic relationships may still partially rely on reduced-form dynamics. For instance, we have allowed

[14] Appendix G shows the historical data decomposition for the output gap.

[15] The inflation gap is defined as actual inflation minus long-run expectations of inflation (PTR), which is negative on average during and following the Great Recession, as seen in Appendix G.

[16] The relationship between r_t^* and the gaps arising from the rule is not as straightforward as that from the channel highlighted by Lopez-Salido et al. (2020) because it is not the level but a quasi difference of the former (i.e., $(1-\rho)r_t^*$) that is a function of the latter.

Figure 4. Estimated r_t^* Comparison across Different Parameter Assumptions

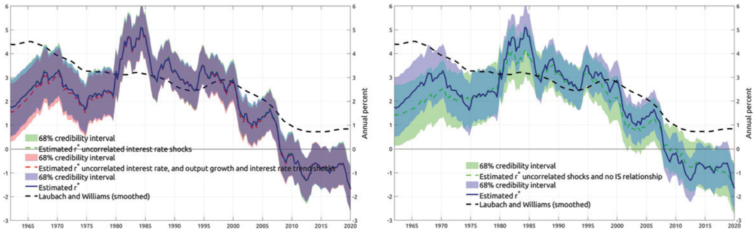

A. Correlation Coefficients B. IS Relationship

for correlation between the shocks to r_t^* and trend output growth, μ_t, in order to link r_t^* to factors such as productivity or population growth. Similarly, we introduce a role for conventional monetary policy by assuming that the shocks to the transitory components of R_t and i_t^{10} are correlated, which is a reduced-form substitute for an explicit modeling of the expectational component of the long-term interest rate. Perhaps more importantly, we introduce an IS-type relationship in which the cyclical component of the long-term real interest rate affects the cyclical component of output. How does each of these features affect the estimate of r_t^* and how does a model fit comparison discriminate among them?

Figure 4A shows a three-way comparison in which our baseline model estimate of r_t^* is contrasted against two other estimates in which the correlation coefficients we previously mentioned are set to zero. As it can be seen, these two restrictions on the correlations of the innovations have negligible effects on the estimated path of r_t^* compared with the baseline path. However, a marginal likelihood comparison across the three specifications indicates that the data strongly prefer a model without correlated interest rate disturbances, but with correlation between output growth and real interest rate trends.[17]

Figure 4B shows a comparison between the estimate of r_t^* obtained with our baseline specification and that of a model that

[17]The baseline model achieves a marginal data density equal to −652.9, the model without correlation between η_t^R and ε_t^{10}, one equal to −631.6, and −677.5 for the model without correlation between η_t^R and ε_t^{10}, and between η_t^μ and $\eta_t^{r^*}$.

Figure 5. Results of the Uncensored Model

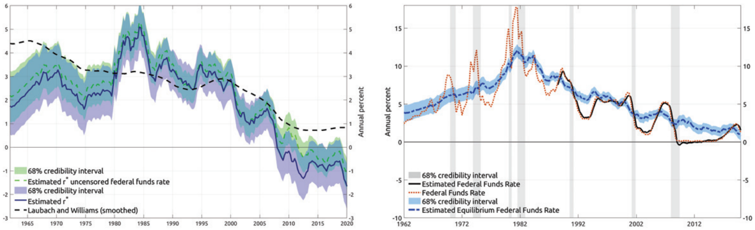

A. r_t^*

B. Observed and Equilibrium Federal Funds Rate

Note: Shaded vertical areas indicate NBER recession periods. Smoothed estimates are reported.

assumes zero correlation between the aforementioned pairs of shocks, like the red line in Figure 4A, as well as the absence of an IS relationship. (We refer to this latter model as the plain model.) Judging by the overlapping of the confidence sets, it is likely that the two estimates of r_t^* may not be different between these two specifications. However, it is noticeable that, on average, the estimate of the baseline model is higher than that of the plain one before the onset of the GFC, whereas the former is lower than the latter after 2008. Moreover, a comparison of the marginal data densities indicates that the data strongly prefer the model with an IS-type relationship.

5.5 The Role of Censoring

In our model specification, we incorporate the fact that the federal funds rate was censored from below in the aftermath of the GFC. As it could be easily foreseeable, ignoring censoring will in all likelihood distort the estimates of the policy rule, including the values of r_t^*. We find that while all the parameters of the model experience changes when we do not incorporate censoring, the reaction of the federal funds rate to the output gap changes substantially, reduced to half of the original estimated coefficient. The rule also becomes more persistent and the shocks are more volatile. Also as expected, the estimate of r_t^* is higher in the specification that ignores censoring of the federal funds rate, as can be seen in Figure 5, which also shows the neutral federal funds rate in this case.

Despite the estimate of r_t^* being within the confidence set of the estimate that considers censoring for most of the sample period, the two estimates likely differ shortly after the onset of the GFC, with the former (ignoring censoring) higher and in positive territory compared with the latter (that takes censoring into account); the difference averages about 1 percentage point between 2011 and 2013. In addition, notice that the estimated neutral federal funds rate is higher than that shown in Figure 2. These results suggest that estimating the natural rate of interest using real interest rate data computed with information from the (censored) federal funds rate, as in LW, would likely overestimate r_t^*.

6. Model Evaluations

In this section, we conduct a formal comparison of our model with other model specifications to gauge what features make our setup beneficial. To that end, we evaluate the out-of-sample forecasting capabilities of each alternative specification because the alternatives do not always include the same set of observable variables, making the marginal data density approach more cumbersome.

6.1 Two Alternative Model Specifications

The approach used so far to treat the output gap as we do any other latent variables, i.e., determined jointly by the model and data, provides flexibility and delivers estimates that reflect the structures of the model from a probabilistic perspective. However, estimation results may, as with any model, be distorted by misspecifications. It is reasonable to assume that differences in the path of the output gap would lead to different dynamics in the natural rate of interest given the tight connections between the output and interest rate gaps implied by our model's IS curve and policy rule. To explore this issue, we estimate a version of the model that includes the output gap series derived from the CBO's estimate of potential output in the data set, in a fashion similar to Johannsen and Mertens (2021). The CBO's estimate of the output gap is a well-recognized measure that is not only model-based but also calculated using a wider set of information than ours as well as economic judgment.

Figure 6. Estimated r_t^* under Alternative Specifications

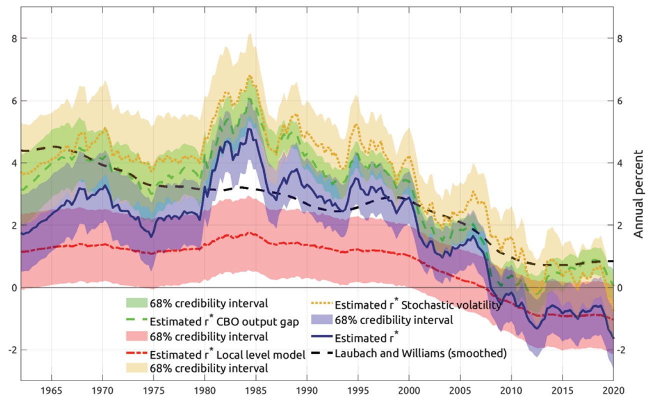

The estimated r_t^* associated with this model specification appears in Figure 6 along with that from our baseline model. Conditioning the estimation of the model on the CBO measure of the output gap yields an estimate of r_t^* about 1 percentage point higher than in our baseline specification, on average. Most of this difference reflects a similar shift in the measures of the output gaps (our estimate of the output gap is about 0.8 percentage point higher, on average, than that of the CBO, as can be seen in Figure 1A). In particular, the estimate of r_t^* using the CBO output gap is positive during the last several years of the sample, averaging about 0.6 percent since 2015. In addition, the posterior mean estimates of the Taylor rule coefficients are close to the upper bound of the 68 percent credible sets of the original specification that did not include the CBO output gap.

Furthermore, the CBO estimate of the output gap entails a more volatile natural rate of unemployment (about four times the volatility of the baseline). For instance, it increases above 8 percent (compared to 6 percent in the baseline) during the Great Recession, whereas it settles at 3.5 percent at the end of our sample, 1 percentage point below the estimate of our baseline model. Finally, the long-run sensitivity of output to the interest rate (given by the IS relationship) is half that of the baseline specification. These results evidence that replacing the output gap estimation is not innocuous in our model.

The other specification we examine has been proposed by Fiorentini et al. (2018): they show how the precision of the estimates of r_t^* in the HLW model deteriorates greatly when the IS and Phillips curves are close to being flat. They propose to estimate r_t^* using information on the ex post real interest rate only, specifying a local-level model in which the trend, r_t^*, follows a unit root and deviations of the real interest rate from the trend are stationary. We implement their proposal by replacing the policy rule specification of our baseline model with the following equation:

$$R_t = r_t^* + \pi_t^* + c_t^R,$$

where c_t^R is a stationary AR(2) process and R_t is still uncensored.[18] In this way, we can assess the effect of assuming a policy rule specification for the federal funds rate on our estimate of r_t^*.

This setup is similar to that in which one obtains the real interest rate by subtracting a measure of long-run inflation expectations from the federal funds rate and uses that information in the local-level model, in the spirit of Lopez-Salido et al. (2020). However, in both Fiorentini et al. (2018) and Lopez-Salido et al. (2020) as well as in HLW, the real interest rate is obtained from a censored nominal interest rate, which could be analogous to having used the specification in this paper that ignores censoring and that delivered a higher estimated r_t^* than when censoring was taken on board—our baseline specification.

The results in Figure 6 indicate a much lower estimate of r_t^* than the baseline specification in this case, except during the last 10 years of the sample, in which both of them average a level close to –0.8 percent. Of note, the r_t^* estimate fluctuates between 1 percent and 2 percent before 2000, when it starts to decline and becomes negative at the same time as our baseline estimate, at the onset of the GFC. In addition to a consistent downward shift in the level of the series, the estimate of r_t^* under the local-level specification also displays a much smoother path compared to that of the baseline model. The characterization of the cyclical dimension of the real

[18]We continue to assume that the shocks to transitory components of both interest rates are correlated as well as the shocks to the output trend growth and r_t^*.

short-term interest rate, c_t^R, as a latent variable, unattached to the rest of the model's variables and explained solely by a single stochastic shock, appears to give the model ample leeway to capture most of the cyclicality observed in the real rate. In contrast, the policy rule in our baseline specification ties its cyclical component to macroeconomic factors (i.e., the inflation and output gaps), linking the fluctuations of the real short-term interest rate at business frequencies to those of other key determinants of the economy. Which specification is preferable? The next section attempts to shed some light on this question by performing a model evaluation exercise.

6.2 Pseudo-out-of-Sample Forecasting Exercises

In order to broadly evaluate the model specifications shown so far, we estimate them and generate projections in a pseudo-real-time forecasting environment. More precisely, we begin with the initial sample spanning the period 1962:Q1 through 2002:Q3, estimate the models and, jumping off from the last quarter of the aforementioned sample, produce one- to four-quarter-ahead forecasts for all the observable variables, using every draw from the posterior distribution of the parameters. We then roll forward the sample by adding one quarter at a time and reestimate the model, producing forecasts of all the observable variables for every posterior draw once again. We continue adding one period at a time until 2019:Q1 to produce the forecasts one to four quarters ahead. Table 1 shows the continuous ranked probability scores (see Gneiting and Raftery 2007) of the one- and four-quarter-ahead forecasts for the unemployment rate, the inflation rate, and the federal funds rate.[19]

The results show that, broadly speaking, out of sample the baseline specification (line 1) outperforms the alternatives considered so far in the paper (lines 2–5 and 8–11). For instance, the model that ignores censoring (lines 2 and 8) forecasts the unemployment and federal funds rates worse than our original specification. The model that uses the CBO output gap (lines 4 and 10) is able to forecast inflation better than our baseline model, but its performance worsens with respect to the unemployment and federal funds rates. Finally,

[19] We consider 1,000 draws from the posterior distribution after burning in 6,000 draws and thinning every 12th draw; that is, we use a total of 18,000 draws.

Table 1. Continuous Ranked Probability Scores

Line	Quarters Ahead	Model	Unemployment	Inflation	Federal Funds
1	One	Baseline	0.13	0.39	0.23
2		Without Shadow Rate	0.14	0.41	0.28
3		Without FFR	0.14	0.58	
4		With CBO Output Gap	0.15	0.36	0.25
5		With Local-Level FFR	0.14	0.45	0.55
6		Baseline with SV	0.12	0.37	0.25
7	Four	Baseline	0.48	0.37	0.41
8		Without Shadow Rate	0.55	0.38	0.50
9		Without FFR	0.54	0.52	
10		With CBO Output Gap	0.51	0.36	0.44
11		With Local-Level FFR	0.52	0.40	0.63
12		Baseline with SV	0.44	0.33	0.44

Note: "FFR" denotes "federal funds rate." "SV" denotes "stochastic volatility." The evaluation window of the forecasts starts in 2002:Q4 and ends in 2020:Q1.

the model that omits a policy rule specification for the federal funds rate—and uses a local-level model in its place—(lines 5 and 11) is overall worse than the baseline specification, and the worst among the alternatives to forecast the federal funds rate.

We also investigate the out-of-sample forecasting performance of a model that does not include interest rates, either short or long, as observable variables (lines 3 and 9). The results show that the ability of the model to forecast inflation deteriorates significantly compared with that of the models that do include an interest rate block.[20]

6.3 A Model with Stochastic Volatility: Parsing the COVID-19 Pandemic Period

Our sample includes episodes of high inflation, output, and interest rate volatilities that could influence how the Durbin and Koopman (2002) simulation smoother parses the information of the data to obtain estimates of the parameters and latent variables of our model, including r_t^*. Up to now, we have assumed a constant variance in the innovations, as it facilitates the comparison with models from the existing literature and allows us to disentangle more easily the role that each assumption of our model specifications plays in our estimate of r_t^*.

However, a growing number of recent additions to this literature have rejected a homoskedastic specification for one that allows for time variation in the variances of the innovations, usually with a stochastic volatility (SV) setup (see Johannsen and Mertens 2021 and Zaman 2021 for a few examples). We now explore the implications of allowing for SV on key aspects of the model's inferences and estimates.

To account for the possibility of time-varying volatility, we specify the variance of each error term in the model as follows:

$$\sigma_t^2 = \exp(h_t),$$
$$h_t = h_{t-1} + \eta_t^h, \quad \eta_t^h \sim \text{i.i.d. } N(0, \sigma_{\eta^h}^2),$$
$$h_0 \sim N(\mu_0, \sigma_0^2).$$

[20]In Appendix J we show real-time estimates of the output gap and r_t^* for the models that include an interest rate block. The results show that the real-time estimate of the baseline model is reasonably close to its smoothed counterpart.

We use the mixture simulator proposed by Kim, Shephard, and Chib (1998) to estimate the parameters and latent states of the model with SV, using the same sample information as in the previous sections.[21]

Figure 6 shows that modeling SV leads to an upward shift in the estimate of r_t^*. One can observe that the magnitude of the difference between the estimates with and without SV is starker during episodes of higher inflation. There is a well-recognized literature (see Stock and Watson 2007, for example) on the estimation of the processes underlying inflation with SV and how the estimates of the variances of the innovations are substantially larger during the inflationary episodes spanning the late 1970s and early 1980s. The nature of our filtering procedure entails that observations during these episodes are given less weight (the signal-to-noise ratio is smaller during these episodes due to the larger variances) than under a structure with smaller (and constant) estimated variances. Framed in terms of the data contributions presented in Section 5.3, the offset from the inflation data on the contributions from the rising 10-year Treasury yield in the late 1970s and early 1980s is now smaller compared with that of the model without time-varying volatility.[22] Consequently, the high inflation and interest rate episodes in the late 1970s and early 1980s are now consistent with higher levels of the r_t^* estimate. All in all, our estimate of r_t^* with SV stands close to −1 percent in early 2020 after hovering slightly above zero in the decade before.

Table 1 (lines 6 and 12) shows the forecasting performance of this baseline model with SV. The addition of SV helps predict the unemployment and inflation rates better than the baseline specification (which was the overall best specification so far), but it worsens the federal funds rate predictions. The differences are rather small

[21] We assume that μ_0 for each shock is equal to the log of the posterior mean estimate of the variance of the respective shock of the model without SV, and that $\sigma_0^2 = 1$ for all the shocks. The prior distribution for $\sigma_{\eta^h}^2$ in each shock is inverse-gamma with mean equal to one-hundredth the value of $\exp(\mu_0)$ and shape parameter equal to 3.

[22] The contribution from the nominal interest rates may also be more muted than in the baseline model during the more volatile episodes but, ultimately, what matters is the relative decline in the contributions, i.e., as long as the reduction from the contribution of the inflation data is larger (in absolute value) than that of the interest rates from allowing for time-varying volatilities, the changes in the estimate of r_t^* from rising rates will be larger.

and hence both versions perform similarly. Because neither of the two specifications (baseline with and without SV) dominates the other, according to this measure of performance, both can be considered equivalently valid representations of the data over the historical sample ending in 2020:Q1 from this perspective.

The recent COVID-19 pandemic, with the exceptional swings observed in key macroeconomic data, constitutes an episode for which simply relying on constant variances based on pre-pandemic samples is in all likelihood misguided. After all, the magnitude of these changes is the primary rationale behind ending the sample with the first quarter of 2020, and a feature that has motivated an all-new literature on how to deal in practice with these recent exceptional movements in the data. Carriero et al. (2022), Schorfheide and Song (2021), and Lenza and Primiceri (2022) are examples of such new literature.

Carriero et al. (2022) note (in a VAR context) that the popular specification of SV may not be entirely congruent with the uniqueness, magnitude, and short-livedness of the variations in the data during the pandemic. They evaluate the model fit of extensions of the SV specification such as an outlier-augmented SV setup (SVO hereafter) and find that it performs better than SV alone or an SV specification that treats the pandemic data as missing, according to in-sample and prediction metrics. However, both Schorfheide and Song (2021) and Lenza and Primiceri (2022) argue that the missing data approach, instead of the outliers treatment of the pandemic data, can also be a valid alternative in empirical work mainly because of its simplicity and because of its adequacy for either forecasting (in the former paper) or parameter estimation (in the latter).[23]

Because the aim of this section is not to determine the best model over some set of candidates, but to explore the latent variables estimates of our proposed model through the pandemic episode, we opt for the missing data approach, using the model specification with SV, as its implementation is straightforward and its computational

[23]Lenza and Primiceri (2022) assume that the pandemic induced a common shift in volatility in a constant-variance Bayesian VAR instead of having one SV process for each of the perturbations of the model, which is a somewhat standard assumption in the literature. Carriero et al. (2022) find that making the outlier common to all the series seems to provide no advantage and that said approach only registers outliers during the COVID-19 pandemic.

cost, marginal. We take advantage of the state-space representation of our model, as the Kalman filter allows us to account for the missing values using those implied by the (random draws of the) dynamics of the model itself.[24]

To parse the pandemic sample with our model, we first start with the baseline specification with SV estimated over the (pre-pandemic) sample previously described and its posterior draws of parameters. Next, to identify the dates to treat as missing data, we assume that any observation that is beyond a threshold factor of 10 of the interquartile range is an outlier, meaning that only real GDP observations and the unemployment rate for 2020:Q2–Q3 are flagged for omission. The pandemic observations for the other data series are not sufficiently unusual at a quarterly frequency to be discarded.[25] Finally, for each set of parameters drawn from the posterior distribution, we hold them fixed and draw the model's latent variables, including the SV processes, with the Durbin and Koopman (2002) simulation smoother and the Kim, Shephard, and Chib (1998) mixture simulator sequentially, using data from 1962:Q1 to 2023:Q1. In that process, we make sure the latent states satisfy the ELB for the federal funds rate starting in the second quarter of 2020 through the fourth quarter of 2021, i.e., that the shadow federal funds rate is below the ELB.[26]

Figure 7A shows that the model's output gap reaches –2.5 percent at the onset of the pandemic in 2020:Q2 and bottoms out at –4.5 percent in the second half of 2020, a much less drastic decline compared with that of the CBO for instance, which falls to as much

---

[24]One disadvantage of the missing observations approach compared with the SVO specification is that the former is unable to account for the possibility of future outliers. However, as our objective is not forecasting but instead parsing the data, the advantages of the SVO approach over the missing data one may not be as significant.

[25]Carriero et al. (2022) carry out their analysis at the threshold factors of 5 and 10 and eventual settle for the former. We picked a factor of 10 rather than 5 because the latter entailed excluding the unemployment rate observation until the end of 2020, which seems at odds with the conventional appraisal of the data over the pandemic. The use of quarterly observations rather than monthly will likely bias toward overomitting information and, as a result, we selected the conservative factor of 10.

[26]The SV estimates for each of the perturbations of the model appear in Appendix I.

Figure 7. Parsing of the COVID-19 Pandemic Period

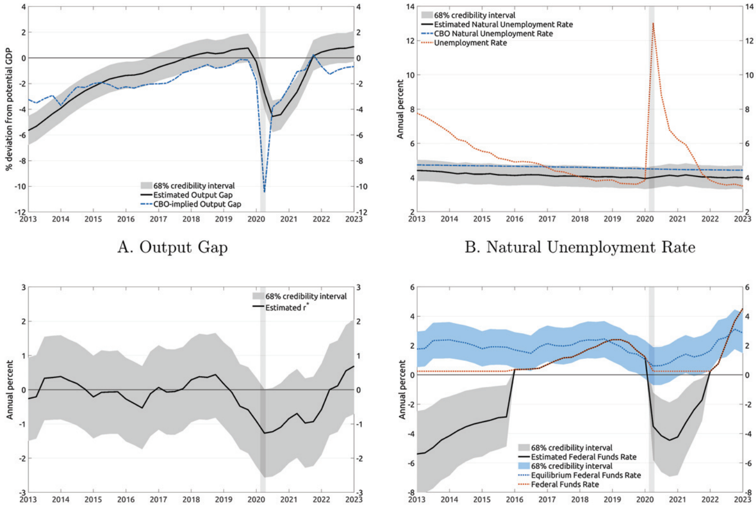

A. Output Gap

B. Natural Unemployment Rate

C. r_t^*

D. Shadow and Equilibirum Federal Funds Rates

Note: Shaded vertical areas indicate NBER recession periods. Smoothed estimates are reported.

as −11 percent. The model does not infer a spike in the natural unemployment rate either (Figure 7B) during these volatile periods. While these results may be expected given the removal of the data with the largest movements, they also indicate how the responses of the other observable variables to the factors corresponding to the pandemic have been unexceptional: none of the movements in inflation and interest rates in mid-2020 indicate a large decline in the cyclical position of the economy. We also notice that our output gap estimate rapidly aligns with that of the CBO in the recovery phase, but has diverged in the last year: our estimate indicates that output is almost 1 percent above potential in early 2023 whereas the CBO estimates that it is about 0.7 percent below.

Figure 7C shows that the estimate of r_t^* has rebounded from its negative level of about –1.25 percent at the onset of the pandemic to about 0.7 percent in early 2023, although the 68 percent credible set still includes zero, as this was the norm during the several years before the pandemic. As a result of the mild response of the output gap as the pandemic unfolds, the decline in the shadow rate is also relatively mild, reaching about –4.5 percent at the end of 2020 (Figure 7D). The equilibrium or trend federal funds rate hovered around 1 percent during the pandemic ELB episode and stands at 2.9 percent in early 2023, indicating that the current stance of monetary policy is contractionary.[27]

7. Conclusion

In this paper, we formulated and estimated a semi-structural model of the U.S. economy that provides measures of the natural rates of unemployment and interest, which can inform the decisions of monetary policymakers. Our model also provides an estimate of the output gap that is roughly consistent with institutional and judgmentally driven estimates, such as those produced by the CBO or the Federal Reserve Board's staff, in contrast to the estimates of LW and HLW.

We note that introducing censoring in the monetary policy rule lowers the estimate of r_t^* compared with a model in which censoring is ignored. This consideration also implies a lower neutral federal funds rate, which is a benchmark recommended by economic theory to evaluate the stance of monetary policy.

We also find that movements in the long-run interest and inflation rates are the most important contributors to the downward secular trend in our r_t^* estimate, especially since the Great Recession. Lastly, an estimation of the model incorporating stochastic volatility shows that r_t^* may have drifted significantly below zero during the COVID-19 pandemic and has increased to 0.7 percent in the recent past, above the pre-pandemic norm, which was close to zero.

[27] Appendix I shows the results for the case in which no observations are omitted, i.e., a straight read from our baseline model with SV. In that case, the response of the output gap fully reflects the swings in the data and output is estimated to have fallen about 20 percent below potential, almost twice as much as the CBO's estimate. Interestingly, and probably because of the quick rebound in real GDP and decline in the unemployment rate, the economic trends of interest in this paper (r_t^* and u_t^*) do not change much with respect to the results in Figure 7.

Appendix A. Literature Review

Laubach and Williams (2003) (LW hereafter) and, subsequently, Holston, Laubach, and Williams (2017) (HLW hereafter) are seminal works on the estimation of natural rates of interest for the U.S. economy and, in the latter case, other advanced economies. One key element of their identification strategy is the relationship between the growth rate of the economy and the real short-term interest rate implied by standard economic theory. Using information on output, the inflation rate, and the short-term interest rate, they document a downward trending estimate of r^*, which in the case of the U.S. economy eventually falls close to zero. Their estimates have become a staple in the economic and policy discussions of r^*, and updates are regularly made publicly available.[28] Nonetheless, numerous studies have sought to improve the LW methodology and estimates.

Lewis and Vazquez-Grande (2019), Beyer and Wieland (2019), Kiley (2020), and Brand and Mazelis (2019) are fairly recent examples of such work. For instance, all four papers use Bayesian methods rather than a multistep procedure likelihood-based estimator to address the pile-up problem that often afflicts classical estimation approaches.

Lewis and Vazquez-Grande (2019) also study the consequences of assuming that the nongrowth component of r^* is first-difference stationary (as in LW) rather than persistent but stationary. They argue that a mixture of permanent and transitory processes to characterize the natural rate of interest is preferable to the original specification of LW. Their estimate is more procyclical and displays less of a secular decline than the one shown in LW and HLW.

Beyer and Wieland (2019) argue that a large degree of uncertainty surrounds the estimates of LW and that their methodology and estimation methods are highly sensitive to the choice made by the econometrician. They note the challenge of simultaneously estimating many unobserved variables in a large state-space model. For instance, they find that the precision of the estimates does not increase even after adding more than one decade of data relative to the original set of LW, which ended in 2002.

[28]See the Federal Reserve Bank of New York webpage "Measuring the Natural Rate of Interest" at https://www.newyorkfed.org/research/policy/rstar.

Kiley (2020) also points out the weak identification of the natural rate of interest in the original LW setup. This observation motivates him to investigate possible ways to improve the identification of r^*. He proposes to add an Okun's law equation to the system and account for the role of additional demand shifters (e.g., asset prices, fiscal policy, and credit conditions) in the IS-curve equation. The addition of credit spreads is one factor that significantly helps with improving the identification of r^*. Following these changes, estimates of r^* are more stable over time and do not exhibit the same kind of gradual secular decline as shown by the LW estimates.

Brand and Mazelis (2019) estimate a semi-structural model of the U.S. economy featuring key elements of the LW model but also a Taylor-type policy rule to better identify r^*. Their estimate of the r^* process for the U.S. is far more volatile than that of LW and falls well below zero following the Great Recession. They do not, however, account explicitly for the presence of the ELB and assume instead that the observed short-term interest rate is what would have prevailed under their rule, even in the absence of the ELB.

The studies discussed so far in this section have adopted the definition of r^* from LW and mostly followed or investigated the robustness of the assumptions of their model. However, economists have also come up with different concepts and methodologies to characterize the stance of monetary policy.

For instance, Christensen and Rudebusch (2019) employ flexible dynamic term structure models and financial data (e.g., inflation-indexed debt) to obtain estimates of the real rate that prevail, on average, between the 5- to 10-year horizon window, once business fluctuations have mostly faded. Their framework allows them to compute an equilibrium rate without having to correctly specify the dynamics of the output gap and inflation. The results show that the natural rate of interest has gradually declined over the past two decades to a level close to zero.

Another paper that computes a longer-run (i.e., five-year horizon) measure of r^* under a flexible approach is Lubik and Matthes (2015). They estimate a time-varying vector autoregressive (TVP-VAR) model, which imposes much fewer theoretical restrictions than LW. Their measure of r^* is the five-year conditional forecast of the observed real rate implied by this model. Although using a different approach, Lubik and Matthes estimate a path of r^* that is roughly consistent with that of LW starting around the mid-1980s. Unsurprisingly, with few restrictions and time-varying coefficients, the degree of uncertainty around their estimates is relatively large.

Cúrdia et al. (2015) argue that policy rules responding to the efficient real interest rate characterize the evolution of the federal funds rate since late 1987 better than traditional monetary policy rules based on estimates of the output gap.[29] They refer to the former as Wicksellian policy rules. It is worth noting that the dynamics of their efficient interest rate—and hence their results—are highly dependent on the model specifications and underlying assumptions.

Del Negro et al. (2017) compare the measure of r^* computed from a low-frequency estimate of the short-term interest rate in a VAR model with common trends to the efficient interest rate in a version of the Federal Reserve Bank of New York DSGE model (see Del Negro, Giannoni, and Schorfheide 2015). The two methodologies deliver fairly consistent views regarding the gradual decline in the short-term real interest rate observed over the past few decades.

There are two papers that are most closely related to ours. The first is Johannsen and Mertens (2021). They propose a flexible time-series approach that decomposes their data as trends and cycles and explicitly accounts for the presence of the ELB by simulating a shadow rate for the periods when the ELB is binding. They also allow for stochastic volatility in the variance of some of the innovations. However, and in contrast to our methodology, they do not identify and infer the output gap based on the structure of their model and the data. Instead, they take the CBO estimate as observed values. The reliance on a reaction function in which the output gap is a significant determinant of the monetary policy rate entails strong identification linkages between the estimate of r^*, the shadow rate, and the output gap. In our paper, we seek to capture the simultaneous directionality of these influences as well as to take into account

[29]The efficient real interest rate in a DSGE model is that which would prevail in an economy in which prices are flexible and desired markups are zero.

the contribution of the uncertainty around the output gap estimate to the uncertainty surrounding the estimate of r_t^*.

The second paper is Zaman (2021). It is a comprehensive study of a semi-structural model of the U.S. economy that shares many features with our paper. The model is estimated by Bayesian methods, includes information from survey data (see below) and specifies the cyclical component of the short-term interest rate using a Taylor-like policy rule. The paper also allows for time variation in some of the parameter estimates, both in the variances of the innovations (like our paper and Johannsen and Mertens 2021) and some regression parameters. The main difference with our paper is that Zaman (2021) does not sample a model-consistent distribution of the shadow rate at the ELB but instead uses the series implied by the model of Wu and Xia (2016) as an observable variable. It is also worth noting that the reliance on survey data in the paper is far more extensive than ours, as its data set may go as far as including the long-run projections of the three-month Treasury bill, real output growth, the unemployment rate, and GDP deflator inflation. Lastly, and in contrast with Zaman's approach, we use information of the 10-year Treasury rate to directly inform our estimate of r_t^*.

Finally, our paper relates to a strand of literature that deals with the ELB and censoring of the federal funds rate in the estimation of dynamic models with structural identification. For instance, Mavroeidis (2021) and Aruoba et al. (2022) propose econometric strategies to account for the censoring of the policy rate below the ELB in the context of a (structural) VAR model. Both carry out their analyses with a canonical three-equation characterization of the U.S. economy. Mavroeidis relies on a maximum likelihood estimator while Aruoba et al. carry out their estimation with Bayesian methods. These papers focus on the identification and estimation of the dynamic coefficients of the econometric system that allow them to investigate possible magnitude changes in the response of conventional monetary policy at the ELB, as well as the effectiveness of unconventional monetary policy relative to its conventional counterpart. Consistent with Johannsen and Mertens (2021), the results indicate that the economy is in general more responsive to monetary policy stimulus when the ELB is binding than when it is not. In contrast, our primary interest is the estimation of key trends rather than the question of whether the dynamic of the response of the economy to a monetary policy shock may change when policy is constrained by the ELB. Another paper worth highlighting is Jones, Kulish, and Morley (2021), who perform their analysis with a full structural DSGE model. In particular, they assume full information rational expectation (FIRE), which likely drives their estimate of the shadow rate being actually above the ELB in the early periods of the 2008 financial crisis. FIRE is a very strong assumption and is valid only to the extent it reflects the data-generating process accurately. Being skeptical of the assumption and more cautious, our model does not embed any kind of structural foresight. In particular, there is no structural identification of explicit future innovations to monetary policy in our setup.

Appendix B. Model in State-Space Form

The benchmark model is as follows:

$$\begin{bmatrix} y_t \\ u_t \\ \pi_t \\ R_t \\ \pi_t^e \\ i_t^{10} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \theta_1 & \theta_2 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}$$

$$\times \begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ z_t \\ c_t \\ c_{t-1} \\ c_{t-2} \\ y_t^* \\ \mu_t \\ u_t^* \\ r_t^* \\ c_t^{10} \\ c_{t-1}^{10} \\ c_{t-2}^{10} \\ \pi_t^* \\ p_t^{10} \end{bmatrix} + \begin{bmatrix} 0 \\ \upsilon_t \\ 0 \\ 0 \\ e_t \\ 0 \end{bmatrix}, \tag{B.1}$$

$$
\begin{bmatrix}
1 & 0 & 0 & 0 & 0 & -\kappa & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\beta_1 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
-\frac{1}{4}(1-\rho)\alpha^{\pi} & -\frac{1}{4}(1-\rho)\alpha^{\pi} & -\frac{1}{4}(1-\rho)\alpha^{\pi} & -\frac{1}{4}(1-\rho)\alpha^{\pi} & 1 & -(1-\rho)\alpha^{y} & 0 & 0 & 0 & 0 & 0 & -(1-\rho) & 0 & 0 & 0 & (1-\rho)\alpha^{\pi} & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}
\begin{bmatrix}
x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ z_t \\ c_t \\ c_{t-1} \\ c_{t-2} \\ y^*_{t-1} \\ \mu_t \\ u^*_{t-1} \\ r^*_t \\ c^{10}_t \\ c^{10}_{t-1} \\ c^{10}_{t-2} \\ \pi^*_t \\ p^{10}_t
\end{bmatrix}
$$

$$
= \begin{bmatrix}
1-\beta_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & \rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & \phi_1 & \phi_2 & 0 & 0 & 0 & 0 & 0 & \lambda_1 & \lambda_2 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \psi_1 & \psi_2 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}
\begin{bmatrix}
x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ z_{t-1} \\ c_{t-1} \\ c_{t-2} \\ c_{t-3} \\ y^*_{t-1} \\ \mu_{t-1} \\ u^*_{t-1} \\ r^*_{t-1} \\ c^{10}_{t-1} \\ c^{10}_{t-2} \\ c^{10}_{t-3} \\ \pi^*_{t-1} \\ p^{10}_{t-1}
\end{bmatrix}
+
\begin{bmatrix}
\eta^{\pi}_t \\ 0 \\ 0 \\ 0 \\ \eta^R_t \\ \varepsilon_t \\ 0 \\ 0 \\ \eta^{y^*}_t \\ \eta^{\mu}_t \\ \eta^{u^*}_t \\ \eta^{r^*}_t \\ \varepsilon^{10}_t \\ 0 \\ 0 \\ \eta^{\pi^*}_t \\ \eta^{p^{10}}_t
\end{bmatrix},
\tag{B.2}
$$

where $\mathrm{corr}(\eta^{r*}_t, \eta^{\mu}_t) = \omega_{\eta^{r^*},\eta^{\mu}}$, and $\mathrm{corr}(\eta^R_t, \varepsilon^{10}_t) = \omega_{\eta^R,\varepsilon^{10}}$.

Appendix C. Data Details

Our sample initially covers the period 1962:Q1 to 2020:Q1. Later on in the paper, we add observations up until 2023:Q1 as we parse the data from the pandemic and its aftermath with a version of our model. The information about each variable appears below:

- Real GDP: Inflation-adjusted value of the goods and services produced by labor and property located in the United States, billions of chained 2012 dollars, seasonally adjusted, annual rate, quarterly frequency from the Federal Reserve Bank of St. Louis FRED database.
- Unemployment rate: Number of unemployed as a percentage of the labor force, seasonally adjusted, monthly frequency from the FRED database, transformed to quarterly frequency by taking the average of the months in the quarter.
- Inflation rate: Annualized quarterly percentage change in the chain-type price index of the personal consumption expenditures excluding food and energy, seasonally adjusted, quarterly frequency from the FRED database.
- Federal funds rate: Effective federal funds rate calculated as a volume-weighted median of overnight federal funds transactions reported in the FR 2420 Report of Selected Money Market Rates, percent, not seasonally adjusted, daily frequency from the FRED database, transformed to quarterly frequency by taking the average of the days in the quarter. We assume a lower bound equal to 0.25 percent that binds between 2009:Q1 and 2015:Q4.
- Ten-year Treasury yield: Yield on the 10-year Treasury security at constant maturity, percent, not seasonally adjusted, daily frequency from the FRED database, transformed to quarterly frequency by taking the average of the days in the quarter.
- Inflation expectations (PTR): This is the Federal Reserve's perceived target rate of inflation used in the FRB/US model (see Board of Governors 2022).
- Board of Governors of the Federal Reserve System output gap estimate: Real-time estimates and projections of the output gap used by the staff of the Board of Governors of the

Federal Reserve System in constructing its Greenbook forecast. Obtained from the Federal Reserve Bank of Philadelphia Greenbook Data Sets.

- CBO potential output: The CBO's estimate of the output the economy would produce with a high rate of use of its capital and labor resources. The data are adjusted to remove the effects of inflation. Obtained from the FRED database.

Appendix D. Gibbs Sampler Details

Let $\Theta_y = \{\theta_1, \theta_2, \beta, \kappa, \rho, \alpha^{\pi}, \alpha^{y}, \sigma^2_{\upsilon}, \sigma^2_{\eta^{\pi}}, \sigma^2_{\eta^{R}}, \sigma^2_{e}\}$ be the parameters of the observation equations and $\Theta_x = \{\phi_1, \phi_2, p^{10}, \psi_1, \psi_2, \sigma^2_{\eta^{y*}}, \sigma^2_{\nu}, \sigma^2_{\varepsilon}, \sigma^2_{\eta^{u*}}, \sigma^2_{\eta^{\pi*}}, \sigma^2_{\eta^{r*}}, \sigma^2_{\varepsilon^{10}}, \sigma^2_{\eta^{p^{10}}}\}$, the parameters of the transition equations. Let $\mathbf{y}_t$ be the vector of variables of the observation equation (B.1) and $\mathbf{x}_t$, the latent variables of the transition equation (B.2). The Gibbs sampler operates as follows:[30] The

[30] Whenever we obtain a posterior draw of the coefficients of the linear regression model

$$Y_t = X_t'\delta + \xi_t, \qquad \xi_t \sim i.i.dN(0, \sigma^2_{\xi}), \qquad t = 1, 2, \ldots, T,$$

we use an independent normal-inverse-gamma posterior distribution with mean

$$\left(\underline{\Sigma}^{-1} + \sum_{t=1}^{T} X_t X_t'/\sigma^2_{\xi}\right)^{-1} \left(\underline{\Sigma}^{-1}\underline{\mu} + \sum_{t=1}^{T} X_t Y_t/\sigma^2_{\xi}\right)$$

and variance

$$\left(\underline{\Sigma}^{-1} + \sum_{t=1}^{T} X_t X_t'/\sigma^2_{\xi}\right)^{-1},$$

with shape coefficient

$$\underline{a}_{\sigma^2_{\xi}} + 0.5 * T$$

and rate coefficient

$$\underline{b}_{\sigma^2_{\xi}} + 0.5 * \hat{\xi}'\hat{\xi},$$

where $\hat{\xi}$ is the vector of residuals conditional on the draw of δ, $\underline{\mu}$ and $\underline{\Sigma}$ are the prior mean and variance, respectively, of the normal prior distribution of δ, whereas $\underline{a}_{\sigma^2_{\xi}}$ and $\underline{b}_{\sigma^2_{\xi}}$ are the prior shape and rate coefficients of the prior inverse-gamma distribution of σ^2_{ξ}.

initialization of the Gibbs sampler consists in setting initial values for Θ_y and Θ_x. Moreover, the observations for the initialization of the shadow rate (R_t) are obtained by drawing from a Tobit model in which $i_t = \max\{R_t, \underline{i}\}$, with $R_t = \rho R_{t-1} + (1-\rho)\left(r^* + \pi_t^* + \alpha^\pi\left(\pi_t - \pi_t^*\right) + \alpha^y c_t\right) + \eta_t^R$ and where r^* is a constant to be estimated, $\pi_t^* = \pi_t^e$ (PTR), and the rest of the regressors are data on the federal funds rate, the PCE core inflation rate, and the CBO's estimate of the output gap. By construction, this initial step yields a shadow rate that is below the ELB during the periods it is binding.

1. Use the Durbin and Koopman (2002) simulator smoother to obtain a random draw of the latent variables, $\{\mathbf{x}_t\}_{t=1}^T$, using the state-space system in Appendix B.

2. Using the simulated values of $Y_t = c_t$ and $X_t = [c_{t-1}, c_{t-2}, c_{t-1}^{10}, c_{t-2}^{10}]'$, sample ϕ_1, ϕ_2, λ_1, λ_2, and $\sigma_{\eta^R}^2$ from a truncated (to ensure covariance stationarity) independent normal-inverse-gamma posterior distribution.

3. Sample y_0^*, μ_0, u_0^*, r_0^*, π_0^*, and p^{10} using a normal distribution with posterior mean $\sigma_{x_0}^2(\bar{x_0}/s_{x_0}^2 + x_1/\sigma_{\eta^x}^2)$ and posterior variance $\sigma_{x_0}^2 = 1/(1/s_{x_0}^2 + 1/\sigma_{\eta^x}^2)$, for $x = y^*, \mu, u^*, r^*$, π^*, and p^{10}, where $\bar{x_0}$ and $s_{x_0}^2$ are the prior mean and variance, respectively.

4. Sample $\sigma_{\eta^x}^2$ for $x = y^*, u^*, \pi^*$, and p^{10} from an inverse-gamma distribution with shape coefficient $\underline{a}_{\sigma_{\eta^x}^2} + 0.5 * T$ and rate coefficient $\underline{b}_{\sigma_{\eta^x}^2} + 0.5 * \hat{\eta}^{x\intercal}\hat{\eta}^x$, where $\hat{\eta}^x$ is the vector of residuals obtained from $x_t - x_{t-1}, t = 1, 2, \ldots, T$, and $\underline{a}_{\sigma_{\eta^x}^2}$ and $\underline{b}_{\sigma_{\eta^x}^2}$ are the prior shape and rate coefficients.

5. Sample $\sigma_{\eta^\mu}^2$, $\sigma_{\eta^{r*}}^2$, and $\omega_{\eta^{r*},\eta^\mu}$ from an inverse-Wishart distribution with scale matrix $\sum_{t=1}^T \hat{v}_t\hat{v}_t' + \nu_0 \times \Sigma_0^{\eta^\mu,\eta^{r*}}$ and degrees of freedom $T + \nu_0$, where (i) ν_0 and $\Sigma_0^{\eta^\mu,\eta^{r*}}$ are the prior degrees of freedom and variance-covariance matrix between η^μ and η^{r*}, respectively, (ii) $\hat{v}_t = [\hat{\eta}_t^\mu, \hat{\eta}_t^{r^*}]'$ is a vector of residuals, and (iii) $\hat{\eta}_t^\mu = \mu_t - \mu_{t-1}$ and $\hat{\eta}_t^{r^*} = r_t^* - r_{t-1}^*$.

6. Using the simulated values of π_t^*, obtain $\hat{e}_t = \pi_t^e - \pi_t^*$ to sample σ_e^2 from an inverse-gamma distribution with shape $\underline{a}_{\sigma_e^2} + 0.5 * T$ and rate $\underline{b}_{\sigma_e^2} + 0.5 * \hat{e}'\hat{e}$, where $\underline{a}_{\sigma_e^2}$ and $\underline{b}_{\sigma_e^2}$ are the prior shape and rate coefficients, respectively.

7. Using the observed and simulated values of $Y_t = u_t - u_t^*$ and $X_t = [c_t, c_{t-1}, c_{t-2}]'$, sample θ_1, θ_2, and σ_υ^2 from an independent normal-inverse-gamma distribution.

8. Using the observed values for $Y_t = \pi_t - \pi_{t-1}$ and observed and simulated values for $X_t = [\pi_t^* - \pi_{t-1}, c_t]'$, sample β, κ, and $\sigma_{\eta^\pi}^2$ from a truncated (to ensure homogeneity and positiveness) independent normal-inverse-gamma posterior distribution.

9. Using the observed and simulated values of $X_t = [R_{t-1}, r_t^* + \pi_t^*, \bar{\pi}_t - \pi_t^*, c_t]'$, generate R_t for t in the set of ELB periods from a truncated (from above at 0.25) normal distribution with mean $X_t'\delta$, where $\delta = [\rho, (1-\rho), (1-\rho)\alpha^\pi, (1-\rho)\alpha^y]'$, and variance $\sigma_{\eta^R}^2$. Set $Y_t = [R_t - r_t^* - \pi_t^*, c_t^{10}]'$, $X_t = [R_{t-1} - r_t^* - \pi_t^*, \bar{\pi}_t - \pi_t^*, c_t]'$, and $Z_t = [c_{t-1}^{10}, c_{t-2}^{10}]'$. Notice that R_t takes the place of i_t only during ELB periods. Set $\Omega = \begin{bmatrix} \sigma_{\eta^R}^2 & \omega_{\eta^R,\varepsilon^{10}}\sigma_{\eta^R}\sigma_{\varepsilon^{10}} \\ \omega_{\eta^R,\varepsilon^{10}}\sigma_{\eta^R}\sigma_{\varepsilon^{10}} & \sigma_{\varepsilon^{10}}^2 \end{bmatrix}$ and $W_t = \begin{bmatrix} X_t & 0 \\ 0 & Z_t \end{bmatrix}$. Draw ρ, $(1-\rho)(\alpha^\pi - 1)$, $(1-\rho)\alpha^y$, ψ_1, and ψ_2 from a truncated normal posterior distribution (to ensure covariance stationarity and the Taylor principle) with mean $(\Sigma_0^{-1} + \sum_{t=1}^T W_t\Omega^{-1}W_t')^{-1}(\Sigma_0^{-1}\delta_0 + \sum_{t=1}^T W_t\Omega^{-1}Y_t)$ and variance $(\Sigma_0^{-1} + \sum_{t=1}^T W_t\Omega^{-1}W_t')^{-1}$, where δ_0 and Σ_0^{-1} are the prior mean and variance, respectively, of the parameters to be drawn.

10. Sample $\sigma_{\eta^R}^2$, $\sigma_{\varepsilon^{10}}^2$, and $\omega_{\eta^R,\varepsilon^{10}}$ from an inverse-Wishart distribution with scale matrix $\sum_{t=1}^T \hat{\upsilon}_t\hat{\upsilon}_t' + \nu_0 \times \Sigma_0^{\eta^R,\varepsilon^{10}}$ and degrees of freedom $T + \nu_0$, where (i) ν_0 and $\Sigma_0^{\eta^R,\varepsilon^{c^10}}$ are the prior degrees of freedom and variance-covariance matrix between η^R and ε^{10}, respectively, (ii) $\hat{\upsilon}_t = [\hat{\eta}_t^R, \hat{\varepsilon}_t^{10}]'$ is a vector of

residuals, and (iii) $\hat{\eta}_t^R$ and $\hat{\varepsilon}_t^{10}$ are the residuals of the shadow and long rate equations, respectively.

11. With the newly generated R_t, initiate a new iteration by going back to step 1.

Appendix E. Prior Distributions

Table E.1 presents the prior distributions and their hyperparameters in the second column. The hyperparameters of the prior distributions associated with output and the unemployment rate are informed by the relatively standard results in the literature of trend-cycle decompositions (see Clark 1989; González-Astudillo and Roberts 2022, for example). With respect to inflation, Basistha and Nelson (2007) estimate the coefficient linked to inflation expectations to be between roughly 0.8 and 0.9, Chan and Grant (2017) estimate a posterior mean close to 0.7, and Blanchard (2016)—in a time-varying setting—estimates a sample average close to 0.6; we take a somewhat more conservative stance and set the prior mean of the persistence coefficient equal to 0.5. We also use the estimates from Blanchard to center our prior for the slope of the Phillips curve at 0.2. The variance of the inflation equation is centered at the estimated value in Basistha and Nelson (2007), whereas that of the inflation trend is centered close to the upper bound of the estimates in Stock and Watson (2007). The standard deviation of the measurement equation of inflation expectations is centered at 0.5 to allow for discrepancies between the data about inflation expectations and the inflation trend; we have not been able to find results in the literature that allow us to better inform our choice.

Regarding the monetary policy rule, we center the prior means of its parameters following the calibration of the FRB/US model (see Brayton, Laubach, and Reifschneider 2014), as well as parameter estimates of an inertial version of the Taylor (1993) rule that take into account the ELB and endogeneity, as in González-Astudillo (2018). The shock to r^* has a variance whose prior distribution is centered at a value close to the estimates in Kiley (2020). For the correlation coefficient between output growth and interest rate trends, we assume an inverse-Wishart prior distribution with 4 degrees of freedom centered at the implied estimate from HLW. In terms of the long interest rate, we choose prior means such that the cycle

Table E.1. Estimates of the Benchmark Model

	Prior Distribution	Posterior Mean	68% Credibility Interval
ϕ_1	N(1.5,1)	1.54	[1.44, 1.64]
ϕ_2	N(–0.6,1)	–0.58	[–0.68, –0.49]
λ_1	N(0.05,1)	–0.50	[–0.66, –0.33]
λ_2	N(0.05,1)	0.15	[–0.02, 0.33]
σ^2_{ε}	IG(2,0.36)	0.15	[0.11, 0.18]
$\sigma^2_{\eta^{y*}}$	IG(2,0.49)	0.29	[0.26, 0.33]
$\sigma^2_{\eta^{\mu}}$	IW(0.03^2,4)	0.03^2	[0.02^2, 0.04^2]
θ_1	N(–0.25,0.5)	–0.35	[–0.39, –0.30]
θ_2	N(–0.25,0.5)	–0.19	[–0.23, –0.15]
$\sigma^2_{\eta^{u*}}$	IG(2,0.04)	0.01	[0.01, 0.02]
σ^2_{υ}	IG(2,0.01)	0.002	[0.002, 0.003]
β	N(0.5,1)	0.26	[0.21, 0.30]
κ	N(0.2,0.5)	0.07	[0.05, 0.09]
$\sigma^2_{\eta^{\pi}}$	IG(2,1)	0.64	[0.58, 0.70]
$\sigma^2_{\eta^{\pi*}}$	IG(2,1)	0.03	[0.03, 0.04]
σ^2_{e}	IG(2,0.25)	0.01	[0.01, 0.01]
ρ	N(0.85,0.1)	0.71	[0.68, 0.75]
α^{π}	N(1.5,0.5)	1.29	[1.13, 1.45]
α^{y}	N(1,0.5)	0.93	[0.80, 1.06]
$\sigma^2_{\eta^{R}}$	IW(1,4)	0.11	[0.09, 0.13]
$\sigma^2_{\eta^{r*}}$	IG(4,0.01)	0.08	[0.06, 0.10]
ψ_1	N(1.5,1)	1.02	[0.90, 1.15]
ψ_2	N(–0.6,1)	–0.21	[–0.33, –0.10]
$\sigma^2_{\varepsilon^{10}}$	IW(1,4)	0.10	[0.09, 0.12]
$\sigma^2_{\eta^{p^{10}}}$	IG(2,0.01)	0.006	[0.003, 0.010]
y^*_0	N(816.1,5)	814.64	[813.21, 816.15]
μ_0	N(1.2,0.5)	0.98	[0.84, 1.11]
u^*_0	N(5.8,1)	6.50	[5.77, 7.21]
r^*_0	N(1.2,1)	1.74	[0.55, 2.92]
π^*_0	N(1.7,1)	1.69	[1.48, 1.89]
p^{10}_0	N(1.4,1)	0.57	[–0.57, 1.69]
$\omega_{\eta^{r*},\eta^{\mu}}$	IW(0.6,4)	0.50	[0.21, 0.75]
$\omega_{\eta^{R},\varepsilon^{10}}$	IW(0.21,4)	0.05	[–0.06, 0.16]

Note: "N" stands for normal distribution, "IG" stands for inverse-gamma distribution, and "IW" stands for inverse-Wishart distribution. In the normal, the first parameter is the mean and the second is the standard deviation. In the inverse-gamma, the first is the shape coefficient, denoted a, and the second is the scale, denoted b; the mean of the distribution is $b/(a-1)$ and the variance is $b^2/((a-1)^2(a-2))$. In the inverse-Wishart, the first parameter is the mean of the distribution of the variance or the correlation coefficient (depending on the parameter) and the second, the degrees of freedom. Strictly speaking, we produce draws of the covariance between shocks.

has a hump shape and an average yield equal to that in the sample; the variance of the disturbance is centered at 1, for the lack of information in the literature. Nevertheless, the hyperparameters of the inverse-gamma prior distributions of the variances are such that only their means are well defined, whereas their variances are not, which allows the estimation to more freely pick up the posterior means of these coefficients. Additionally, the means of the prior distributions of the initial values of the nonstationary latent factors are set in accordance with the initial values of the relevant variables in the sample (we use the term premium series from Adrian, Crump, and Moench 2013 to initialize p_t^{10}). Finally, the correlation coefficient between trend output growth and r^* is centered at 0.6, derived from the results in Laubach and Williams (2003), whereas the correlation coefficient between the shadow interest rate and the cycle of the 10-year Treasury yield is centered at 0.21, which is the correlation coefficient between the policy news shock in Nakamura and Steinsson (2018) and the change in the nominal yield of the zero-coupon 10-year Treasury bond.

Appendix F. Parameter Diagnostics

Figures F.1, F.2, and F.3 show the prior and posterior distributions of the parameters of the benchmark model.

Figure F.1. Prior and Posterior Distributions: Conditional Mean Parameters

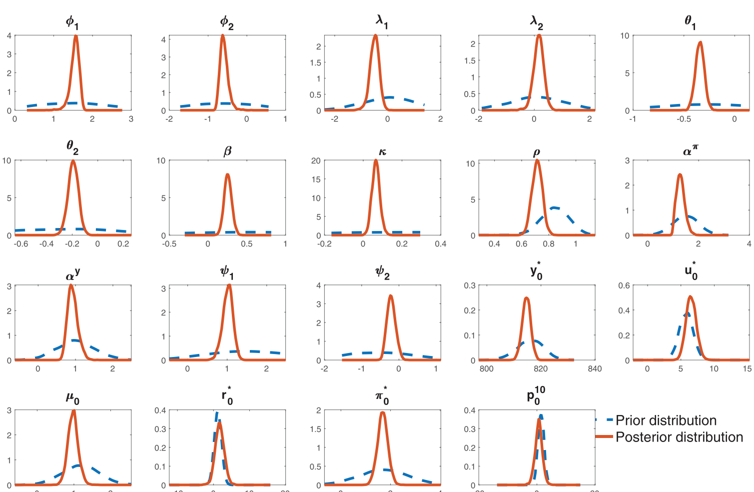

Figure F.2. Prior and Posterior Distributions: Variance Parameters

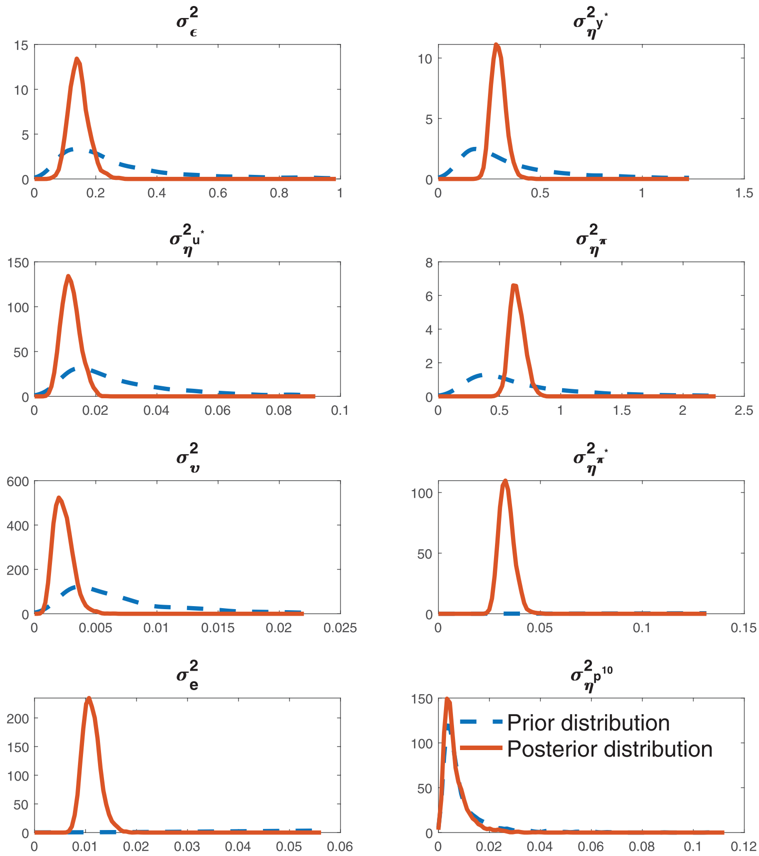

Figure F.3. Prior and Posterior Distributions: Variance and Correlation Parameters

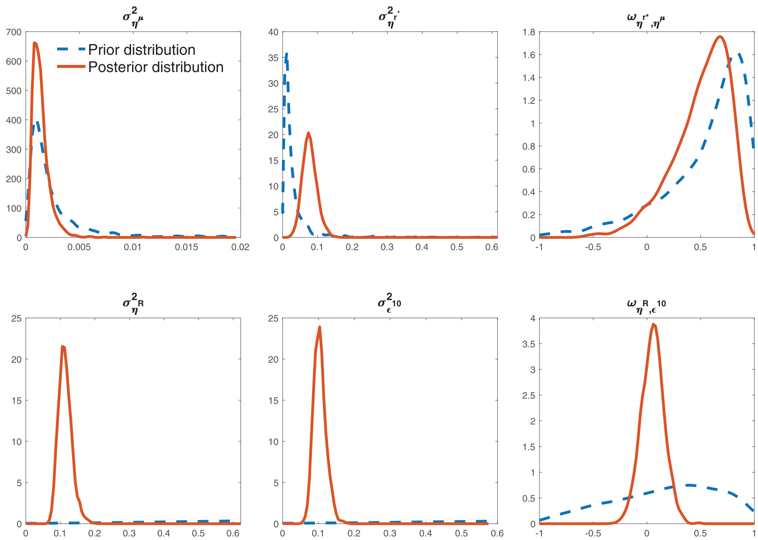

Appendix G. Data Decomposition

This appendix shows the inflation gap, defined as the difference between actual and trend inflation, in Figure G.1 and the historical data decomposition of the output gap in Figure G.2.

Figure G.1. Inflation Gap

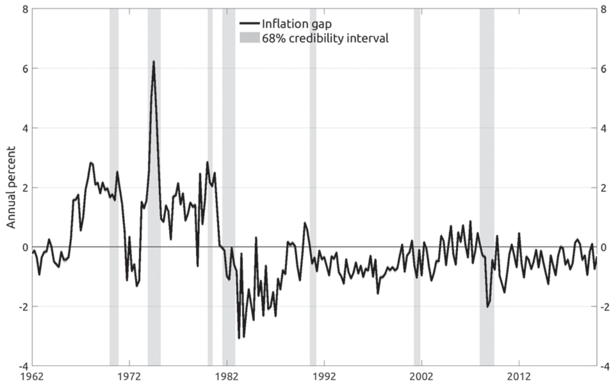

Note: Shaded vertical areas indicate NBER recession periods. Smoothed estimates are reported. The inflation gap is defined as the posterior mean of $\pi_t - \pi_t^*$.

Figure G.2. Historical Data Decomposition of the Output Gap

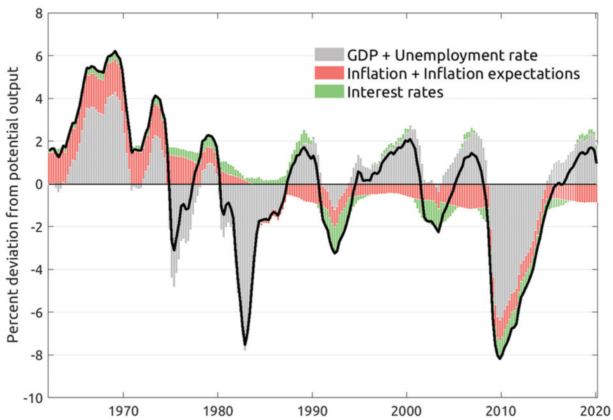

Note: The contributions of GDP and the unemployment rate (GDP + Unemployment rate) have been added together. The same is true for the inflation rate and PTR (Inflation + Inflation expectations), and for the federal funds rate and the 10-year Treasury yield (Interest rates).

Appendix H. Impulse Response Analysis

Beside the structural dynamics and constraints imposed on some of the coefficients, the correlation between the policy equation innovation, $\eta_t^{R^*}$, and that of the long-term interest rate gap, ε_t^{10}, entails an influence of the former on the rest of the economy. As explained in Section 2.4, without this feature and in the absence of any explicit role for expectations, policy rate innovations would have no bearing on real activity and inflation. The correlated innovations allow us to mimic the effects of a conventional monetary policy shock.

The impulse response functions of an increase of 100 basis points (bps) in the federal funds rate are depicted in Figure H.1A. Our model predicts that the 10-year Treasury yield would increase 5 bps on impact and would then decline, causing the output gap to decline by almost 14 bps at the trough with the inflation rate declining about 3 bps. The sizes of these responses are similar to those obtained in the FRB/US model with model-consistent expectations (see Brayton, Laubach, and Reifschneider 2014; Laforte and Roberts 2014).

The structure of the model allows us to treat and interpret the cyclical component of the 10-year Treasury yield as a proxy for forward guidance or asset purchases by the central bank (i.e., unconventional monetary policy). Given a shock of 100 bps to this cyclical component of long rate, output declines 2.5 percentage points (pp) at the trough and inflation, 0.6 pp, followed by a decline in the federal funds rate of 2.5 pp, as shown in Figure H.1B. These effects are much larger than those obtained with the FRB/US model, for instance, after an increase of 100 bps in term premiums as shown in Laforte and Roberts (2014).

The structural character of our model is, however, limited. For instance, a direct innovation to the inflation process will have no repercussion on the rest of the economy for reasons explained earlier. Nonetheless, it is possible to gauge the dynamics of our model from its "multipliers," i.e., the magnitude of the response of a given process/variable to a change in a particular economic factor (which is not necessarily a fundamental shock of the model), a practice well-established in the business of professional forecasting. For instance, it is common to be interested in how much inflation will react to a change in economic conditions, like in the output gap.

Figure H.1. Impulse-Response Functions

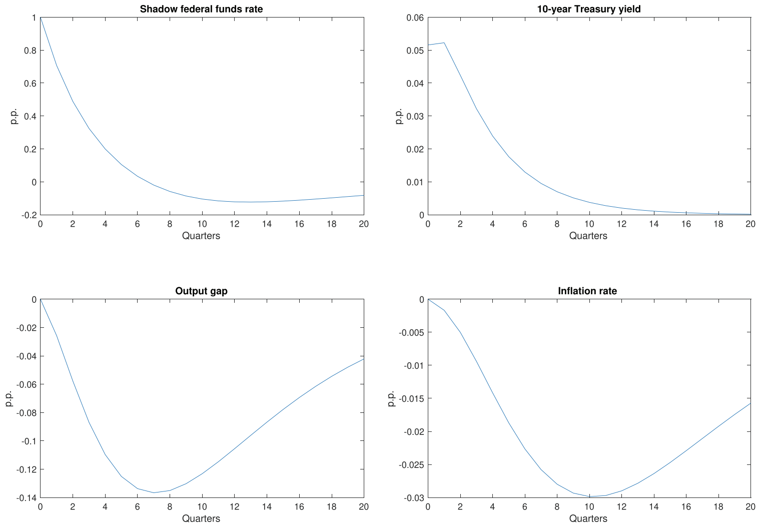

A. Shock to the Federal Funds Rate

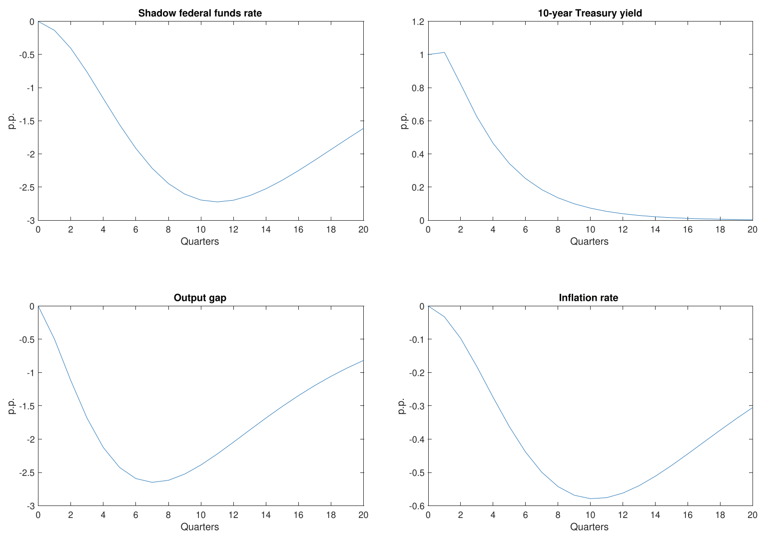

B. Shock to the Long-Term Interest Rate Cycle

(*continued*)

Figure H.1. (Continued)

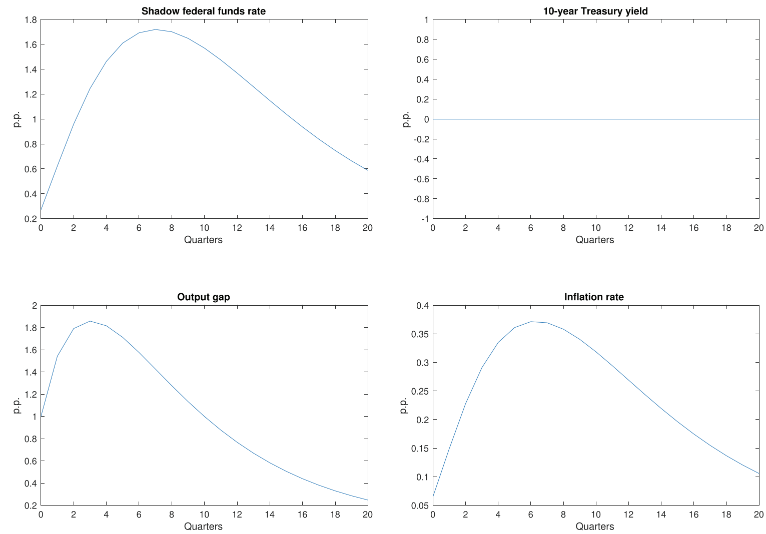

C. Shock to the Output Gap

We can perform such an exercise with our model. Figure H.1C shows that inflation will have a peak increase of about 0.4 pp and the federal funds rate would peak at around 1.7 pp, following a 1 percent increase of output above potential. This is because there is no connection between the 10-year Treasury yield and the federal funds rate beyond the correlation between their shocks, and as explained earlier, the former does not move despite the increase in the latter.

Appendix I. Stochastic Volatility Results

This section presents the results of the estimation that includes stochastic volatility. Figure I.1 shows the estimated stochastic volatility process through 2023:Q1 in which the missing observations approach was implemented.

Figure I.2A shows that taking the most extreme observations at face value rather than as being missing, our estimate of the output gap reached about –21 percent in the second quarter of 2020 compared with the CBO's –11 percent, whereas the natural unemployment rate (Figure I.2B) did not suffer any significant break despite

Figure I.1. Estimated Stochastic Volatilities for Each Shock

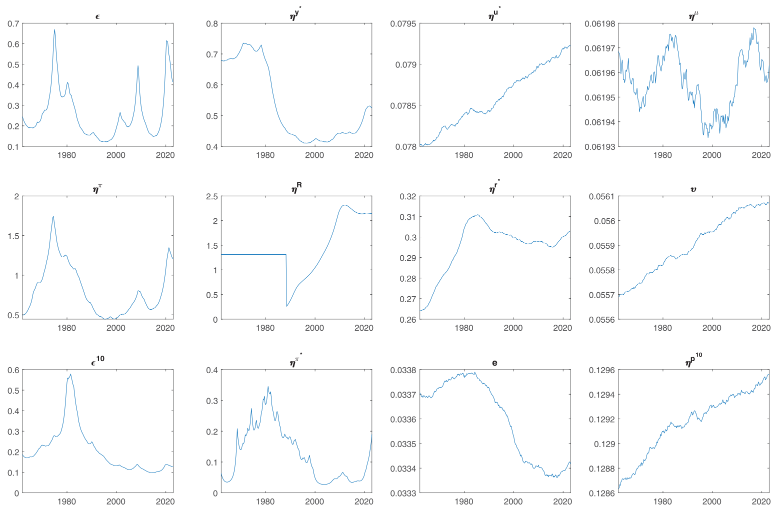

Figure I.2. Parsing of the COVID-19 Pandemic Period

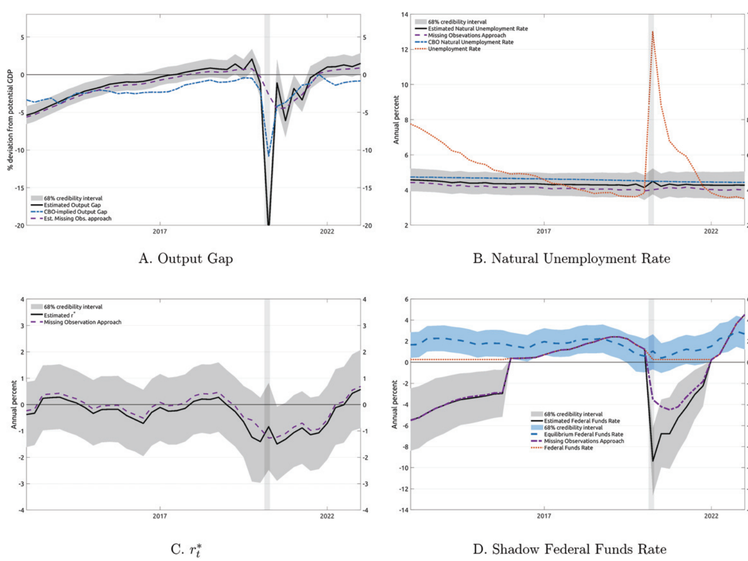

Note: Shaded vertical areas indicate NBER recession periods. Smoothed estimates are reported.

the increase in the actual unemployment rate (hence the massive slack estimate, given the Okun's law relationship in our model); u^* is estimated to be 4.2 percent at the end of the sample.

Figure I.2D shows the federal funds rate along with its shadow and equilibrium counterparts while Figure I.2C shows r_t^*. The results indicate that the equilibrium policy rate was below 1 percent during the early stage of the pandemic period and has increased since then, standing at 2.8 percent in early 2023. Of note, the estimates of r_t^* with and without missing observations are practically identical.

Appendix J. Pseudo-Real-Time Estimates of the Output Gap and r_t^*

Figures J.1 and J.2 show the pseudo-real-time mean estimates of the output gap and r_t^* for five of our models. For each model, the estimated value shown in period t corresponds to the estimate at the end of the sample for the vintage whose date of the last observation is period t, i.e., the value obtained from conditioning solely on data through period t. For example, the output gap estimate in 2017:Q1 is the last implied value by the model estimated with the sample through 2017:Q1. Estimates are calculated for each model

Figure J.1. Pseudo-Real-Time Estimates of the Output Gap across Models

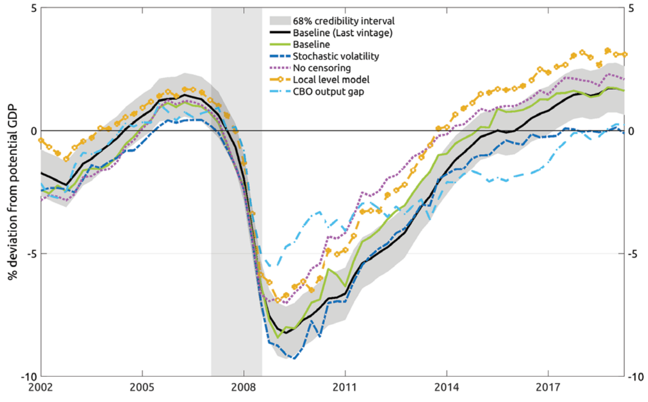

Note: Shaded vertical area indicates NBER recession period.

Figure J.2. Pseudo-Real-Time Estimates of r_t^* across Models

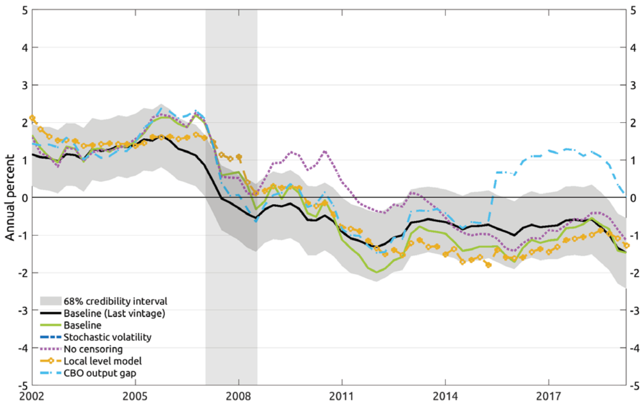

Note: Shaded vertical area indicates NBER recession period.

over their respective posterior distribution of the parameters. The gray area shows the 68 percent credibility interval of the baseline model estimated with the whole sample through 2020:Q1.

References

Adrian, T., R. K. Crump, and E. Moench. 2013. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* 110: 110–38.

Andrle, M. 2013. "Understanding DSGE Filters in Forecasting and Policy Analysis." IMF Working Paper No. 2013/098.

Aruoba, S. B., M. Mlikota, F. Schorfheide, and S. Villalvazo. 2022. "SVARs with Occasionally-Binding Constraints." *Journal of Econometrics* 231: 477–99.

Aruoba, S. B., and F. Schorfheide. 2011. "Sticky Prices versus Monetary Frictions: An Estimation of Policy Trade-Offs." *American Economic Journal: Macroeconomics* 3 (1): 60–90.

Ascari, G., and A. M. Sbordone. 2014. "The Macroeconomics of Trend Inflation." *Journal of Economic Literature* 52 (3): 679–739.

Basistha, A., and C. R. Nelson. 2007. "New Measures of the Output Gap Based on the Forward-Looking New Keynesian Phillips Curve." *Journal of Monetary Economics* 54 (2): 498–511.

Bauer, M. D., and G. D. Rudebusch. 2016. "Monetary Policy Expectations at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (7): 1439–65.

———. 2020. "Interest Rates under Falling Stars." *American Economic Review* 110 (5): 1316–54.

Beyer, R. C. M., and V. Wieland. 2019. "Instability, Imprecision and Inconsistent Use of Equilibrium Real Interest Rate Estimates." *Journal of International Money and Finance* 94 (June): 1–14.

Blanchard, O. 2016. "The US Phillips Curve: Back to the 60s?" Policy Brief No. 16-1, Peterson Institute for International Economics.

Blanchard, O., E. Cerutti, and L. Summers. 2015. "Inflation and Activity—Two Explorations and Their Monetary Policy Implications." Technical Report, National Bureau of Economic Research.

Board of Governors of the Federal Reserve System. 2018. "Policy Rules and How Policymakers Use Them." Accessed May 20, 2020. https://www.federalreserve.gov/monetarypolicy/policy-rules-and-how-policymakers-use-them.htm.

———. 2022. "Data for FRB/US Model." Accessed April 29, 2022.

Brand, C., and F. Mazelis. 2019. "Taylor-Rule Consistent Estimates of the Natural Rate of Interest." ECB Working Paper No. 2257.

Brayton, F., T. Laubach, and D. L. Reifschneider. 2014. "The FRB/US Model: A Tool for Macroeconomic Policy Analysis." FEDS Notes, Board of Governors of the Federal Reserve System.

Bullard, J. 2018. "Modernizing Monetary Policy Rules." Presentation to the Economic Club of Memphis, Memphis, TN, October 18, 2018. Accessed January 24, 2020.

Carriero, A., T. E. Clark, M. Marcellino, and E. Mertens. 2022. "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility." Forthcoming in Review of Economics and Statistics. https://doi.org/10.1162/rest_a_01213.

———. 2023. "Shadow-Rate VARs." Technical Report No. 14, Deutsche Bundesbank.

Chan, J., and A. Grant. 2017. "Measuring the Output Gap Using Stochastic Model Specification Search." CAMA Working Paper

No. 2/2017, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, Australian National University.

Chib, S. 1992. "Bayes Inference in the Tobit Censored Regression Model." *Journal of Econometrics* 51 (1–2): 79–99.

———. 1995. "Marginal Likelihood from the Gibbs Output." *Journal of the American Statistical Association* 90 (432): 1313–21.

Christensen, J. H. E., and G. D. Rudebusch. 2019. "A New Normal for Interest Rates? Evidence from Inflation-Indexed Debt." *Review of Economics and Statistics* 101 (5): 933–49.

Chung, H., C. Fuentes-Albero, M. Paustian, and D. Pfajfar. 2021. "Latent Variables Analysis in Structural Models: A New Decomposition of the Kalman Smoother." *Journal of Economic Dynamics and Control* 125: Article 104097.

Clark, P. K. 1989. "Trend Reversion in Real Output and Unemployment." *Journal of Econometrics* 40 (1): 15–32.

Cochrane, J. H., and M. Piazzesi. 2002. "The Fed and Interest Rates—A High-Frequency Identification." *American Economic Review* 92 (2): 90–95.

Cogley, T., and T. Sargent. 2015. "Measuring Price-Level Uncertainty and Instability in the United States, 1850–2012." *Review of Economics and Statistics* 97 (4): 827–38.

Cúrdia, V., A. Ferrero, G. C. Ng, and A. Tambalotti. 2015. "Has U.S. Monetary Policy Tracked the Efficient Interest Rate?" *Journal of Monetary Economics* 70 (March): 72–83.

Del Negro, M., D. Giannone, M. Giannoni, and A. Tambalotti. 2017. "Safety, Liquidity, and the Natural Rate of Interest." *Brookings Papers on Economic Activity* 48 (Spring): 235–316.

Del Negro, M., M. Giannoni, and F. Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics* 7 (1): 168–96.

Durbin, J., and S. J. Koopman. 2002. "A Simple and Efficient Simulation Smoother for State Space Time Series Analysis." *Biometrika* 89 (3): 603–16.

Fiorentini, G., A. Galesi, G. Pérez-Quirés, and E. Sentana. 2018. "The Rise and Fall of the Natural Interest Rate." CEPR Discussion Paper No. 13042.

Gneiting, T., and A. E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102 (477): 359–78.

González-Astudillo, M. 2018. "Identifying the Stance of Monetary Policy at the Zero Lower Bound: A Markov-Switching Estimation Exploiting Monetary-Fiscal Policy Interdependence." *Journal of Money, Credit and Banking* 50 (1): 115–54.

González-Astudillo, M., and J. M. Roberts. 2022. "When Are Trend-cycle Decompositions of GDP Reliable?" *Empirical Economics* 62 (5): 2417–60.

Holston, K., T. Laubach, and J. Williams. 2017. "Measuring the Natural Rate of Interest: International Trends and Determinants." *Journal of International Economics* 108 (S1): S59–S75.

Johannsen, B. K., and E. Mertens. 2021. "A Time-Series Model of Interest Rates with the Effective Lower Bound." *Journal of Money, Credit and Banking* 53 (5): 1005–46.

Jones, C. J., M. Kulish, and J. Morley. 2021. "A Structural Measure of the Shadow Federal Funds Rate." Finance and Economics Discussion Series No. 2021-064, Board of Governors of the Federal Reserve System.

Kiley, M. T. 2020. "What Can the Data Tell Us about the Equilibrium Real Interest Rate?" *International Journal of Central Banking* 16 (3): 181–209.

Kim, S., N. Shephard, and S. Chib. 1998. "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models." *Review of Economic Studies* 65 (3): 361–93.

Koopman, S. J., and A. Harvey. 2003. "Computing Observation Weights for Signal Extraction and Filtering." *Journal of Economic Dynamics and Control* 27 (7): 1317–33.

Laforte, J.-P., and J. M. Roberts. 2014. "November 2014 Update of the FRB/US Model." Technical Report, Board of Governors of the Federal Reserve System.

Laubach, T., and J. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 85 (4): 1063–70.

Lenza, M., and G. E. Primiceri. 2022. "How to Estimate a Vector Autoregression after March 2020." *Journal of Applied Econometrics* 37 (4): 688–99.

Lewis, K. F., and F. Vazquez-Grande. 2019. "Measuring the Natural Rate of Interest: A Note on Transitory Shocks." *Journal of Applied Econometrics* 34 (3): 425–36.

Lopez-Salido, D., G. Sanz-Maldonado, C. Schippits, and M. Wei. 2020. "Measuring the Natural Rate of Interest: The Role of

Inflation Expectations." FEDS Notes, Board of Governors of the Federal Reserve System, June 16.

Lubik, T. A., and C. Matthes. 2015. "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches." Economic Brief No. 15-10, Federal Reserve Bank of Richmond.

Mavroeidis, S. 2021. "Identification at the Zero Lower Bound." *Econometrica* 89 (6): 2855–85.

Mertens, E. 2016. "Measuring the Level and Uncertainty of Trend Inflation." *Review of Economics and Statistics* 98 (5): 950–67.

Morley, J. C., C. R. Nelson, and E. Zivot. 2003. "Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?" *Review of Economics and Statistics* 85 (2): 235–43.

Nakamura, E., and J. Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–330.

Orphanides, A., and J. Williams. 2002. "Robust Monetary Policy Rules with Unknown Natural Rates." *Brookings Papers on Economic Activity* 2002 (2): 63–146.

Powell, J. 2019. "Building on the Gains from the Long Expansion." Presentation at the Annual Meeting of the Greater Providence Chamber of Commerce, Providence, Rhode Island. Accessed February 18, 2020.

Roberts, J. 2018. "An Estimate of the Long-Term Neutral Rate of Interest." FEDS Notes, Board of Governors of the Federal Reserve System, September 5.

Sander, N. 2013. "Fresh Perspectives on Unobservable Variables: Data Decomposition of the Kalman Smoother." Analytical Notes Series No. 2013/09, Reserve Bank of New Zealand.

Schorfheide, F., and D. Song. 2021. "Real-Time Forecasting with a (Standard) Mixed-Frequency VAR during a Pandemic." Technical Report, National Bureau of Economic Research.

Smets, F., and R. Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Stock, J. H., and M. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (s1): 3–33.

Summers, L. H., and L. Rachel. 2019. "On Falling Neutral Real Rates, Fiscal Policy and the Risk of Secular Stagnation." *Brookings Papers on Economic Activity*, BPEA Conference Drafts, March 7–8.

Taylor, J. 1993. "Discretion Versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39 (December): 195–214.

———. 1999. "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, ed. J. B. Taylor, 319–48. Chicago: University of Chicago Press.

Williams, P., Y. Abdih, and E. Kopp. 2020. "Reading the Stars." IMF Working Paper No. 2020/136.

Wu, J. C., and F. D. Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (2–3): 253–91.

Yellen, J. 2017. "The Economic Outlook and the Conduct of Monetary Policy." Presentation at Stanford Institute for Economic Policy Research, Stanford, California, January 19. Accessed: January 24, 2020.

Zaman, S. 2021. "A Unified Framework to Estimate Macroeconomic Stars." Working Paper No. 21-23R, Federal Reserve Bank of Cleveland.

Inflation Expectations Anchoring: New Insights from Microevidence of a Survey at High Frequency and of Distributions*

Nikos Apokoritis,[a] Gabriele Galati,[b] and
Richhild Moessner[c]
[a]IRI
[b]De Nederlandsche Bank
[c]University of Heidelberg, Bank for International Settlements, and CESifo

We analyze the anchoring of long-term euro-area inflation expectations in the decade following the Global Financial Crisis, by exploring a new weekly survey, conducted among around 25 macroanalysts from 2010 to 2018. We perform a battery of tests on level expectations from the weekly survey and measures based on the distribution of inflation expectations from a quarterly survey. These include measures of uncertainty, the probability of expected long-term inflation lying between 1.5 percent and 2.5 percent, and deflation risk. We find that long-term euro-area inflation expectations remained broadly anchored to the European Central Bank's inflation aim.

JEL Codes: E31, E58, F62.

1. Introduction

This paper provides new evidence on the anchoring of inflation expectations of professionals using a new data set on short- and long-term inflation expectations in the euro area, which is based on a new survey at weekly frequency.

*The views expressed are those of the authors and do not necessarily reflect those of the Bank for International Settlements, De Nederlandsche Bank, or the Eurosystem. We would like to thank Maurice Bun, Sebastiano Manzan, an anonymous referee, and participants at the Bank of England CCBS Chief Economists' Workshop 2019 and the Swiss Society of Economics and Statistics Annual Conference 2021 for useful comments and discussions. Author e-mails: nikosapok@gmail.com, e.b.g.galati@dnb.nl, richhild.moessner@uni-heidelberg.de.

Inflation expectations play a key role in macroeconomic models and monetary policy (Bernanke 2007, 2022; Williams 2022). They are carefully monitored by central banks to gauge how private agents perceive the credibility of monetary policy in pursuing price stability. In the decade following the Global Financial Crisis, the credibility of monetary authorities in advanced economies was challenged by persistently low inflation. With inflation stuck at low levels and policy rates close to their effective lower bound, there were growing concerns that long-term inflation expectations would fall below central banks' inflation targets, thereby affecting the effectiveness of monetary policy (Schnabel 2021). When central banks reviewed their monetary strategies in those years, efforts to better anchor inflation expectations therefore played a significant role (Powell 2020; European Central Bank 2021). Since mid-2021, instead, global inflation increased persistently, raising concerns about a de-anchoring of inflation expectations on the upside.

A key question in the policy debate and the research literature is therefore whether inflation expectations have been firmly anchored to central banks' price inflation targets (e.g., Corsello, Neri, and Tagliabracci 2019; Moessner and Takáts 2020; Bems et al. 2021; Goel and Tsatsaronis 2022). This is particularly the case for the euro area, where significant cross-country differences in wage and price setting make the coordinating role of a nominal anchor more important (Cœuré 2019).

In the literature, the concept of anchored inflation expectations refers to long-term expectations and is defined in terms of several conditions (Kumar et al. 2015; Neri et al. 2022). First, average expectations should be close to the central bank's target ("level anchoring"). Second, long-term expectations should not co-move with changes in actual inflation, inflation surprises, or short-term expectations ("shock anchoring"). Third, expectations should not be overly dispersed among individuals. Fourth, agents should be fairly confident about their best guess of future inflation and have little uncertainty about the long term. And finally, agents should not attach a large weight to extreme inflation outcomes in the future. According to this view, a full picture of anchoring of expectations would therefore involve information also on the higher moments of their distribution.

Expectations of different types of agents—market participants, professional forecasters, firms, and households—matter for the transmission of monetary policy to the economy. Over the past years, also because of an increasing availability of data, inflation expectations held by firms and households have come to play a bigger role in both research and policy (Candia, Coibion, and Gorodnichenko 2021; Adrian 2022; D'Acunto, Malmendier, and Weber 2022; Neri et al. 2022; Weber et al. 2022). Nevertheless, central banks are, in practice, still focusing on inflation expectations held by professional forecasters and financial market participants (European Central Bank 2021). One reason is that expectations of financial market participants play an important role in the monetary transmission mechanism, since they are a driver or financial prices and hence financing conditions for firms and households. Another reason is that expectations of professional forecasters can be an input in wage negotiations and firms' price-setting decisions (see, e.g., Conflitti and Zizza 2021).

We shed new light on the behavior of short- and long-term euro-area inflation expectations between July 2010 and December 2018 by using microevidence from a new type of survey at weekly frequency. This survey has been conducted since July 2010 among economists, financial analysts, and statisticians at De Nederlandsche Bank (DNB, the Dutch central bank). Participants answer every week on Monday a questionnaire about their short- and long-term expectations of euro-area Harmonised Index of Consumer Prices (HICP) inflation.

Our survey has two main advantages with respect to existing surveys of professionals. First, the weekly frequency of our survey is unique since surveys of professional forecasters' expectations are typically conducted at monthly or quarterly frequency. The higher frequency of our survey allows a richer characterization of the anchoring of inflation expectations by using methods that in the literature have been applied to high-frequency data on market-based expectations measures. In particular, the high frequency allows an analysis of the reaction of expectations to news about inflation in the euro area. In this respect, our paper is related to research that exploits variation across survey panelists in the exact survey dates in monthly or quarterly surveys to investigate the effect of macroeconomic news and monetary policy decisions on expectations. These papers typically examine expectations of households and focus on the United States.[1] To our knowledge, the paper by Bottone and Rosolia (2019), which examines the response of Italian firms' expectations of inflation in Italy to monetary policy shocks, is the only paper focusing on the euro area using this approach. In this respect, our paper is also related to that of Clements (2012); however, he uses a different approach than ours and relies on low-frequency survey data. By contrast, a large literature exists on whether financial market expectations pay attention to data releases, due to the availability of financial market data at high frequency (daily and intraday) (see, e.g., Fleming and Remolona 1997; Gürkaynak, Levin, and Swanson 2010; Beechey, Johannsen, and Levin 2011). Due to the high (weekly) frequency of the DNB survey, it allows us to study this question also for survey expectations.

Second, survey participants also answer once per quarter questions about the entire distribution of their inflation expectations. Only a few surveys of professional forecasters provide information about the probability distribution of individuals' inflation expectations, including the Survey of Professional Forecasters for the euro area (e.g., Rich and Tracy 2018), the Bank of England survey of external forecasters (Boero, Smith, and Wallis 2008), and the Survey of Professional Forecasters (D'Amico and Orphanides 2008) and the Federal Reserve Bank of New York Survey of Consumer Expectations (Bruine de Bruin et al. 2011) for the United States.

We use several methods to study whether long-term euro-area inflation expectations of DNB survey respondents have been well anchored, in line with the different conditions used in the literature to define anchoring.

To assess level anchoring of DNB survey inflation expectations, we investigate whether the European Central Bank's (ECB's) inflation aim has acted as a focal point for expectations. As an alternative focal point, we also test the role of Consensus survey inflation

[1] These papers on expectations of U.S. households use data from the Federal Reserve Bank of New York's Survey of Consumer Expectations (De Fiore, Lombardi, and Schuffels 2019; Binder, Campbell, and Ryngaert 2022), a Gallop survey (Lewis, Makridis, and Mertens 2019), or an ad hoc survey (Lamla and Vinogradov 2019).

expectations, which are included in the information set available to DNB survey participants.

To assess shock anchoring of long-term DNB survey expectations, we test whether they responded to data releases on inflation or to inflation data surprises. The response of long-term inflation expectations to macroeconomic data surprises is a common measure for the anchoring of inflation expectations (Gürkaynak et al. 2007; Beechey, Johannsen, and Levin 2011). If long-term expectations are well anchored, they should not respond to data surprises.

We also study whether long-term inflation expectations of DNB survey respondents have been shock anchored by investigating whether changes in long-term DNB survey expectations responded to changes in short-term DNB survey expectations. Such an approach has been considered, e.g., in Buono and Formai (2018). We also study whether there has been heterogeneity across survey respondents in these reactions. Heterogeneity in inflation expectations formation may matter for the anchoring of inflation expectations. Busetti et al. (2017) find that under heterogeneity in inflation expectations formation, a sequence of negative shocks may lead inflation to deviate from target and reinforce a de-anchoring of expectations.

Furthermore, we study the distribution of inflation expectations, and consider two measures of the anchoring of long-term inflation expectations based on the full distribution from the quarterly DNB survey, namely uncertainty and the effect of short-term deflation risk on long-term deflation risk.

We consider uncertainty about long-term inflation expectations as a distributions-based measure of the anchoring of long-term inflation expectations (Dovern and Kenny 2020). Moreover, we consider the survey-based probability of future inflation being in a certain range that is consistent with the inflation target as a measure of anchoring, in particular the probability of expected long-term inflation lying between 1.5 percent and 2.5 percent, as proposed by Grishchenko, Mouabbi, and Renne (2019). Relatedly, Mehrotra and Yetman (2018) consider the precision around forecasts of the level of inflation as a measure of the anchoring of inflation expectations.

Second, we consider the effect of short-term on long-term deflation risk from the DNB survey as a measure of the anchoring of long-term inflation expectations. A related measure has been applied to deflation risk derived from market-based inflation options by Galati et al. (2018), who consider Granger causality between short-term and long-term deflation risk. A related measure is also presented by Antunes (2015) and Natoli and Sigalotti (2018), who analyze the tail co-movement between the moments of short- and long-term distributions of inflation expectations derived from market-based inflation options. Differently from these two papers, we investigate the co-movement of short-term and long-term deflation risk using deflation risk derived from survey-based distributions of inflation expectations, rather than using market-based measures of deflation risk.

Using the weekly survey, almost all the tests we conducted suggest that over the period 2010–18, long-term inflation expectations remained well anchored to the ECB's inflation aim, which has acted as a focal point. By contrast, we find no evidence that professional forecasts (reported by Consensus Economics) acted as focal points. But for one of the approaches we follow, namely tests of the reaction of long-term inflation expectations to short-term expectations, there are subtle signs of long-term inflation expectations not being perfectly well anchored, in line with the conclusions in ECB (2021). We also find that, notwithstanding the relative homogeneity of the sampled population, there is some evidence of heterogeneity in the anchoring of long-term inflation expectations.

Tests that use measures based on the distribution of inflation expectations—uncertainty based on the full distribution, the probability of expected long-term inflation lying between 1.5 percent and 2.5 percent, and the effect of short-term deflation risk on long-term deflation risk—confirm that long-term inflation expectations were well anchored and became better anchored at the end of the sample period in 2018 compared with the start of the sample period in 2011.

The remainder of the paper is organized as follows. Section 2 introduces the DNB inflation expectations survey. Section 3 presents the method and Section 4 the results. Finally, Section 5 concludes.

2. The DNB Inflation Expectations Survey

Since July 2010, participants in the DNB inflation expectations survey answer a questionnaire about their short- and long-term expectations of euro-area HICP inflation every week on Monday.

In addition, participants are asked questions about the distribution of their inflation expectations once per quarter.

The survey panel consists of around 25 economists, financial analysts, and statisticians employed by DNB per week. In order to deal with panel attrition, new participants are added to the survey panel to replace participants with similar characteristics that left the panel. The participants in this survey have a background that is comparable to that of respondents to the ECB's Survey of Professional Forecasters (SPF), who are experts employed by financial or non-financial institutions, such as economic research institutions (Garcia 2003). In line with other surveys of inflation expectations, participants and their answers are treated anonymously, to encourage participants to submit their input without any concern about forecast errors. Panelists generally receive an e-mail on Monday morning with three questions on their short- and long-term inflation expectations, and generally answer the e-mail within that day. An example of this e-mail is provided in the appendix.

The survey has two novel features compared with existing surveys. First, the weekly frequency is higher than the frequency of other surveys of professional forecasters, which typically ranges from monthly to semi-annual. Secondly, participants in our survey are provided with common information sets. In particular, together with the questionnaire, participants receive each week an update of relevant data related to inflation in the euro area. This background information includes data releases on inflation for the euro area as a whole and for six euro-area member countries (Germany, France, Italy, Spain, the Netherlands, and Belgium) that were published during the previous week, a table with the latest Consensus forecasts for euro-area HICP inflation, and a graph with current and past actual euro-area HICP inflation.

The quarterly information on the distribution of expectations allows for tracking changes in the higher moments of expectations—in particular, uncertainty—over time.

The combination of a homogeneous set of participants, a common information set, and a high frequency allows us to focus on mechanisms of expectations formation and their heterogeneity since the Global Financial Crisis, a period characterized by high uncertainty. In particular, we can study more carefully some aspects of expectation formation, such as how inflation expectations depend on

Figure 1. Euro-Area Inflation Expectations from DNB Survey, in Percent

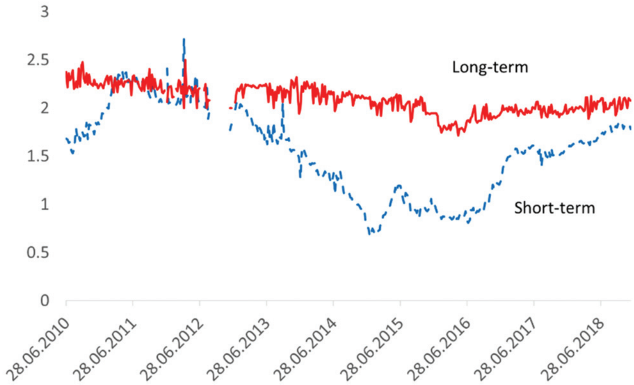

Note: Mean long-term and short-term euro-area inflation expectations from DNB survey.

realized inflation data and surprises; whether and how the anchoring of expectations changes with a crisis; and the role of focal points, such as the ECB's inflation aim or professional forecasters' inflation predictions.

Short-term and long-term mean DNB survey expectations are shown in Figure 1. For long-term DNB survey expectations, these are the direct survey responses. For short-term DNB survey expectations, we interpolate between the current-year, π_{it}^{c}, and next-year, π_{it}^{n}, survey responses, in order to obtain a constant-horizon short-term expectation, π_{it}^{ST}, according to

$$\pi_{it}^{ST} = \left(1 - \frac{(m-1)}{11}\right)\pi_{it}^{c} + \frac{(m-1)}{11}\pi_{\mathrm{it}}^{n} \tag{1}$$

with $m = 1, \ldots, 12$, and $m = 1$ for January (this is when the survey expectations for the current year and the next year each switch to the following year), $m = 2$ for February, etc.[2] π_{it}^{ST} is referred to as short-term DNB survey expectations in the remainder of this paper.

[2]This is the most commonly used approach in the literature for approximating fixed-horizon forecasts using fixed-event forecasts (e.g., Gerlach 2007; Dovern, Fritsche, and Slacalek 2012; Siklos 2013). For an alternative approach constructing optimal weights, see Knüppel and Vladu (2016).

Figure 2. Long-Term Survey-Based Euro-Area Inflation Expectations, in Percent

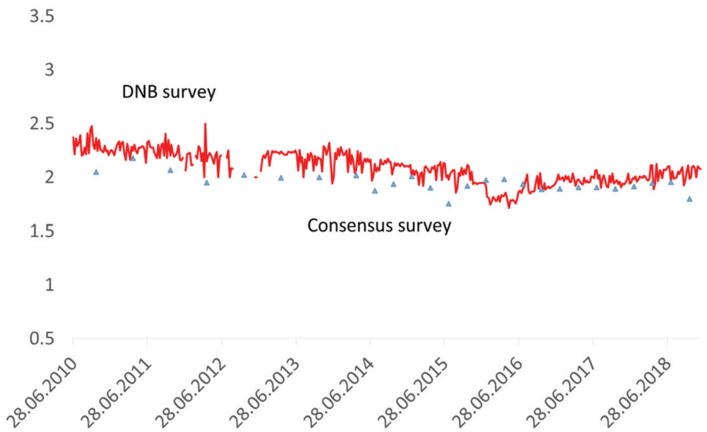

Note: Mean long-term inflation expectations from DNB survey and mean long-term (5 to 10 years ahead) inflation expectations from Consensus Economics surveys.

There is no consensus in the literature on the process through which agents form inflation expectations. Commonly used measures extracted from surveys or financial markets do not provide a uniform answer. In the euro area, for example, there is a visible difference in the level and variance between these two types of measures (see Figures 2 and 3). Survey-based measures are usually quite persistent, while financial-market-based measures are typically quite volatile. Both survey-based measures and financial-market-based measures of inflation expectations have advantages and drawbacks (for an overview, see ECB 2021; Neri et al. 2022). A main disadvantage of survey-based measures is that they are usually only available at low frequency. A main disadvantage of financial-market-based measures is that they are usually affected by risk and liquidity premia.

One caveat about the setup of this survey is that its external validity depends on how representative the participants are of the general population of macroanalysts. In particular, our results may be biased if there are incentive issues for employees of the central bank in our panel despite the anonymous character of the survey. A comparison of long-term DNB survey inflation expectations and

Figure 3. Long-Term Market-Based Euro-Area Inflation Expectations, in Percent

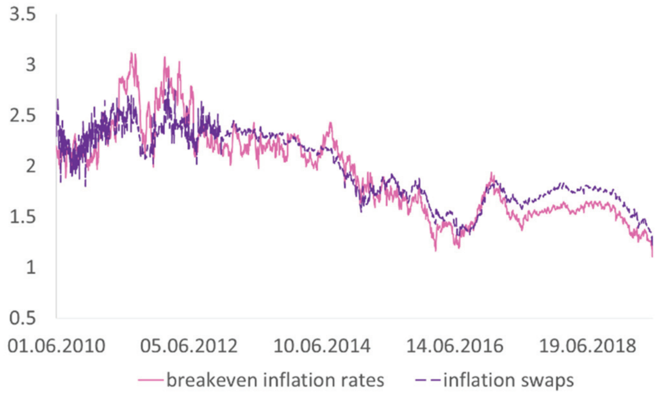

Note: Long-term market-based euro-area inflation expectations; inflation swap rates derived from euro inflation swaps; breakeven inflation rates derived from nominal and index-linked government bonds, average for France and Germany; five-year rates five years ahead.

long-term inflation expectations based on Consensus surveys suggests that DNB survey respondents are representative of professional macroanalysts. Figure 2 shows that long-term DNB survey expectations lie in a range similar to that of long-term Consensus survey-based expectations. This similarly holds for short-term inflation expectations.[3]

Long-term inflation expectations from our survey instead differ visibly from those based on financial market prices, namely breakeven inflation rates based on government bond yields, and forward inflation rates based on inflation swaps, which are shown in Figure 3. In both cases we show five-year/five-year forward inflation rates commonly used as a measure of monetary policy credibility. This is in line with the literature, which commonly finds significant differences in expectations measures extracted from surveys and financial market prices.

In addition to the weekly questions about their inflation expectations, participants in the DNB survey are asked questions about

[3]This is available from the authors upon request.

Figure 4. Examples of Distributions of Long-Term Inflation Expectations from DNB Survey, in Percent

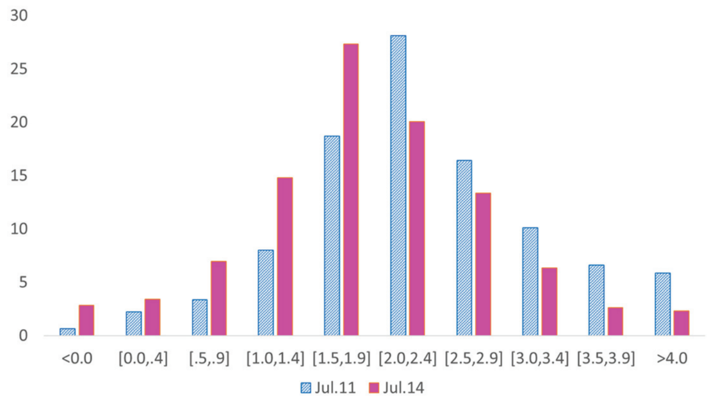

Note: Frequencies of the aggregate full distributions of expected long-term inflation rates per inflation interval from DNB surveys of July 2011 and July 2014, calculated from individual survey responses according to Equation (13); inflation interval in percent shown on x-axis.

the distribution of their short-term and long-term euro-area inflation expectations once per quarter. Survey respondents are asked to assign probabilities to $J = 10$ intervals $j, j = 1, \ldots, J$. These intervals are defined as <0.0, [0.0,.5[, [.5,1.0[, [1.0,1.5[, [1.5,2.0[, [2.0,2.5[, [2.5,3.0[, [3.0,3.5[, [3.5,4.0], and >4.0, in percent, where [,] denotes a closed interval and [, [denotes an interval closed on the left and open on the right. The frequency assigned by respondent i to interval j at horizon h and time t is denoted by $f_{it}^{j,h}$, where $h = LT$ or $h = ST$ for the long-term or short-term horizon, respectively. Examples of DNB survey responses for the distribution of long-term inflation expectations are shown in Figure 4. Figure 5 shows the mean of short-term and long-term deflation risk from the DNB survey over the sample period. Long-term deflation risk, dr_{it}^{LT}, is obtained directly from survey responses for the interval $j = 1$. Short-term deflation risk at a constant horizon, dr_{it}^{ST}, is obtained by interpolating between survey responses for current-year deflation risk, dr_{it}^{c}, and next-year deflation risk, dr_{it}^{n}, according to

$$dr_{it}^{ST} = \left(1 - \frac{(q-1)}{3}\right) dr_{it}^{c} + \frac{(q-1)}{3} dr_{it}^{n} \tag{2}$$

Figure 5. Euro-Area Deflation Risk from DNB Survey

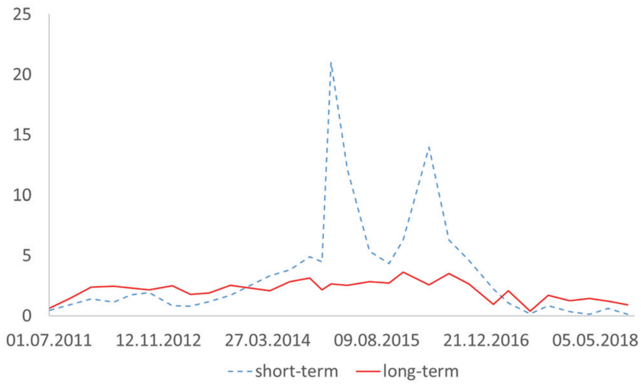

Note: Mean long-term and short-term euro-area deflation expectations from DNB survey, calculated from individual survey responses according to Equation (13).

with $q = 1, \ldots, 4$, and $q = 1$ for the first quarter, $q = 2$ for the second quarter, and so on.

3. Method

We analyze inflation expectations formation by means of panel data estimation over the period June 28, 2010 to December 10, 2018 and with around 25 respondents per week, using weekly data.

We test whether long-term DNB survey expectations are well anchored or not, by verifying whether the different conditions hold that are used in the literature to characterize anchoring. These include average expectations being close to the central bank's target ("level anchoring"); long-term expectations not co-moving with changes in actual inflation, inflation surprises, or short-term expectations ("shock anchoring"); expectations not being overly dispersed between individuals; agents being fairly confident about their best guess of future inflation and having little uncertainty about inflation in the long term; and agents not attaching a large weight to extreme inflation outcomes in the future.

We assess level anchoring by testing whether the ECB's inflation aim of close to but below 2 percent has acted as a focal point for long-term DNB survey expectations, by estimating

Table 1. Role of the ECB's Inflation Aim

Dependent Variable: π^{LT}	**Full Sample Period**	**Including Euro-Area Crisis[1]**	**Post-Euro-Area Crisis[2]**
Constant	2.097***	2.232***	2.007***
Wald Test of Const = 2 (p-value)	0.000	0.000	0.3279
No. of Observations	8,821	3,530	5,291

[1]Including euro-area sovereign debt crisis, June 28, 2010–December 31, 2013.
[2]Post-euro-area sovereign debt crisis, January 6, 2014–December 10, 2018. Pooled OLS regression; robust standard errors.
Note: ***, **, and * represent significance at the 1 percent, 5 percent, and 10 percent levels. Sample period: June 28, 2010–December 10, 2018, weekly data.

$$\pi_{it}^{LT} = c + \varepsilon_{it} \tag{3}$$

using pooled ordinary least squares (OLS) regression with robust standard errors. The results are shown in Table 1.

Consensus survey expectations could act as an alternative focal point for the formation of inflation expectations of DNB survey respondents. To test this hypothesis, we test whether changes in long-term Consensus survey inflation expectations affect changes in long-term DNB survey inflation expectations,

$$\triangle\pi_{it}^{LT} = \alpha_i + \beta\triangle\pi_t^{Cons,LT} + \varepsilon_{it}, \tag{4}$$

where $\pi_t^{Cons,LT}$ are long-term Consensus survey inflation expectations available at the time of the DNB survey in week t. We also include survey individual fixed effects (α_i) to control for any observed or unobserved time-invariant heterogeneity among survey respondents. We use fixed-effects within-group panel estimation. We also estimate Equation (4) for changes in short-term Consensus survey inflation expectations. The results are shown in Table 2.[4]

[4]In principle we could pool Equations (3) and (4) and test the joint hypothesis that the ECB's inflation aim acts as a focal point while survey participants do not react to changes in Consensus Forecasts. In practice, there is little variation in Consensus Forecasts' long-term expectations, and as a consequence, this joint test would have low power to identify any impact of the latter.

Table 2. Effects of Changes in Consensus Survey on Changes in Long-Term DNB Survey Inflation Expectations

Dependent Variable: $\Delta\pi^{LT}$		
$\triangle\pi^{Cons,ST}$	0.0552	—
$\triangle\pi^{Cons,LT}$	—	0.0623
No. of Observations	7,266	7,266

Note: ***, **, and * represent significance at the 1 percent, 5 percent, and 10 percent levels. Sample period: June 28, 2010–December 10, 2018, weekly changes. Fixed-effects within-group panel regression; robust standard errors. Using latest available Consensus survey.

To test for shock anchoring of inflation expectations, we first estimate whether long-term inflation expectations respond to changes in inflation,

$$\triangle\pi_{it}^{LT} = \alpha_i + \beta\triangle\pi_t + \varepsilon_{it}. \tag{5}$$

Here, $\triangle\pi_{it}^{LT}$ are weekly changes in long-term DNB survey expectations of respondent i in week t, and $\triangle\pi_t$ are weekly changes in euro-area HICP inflation (for the weeks in which new HICP inflation data are released, and zero otherwise). The hypothesis is that if long-term expectations are well anchored, they should be unresponsive to short-term developments in actual inflation, hence the estimate of β should not be significantly different from 0. Here and in the following regressions we again include survey individual fixed effects to control for any observed or unobserved time-invariant heterogeneity among survey respondents, and use fixed-effects within-group panel estimation. We use robust standard errors in this and all other regressions in this paper. We also estimate another variant of Equation (5) where we replace weekly changes in euro-area HICP inflation by weekly changes in the flash estimate of euro-area HICP inflation, $\triangle\pi_t^{flash}$. The results are shown in columns 1 and 2 of Table 3.

As a variant of the previous test, we also verify whether long-term DNB survey expectations respond to surprises in inflation, as

Table 3. Effects of Changes in Inflation, Inflation Surprises, and Short-Term DNB Survey Expectations on Changes in Long-Term DNB Survey Inflation Expectations

	(1)	(2)	(3)	(4)	(5)
Dependent Variable	$\boldsymbol{\Delta\pi^{LT}}$	$\boldsymbol{\Delta\pi^{LT}}$	$\boldsymbol{\Delta\pi^{LT}}$	$\boldsymbol{\Delta\pi^{LT}}$	$\boldsymbol{\Delta\pi^{LT}}$
$\Delta\pi$	0.0013	—	—	—	—
$\Delta\pi^{flash}$	—	0.0064	—	—	—
π^{sur}	—	—	–0.076	—	—
$\pi^{flash,sur}$	—	—	—	0.018	—
$\Delta\pi^{ST}$	—	—	—	—	0.0755*
No. of Observations	7,266	7,266	1,761	1,656	7,266

Note: ***, **, and * represent significance at the 1 percent, 5 percent, and 10 percent levels. Sample period: June 28, 2010–December 10, 2018, weekly changes. Fixed-effects within-group panel regression; robust standard errors. Inflation surprises relative to median Bloomberg survey expectations.

measured by actual euro-area HICP inflation minus median Bloomberg survey expectations, π_t^{sur}, according to

$$\triangle\pi_{it}^{LT} = \alpha_i + \beta\pi_t^{sur} + \varepsilon_{it}. \tag{6}$$

This empirical specification is similar to that typically used in the empirical literature on inflation expectations anchoring that relies on high-frequency market-based measures of inflation expectations. We also estimate Equation (6) when replacing surprises in euro-area HICP inflation with surprises in the flash estimate of euro-area HICP inflation, $\pi_t^{flash,sur}$, since there is evidence that flash data releases for inflation have a bigger impact on financial-market-based inflation expectations compared with the final data releases (Garcia and Werner 2018). The results are shown in columns 3 and 4 of Table 3.

As a further test of whether long-term DNB survey expectations are well anchored in the sense of shock anchoring, we also estimate whether they respond to changes in short-term DNB survey expectations,

$$\triangle\pi_{it}^{LT} = \alpha_i + \beta\triangle\pi_{it}^{ST} + \varepsilon_{it}, \tag{7}$$

where $\triangle\pi_{it}^{ST}$ are weekly changes in short-term DNB survey expectations, again using fixed-effects within-group panel estimation. This is a common test of expectations anchoring in the literature. The hypothesis here is that if long-term inflation expectations are well anchored to the central bank's inflation target, they should be unresponsive to changes in short-term inflation expectations, which reflect changing views of the short-term economic outlook. The results are shown in column 5 of Table 3.

To assess whether expectations anchoring is dispersed between individuals, we first consider possible heterogeneity in whether the ECB's inflation aim of close to but below 2 percent has acted as a focal point for long-term DNB survey expectations, by allowing the intercept in the regression of long-term expectations to vary by respondent,

$$\pi_{it}^{LT} = c_i + \varepsilon_{it}. \tag{8}$$

Moreover, we study possible heterogeneity in the anchoring of long-term DNB survey expectations by allowing the coefficient of changes in long-term expectations on changes in HICP inflation to vary by respondent,

$$\triangle\pi_{it}^{LT} = \alpha_i + \beta_i \triangle\pi_t + \varepsilon_{it}. \tag{9}$$

We also study possible heterogeneity in the response of long-term DNB survey expectations to changes in short-term DNB survey expectations, by allowing the coefficient on changes in short-term DNB survey expectations to vary by respondent,

$$\triangle\pi_{it}^{LT} = \alpha_i + \beta_{\mathrm{i}} \triangle\pi_{it}^{ST} + \varepsilon_{it}. \tag{10}$$

The regressions of Equations (8), (9), and (10) all use fixed-effects within-group panel estimation.

Our survey allows us to also assess the role of demographic characteristics in the anchoring properties of long-term inflation expectations, considering age and gender on which we have information in the DNB survey. To do so, we rerun the regressions of Equations (3), (5), and (7) separately for women and men, as well as separately for the group of younger respondents (below 40 years of age) and of older respondents (40 years of age or above).

Furthermore, we study the anchoring of long-term inflation expectations by considering measures based on the full distribution and the second moments of the distribution. We consider disagreement between respondents (the second moment of the distribution of different agents' levels expectations), as well as average individual uncertainty from the quarterly DNB survey of individuals' expected distributions. The underlying idea is that changes in the higher moments of the distribution of long-term inflation expectations could foreshadow changes in the anchoring of the mean of the distribution.

For their expectations to be well anchored, there should be little disagreement between agents. Our survey allows us to determine disagreement between individual respondents as well as average individual uncertainty about long-term expected inflation. We calculate disagreement between individual respondents about long-term expected inflation as the standard deviation of individual respondents' expected levels of inflation in the long term, from the weekly surveys of levels.

Another condition for well-anchored expectations is that agents should be fairly confident about their best guess of future inflation and have little uncertainty about inflation in the long term. As argued by Kumar et al. (2015), the idea is that agents should perceive little risk of either high or low inflation in the future, and hence consider the range of possible outcomes for inflation to be limited. Importantly, this condition—and the concept of anchored expectations more generally—refers to expectations over the long term, over which unpredictable shocks and consequent short- to medium-term deviations from the inflation target have faded.

We determine average individual uncertainty about long-term expected inflation at time t from the average of individuals' interquartile range of their expected probability distribution of inflation in the long term, using the quarterly survey of distributions at time t. To determine the interquartile range of individual i at time t, we first calculate the expected cumulative distribution of individual i at time $t, cdf_{it}^{j,LT}$, from the frequencies assigned by respondent i to the 10 intervals at time t for the long-term horizon, $f_{it}^{k,LT}$. From this cumulative distribution we determine the first quartile, $\mathrm{Q1}_{it}^{LT}$, as the midpoint of the inflation interval j in which the cumulative distribution first reaches 0.25. That is, we assume that the probability mass in each interval is concentrated at its midpoint. For the open intervals at either end of the distribution, we truncate the distribution by assuming that the interval has the same size as the other intervals, 0.5 percentage point (pp). Both these assumptions are based on D'Amico and Orphanides (2008). Similarly, we determine the third quartile, $\mathrm{Q3}_{it}^{LT}$, as the midpoint of the inflation interval in which the cumulative distribution first reaches 0.75. The interquartile range of the expected distribution of long-term inflation of individual i is then given by $iqr_{it}^{LT} = \mathrm{Q3}_{\mathrm{it}}^{LT} - \mathrm{Q1}_{it}^{LT}$. The average individual interquartile range at time t, iqr_t^{LT}, is then calculated as the average of the interquartile range over all N respondents. Average individual uncertainty about long-term expected inflation, $unc_t^{indiv,LT}$, is then calculated as the average individual interquartile range iqr_t^{LT} divided by 1.35 to make this measure more comparable to the standard deviation used as a measure for disagreement, since for a normal probability distribution the standard deviation equals the interquartile range divided by 1.35. Average individual uncertainty about long-term expected inflation is then given by

$$unc_t^{indiv,LT} = \frac{1}{1.35}\frac{1}{N}\sum_{i=1}^{N} iqr_{it}^{LT}. \tag{11}$$

Next, we study the anchoring of long-term inflation expectations by considering a measure based on the full aggregate distribution of inflation expectations from the quarterly DNB survey, namely the probability of future euro-area inflation being in a range that is consistent with the inflation target as a measure of anchoring. For expectations to be well anchored, agents should not attach a large weight to extreme inflation outcomes in the future. We therefore consider the survey-based probability of future euro-area inflation being in a certain range that is consistent with the inflation target as a measure of anchoring—in particular, the probability of expected long-term inflation lying between 1.5 percent and 2.5 percent (as in Grishchenko, Mouabbi, and Renne 2019). This probability, ptr_t^{LT}, is calculated as the sum of the frequencies assigned in the aggregated histogram at the long-term horizon at time t to inflation being in the two intervals $j = 5$ and $j = 6$, which together make up the interval between 1.5 percent and 2.5 percent, according to

$$ptr_t^{LT} = \sum_{j=5}^{6} f_t^{j,h}. \tag{12}$$

Here, the frequency of the aggregate histogram at time t in each interval j, $f_t^{j,h}$, is calculated according to (see Krueger and Nolte 2016)

$$f_t^{j,h} = \frac{1}{N} \sum_{i=1}^{N} f_{it}^{j,h}, \tag{13}$$

where N is the number of respondents to the survey questions about the distribution of inflation expectations. That is, we construct a histogram of the aggregate distribution of inflation expectations by a linear combination of the histograms of the individual distributions, with equal weights.

Finally, we consider information from the tails of individuals' expected distributions on the anchoring of long-term inflation expectations. To test whether long-term DNB survey expectations are well anchored or not, we estimate whether changes in long-term deflation risk respond to changes in short-term deflation risk derived from the DNB survey,

$$\triangle dr_{it}^{LT} = \alpha_i + \beta \triangle dr_{it}^{ST} + \varepsilon_{it}, \tag{14}$$

where $\triangle dr_{it}^{LT}$ are quarterly changes in long-term deflation risk, and $\triangle dr_{it}^{ST}$ are quarterly changes in short-term deflation risk, again using fixed-effects within-group panel estimation. The results are shown in Table 7.

4. Results

Overall, most but not all empirical tests suggest that over the period 2010–18, inflation expectations measured by our DNB survey have remained well anchored to the ECB's inflation aim. Our main findings are presented in this section.

First, we find evidence of well-anchored long-term inflation expectations based on the level-anchoring condition. The ECB's inflation aim has acted as a focal point for long-term DNB survey expectations, especially after the euro-area sovereign debt crisis, where we cannot reject that the mean of long-term DNB survey expectations equals 2 percent based on Equation (3) (Table 1). This is the case even though the mean short-term DNB inflation expectations were well below 2 percent after the euro-area crisis, at around 1.25 percent. But in the period including the euro-area sovereign debt crisis, mean long-term DNB survey expectations were slightly (around 25 basis points) above 2 percent (Table 1).

We find that Consensus surveys, which are provided to survey respondents as part of a common information set, do not act as focal points for long-term DNB survey expectations. There are no significant reactions of changes in long-term DNB survey expectations to changes in either long-term or short-term Consensus survey expectations based on Equation (4) (Table 2).

Second, tests for shock anchoring show some subtle signs of not perfectly well-anchored long-term inflation expectations for the group of survey respondents as a whole. There are no significant reactions of changes in long-term DNB survey expectations to changes in inflation, or in the flash estimate of inflation based on Equation (5) (Table 3, columns 1 and 2). Similarly, there are no significant reactions of changes in long-term DNB survey expectations to surprises in inflation, or in the flash estimate of inflation using Equation (6) (Table 3, columns 3 and 4). However, the coefficient of changes in long-term DNB survey expectations on changes in short-term DNB survey expectations is statistically significant, although only at the 10 percent significance level, and economically small (with a value of around 0.08) using Equation (7) (Table 3, column 5). This is consistent with results on subtle signs of a change in the anchoring properties of long-term inflation expectations found in other papers for the euro area (see ECB 2021).

Third, we find evidence that notwithstanding a fairly homogenous panel of survey participants and a common information set, there is heterogeneity across survey participants in their expected level of future inflation and the responsiveness of their expectations to shocks. The results for the individual intercepts c_i of Equation (8), which provide a measure of level anchoring, are shown as a histogram in Figure 6. We can see that there is some heterogeneity in this intercept. We therefore find some evidence of heterogeneity across survey respondents on whether the ECB's inflation aim of close to but below 2 percent has acted as focal point for long-term DNB survey expectations.

Figure 6. Heterogeneity in ECB's Inflation Aim Acting as Focal Point for Long-Term DNB Survey Inflation Expectations

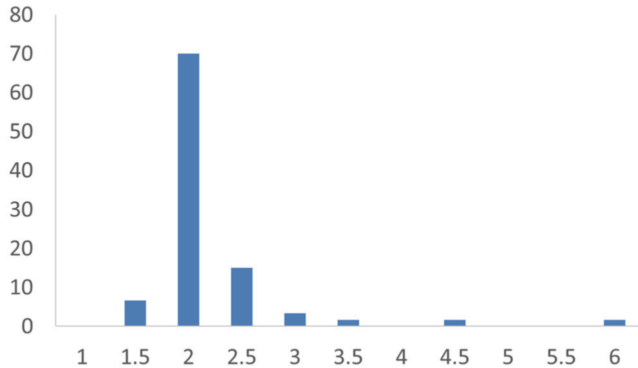

Note: Histogram of individuals' constant term c_i from Equation (8), for bins of width 0.5 pp with midpoint of bin shown on x-axis (in percent), from regression over full sample period of June 28, 2010–December 10, 2018.

The results for the coefficients β_i in estimates of Equation (9) suggest that there is also some heterogeneity in the response of long-term DNB inflation expectations to inflation (see Figure 7). Similarly, the results for the coefficients β_i in estimates of Equation (10) suggest that there is also some heterogeneity in the response of long-term to short-term DNB inflation expectations (see Figure 8). We therefore find some evidence of heterogeneity in the shock-anchoring properties across survey respondents on these measures.

Fourth, heterogeneity across survey participants in the anchoring of long-term inflation expectations can in part be explained by demographic characteristics. The results of Equation (3) estimated separately for women and men are shown in Table 4 (columns 1 and 2). The intercept is significantly above the ECB's inflation aim of 2 percent for both women and men, but it is slightly higher for men. This also suggests that within this group of professionals, long-term inflation expectations of women are slightly better anchored than those of men. The results of Equation (3) estimated separately for the group of younger and older respondents are also shown in Table 4 (columns 3 and 4). The intercept is slightly below the ECB's

Figure 7. Heterogeneity in Effect of Changes in Inflation on Changes in Long-Term DNB Survey Inflation Expectations

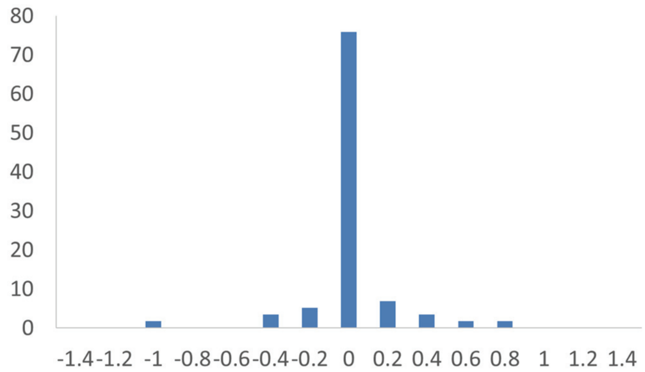

Note: Histogram of individuals' coefficient β_i from Equation (9), for bins of width 0.2 with midpoint of bin shown on x-axis, from regression over full sample period of June 28, 2010–December 10, 2018.

Figure 8. Heterogeneity in Effect of Changes in Short-Term on Changes in Long-Term DNB Survey Inflation Expectations

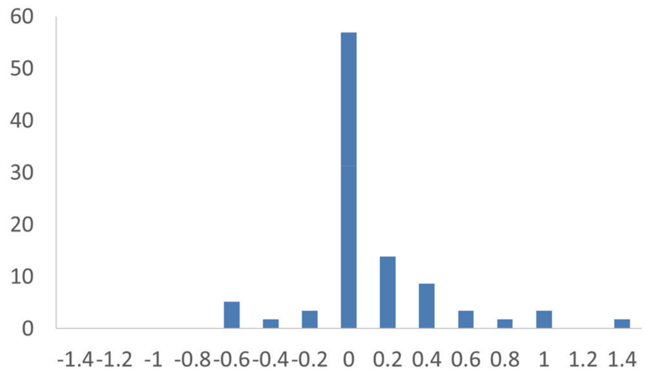

Note: Histogram of individuals' coefficient β_i from Equation (10), for bins of width 0.2 with midpoint of bin shown on x-axis, from regression over full sample period of June 28, 2010–December 10, 2018.

Table 4. Demographic Characteristics and Role of ECB's Inflation Aim

	(1)	(2)	(3)	(4)
Demographic Characteristics	Female	Male	Young	Old
Dependent Variable	π^{LT}	π^{LT}	π^{LT}	π^{LT}
Constant	2.0519***	2.1053***	1.9645***	2.1784***
Wald Test of Const=2 (p-value)	0.0000	0.0000	0.0000	0.0000
No. of Observations	1,402	7,419	3,364	5,457

Note: ***, **, and * represent significance at the 1 percent, 5 percent, and 10 percent levels. Sample period: June 28, 2010–December 10, 2018, weekly data. Pooled OLS regression; robust standard errors.

Table 5. Role of Demographic Characteristics for Effects of Changes in Inflation on Changes in Long-Term DNB Survey Inflation Expectations

	(1)	(2)	(3)	(4)
Demographic Characteristics	Female	Male	Young	Old
Dependent Variable	$\Delta\pi^{LT}$	$\Delta\pi^{LT}$	$\Delta\pi^{LT}$	$\Delta\pi^{LT}$
$\Delta\pi$	–0.0237	0.0079	0.0220	–0.0010
No. of Observations	1,092	6,174	2,688	4,578

Note: ***, **, and * represent significance at the 1 percent, 5 percent, and 10 percent levels. Sample period: June 28, 2010–December 10, 2018, weekly data. Fixed-effects within-group panel; robust standard errors.

inflation aim of 2 percent for younger respondents, but it is slightly above 2 percent for older respondents, and both results are significant. This also suggests that within this group of professionals, long-term inflation expectations of younger respondents are slightly better anchored than those of older ones.

The results of Equation (5) estimated separately for women and men are shown in Table 5 (columns 1 and 2). Those estimated separately for the group of younger and older respondents are also shown in Table 5 (columns 3 and 4). The coefficient for the effects of changes in HICP inflation on changes in long-term expectations is insignificant for all the four different demographic groups, as in the sample

Table 6. Role of Demographic Characteristics for Effects of Changes in Short-Term on Changes in Long-Term DNB Survey Inflation Expectations

	(1)	(2)	(3)	(4)
Demographic Characteristics	Female	Male	Young	Old
Dependent Variable	$\Delta\pi^{LT}$	$\Delta\pi^{LT}$	$\Delta\pi^{LT}$	$\Delta\pi^{LT}$
$\Delta\pi^{ST}$	0.0251	0.0861***	0.0832	0.0708**
No. of Observations	1,092	6,174	2,688	4,578

Note: ***, **, and * represent significance at the 1 percent, 5 percent, and 10 percent levels. Sample period: June 28, 2010–December 10, 2018, weekly changes. Fixed-effects within-group panel regression; robust standard errors.

as a whole shown in Table 3. We therefore find that these demographic characteristics do not affect the anchoring properties on this measure.

The results of Equation (7) estimated separately for women and men are shown in Table 6 (columns 1 and 2). The coefficient for the effects of short-term expectations on long-term expectations is larger and more significant for men than for women. We therefore find that on this anchoring measure, the expectations of women are better anchored than those of men within this group of professionals. Household surveys tend to find that women's inflation expectations are less well anchored than those of men (see, e.g., Galati, Moessner, and van Rooij 2021 for euro-area expectations). The difference is likely to arise since we are considering a group of professionals, rather than households representative of the whole population. Moreover, most household surveys consider short- or medium-term inflation expectations rather than long-term expectations. The results of Equation (7) estimated separately for the group of younger and older respondents are also shown in Table 6 (columns 3 and 4). The coefficient for the effects of short-term on long-term expectations is of similar magnitude for older and younger respondents, but it is only significant for older respondents. This suggests that on this anchoring measure, the expectations of older respondents are slightly less well anchored than those of younger ones within this group of professionals.

Figure 9. Average Individual Uncertainty and Disagreement for Long-Term DNB Survey Inflation Expectations

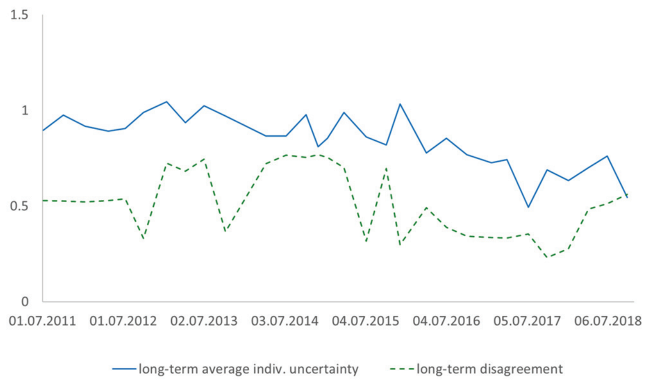

Note: Average individual uncertainty shown is $unc_t^{indiv,LT}$ of Equation (11) from the quarterly DNB survey of distributions of long-term expected inflation. Disagreement is the standard deviation of individuals' expected levels from the weekly DNB survey of long-term levels available closest to the time of the quarterly distributions survey.

Fifth, there is evidence that the patterns of disagreement between respondents and of individual uncertainty about future inflation show some differences. Figure 9 shows the time series of average individual uncertainty calculated from Equation (11), $unc_t^{indiv,LT}$, as well as disagreement, for long-term DNB survey inflation expectations. Disagreement is the standard deviation of individuals' expected levels from the weekly survey of levels available closest to the time of the quarterly distributions survey. We can see that average individual uncertainty has fallen over the sample period. This measure therefore points to long-term inflation expectations having become better anchored over the sample period. Disagreement shows a slightly different pattern, rising toward the end of the sample period.

Next, we find evidence that over time agents have tended to attach a lower weight to extreme inflation outcomes in the future. One way to see this is by tracking the survey-based probability of future euro-area inflation being in a certain range that is consistent

Figure 10. Probability of Expected Long-Term Inflation Lying between 1.5 Percent and 2.5 Percent from DNB Survey

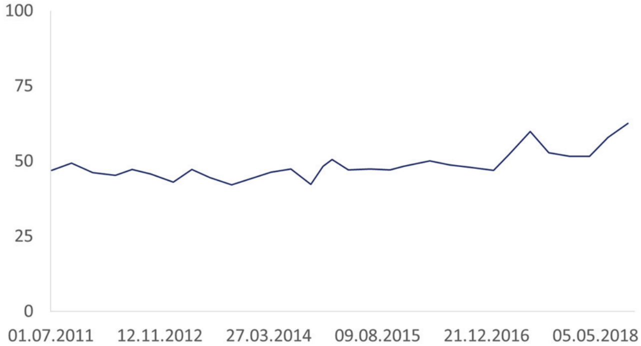

Note: Mean expected probability of inflation lying between 1.5 percent and 2.5 percent in the long term, calculated from individual survey responses to DNB survey according to Equation (12).

with the inflation target as a measure of anchoring. Figure 10 shows the time series of the probability ptr_t^{LT} of expected long-term inflation lying between 1.5 percent and 2.5 percent derived from the DNB survey according to Equation (12). We can see that this probability has increased slightly over the sample period. This measure therefore also points to long-term inflation expectations having become better anchored over the sample period.

Finally, we present information from the tails of individuals' expected distributions on the anchoring of long-term inflation expectations using Equation (14). We find that changes in short-term deflation risk have no significant effect on changes in long-term deflation risk from the DNB survey, which also suggests that long-term euro-area inflation expectations have remained well anchored (Table 7).

5. Conclusions

We shed new light on the behavior of short- and long-term euro-area inflation expectations between 2010 and 2018 by using microevidence from a new type of survey at high (weekly) frequency. These

Table 7. Effects of Changes in Short-Term DNB Survey Deflation Risk on Changes in Long-Term DNB Survey Deflation Risk

Dependent Variable: Δdr^{LT}	
Δdr^{ST}	0.009
No. of Observations	369

Note: ***, **, and * represent significance at the 1 percent, 5 percent, and 10 percent levels. Sample period: 2011:Q3–2018:Q4, quarterly changes. Fixed-effects within-group panel regression; robust standard errors.

data allow us to shed new light on the different dynamics of professional forecasters' inflation expectations, as reflected in survey measures of inflation expectations and market-based measures. A caveat is that the external validity of our setup depends on the representativeness of our sample of DNB survey participants for the general population of macroanalysts. Descriptive evidence suggests that this is indeed the case.

We run a battery of tests of anchoring of long-term inflation expectations to the ECB's inflation aim. In the literature, some of these tests have so far been applied only to market-based measures of inflation expectations. Overall, we find at most subtle signs of inflation expectations that are not firmly anchored.

We find that in the sense of level anchoring, long-term inflation expectations remained well anchored to the ECB's inflation aim, which has acted as a focal point. By contrast, we find no evidence that professional forecasts (reported by Consensus Economics) acted as focal points.

But when we look at tests for shock anchoring, we detect some subtle signs of long-term inflation expectations not being perfectly well anchored. This shows that subtle changes in the anchoring of inflation expectations by professionals can be detected by using survey-based measures at a weekly frequency. These changes are much more nuanced than those found in empirical exercises that rely on market-based measures of inflation expectations. We also find that notwithstanding a fairly homogenous panel of survey participants and a common information set, there is heterogeneity across survey participants in their expected level of future inflation and the responsiveness of their expectations to shocks.

Using measures based on the full distribution of inflation expectations, namely average individual uncertainty based on the full expected distribution, the probability of expected long-term inflation lying between 1.5 percent and 2.5 percent, and the effect of short-term on long-term deflation risk, we find that long-term inflation expectations remained well anchored and became better anchored at the end of the sample period in 2018 compared with the start of the sample period in 2011.

Appendix. Example of Survey E-mail

Weekly Inflation Expectations Survey, October 15, 2018

Dear survey participant,

Please find attached the updated background information on euro area inflation.

Please send us your answers to the questions below by Monday 5pm.

1. What HICP inflation do you expect for the euro area for the whole calendar year 2019?

2. What HICP inflation do you expect for the euro area for the whole calendar year 2020?

3. What HICP inflation do you expect for the euro area for the whole calendar year 2028?

Background Information

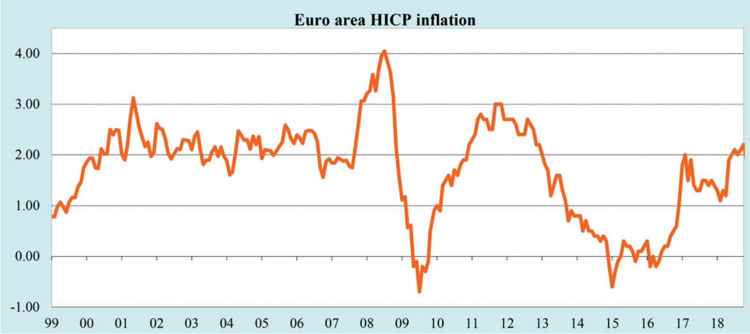

Consensus Forecast Euro-Area Inflation (% change from previous calendar year)

	2015	2016	2017	2018	2019	5–10 Years Ahead
Jan. 16		x.x	x.x			
Feb. 16		x.x	x.x			
Mar. 16		x.x	x.x			
Apr. 16		x.x	x.x			x.x
May 16		x.x	x.x			
Jun. 16		x.x	x.x			
Jul. 16		x.x	x.x			
Aug. 16		x.x	x.x			
Sep. 16		x.x	x.x			
Oct. 16		x.x	x.x			x.x
Nov. 16		x.x	x.x			
Dec. 16		x.x	x.x			
Jan. 17			x.x	x.x		
Feb. 17			x.x	x.x		
Mar. 17			x.x	x.x		
Apr. 17			x.x	x.x		x.x
May 17			x.x	x.x		
Jun. 17			x.x	x.x		
Jul. 17			x.x	x.x		
Aug. 17			x.x	x.x		
Sep. 17			x.x	x.x		
Oct. 17			x.x	x.x		x.x
Nov. 17			x.x	x.x		
Dec. 17			x.x	x.x		
Jan. 18				x.x	x.x	
Feb. 18				x.x	x.x	
Mar. 18				x.x	x.x	
Apr. 18				x.x	x.x	x.x
May 18				x.x	x.x	
Jun. 18				x.x	x.x	
Jul. 18				x.x	x.x	
Aug. 18				x.x	x.x	
Sep. 18				x.x	x.x	

Note: Numbers for Consensus forecasts in this table, which were provided to survey respondents, have been replaced by “x.x” in this paper for license reasons.

Consumer Prices

	Jul.	Aug.	Sep.
Germany (Final)			
%m/m, nsa	0.3	0.1	0.4
%oya, nsa	2.0	2.0	2.3
HICP (%oya)	2.1	1.9	2.2
France (Final)			
%m/m, nsa	–0.1	0.5	–0.2
Index ex Tobacco, na	2.96	3.48	3.25
%oya, nsa	2.3	2.3	2.2
HICP (%oya)	2.6	2.6	2.5
Spain (Final)			
%m/m, nsa	–0.7	0.1	0.2
%oya, nsa	2.2	2.2	2.3
HICP (%oya)	2.3	2.2	2.3

References

Adrian, T. 2022. "Are Household Inflation Expectations De-anchoring?" Speech at the European Banking Institute, Frankfurt am Main, May 17.

Antunes, A. 2015. "Co-movement of Revisions in Short- and Long-Term Inflation Expectations." *Bank of Portugal Economic Studies* 1 (May): 1–20.

Beechey, M., B. Johannsen, and A. Levin. 2011. "Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area than in the United States?" *American Economic Journal: Macroeconomics* 3 (2): 104–29.

Bems, R., F. Caselli, F. Grigoli, and B. Gruss. 2021. "Expectations' Anchoring and Inflation Persistence." *Journal of International Economics* 132 (September): Article 103516.

Bernanke, B. 2007. "Inflation Expectations and Inflation Forecasting." Remarks at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts, July 10.

———. 2022. “Inflation Expectations and Monetary Policy.” Keynote speech at the Brookings Institution Conference on Inflation Expectations: Determinants and Consequence, May 19.

Binder, C., J. Campbell, and J. Ryngaert. 2022. “Consumer Inflation Expectations: Daily Dynamics.” Mimeo, December 15.

Boero, G., J. Smith, and K. Wallis. 2008. “Uncertainty and Disagreement in Economic Prediction: The Bank of England Survey of External Forecasters.” *Economic Journal* 118 (530): 1107–27.

Bottone, M., and A. Rosolia. 2019. “Monetary Policy, Firms’ Inflation Expectations and Prices: Causal Evidence from Firm-Level Data.” Bank of Italy Working Paper No. 1218.

Bruine de Bruin, W., C. Manski, G. Topa, and W. van der Klaauw. 2011. “Measuring Consumer Uncertainty About Future Inflation.” *Journal of Applied Econometrics* 26 (3): 454–78.

Buono, I., and S. Formai. 2018. “The Evolution of the Anchoring of Inflation Expectations.” *Journal of Macroeconomics* 57 (3): 39–54.

Busetti, F., D. Delle Monache, A. Gerali, and A. Locarno. 2017. “Trust, but Verify. De-anchoring of Inflation Expectations under Learning and Heterogeneity.” ECB Working Paper No. 1994.

Candia, B., O. Coibion, and Y. Gorodnichenko. 2021. “The Inflation Expectations of US Firms: Evidence from a New Survey.” NBER Working Paper No. 28836, May.

Clements, M. 2012. “Do Professional Forecasters Pay Attention to Data Releases?” *International Journal of Forecasting* 28 (2): 297–308.

Cœuré, B. 2019. “Inflation Expectations and the Conduct of Monetary Policy.” Speech at an event organized by the SAFE Policy Center, Frankfurt am Main, July 11.

Conflitti, C., and R. Zizza. 2021. “What’s Behind Firms’ Inflation Forecasts?” *Empirical Economics* 5 (November): 2449–75.

Corsello, F., S. Neri, and A. Tagliabracci. 2019. “Anchored or De-anchored? That is the Question.” Bank of Italy Occasional Paper No. 516.

D’Acunto, F., U. Malmendier, and M. Weber. 2022. “What Do the Data Tell Us About Inflation Expectations?” NBER Working Paper No. 29825, March.

D'Amico, S., and A. Orphanides. 2008. "Uncertainty and Disagreement in Economic Forecasting." Finance and Economics Discussion Series No. 2008-56, Board of Governors of the Federal Reserve System.

De Fiore, F., M. Lombardi, and J. Schuffels. 2019. "Are Households Indifferent to Monetary Policy Announcements?" BIS Working Paper No. 956, August.

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

Dovern, J., and G. Kenny. 2020. "Anchoring Inflation Expectations in Unconventional Times: Micro Evidence for the Euro Area." *International Journal of Central Banking* 16 (5): 309–47.

European Central Bank. 2021. "Inflation Expectations and Their Role in Eurosystem Forecasting." Occasional Paper No. 264, September.

Fleming, M., and E. Remolona. 1997. "What Moves the Bond Market?" *Economic Policy Review* (Federal Reserve Bank of New York) 3 (4): 31–50.

Galati, G., Z. Gorgi, R. Moessner, and C. Zhou. 2018. "Deflation Risk in the Euro Area and Central Bank Credibility." *Economics Letters* 167 (June): 124–26.

Galati, G., R. Moessner, and M. van Rooij. 2021. "Anchoring of Consumers' Long-Term Euro Area Inflation Expectations during the Pandemic." DNB Working Paper No. 715.

Garcia, J. 2003. "An Introduction to the ECB's Survey of Professional Forecasters." ECB Occasional Paper No. 8, September.

Garcia, J. A., and S. Werner. 2018. "Inflation News and Euro Area Inflation Expectations." IMF Working Paper No. 18/167.

Gerlach, S. 2007. "Interest Rate Setting by the ECB, 1999-2006: Words and Deeds." *International Journal of Central Banking* 3 (3): 1–46.

Goel, T., and K. Tsatsaronis. 2022. "Anchoring of Inflation Expectations: Has Past Progress Paid Off?" BIS Bulletin No. 51, March 17.

Grishchenko, O., S. Mouabbi, and J.-P. Renne. 2019. "Measuring Inflation Anchoring and Uncertainty: A U.S. and Euro Area Comparison." *Journal of Money, Credit and Banking* 51 (5): 1053–96.

Gürkaynak, R., A. Levin, A. Marder, and E. Swanson. 2007. "Inflation Targeting and the Anchoring of Inflation Expectations in the Western Hemisphere." *FRBSF Economic Review* (Federal Reserve Bank of San Francisco) 2007: 25–47.

Gürkaynak, R., A. Levin, and E. Swanson. 2010. "Does Inflation Targeting Anchor Long-Term Inflation Expectations?" *Journal of the European Economic Association* 8 (6): 1208–42.

Knüppel, M., and A. Vladu. 2016. "Approximating Fixed-Horizon Forecasts Using Fixed-Event Forecasts." Deutsche Bundesbank Discussion Paper No. 28/2016.

Krueger, F., and I. Nolte. 2016. "Disagreement versus Uncertainty: Evidence from Distribution Forecasts." *Journal of Banking and Finance* 72 (Supplement): S172–S186.

Kumar, S., H. Afrouzi, O. Coibion, and Y. Gorodnichenko. 2015. "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand." *Brookings Papers on Economic Activity* 46 (Fall): 151–225.

Lamla, M., and D. Vinogradov. 2019. "Central Bank Announcements: Big News for Little People?" *Journal of Monetary Economics* 108 (December): 21–38.

Lewis, D., C. Makridis, and K. Mertens. 2019. "Do Monetary Policy Announcements Shift Household Expectations?" Federal Reserve Bank of New York Staff Report No. 897.

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): 55–102.

Moessner, R., and E. Takáts. 2020. "How Well-Anchored Are Long-Term Inflation Expectations?" BIS Working Paper No. 869, June.

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.

Neri, S., G. Bulligan, S. Cecchetti, F. Corsello, A. Papetti, M. Riggi, C. Rondinelli, and A. Tagliabracci. 2022. "On the Anchoring of Inflation Expectations in the Euro Area." Bank of Italy Occasional Paper No. 712, September.

Powell. 2020. "New Economic Challenges and the Fed's Monetary Policy Review." Speech at the "Navigating the Decade Ahead:

Implications for Monetary Policy" economic policy symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 27.

Rich, R., and J. Tracy. 2018. "A Closer Look at the Behavior of Uncertainty and Disagreement: Micro Evidence from the Euro Area." Federal Reserve Bank of Dallas Working Paper No. 1811.

Schnabel, I. 2021. "New Narratives on Monetary Policy — The Spectre of Inflation." Speech at the 148th Baden-Baden Entrepreneurs' Talk, Frankfurt am Main, September 13.

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

Weber, M., F. D'Acunto, Y. Gorodnichenko, and O. Coibion. 2022. "Reality Check: How People Form Inflation Expectations and Why You Should Care." Becker Friedman Institute Working Paper No. 2022-40.

Williams, J. 2022. "A Steady Anchor in a Stormy Sea." Remarks at the SNB-FRB-BIS High-Level Conference on Global Risk, Uncertainty, and Volatility, Zurich, Switzerland, September 9.

Do Buffer Requirements for European Systemically Important Banks Make Them Less Systemic?*

Carmen Broto, Luis Fernández Lafuerza,
and Mariya Melnychuk
Banco de España

With a panel data model for a sample of listed European banks, we demonstrate that capital requirements for systemically important institutions (SIIs) effectively reduce the perceived systemic risk of these institutions, which we proxy with the SRISK indicator in Brownlees and Engle (2017). We also study the impact of the adjustment mechanisms that banks use to comply with SII requirements. The results show that banks mainly respond to higher SII buffers by increasing their equity. Once we control for the options SIIs employ to fulfill these requirements and SII characteristics, we find a residual effect of having SII status.

JEL Codes: C54, E58, G21, G32.

1. Introduction

The Global Financial Crisis (GFC) made systemic risk a central topic of research and policy. Systemic risk can be analyzed either in its time/cyclical dimension or in its cross-sectional/structural dimension (see European Systemic Risk Board 2013).[1] In this paper

*We thank Ángel Estrada, David Martínez-Miera, Javier Mencía, Carlos Pérez Montes, and seminar participants at the Banco de España for their helpful comments. The opinions expressed in this paper are solely the responsibility of the authors and do not represent the views of the Banco de España. Author e-mails: carmen.broto@bde.es, luisg.fernandez@bde.es, mariya.melnychuk@bde.es@bde.es.

[1]Whereas the time dimension is related to the buildup of risks over time and the procyclical accumulation of financial vulnerabilities, the structural dimension of systemic risk focuses on how a specific shock to the financial system can spread and become systemic (International Monetary Fund 2011).

we focus on this second dimension of systemic risk, specifically on systemically important institutions (SIIs). The GFC evidenced that the failure of these large and complex banks could spill over into the whole financial sector and also harm the real economy. For this reason, these SIIs can be considered to be "too big to fail" and could engage in moral hazard behavior (see Stern and Feldman 2004), so that during boom periods these institutions could have incentives to take excessive risks, as they expect to receive support during crisis episodes. These SII characteristics justify the adoption of specific policy measures.

To address this competitive advantage of SIIs and the associated risk that they create, the Basel Committee on Banking Supervision (BCBS) launched in 2011 its framework for dealing with systemically important banks, with new additional capital requirements with a macroprudential focus (see BCBS 2011).[2] The rationale behind these additional capital buffers for SIIs was precisely to account for the negative externalities stemming from their size and interconnectedness, as well as to increase their resilience and loss-absorbing capacity. Namely, there are two possible structural buffers to address SIIs' particularities: (i) the capital requirements for global systemically important institutions (G-SIIs), which are systemically relevant institutions at the global level,[3] and (ii) the requirements for other systemically important institutions (O-SIIs), which are institutions that are more likely to create risks to financial stability at the national level.[4] In Europe, the region that constitutes the focus of our analysis, G-SIIs are also O-SIIs, and the higher of the two buffers is applied.[5] Additionally, under the CRD IV, the systemic risk buffer

[2]The BCBS framework was implemented in the European Union (EU) with the transposition of Capital Requirements Directive 2013/36/EU (CRD IV), which entered into force in 2013.

[3]Since 2011 the Financial Stability Board (FSB) identifies the list of G-SIIs annually in consultation with the BCBS (see BCBS 2013). The G-SII buffers were first activated in 2016.

[4]Since 2014, O-SIIs are annually selected in accordance with the European Banking Authority (EBA) guidelines (see EBA 2014). These lists and corresponding buffers are revised annually by the national regulatory authorities and communicated to the ESRB, and also submitted to and disclosed by the EBA. O-SII buffers became active in 2016.

[5]The criteria for identifying SIIs, both G-SIIs and O-SIIs, follow an indicator-based measurement approach that takes into account five dimensions of systemic

(SyRB) aims to tackle systemic risks of a long-term structural and non-cyclical nature that are not covered by the CRR.[6] Henceforth, we use the term "SII buffer" to refer to the buffer applicable to SIIs, which is a combination of the O-SII, G-SII, and SyRB capital requirements, depending on the institution and country.[7]

Most of the literature on the effect of higher capital requirements on banks' performance analyzes their impact in general, not that of the SII buffer in particular. These papers tend to conclude that under tighter capital requirements, banks reduce their risk-weighted assets and cut lending in the short run—see, for instance Aiyar, Calomiris, and Wieladek (2014, 2016), Bridges et al. (2014), Gropp et al. (2019), and Mayordomo and Rodríguez-Moreno (2020), among others. However, in the long run capital buffers smooth credit supply cycles and have a positive effect on firm-level aggregate financing and performance (see Drehmann and Gambacorta 2012 and Jiménez et al. 2017).

There is little empirical evidence on the specific impact of SII capital buffers. For instance, there is some literature on the effect of the activation of SII buffers on lending—see Cappelletti et al. (2019, 2020).[8] Additionally, a few studies analyze the impact of SII buffers on banks' solvency and, separately, on the financial markets' response. For instance, Dautović (2020) concludes that an increase in SII buffers was associated with increases in both common equity

importance, namely size, interconnectedness, substitutability, complexity, and cross-jurisdictional activity.

[6]CRR: Capital Requirements Regulation (EU) No. 575/2013. See Article 133 of the CRD IV for further details. Unlike the G-SII and O-SII buffers, the SyRB is an EU instrument beyond the Basel III Framework. It aims to address the risks stemming from structural features that have the potential to amplify shocks and losses such as high indebtedness, interconnectedness, or exposure to common shocks, among others. CRD IV sets out the rules to accumulate this buffer with the G-SII and O-SII buffer rates.

[7]The BCBS calls these two capital requirements the global systemically important bank (G-SIB) buffer and the domestic systemically important bank (D-SIB) buffer, instead of the EU denomination (i.e., G-SII for G-SIB and O-SII for D-SIB)—see BCBS (2012). In this paper we use the latter.

[8]Cappelletti et al. (2019) find that O-SIIs reduce lending to household and financial sectors in the short term, while in the medium term the effect is much smaller and heterogeneous. However, Cappelletti et al. (2020) suggest that O-SIIs curtail lending to credit institutions the most, leaving loan supply to non-financial corporations (NFCs) almost unchanged.

tier 1 (CET1) capital levels and the average risk weights of the asset portfolio.[9] Regarding the impact of SII announcements on financial markets, the empirical evidence suggests that higher capital requirements lead to lower stock prices and credit default swap (CDS) spread increases, although this market response is temporary—see Andrieş et al. (2020) and Gündüz (2020).

The literature on the effect of capital requirements on systemic risk is even scarcer. As far as we know, only Bostandzic et al. (2022) analyze the impact of higher capital requirements on a set of systemic risk measures. These authors use the 2011 EBA capital exercise to conclude that one-off capital increases deteriorate a set of market-based measures of systemic risk. However, SII buffers, which were gradually phased in from 2014, are out of the scope of this analysis. That is, to date, the effectiveness of SII buffers at lowering their systemic risk is still an open question, although this instrument was designed to address the systemic risk posed by these large and interconnected institutions.

Our paper has a dual objective. First, we analyze whether higher capital buffers for SIIs reduce their contribution to systemic risk. For this purpose, we fit a panel data model with fixed effects for all listed European banks, be they SIIs or not. The dependent variable is an indicator that quantifies systemic risk, namely the SRISK indicator in Brownlees and Engle (2017). This metric can be easily computed with publicly available bank- and market-based data. Like other market-based measures, SRISK is available at high frequencies and can be calculated for listed institutions only.[10]

Second, we analyze the impact of the adjustment mechanisms that banks employ to comply with SII buffers on their contribution to systemic risk. Understanding which of these mechanisms dominates banks' behavior toward increases in SII capital requirements is central to evaluating the implications of SII buffers. Broadly speaking, in response to higher capital requirements, banks have four

[9]The findings in Dautović (2020) suggest that banks comply with the regulation by raising equity capital, but at the same time reallocate their portfolio toward riskier assets, thus the overall net effect on solvency is unclear.

[10]Since the GFC, the literature on market-based measures to gauge systemic risk has grown. See Bisias et al. (2012) and Benoit et al. (2017) for two surveys.

main options at their disposal (see Bank for International Settlements 2012, Cohen and Scatigna 2016, and Braouezec and Kiani 2021).[11] Namely, a bank can either (i) issue new equity; (ii) increase its retained earnings; (iii) run down its assets; or (iv) reduce its risk-weighted assets. However, these alternatives have associated costs. For instance, new share issuance might be very costly for banks in the context of EU banks' historically low valuations, especially after the GFC. On the other hand, the retained earnings strategy would be more favorable from a regulator's perspective, but, as in the case of lowering dividends, it might take many years to increase the capital ratio and might lead to negative reaction from investors. Also, if banks' response to the higher requirements is to run down their assets, lending to the real sector could be negatively affected (see Gropp et al. 2019). Finally, shifting the composition of assets toward lower risk-weighted exposures could decrease expected profitability (see Bostandzic et al. 2022).

According to our results, SII buffers do decrease European banks' contribution to systemic risk in the medium term. Furthermore, we find that this effect is partially driven by the increase in banks' equity, and, contrary to Dautović (2020), we do not find evidence that banks take more risks. From a financial stability perspective, this is an important implication that suggests that banks respond to SII buffers as intended. Finally, once we control for the adjustment mechanisms that banks use to comply with the SII buffer, the residual effect of having SII status on perceived systemic risk is still negative and significant. This outcome implies that being an SII provides a positive signal to markets, which further reduces the institution's contribution to systemic risk.

The remainder of the paper is organized as follows. Section 2 describes how we quantify the contribution of a bank to systemic risk by means of the SRISK measure. Section 3 details our data set. Section 4 then describes the empirical model to analyze the relationship between buffers and systemic risk, and Section 5 summarizes the main results. Finally, Section 6 contains our conclusions.

[11] These four possible responses to higher requirements entail the assumption of a constant score, which is the indicator-based measurement that represents its systemic riskiness.

2. SRISK: A Systemic Risk Indicator for Banks

The concept of systemic risk is very complex to capture in a unique framework (see Hansen 2014). In this paper we focus on market-based metrics of systemic distress that allow us to explore the systemic importance of individual banks. Since the GFC, the literature that analyzes such metrics has significantly increased. Broadly speaking, these indicators can be classified into two groups. The first one consists of those metrics that are purely market based, such as the conditional autoregressive value-at-risk (VaR) in Engle and Manganelli (2004), and the ΔCoVaR in Adrian and Brunnermeier (2016), among others. The second set of indicators comprise those metrics that use balance sheet data in addition to market information, such as the marginal expected shortfall (MES) and the systemic expected shortfall (SES) in Acharya et al. (2017) and the SRISK in Brownlees and Engle (2017). We focus on this second type of metrics that associate systemic risk with the capital shortfall of the financial system conditional on the materialization of a systemic event.[12]

More specifically, our dependent variable is systemic risk as proxied by the SRISK indicator in Brownlees and Engle (2017). SRISK is inspired by the SES index in Acharya et al. (2017). Thus, SRISK associates the systemic risk contribution of an institution i with its expected capital shortfall conditional on a severe market downturn. The capital shortfall in t, CS_{it}, is the difference between the market value of equity and a prudential fraction k of the market value of the institution's assets, that is,

$$CS_{it} = k(D_{it} + MV_{it}) - MV_{it}, \tag{1}$$

where D is the book value of total liabilities, MV is the market value of equity, and k is the prudential capital ratio, which is the percentage of total assets that the financial institution holds as reserves

[12]While ΔCoVaR measures the VaR of the financial system conditional on an event affecting a specific bank, SRISK and MES are conditioned by a shock throughout the entire system. Accordingly, the direction of ΔCoVaR is from individual distress to the system, while MES and SRISK measure how much a given financial institution is undercapitalized when the whole financial system is undercapitalized.

because of regulation or prudential management.[13] Brownlees and Engle (2017) define SRISK as the conditional expectation of the future CS in the case of a systemic event, i.e., how much an institution's equity drops below a given fraction of its assets, when there is a crisis affecting the whole financial system, that is,

$$SRISK_{it} = E_t \left[CS_{i,t+h} \mid R_{m,t+1:t+h} < C \right], \tag{2}$$

where $R_{m,t+1:t+h}$ represents the market return between $t + 1$ and $t + h$, and C a threshold of market decline over time horizon h, defining a crisis, so that $R_{m,t+1:t+h} < C$ corresponds to the systemic event.

If we assume, like in Brownlees and Engle (2017), that the book value of the bank's liabilities remains fixed during the hypothetical systemic event, this expected capital shortfall can be expressed in terms of the firm equity return conditional on the systemic event, the long-run marginal expected shortfall (LRMES), that is,

$$SRISK_{it} = k(D_{it} + MV_{it}(1 - LRMES_{it})) - MV_{it}(1 - LRMES_{it}), \tag{3}$$

where $LRMES$ denotes the expected drop in the equity value of an institution i when the market falls below a threshold C within time horizon h,

$$LRMES_{it} = -E_t \left[R_{i,t+1:t+h} \mid R_{m,t+1:t+h} < C \right]. \tag{4}$$

$LRMES$, as defined in (4) is non-observable. Following Brownlees and Engle (2017), we estimate $LRMES$ with a dynamic conditional correlation (DCC) generalized autoregressive conditional heteroskedasticity (GARCH) model (see Engle 2002, 2009). For further details on the $LRMES$ estimation, see Appendix A.

SRISK has at least four properties that make this indicator an appropriate choice to measure systemic risk. First, it explicitly depends on the size and the degree of leverage of an institution. Second, SRISK can be easily computed with publicly available data.

[13] Brownlees and Engle (2017) call "quasi assets" the sum of book value of liabilities, D, and market value of equity, MV.

Third, as this indicator is also based on market data, it can be available at high frequencies, so that sudden shifts in systemic risk can be detected quickly.[14] Finally, SRISK is a forward-looking measure, as it signals the degree of systemic risk that has not yet materialized, but such risk could lead to economic losses in the event of a severe financial market downturn. However, like other market-based measures, SRISK can only be calculated for listed institutions.[15] Despite its limitations, SRISK is broadly used for empirical purposes by both policymakers and academics.[16]

3. Data and Variables

To test the implications in terms of systemicity of the implementation of SII buffers, we analyze a panel data set of listed banks from 24 European countries. The data set is quarterly and the sample period runs from 2008:Q1 to 2021:Q3.[17] The inclusion of 2020 data in the sample poses sizable challenges, given the sharp changes in a number of variables from the onset of the pandemic that considerably affect the estimates. Our approach to address this issue is to analyze the data set for the complete sample and also for two subsamples: the one that runs from 2008:Q1 to 2019:Q4 and the subsample that corresponds to the pandemic period, from 2020:Q1 to 2021:Q3. As the pandemic represents an exogenous shock independent of the financial cycle, focusing on the first subsample allows us to disentangle the effect of SII buffers in normal times. Specifically, the outbreak of COVID-19 led to an abrupt decrease in banks' market valuations,

[14] While the value of debt is usually available at quarterly frequencies, the market value and the LRMES can be updated daily, which allows us to capture short-term dynamics.

[15] Another limitation of SRISK is that it only reflects the markets' perception on an institution, so that this measure does not allow us to disentangle different risk factors (e.g., contagion, liquidity, solvency, funding, fire sales, etc.). That is, it is less informative than a fully fledged stress test.

[16] See, for instance, Tavolaro and Visnovsky (2014), Grinderslev and Kristiansen (2016), Coleman, LaPlante, and Rubtsov (2018), Engle and Ruan (2019), Bats and Houben (2020), or Brownlees et al. (2020) for some empirical works based on the SRISK indicator.

[17] To minimize the data gaps, especially at the beginning of the sample period where only annual or half-yearly data are available for some banks, we linearly interpolate the missing data to proxy quarterly series.

which held SRISK at historically high levels in 2020:Q1, close to the levels during the 2012 European sovereign debt crisis and above the levels of the 2008–09 Global Financial Crisis.[18] Also, the severity of the shock and the lack of alternative buffers led several countries to release the SII requirements in full or in part in 2020:Q2 as an immediate alternative available to ease the regulatory pressure on their credit institutions.[19]

Our total sample consists of 168 different banks. As Figure 1 shows, since 2008:Q1, when the number of banks was 82, this amount has gradually increased to reach a maximum of 158 banks in 2020:Q1. Subsequently, our sample decreases to 127 banks in 2021:Q3.[20] This sample is fairly representative and accounts for about 80 percent of total EU banks' assets.[21] Our sample consists of all publicly traded European banks reported by Refinitiv Datastream. More specifically, we compare the group of 14 and 52 banks in our sample that have been classified as G-SIIs and O-SIIs, respectively, at any time, and a control group of 102 banks that have never become an SII.[22] We assume that a bank's country is that of its primary listing where its stock is traded. Appendix B details the complete list of banks.

The dependent variable is the systemic risk of each bank as proxied by the SRISK indicator in Brownlees and Engle (2017), which we call $SRISK$, as shown in (3). To compute $SRISK$ we exploit both balance sheet and market data. Regarding balance sheet data, we use total liabilities at the consolidated level. Market data are also at the consolidated level and consist of the market value (MV)

[18]This increase in SRISK was the result of the higher uncertainty around the course of the pandemic that resulted in a sharp decline in stock market performance.

[19]Specifically, Estonia, Finland, Hungary, the Netherlands, and Poland fully or partially released their SII buffers in 2020:Q2. In addition, Cyprus, Greece, Lithuania, Malta, and Portugal postponed the phasing-in of planned O-SII buffers increases by one year (see ESRB 2021).

[20]Not all the banks in the sample continue over the entire period, either because of failures or mergers and acquisitions. This fact explains the gap between the total number of different banks and its peak reached in one quarter.

[21]This evidence is based on the total consolidated assets in 2020 (European Central Bank Statistical Data Warehouse).

[22]Out of these 66 SIIs, 50 are always SIIs (either O-SIIs or G-SIIs) throughout the sample period, while 16 banks have changed their status at any time.

Figure 1. Mean *SRISK* and Number of Banks

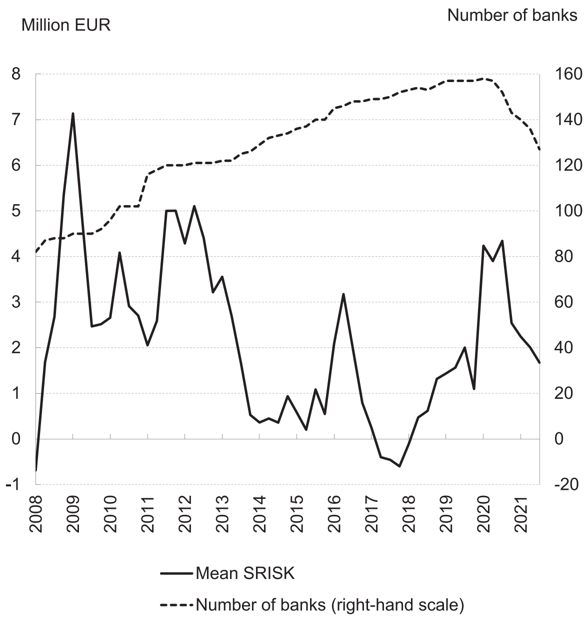

Note: Mean *SRISK* and number of banks (right-hand scale) for a panel of 168 European listed banks.

of each bank and the STOXX Europe 600 stock market index. We chose this index as the market index required to calculate LRMES in (4), as shown in Appendix A. We assume that the parameters to compute $SRISK$ in (4) are $k = 4.5\%$ for the capital requirement, $C = 10$ for the market decline threshold, and $h = 22$ business days for the period over which the hypothetical market decline occurs.[23] We calculate $SRISK$ for each listed bank with our own codes.[24]

[23]After several robustness tests, we assume k to be equal to the minimum CET1 ratio in accordance with the Basel III minimum own funds requirement (Pillar 1). In our specification, C and h have the same values as in Brownlees and Engle (2017), while they assume $k = 5.5$.

[24]Our MATLAB codes to compute $SRISK$ are available upon request. Alternatively, $SRISK$ data could be directly obtained from the Volatility Laboratory (V-Lab) (https://vlab.stern.nyu.edu). However, this source does not contain information about all individual listed banks in Europe, and the use of our own codes allows us to better control for the parameters of the indicator.

In line with Bostandzic et al. (2022), we do not restrict $SRISK$ to being positive, that is, it allows us to capture both capital shortfalls and surpluses. Figure 1 depicts the mean $SRISK$ throughout the sample. The average perceived systemic risk of banks as proxied by $SRISK$ increased in 2008 with the GFC, in 2012, coinciding with the European sovereign debt crisis, and also at the beginning of 2020 with the onset of the pandemic.

Our main explanatory variable is the SII buffer rate applied to each systemic bank, which we denote as S_BUF. This variable allows us to disentangle whether SII requirements do influence the systemic importance of the bank. In the EU, SII buffer requirements are set at the individual bank level. We obtain the SII buffer rates from the ESRB website. To calculate S_BUF we account for all the possible combinations to set the SII buffer at the domestic level. Thus, this capital requirement is usually the higher of the G-SII buffer, the O-SII buffer, and the SyRB, although there are some exceptions in certain jurisdictions.[25]

We also analyze whether merely designating the bank as an SII creates a signaling effect, which could be due to the implicit government guarantees in the event of distress, irrespective of the S_BUF level. For this, we define three dummy variables based on the assignment of the SII status by the competent authority—namely, the FSB for G-SIIs and the EBA for O-SIIs.[26] The first one, SII_STAT, takes into account the fact of having SII status. It is a step variable that takes the value of 1 once the EBA or the FSB identifies the bank as an SII and it is equal to 1 while the institution remains on the

[25]For instance, the O-SII buffer was not activated in Denmark or the Czech Republic until the end of 2019, while the SyRB was applied to SIIs in both countries before that date. In Bulgaria, Estonia, and Slovakia, the SyRB is cumulative, and the higher of the O-SII and G-SII buffers are set. Finally, in Bulgaria, Croatia, Estonia, Norway, and Poland all banks—not just SIIs—are subject to the SyRB. For more details, see the annual notifications available on the EBA website and the overview of national macroprudential and capital-based measures updated quarterly by the ESRB.

[26]Our definitions are based on the dates when the SII status is effective. The first G-SII list took effect in January 2012, and since then it is updated annually. The EBA published the O-SII list for the first time on April 25, 2016. However, most competent authorities began to assign the O-SII status in late 2015 (some of them in 2014). In the computation of S_BUF we have checked all notifications that are available on the ESRB website prior to the first EBA list to account for such cases.

SII list. This variable allows us to analyze whether designation as an SII has an impact on the systemicity of the bank, regardless of the buffer level. That is, for bank i in period t we define SII_STAT as

$$SII\_STAT_{it} = \begin{cases} 1 & \text{if designated SII in} \quad t \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

The second dummy variable, SII_IN, is 1 only in the quarter when the bank becomes an SII, while the third one, SII_OUT, takes the value of 1 only when the institution loses SII status. Both SII_IN and SII_OUT allow us to quantify the immediate market reaction after the announcements themselves, which is in line with the empirical approach in Bekaert and Breckenfelder (2019), Andrieş et al. (2020), and Gündüz (2020) to analyze the market reaction to the disclosure of the list of O-SIIs by the EBA. Both binary variables are expressed as follows:

$$SII\_IN_{it} = \begin{cases} 1 & \text{if} & SII\_STAT_{it} = 1 \\ & \text{and} & SII\_STAT_{it-1} = 0 \\ 0 & \text{otherwise,} & \end{cases} \tag{6}$$

$$SII\_OUT_{it} = \begin{cases} 1 & \text{if} & SII\_STAT_{it} = 0 \\ & \text{and} & SII\_STAT_{it-1} = 1 \\ 0 & \text{otherwise.} & \end{cases} \tag{7}$$

We also explore whether the different adjustment mechanisms to higher capital requirements affect $SRISK$ once the SII buffer is implemented. Following Cohen and Scatigna (2016) and Braouezec and Kiani (2021), among others, we explore the impact on SRISK of four transmission channels, namely (i) total equity; (ii) retained earnings; (iii) new share issuances; and (iv) the risk-weighted density (RWD).

To check whether SIIs and non-SIIs follow different patterns, Table 1 reports some summary statistics for $SRISK$, the SII buffer level, S_BUF, as well as the four transmission channels for both SIIs and non-SIIs. We analyze the full sample and before and after the pandemic. As expected, the average $SRISK$ is higher for SIIs than for non-SIIs in the entire sample period and the two subperiods. Since the onset of the COVID-19 crisis, SIIs' average $SRISK$

Table 1. Descriptive Statistics of *SRISK*, the SII Buffer Level (*S_BUF*), and Four Transmission Channels

	SIIs				Non-SIIs			
	Mean	**SD**	**Min.**	**Max.**	**Mean**	**SD**	**Min.**	**Max**
A. Full Sample								
SRISK	5,861	18,335	–66,560	90,926	640	6,474	–29,177	85,537
S_BUF	1.2	1.3	0.0	6.5	—	—	—	—
Total Equity	29,880	35,824	137	183,054	4,178	10,627	8	127,892
Retained Earnings	14,724	22,043	–11,715	133,314	2,006	5,186	–21,178	86,138
Risk-Weighted Density (RWD)	42.2	16.1	15.5	86.0	56.3	17.6	3.5	98.3
Equity Issuances	5,807	7,338	0	31,694	546	1871	0	21,594
B. 2008:Q1–2019:Q4								
SRISK	5,011	18,494	–66,560	79,725	712	6,874	–29,177	85,537
S_BUF	1.1	1.3	0.0	5.0	—	—	—	—
Total Equity	32,000	37,100	137	183,000	4,561	11,200	8	128,000
Retained Earnings	15,700	22,900	–11,700	133,000	2,132	5,439	–21,200	86,100
Risk-Weighted Density (RWD)	42.7	16.6	15.5	86.0	57.1	17.6	4.3	98.3
Equity Issuances	6,085	7,301	0	31,694	632	2,062	0.00	21,595
C. 2020:Q1–2021:Q3								
SRISK	8,345	17,653	–14,577	90,926	113	1,748	–6,788	16,780
S_BUF	1.5	1.2	0	6.5	—	—	—	—
Total Equity	23,602	31,145	207	180,427	1,408	2,572	20	20,452
Retained Earnings	11,885	18,964	–11,664	133,314	929	1,635	–1,136	11,873
Risk-Weighted Density (RWD)	40.8	14.5	17.1	77.5	50.4	16.1	3.5	96.6
Equity Issuances	5,080	7,604	0	31,694	290	609	0	3,767

Note: SIIs stands for "systemically important institutions," *SRISK* is the systemic risk of each listed bank as measured by the SRISK indicator in Brownlees and Engle (2017), and *S_BUF* denotes the SII buffer level.

for the pandemic period is much higher than during the first subperiod, while for non-SIIs, this statistic decreased significantly. In other words, during the pandemic, SIIs would have been penalized in terms of systemicity, as measured by $SRISK$, given the drop in their stock prices and market valuations that pushed up the average $SRISK$ from 2020:Q1. Regarding S_BUF, the buffer requirements for SIIs have remained relatively stable in the subsamples. Although the average S_BUF during the pandemic is higher than that of the first subsample, there have been several releases during this last period. Finally, as expected, total equity and equity issuances are, on average, greater for SIIs. Conversely, non-SIIs hold larger shares of risk-weighted assets.[27]

Finally, for the robustness of our results, we also use a set of bank-specific and country-level variables as controls. Specifically, the bank variables consist of the total assets and the return on equity (RoE). Country-level variables allow us to control for the unobserved heterogeneity across countries and comprise (i) macroprudential buffers that are common at national level, namely the countercyclical capital buffer (CCyB),[28] the SyRB, and the capital conservation buffer (CCoB);[29] (ii) real GDP per capita; and (iii) the sovereign CDS spread. See Table 2 for the complete list of variables and data sources.

4. Methodological Approach

The baseline linear panel data model is described by the following expression:

[27]This might be due to different management practices at SIIs as well as the more intense use of internal models when calculating risk-weighted assets.

[28]The total CCyB of a given bank is the average of the CCyB across all countries, weighted by its exposures of the bank in each country. Due to a lack of country-exposures data, we abstract from this complication and only control for the level of the CCyB in its primary listing country.

[29]The adoption of the CCoB was completed in 2015 for the countries with phase-in arrangements, so that since 2015 the CCoB level is 2.5 percent in all jurisdictions. Although this control variable is constant since that date, we consider it given its different dynamics across countries during the phase-in period. The SyRB is non-zero for non-SIIs where this buffer is applied to all banks at the country level.

Table 2. Variable Definitions and Data Sources

Variable	Description	Source
SRISK	Systemic risk contribution of individual banks as proxied by the SRISK indicator (Brownlees and Engle 2017)	Own calculations
Total Liabilities	Quarterly total liabilities at consolidated level, million EUR	S&P Capital IQ
Market Value	Market capitalization, which is the product of share price and the number of ordinary shares in issue, for the last day of each period, million EUR	Refinitiv Datastream
STOXX Europe 600	Stock market index for the last day of each period	Refinitiv Datastream
	SII Characteristics	
S_BUF	Level of structural buffer for SIIs, which is the combination of the G-SII and O-SII buffers and SyRB, for the last day of each period, %	European Systemic Risk Board (ESRB)
SII_STAT	Dummy variable that takes the value of 1 once the bank is designated as an SII in that period and 0 otherwise	European Banking Authority (EBA)
SII_IN	Dummy variable that takes the value of 1 when the bank is designated as an SII and 0 otherwise	European Banking Authority (EBA)
SII_OUT	Dummy variable that takes the value of 1 when the bank stops being designated as an SII and 0 otherwise	European Banking Authority (EBA)

(*continued*)

Table 2. (Continued)

Variable	**Description**	**Source**
	Bank Variables	
Total Equity	Includes par value, paid-in capital, retained earnings, and other adjustments to equity, million EUR	S&P Capital IQ
Retained Earnings	Accumulated net income and dividends, million EUR	S&P Capital IQ
Risk-Weighted Density (RWD)	Total risk-weighted assets, as defined by the latest regulatory and supervisory guidelines, as a percentage of total assets, %	S&P Capital IQ
Equity Issuances	Current value of the share issuance deal, thousand EUR	Dealogic
Total Assets	All assets owned by the company for the last day of each period, as carried on the balance sheet, million EUR	S&P Capital IQ
Return on Equity (RoE)	Net income as a percentage of average equity, %	S&P Capital IQ
	Country Variables	
CCyB	Level of countercyclical capital buffer for the last day of each period, %	European System Risk Board (ESRB)
CCoB	Level of capital conservation buffer, %	European System Risk Board (ESRB)
SyRB	Level of systemic risk buffer if applied to non-SIIs for the last day of each period, %	European System Risk Board (ESRB)
GDP	GDP per capita in purchasing power standards (PPS) with respect to the EU27 average. Average quarter at 2010 = 100	Eurostat
CDS	Sovereign credit default swap (CDS) spreads in basis points	Refinitiv Datastream

$$SRISK_{it} = \alpha_i + T_t + \gamma S\_BUF_{it-1} + \sum_j \beta_j X_{j,it-1} + \sum_k \delta_k Z_{k,it-1} + \varepsilon_{it}, \tag{8}$$

where for all banks $i = 1, \ldots, N$ and periods $t = 1, \ldots, T$, the main explanatory variable is S_BUF to quantify the effect of the introduction of SII capital buffers on a panel of listed banks. The key coefficient in (8) is γ, which can be interpreted as the average effect of a 1 percent increase in SII capital requirements on $SRISK$. Therefore, a negative estimate of γ would suggest that higher SII requirements would lead to a lower contribution to systemic risk as proxied by $SRISK$. Apart from the bank and time dummies, the model includes bank-specific control variables, X_{it}, and country-level variables, Z_{it}, as described in the previous section. We fit the model for the full sample, and also for the two subsamples to characterize the impact of the pandemic on the data set.

Second, we also fit different specifications of the baseline model in (8), replacing S_BUF with the three alternative dummy variables based on the assignment of SII status defined in expressions (5) to (7). Namely, we use the step variable SII_STAT in (5) as an explanatory variable to study whether having SII status influences on the bank systemicity regardless of the SII capital requirement level. Thus, a negative estimate of this coefficient would suggest that being an SII lowers a bank's contribution to systemic risk as proxied by $SRISK$. In other words, being an SII might represent a signaling effect regardless of the buffer level. Further, we explore the possibility that there could be an immediate market response on the announcement of a bank's designation as an SII related to the market perception of its contribution to systemic risk. To this end, we also modify the baseline model in (8) by replacing S_BUF with SII_IN and SII_OUT, as defined in (6) and (7). A positive (negative) estimate of the SII_IN coefficient would indicate that the designation as an SII would immediately increase (decrease) the systemic nature of the bank.

We further study the effect of being identified as an SII over time via local projections (see Jordà 2005). Thus, for quarters $q = 0, 1, 2, \ldots, 12$ we fit the baseline model that considers SII_IN as an explanatory variable instead of S_BUF as follows:

$$SRISK_{it+q} = \alpha_i^q + T_t^q + \lambda^q SII\_IN_{it-1} + \sum_j \beta_j^q X_{j,it-1} + \sum_k \delta_k^q Z_{k,it-1} + \varepsilon_{it+q}. \tag{9}$$

Next, we increase the number of drivers in (8) with the four main options that banks have at their disposal to comply with SII buffers. This specification allows us to disentangle which one dominates in a bank's response in terms of lower systemicity to higher capital requirements. Namely, we explore the impact of four alternative variables entailing changes to the capital structure on $SRISK$: (i) total equity; (ii) retained earnings; (iii) new share issuances; and (iv) risk-weighted density.

$$SRISK_{it} = \alpha_i + T_t + \gamma S\_BUF_{it-1} + \sum_{l=1}^{4} \omega_l C_{l,it-1} + \sum_j \beta_j X_{j,it-1} + \sum_k \delta_k Z_{k,it-1} + \varepsilon_{it}, \tag{10}$$

where $\{C_{l,it}\}_{l=1}^4$ denotes the four different channels.

Finally, we check whether decisions by SII banks to comply with capital requirements do have an impact on their systemicity. For this purpose, we also fit the model in (9) with interactions of the variables related to a bank's capital structure and SII_STAT, which is given by

$$SRISK_{it} = \alpha_i + T_t + \gamma S\_BUF_{it-1} + \sum_{l=1}^{4} \omega_l C_{l,it-1} + \sum_{l=1}^{4} \lambda_l (C_{l,it-1} \times SII\_STAT_{it-1}) + \sum_j \beta_j X_{j,it-1} + \sum_k \delta_k Z_{k,it-1} + \varepsilon_{it}. \tag{11}$$

This last specification allows us to test for the null hypothesis that the influence of these variables on the systemicity is independent of the SII status. That is, for all the bank capital-related variables l it is possible to test for the following null:

$$H_0 : \lambda_l = 0. \tag{12}$$

Model (11) also allows us to quantify the residual impact of S_BUF on $SRISK$ once we control for capital-related variables. This residual impact of S_BUF could be interpreted as the effect of the SII buffer level itself on a bank's contribution to systemic risk. Finally, we replace S_BUF with SII_STAT in (11) to analyze the significance of having SII status once we also consider all the feasible bank choices to fulfill this capital requirement. This effect could be related to a positive sign for the markets of having SII buffers once we control for bank balance sheet variables.

We estimate this linear fixed-effects panel data model with standard errors robust to serial correlation (clustered at bank level) and heteroskedasticity. Another challenge of the analysis is the possibility of endogeneity problems as a result of reverse causality and omitted variables. Reverse causality could be a concern when analyzing the link between $SRISK$ and S_BUF, as a two-way causality relationship could be feasible. For instance, the national authorities could increase the SII buffer to address a bank's higher systemicity. On the other hand, higher capital requirements for SIIs are likely to influence a bank's systemic nature. This latter direction of causality is precisely the focus of our analysis, and, to minimize the effect of the former, the main variables of interest—S_BUF, SII_IN, and SII_OUT—are lagged one period in specifications (8) to (11). Also, all explanatory variables are lagged one period to limit simultaneity bias. Finally, regarding a possible omitted-variable bias, we consider that our set of explanatory variables contains a sufficient number of relevant drivers to analyze of $SRISK$ and, therefore, we consider that our model is not poorly specified.

5. Results

5.1 Baseline Model: Some Initial Results

Table 3 reports the estimates of the baseline model in (8) for the total sample (panel A), as well as for the subsample before and after the onset of the COVID-19 pandemic (panels B and C, respectively).

Table 3. Estimates of the Baseline Model for the Total Sample and for the Pre- and Post-Pandemic Period

	Full Sample			2008:Q1–2019:Q4			2020:Q1–2021:Q3		
S_BUF	–79.40			–645.5***			–27.16		
	(276.7)			(223.4)			(202.6)		
SII_STAT		–1,303**			–1,767***			756.4	
		(518.3)			(506.8)			(476.0)	
SII_IN			2,120***			2,673***			–486.9
			(597.7)			(697.3)			(654.4)
SII_OUT			–214.3			–420.5			–1225**
			(460.3)			(598.8)			(596.4)
Total Assets	0.039***	0.039***	0.039***	0.034***	0.034***	0.034***	–0.019	–0.019	–0.019
	(0.003)	(0.003)	(0.003)	(0.005)	(0.005)	(0.005)	(0.011)	(0.011)	(0.011)
RoE	–8.10***	–8.14***	–8.08***	–7.45**	–7.21**	–7.12**	–1.36	–1.43	–1.34
	(3.01)	(3.01)	(2.99)	(2.99)	(2.91)	(2.89)	(4.35)	(4.33)	(4.35)
CCyB	–499.0**	–513.6**	–522.7*	–439.9*	–679.7***	–722.3**	–411.0	–396.0	–396.8
	(236.1)	(254.0)	(284.2)	(232.7)	(258.7)	(308.2)	(427.2)	(429.9)	(431.1)
SyRB	318.5***	167.9	355.2***	388.3***	413.1***	688.2***	243.8	263.4*	266.3
	(100.6)	(115.3)	(130.5)	(122.3)	(132.0)	(160.3)	(153.1)	(155.9)	(166.4)
CCoB	636.3***	653.9***	617.4***	421.3**	304.1	235.7	—	—	—
	(208.1)	(228.0)	(234.8)	(198.1)	(198.3)	(211.0)	—	—	—
GDPpc	–76.04***	–71.69***	–78.45***	–39.95	–56.64**	–65.93***	–114.6*	–114.1*	–113.5*
	(26.27)	(24.02)	(23.43)	(25.14)	(25.41)	(23.47)	(68.16)	(68.08)	(68.20)
CDS	0.003	–0.013	–0.002	–0.003	–0.043	–0.037	–4.197	–4.307	–4.145
	(0.032)	(0.036)	(0.032)	(0.032)	(0.034)	(0.031)	(10.12)	(10.16)	(10.13)
N	6,487	6,487	6,487	5,501	5,501	5,501	986	986	986
R^2	0.396	0.400	0.397	0.381	0.385	0.380	0.177	0.177	0.177

Note: Dependent variable: systemic risk contribution of individual banks as proxied by the SRISK indicator (see Brownlees and Engle 2017); all explanatory variables are lagged one period; *S_BUF*: SII buffer level; *SII_STAT*: binary dummy, *SII_STAT* = 1 while a bank has SII status; *SII_IN*: binary dummy, *SII_IN* = 1 only in the quarter when a bank becomes an SII; *SII_OUT*: binary dummy, *SII_OUT* = 1 only in the quarter when the bank stops being an SII; RoE: return on equity; CCyB: countercyclical capital buffer; CCoB: capital conservation buffer; GDPpc: GDP per capita; CDS: credit default swap spreads. See Table 2 for a complete description of the explanatory variables and data sources. Standard errors are shown in parentheses. Standard errors are robust to heteroskedasticity and serial correlation. Intercept, time, and bank fixed effects are included but not reported. ***, **, and * refer to significance at 1 percent, 5 percent, and 10 percent level, respectively.

First and most importantly, before the pandemic the increase in the SII buffer level has a negative effect on a bank's contribution to systemic risk, as signaled by the negative and significant estimate of S_BUF. This evidence suggests that higher SII buffers are associated with lower systemic risk. Given that the main objective of SII buffers is to address the systemic riskiness of SIIs, this result means that these capital requirements work as expected over this sample period.[30]

Conversely, during the pandemic the estimate of S_BUF becomes non-significant. This lack of significance also holds for the full sample. The particular dynamics of S_BUF and $SRISK$ during the pandemic explain this result. Thus, during the pandemic the estimate for S_BUF reflects the fact that the released buffers for SIIs in some countries (lower S_BUF) were followed by a reduction in the banks' contribution to systemic risk (lower SRISK) after its peak in 2020:Q1. Therefore, the negative link between the SII buffer level and a bank's contribution to systemic risk identified for the preceding sample does not hold during the pandemic. In fact, this link between S_BUF and $SRISK$ is reversed in those countries that released their SII buffers, so that the lower S_BUF preceded $SRISK$ drops. However, this temporal positive relationship between the two variables does not entail a causality link between them, as the lower $SRISK$ results from the market's correction after the abnormal pattern of $SRISK$ in 2020:Q1.[31] Finally, as the link between S_BUF and $SRISK$ changes during the pandemic, time fixed effects are

[30]According to the estimated coefficient for S_BUF for this subsample, a 1 percentage point (pp) increase in the SII buffer level is associated with a €645.5 million reduction in SRISK. Alternatively, we have fitted the baseline model in (8) with $SRISK$ expressed in logarithms. According to the results, a 1 pp increase in S_BUF is associated with a reduction of 16.4 percent in $SRISK$. The estimates for other coefficients are in line with those presented in Table 3 and are available upon request.

[31]As a robustness check, we have also fit Equation (8) for the full sample with a dummy variable that equals 1 during the COVID period and its interaction with our main explanatory variable, S_BUF. The results are in line with those for separate subsamples presented in Table 3 and are available in Appendix C. As explained, the link between S_BUF and $SRISK$ becomes positive during the pandemic, without involving a causality link between both variables. Besides, the great variety of policy responses implemented by authorities could have also affected in other regressors, which is out of the scope of this article. All in all, we consider it more appropriate to fit the model by subsamples, as in Table 3.

not enough to fully capture the impact of the COVID shock on the variables.

Table 3 also shows that having SII status, irrespective of the buffer level, lowers banks' contribution to systemic risk, as shown by the negative and significant estimates of SII_STAT. This result is to some extent related to Bekaert and Breckenfelder (2019) and Vogel (2020), who find that being designated an SII has different consequences for banks.[32] This outcome holds for the full sample and for the pre-pandemic period, but it becomes non-significant during the pandemic. That is, after the onset of the coronavirus crisis, the contribution of banks to systemic risk was independent of their SII status.[33] From a policy perspective, this significance of having SII status determine the contribution of a bank to systemic risk means that having SII buffers can be interpreted as a signal for the markets of the commitment of these banks to increasing their resilience.

The estimates of SII_IN in Table 3 suggest that the designation as an SII immediately leads to an increase in systemic risk. Conversely, as evidenced by the coefficient of SII_OUT, once the bank ceases to be an SII, its contribution to systemic risk in the next period diminishes. This result is in line with Andrieş et al. (2020) and Gündüz (2020), who find that the initial market reaction to the SII designation tends to be negative given certain stigma effects related to tighter regulation and lower profitability once the requirement is set. To further explore the effect of the designation as an SII over time, we fit model (10), which is inspired by the local projections method (see Jordà 2005). As stated in (10), we consider the baseline equation in (8), and recursively run a set of regressions for the lead dependent variable up to 12 quarters ahead. Figure 2 shows the estimated coefficients of SII_IN and their corresponding 95 percent confidence intervals for the 12 quarters. The initial impact of a bank's SII designation is positive, that is, it is associated with an increase in $SRISK$. However, this effect decreases quickly over

[32]Vogel (2020) concludes that being labeled an O-SII brings a funding cost advantage for deposits, while Bekaert and Breckenfelder (2019) document that bond prices are higher for those banks belonging to the O-SII list.

[33]The subsample after the onset of the pandemic is short and has few new SII designations, which implies a low variation of SII_STAT. This fact complicates the identification, so that the results for this subsample should be interpreted with caution.

Figure 2. Local Projection Estimator of *SII_IN* for 12 Quarters

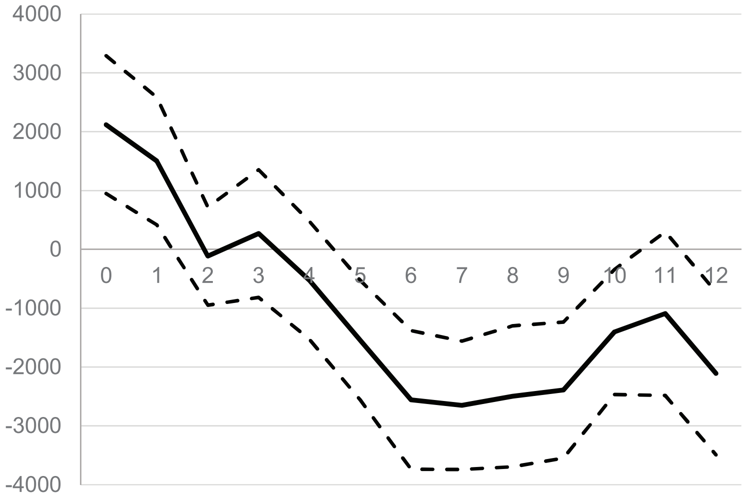

Note: The unbroken line depicts the local projection estimator for *SII_IN*, a dummy variable that is 1 only in the period when the bank becomes an SII, during 12 quarters for model (9). The broken lines are the 95 percent confidence intervals.

time and becomes negative four quarters later in line with the negative coefficients for S_BUF and SII_STAT. This result suggests that banks adapt to the new capital requirements over time, so that the institutions' contribution to systemic risk eventually decreases.

Finally, as a robustness check, we have also explored whether the estimate of γ in (8) is different for banks of both the peripheral and the core countries. Our results suggest that the country location of the bank does not influence their systemicity once the SII buffer is set.[34]

[34]To this end, we have built two indicator variables. The first one is equal to 1 if a bank's home country is Greece, Italy, Portugal, or Spain, and 0 otherwise, while the second one equals 1 if a bank's jurisdiction is Austria, Belgium, France, Germany, Luxembourg, or the Netherlands, and 0 otherwise. We have interacted these two new variables with S_BUF in (8), as well as with SII_IN and SII_OUT. The estimates, which are available upon request, are not significant or conclusive.

5.2 Additional Insights into the Drivers of the Systemic Risk Driven by Banks

Next, we analyze the potential impact of the adjustment mechanisms that banks use to comply with SII buffer requirements on their contribution to systemic risk. This approach allows us to disentangle which option dominates in banks' response and contributes the most to decreasing their systemicity due to higher capital requirements. As the link between $SRISK$ and S_BUF during the pandemic crisis follows an abnormal pattern, hereafter we focus on the pre-pandemic sample up to 2019:Q4.[35]

Table 4 reports the estimates for model (10), which extends the baseline model (8) by adding the main options for capital adjustment as regressors. The results indicate that higher equity, especially in the form of retained earnings, leads to a lower bank contribution to systemic risk. Also, the coefficient of S_BUF decreases once these variables are included. "De-risking," i.e., a lower risk-weighted density, also diminishes a bank's systemic impact, although this estimate is less significant. Next, we fit the model in (11), which includes the interactions of SII_STAT with banks' capital adjustment options, to distinguish their impact for SIIs and non-SIIs. Table 5 reports the results, which show that the estimated coefficients of the interaction terms are mostly significant, while the estimates for the non-SII group are not, except for the risk-weighted density. This outcome indicates that the impact of these options on a bank's contribution to systemic risk depends on having an SII status.

Specifically, the results in Table 5 indicate that the market perception of systemic risk improves when an SII bank increases its equity.[36] However, the two main drivers of this effect, retained earnings and equity issuances, work in opposite directions. Thus, the increase in SIIs' retained earnings is generally positively perceived by the markets, as signaled by the negative coefficient of its interaction with SII_STAT. One possible interpretation is that in this

[35] This is mainly because the pandemic crisis represents an exogenous shock to banks' market valuations and there was a massive release and other prudential changes to SII buffers in some countries, as explained in the previous subsection.

[36] This finding is in line with Dautović (2020), who suggested that a phased-in increase in capital requirements raises the CET1 capital ratio, thereby improving resilience and loss-absorbing capacity.

Table 4. Estimates of the Baseline Model with Bank Variables Involving Changes to the Capital Structure

	M1	M2	M3	M4	M5	M6	M7	M8
S_BUF	–668.1***	–507.0**	–569.0***	–599.1***	–672.7***	–516.8***	–534.4***	–541.3***
	(228.0)	(196.3)	(193.0)	(217.7)	(228.3)	(194.1)	(191.9)	(190.4)
Total Equity		–0.210**				–0.216**		
		(0.0856)				(0.0858)		
Retained Earnings			–0.220***				–0.203***	–0.204***
			(0.0569)				(0.0657)	(0.0658)
Equity Issuances				–0.196			–0.120	–0.126
				(0.121)			(0.134)	(0.136)
RWD					6.724	21.18*		13.61
					(11.75)	(11.90)		(11.91)
Total Assets	0.035***	0.041***	0.039***	0.034***	0.035***	0.041***	0.038***	0.038***
	(0.005)	(0.007)	(0.004)	(0.005)	(0.005)	(0.007)	(0.005)	(0.005)
RoE	–7.50**	–6.20**	–5.99**	–7.63**	–7.46**	–6.02**	–6.18**	–6.09**
	(3.04)	(2.92)	(2.51)	(3.14)	(3.02)	(2.85)	(2.62)	(2.58)
CCyB	–465.8*	–592.9**	–470.2*	–606.9**	–470.9*	–612.7**	–556.1**	–571.1**
	(244.6)	(234.5)	(245.6)	(239.4)	(246.8)	(235.2)	(243.6)	(247.9)
SyRB	388.6***	371.4***	328.6**	384.7***	386.7***	364.8***	330.8**	326.6**
	(127.5)	(138.4)	(137.2)	(125.6)	(126.8)	(139.2)	(135.2)	(134.6)
CCoB	447.7**	270.2	401.9*	389.9*	453.0**	281.8	370.1*	379.0*
	(210.2)	(196.4)	(210.0)	(201.1)	(208.9)	(195.3)	(202.9)	(202.1)
GDPpc	–40.37	–40.65*	–24.39	–55.29**	–40.36	–40.62*	–34.72	–35.15
	(26.51)	(23.23)	(23.50)	(27.62)	(26.39)	(22.77)	(27.69)	(27.82)
CDS	–0.003	–0.016	–0.026	–0.012	–0.001	–0.009	–0.029	–0.025
	(0.032)	(0.029)	(0.027)	(0.032)	(0.033)	(0.030)	(0.028)	(0.029)
N	5,103	5,103	5,103	5,103	5,103	5,103	5,103	5,103
R^2	0.381	0.403	0.398	0.386	0.381	0.404	0.400	0.401

Note: Dependent variable: systemic risk contribution of individual banks as proxied by the SRISK indicator (see Brownlees and Engle 2017); all explanatory variables are lagged one period; *S_BUF*: SII buffer level; RWD: risk-weighted density; RoE: return on equity; CCyB: countercyclical capital buffer; CCoB: capital conservation buffer; GDPpc: GDP per capita; CDS: credit default swap spreads. See Table 2 for a complete description of explanatory variables and data sources. Standard errors are shown in parentheses. Standard errors are robust to heteroskedasticity and serial correlation. Intercept, time, and bank fixed effects are included but not reported. The sample period runs from 2008:Q1 to 2019:Q4. ***, **, and * refer to significance at 1 percent, 5 percent, and 10 percent level, respectively.

Table 5. Estimates of the Baseline Model with Bank Variables Involving Changes to the Capital Structure and Interactions with SII Status

	M1	M2	M3	M4	M5	M6	M7	M8
S_BUF	–668.1***	–413.1**	–307.8*	–577.6***	–653.9**	–651.9***	–197.0	–204.8
	(228.0)	(180.4)	(162.2)	(206.4)	(257.3)	(207.6)	(172.6)	(230.8)
Total Equity		–0.087				–0.078		
		(0.100)				(0.100)		
Total Equity × *SII_STAT*		–0.069**				–0.079**		
		(0.031)				(0.034)		
Retained Earnings			0.0325				0.159	0.159
			(0.0732)				(0.104)	(0.104)
Retained Earnings × *SII_STAT*			–0.191***				–0.345***	–0.345***
			(0.047)				(0.073)	(0.073)
Equity Issuances				–0.0362			–0.442	–0.447
				(0.342)			(0.294)	(0.299)
Equity Issuances × *SII_STAT*				–0.150			0.555**	0.553**
				(0.246)			(0.238)	(0.246)
RWD					7.084	19.59*		14.93
					(11.44)	(10.60)		(11.84)
RWD × *SII_STAT*					–1.539	19.82*		0.053
					(9.479)	(10.31)		(9.83)
N	5,103	5,103	5,103	5,103	5,103	5,103	5,103	5,103
R^2	0.381	0.418	0.428	0.374	0.381	0.421	0.445	0.445

Note: Dependent variable: systemic risk contribution of individual banks as proxied by the SRISK indicator (see Brownlees and Engle 2017); all explanatory variables are lagged one period; *S_BUF*: SII buffer level; *SII_STAT*: binary dummy, *SII_STAT* = 1 while a bank has SII status; RWD: risk-weighted density; RoE: return on equity; CCyB: countercyclical capital buffer; CCoB: capital conservation buffer; GDPpc: GDP per capita; CDS: credit default swap spreads. See Table 2 for a complete description of explanatory variables and data sources. Standard errors are shown in parentheses. Standard errors are robust to heteroskedasticity and serial correlation. Intercept, time, and bank fixed effects are included but not reported. The sample period runs from 2008:Q1 to 2019:Q4. ***, **, and * refer to significance at 1 percent, 5 percent, and 10 percent level, respectively.

way banks could thus seek to improve their profits, for instance by increasing net interest income or reducing overall operating expenses (see Cohen and Scatigna 2016). Conversely, issuing new equity could be perceived as a less attractive option because of its direct diluting effect on the market value of the existing shares and the uncertainty related to EU banks' post-GFC low valuations. Finally, we do not find a strong link between the risk-weighted density and the level of SRISK, in either SIIs or non-SIIs.[37] Indeed, there is very little empirical evidence of a trade-off between the rise in the risk-weighted density to compensate for the increase in equity (see, for instance, Gropp et al. 2019).

Once we consider the differential impact of these variables for the sample of SIIs, the effect of S_BUF on $SRISK$ substantially decreases and even becomes non-significant in those model specifications that include the interaction of SII_STAT with retained earnings and equity issuances as regressors. In other words, once we control for these drivers, the impact of the SII buffer level on a bank's systemic nature disappears. This indicates that an important part of the decrease in $SRISK$ associated with increases in S_BUF is mediated by increases in equity, particularly via retained earnings.

Finally, Table 6 reports the estimates of model (11) with SII_STAT instead of S_BUF as the main explanatory variable. This approach allows us to disentangle the impact of having SII status on banks' contribution to systemic risk in Table 3 from banks' decisions to adjust their capital to comply with SII requirements.[38] That is, we aim to quantify the residual impact of having SII status once we control for these capital-related variables. Contrary to the results in Table 5 for the SII buffer level, the estimate of SII_STAT is still significant once we include all adjustment options—namely, retained earnings, equity issuances, and share of risk-weighted assets. This result suggests that having SII status is itself positively perceived by markets and decreases the contribution to the systemic risk.

[37] This outcome is to some extent contrary to Dautović (2020), who finds that being an SII is associated with potentially higher risk-taking on average.

[38] Estimates of the interactions of SII_STAT with the different bank-related variables are relatively similar to those reported in Table 6. The main difference is that in Table 6 the link between the proportion of risk-weighted assets and $SRISK$ only holds for SIIs.

Table 6. Estimates of the Baseline Model with Bank Variables Involving Changes to the Capital Structure and Interactions with SII Status

	M1	M2	M3	M4	M5	M6	M7	M8
SII_STAT	–1,873***	–239.5	–92.47	–1,433**	–6,932***	–3,870**	–732.0*	–3,320***
	(512.3)	(589.3)	(456.5)	(630.1)	(1,534)	(1,824)	(420.3)	(1,216)
Total Equity		–0.0945				–0.115		
		(0.103)				(0.108)		
Total Equity × *SII_STAT*		–0.0674*				–0.0456		
		(0.0368)				(0.0415)		
Retained Earnings			0.0327				0.146	0.126
			(0.0795)				(0.105)	(0.109)
Retained Earnings × *SII_STAT*			–0.194***				–0.339***	–0.311***
			(0.0527)				(0.0715)	(0.0782)
Equity Issuances				–0.0923			–0.462	–0.461
				(0.343)			(0.285)	(0.293)
Equity Issuances × *SII_STAT*				–0.0584			0.594**	0.600**
				(0.267)			(0.229)	(0.234)
RWD					–1.020	14.65		9.479
					(10.33)	(10.71)		(11.24)
RWD × *SII_STAT*					110.5***	69.35**		50.06**
					(26.86)	(28.63)		(21.06)
N	5,103	5,103	5,103	5,103	5,103	5,103	5,103	5,103
R^2	0.387	0.416	0.427	0.390	0.404	0.423	0.446	0.449

Note: Dependent variable: systemic risk contribution of individual banks as proxied by the SRISK indicator (see Brownlees and Engle 2017); all explanatory variables are lagged one period; *S_BUF*: SII buffer level; *SII_STAT*: binary dummy, *SII_STAT* = 1 while a bank has SII status; RWD: risk-weighted density; RoE: return on equity; CCyB: countercyclical capital buffer; CCoB: capital conservation buffer; GDPpc: GDP per capita; CDS: credit default swap spreads. See Table 2 for a complete description of explanatory variables and data sources. Standard errors are shown in parentheses. Standard errors are robust to heteroskedasticity and serial correlation. Intercept, time, and bank fixed effects are included but not reported. The sample period runs from 2008:Q1 to 2019:Q4. ***, **, and * refer to significance at 1 percent, 5 percent, and 10 percent level, respectively.

6. Conclusions

In this paper we investigate whether SII buffers are effective at lowering the contribution of these large and complex banks to systemic risk. We also analyze what the possible drivers of the banks' systemicity adjustment are. We proxy banks' perceived systemic risk with the measure SRISK in Brownlees and Engle (2017), the dependent variable of our empirical analysis. This is a broadly used metric of systemic risk that can be easily computed with bank- and market-based data. Then, we fit a number of fixed-effects panel data models to analyze the link between the SII buffer level and having SII status, and the SRISK indicator for a sample of listed European banks from 2008:Q1 to 2021:Q3.

According to our results, there is a negative relationship between the SII buffer level and banks' contribution to systemic risk. Therefore, higher capital requirements for systemic banks achieve the goal sought by regulators, as they lead to a decrease in perceived systemic risk. Furthermore, being designated as an SII also decreases banks' contribution to systemic risk, but this effect is time sensitive. The short-term impact of SII designation on SRISK appears to be positive (i.e., it increases SRISK), potentially due to a market stigma effect, while the medium-term effect, once the bank has had the time to adapt to the higher requirements, turns negative. We then control for the main options banks use to comply with higher SII requirements to further analyze the determinants of this perceived lower systemic risk. The results indicate that an increase in banks' equity through retained earnings is the main driver of this effect. Finally, once we control for these bank-based drivers, the residual effect of having SII status on perceived systemic risk is still negative and significant. This outcome means that being an SII provides a positive signal to markets by further decreasing its contribution to systemic risk.

Our results have important financial stability implications, in particular regarding the discussion of the role of SII buffers as an effective instrument for increasing these banks' resilience and for reducing their need for government interventions. Further research to fully understand the impact of SII buffers would be needed to guide policy responses to address the "too big to fail" status of SIIs. For instance, this paper does not address buffer calibration.

Moreover, impact analysis of SII buffers on different measures of systemic risk, not only the SRISK indicator, would be useful to provide a more holistic view on the implications of SII buffers.

Appendix A. Estimation of the Long-Run Marginal Expected Shortfall (LRMES)

This appendix describes the procedure to estimate the firm equity return conditional on the systemic event, $LRMES$, as defined in (4). $LRMES$ is non-observable and, in line with Brownlees and Engle (2017), we calculate this quantity using a DCC-GARCH model (see Engle 2002, 2009). Following Brownlees and Engle (2017), we denote the logarithmic returns of bank i and market m as $r_{it} = \log(1+R_{it})$ and $r_{mt} = \log(1+R_{mt})$. Conditional on the information set available at $t-1$, I_{t-1}, both variables are jointly distributed and follow an unspecified distribution D with zero mean and time-varying variance and covariance matrix,

$$\left[\begin{array}{c} r_{it} \\ r_{mt} \end{array} \right] \Bigg| I_{t-1} \sim D\left(\mathbf{0}, \left[\begin{array}{cc} \sigma_{it}^2 & \rho_{it}\sigma_{it}\sigma_{mt} \\ \rho_{it}\sigma_{it}\sigma_{mt} & \sigma_{mt}^2 \end{array} \right]\right). \tag{A.1}$$

The time-varying volatilities are assumed to follow a GJR-GARCH model (Glosten, Jagannathan, and Runkle 1993) model as follows:

$$\sigma_{it}^2 = \omega_i + (\alpha_i + \gamma_i I_{it-1}^{-}) r_{it}^2 + \beta_i \sigma_{it-1}^2 \tag{A.2}$$

$$\sigma_{mt}^2 = \omega_m + (\alpha_m + \gamma_m I_{mt-1}^{-}) r_{mt}^2 + \beta_m \sigma_{mt-1}^2, \tag{A.3}$$

where $I_{it}^{-} = 1$ if $r_{it} < 0$ and $I_{mt}^{-} = 1$ if $r_{mt} < 0$. Next, like in Brownlees and Engle (2017), we define the standardized log returns as $\epsilon_{jt} = \frac{r_{jt}}{\sigma_{jt}}$, while their correlation is given by

$$Corr\left(\begin{array}{c} \epsilon_{it} \\ \epsilon_{mt} \end{array} \right) = \left[\begin{array}{cc} 1 & \rho_{it} \\ \rho_{it} & 1 \end{array} \right] = \text{diag}(Q_{it})^{-1/2} Q_{it} \text{diag}(Q_{it})^{-1/2}, \tag{A.4}$$

where Q_{it} is the pseudo-correlation matrix with the following expression:

$$Q_{it} = (1 - \alpha_{ci} - \beta_{ci}) S_i + \alpha_{ci} \left[\begin{array}{c} \epsilon_{it-1} \\ \epsilon_{mt-1} \end{array} \right] \left[\begin{array}{c} \epsilon_{it-1} \\ \epsilon_{mt-1} \end{array} \right]' + \beta_i Q_{it-1}, \tag{A.5}$$

where S_i is the the unconditional correlation matrix between bank and market-adjusted returns, r_i and r_m, respectively. Using quasi-maximum likelihood, we can estimate the parameters ($\omega_{V,i}$, $\omega_{V,m}$, $\alpha_{V,i}$, $\alpha_{V,m}$, $\gamma_{V,i}$, $\gamma_{V,m}$, $\beta_{V,i}$, $\beta_{V,m}$, α_C, β_C) as well as the time-varying volatilities $\{\sigma_{i,t},\sigma_{m,t}\}_{t=1,\ldots,T}$ and correlations $\{\rho_{i,t}\}_{t=1,\ldots,T}$. The LRMES can, then, be calculated via simulations. For this purpose, we first compute the standardized innovations,

$$\epsilon_{m,t} = \frac{r_{m,t}}{\sigma_{mt}}, \text{ and } \xi_{i,t} = \left(\frac{r_{i,t}}{\sigma i, t} - \rho_{i,t}\epsilon_{m,t}\right) / \sqrt{1 - \rho_{i,t}^2}, \tag{A.6}$$

for $t = 1, \ldots, T$. To generate joint paths of $\{R_{i,t+l}, R_{m,t+l}\}_{l=1,\ldots,h}$, we first sample with replacement h pairs of standardized returns $\{\epsilon_{m,k}, \xi_{i,k}\}_{k=1,\ldots,h}$. Starting with the estimated $\sigma_{i,T}$, $\sigma_{m,T}$, $\rho_{i,T}$, we can compute $\sigma_{i,T+1}$, $\sigma_{m,T+1}$, $\rho_{i,T+1}$ using expressions from (A.2) to (A.6), and $r_{i,T+1}, r_{m,T+1}$ using the sampled $\{\epsilon_{m,1}, \xi_{i,1}\}$. Iterating, we obtain a simulated sample $\{R_{i,t+l}, R_{m,t+l}\}_{l=1,\ldots,h}$. The LRMES is, then, simply calculated as the average of $R_{i,t+h}$ over paths in which $R_{m,t+h} < C$.

Appendix B. Sample of Listed Banks

- AUSTRIA: Bank für Tirol und Vorarlberg; BAWAG Group; BKS Bank; Erste Group Bank; Oberbank; Raiffeisen Bank International; Volksbank Vorarlberg.
- BELGIUM: Dexia; KBC Group.
- BULGARIA: Bulgarian American Credit Bank; Central Cooperative Bank; First Investment Bank; Texim Bank.
- CROATIA: Privredna banka Zagreb.
- CYPRUS: Bank Cyprus Holdings Public; Hellenic Bank Public.
- CZECH REPUBLIC: Komercní banka; MONETA Money Bank.
- DENMARK: BankNordik; Danske Andelskassers Bank; Danske Bank; Den Jyske Sparekasse; Djurslands Bank; Fynske Bank; GrønlandsBANKEN; Hvidbjerg Bank; Jutlander Bank; Jyske Bank; Kreditbanken; Lån & Spar Bank; Lollands Bank; Møns Bank; Nordfyns Bank; Ringkjøbing Landbobank;

Salling Bank; Skjern Bank; Spar Nord Bank; Sparekassen SjællandFyn; Sydbank; Totalbanken; Vestjysk Bank.

- ESTONIA: AS LHV Group.
- FINLAND: Ålandsbanken; Evli Pankki Oyj; Nordea Bank Abp; Oma Säästöpankki Oyj.
- FRANCE: BNP Paribas; CRCAM de Toulouse 31; CRCAM Paris et IDF; CRCAM d'Ille-et-Villaine; CRCAM du Morbihan; CRCAM de Nord de France; CRCAM Brie Picardie; CRCAM du Languedoc; CRCAM Atlantique Vendee; Crédit Agricole; Natixis; Société Générale.
- GERMANY: Aareal Bank; Comdirect bank; Commerzbank; Deutsche Bank; Deutsche Pfandbriefbk; ProCredit Holding.
- GREECE: Alpha Bank; Attica Bank; Eurobank Ergasias; National Bank Greece; Piraeus Financial Holdings.
- HUNGARY: OTP Bank; Takarék Jelzálogbank Nyrt.
- ITALY: Banca Carige; Banca Finnat Euramerica; Banca Generali; Banca Monte dei Paschi di Siena; Banca Popolare di Milano; Banca Popolare di Sondrio; Banca Profilo; Banca Sistema; Banco BPM Società per Azioni; Banco di Desio e della Brianza; Banco di Sardegna; BPER Banca; Credito Emiliano; FinecoBank; Banca Fineco; Intesa Sanpaolo; Mediobanca Banca di Credito Finanziario; UniCredit; Unione di Banche Italiane.
- LITHUANIA: AB Siauliu Bankas.
- NETHERLANDS: ABN AMRO Bank; ING Groep; Van Lanschot Kempen.
- NORWAY: Aurskog Sparebank; DNB ASA; Høland og Setskog Sparebank; Instabank; Jæren Sparebank; Komplett Bank; Melhus Sparebank; Norwegian Finans Holding; Sandnes Sparebank; Sbanken; Skue Sparebank; Sogn Sparebank; SpareBank 1; SpareBank 1 Helgeland; SpareBank 1 Nord-Norge; Sparebank 1 Nordvest; SpareBank 1 Østfold Akershus; SpareBank 1 Østlandet; SpareBank 1 Ringerike Hadeland; SpareBank 1 SMN; SpareBank 1 SRBank; SpareBank 1 Telemark; Sparebanken Møre; Sparebanken Øst; Sparebanken Sør; Sparebanken Vest; Totens Sparebank; Voss Veksel og Landmandsbank.
- POLAND: Alior Bank; Bank Handlowy w Warszawie; Bank Millennium; Bank Ochrony Srodowiska; Bank Polska Kasa

Opieki; BNP Paribas Bank Polska; Getin Holding; Getin Noble Bank; ING Bank Slaski; mBank; Powszechna Kasa Oszczednosci Bank Polski; Santander Bank Polska.

- PORTUGAL: Banco BPI; Banco Comercial Português; Banco Espírito Santo.
- ROMANIA: Banca Transilvania; BRD Groupe Société Générale.
- SPAIN: Banco Bilbao Vizcaya Argentaria; Banco de Sabadell; Banco de Valencia; Banco Popular Español; Banco Santander; Bankia; Bankinter; CaixaBank; Liberbank; Unicaja Banco.
- SLOVAKIA: OTP Banka Slovensko; Vseobecna uverova banka.
- SWEDEN: Avanza Bank Holding; Collector; Handelsbanken; Skandinaviska Enskilda Banken; Swedbank; TF Bank.
- UNITED KINGDOM: Barclays; HSBC Holdings; Lloyds Banking Group; Metro Bank; NatWest Group; Standard Chartered.

Appendix C. Robustness Exercise: Impact of the Pandemic

Table C.1 shows estimates of the baseline model for the full sample. S_BUF and SII_STAT are interacted with a dummy indicator, $COVID$, that is 1 during the pandemic period and 0 otherwise.

Table C.1. Estimates of Baseline Model for Full Sample

	Full Sample			
	MI	**M2**	**M3**	**M4**
S_BUF	–79.40 (276.7)	–620.5*** (222.4)		
S_BUF × *COVID*		1,642*** (492.8)		
SII_STAT			–1,303** (518.3)	–2,406*** (609.2)
SII_STAT × *COVID*				5,223*** (1,012)
Total Assets	0.039*** (0.003)	0.038*** (0.003)	0.039*** (0.003)	0.038*** (0.003)
RoE	–8.10*** (3.01)	–8.14*** (2.98)	–8.14*** (3.01)	–7.79*** (2.93)
CCyB	–499.0** (236.1)	–448.6** (191.6)	–513.6** (254.0)	–673.6*** (213.6)
SyRB	318.5*** (100.6)	317.7*** (115.3)	167.9 (115.3)	347.6*** (101.2)
CCoB	636.3*** (208.1)	704.1*** (203.8)	653.9*** (228.0)	489.7*** (185.0)
GDPpc	–76.04*** (26.27)	–75.45*** (25.77)	–71.69*** (24.02)	–74.92*** (22.75)
CDS	0.003 (0.032)	0.005 (0.031)	–0.013 (0.036)	0.001 (0.032)
N	6,487	6,487	6,487	6,487
R^2	0.396	0.409	0.400	0.430

Note: Dependent variable: systemic risk contribution of individual banks as proxied by the SRISK indicator (see Brownlees and Engle 2017); all explanatory variables are lagged one period; *S_BUF*: SII buffer level; *SII_STAT*: binary dummy, $SII\_STAT = 1$ while a bank has SII status; *COVID*: binary dummy, $COVID = 1$ from 2020:Q1 to 2021:Q3. See Table 2 for a complete description of the explanatory variables and data sources. Standard errors are shown in parentheses. Standard errors are robust to heteroskedasticity and serial correlation. Intercept, time, and bank fixed effects are included but not reported. ***, **, and * refer to significance at 1 percent, 5 percent, and 10 percent level, respectively.

References

Acharya, V., L. Pedersen, T. Philippon, and M. Richardson. 2017. "Measuring Systemic Risk." *Review of Financial Studies* 30 (1): 2–47.

Adrian, T., and M. K. Brunnermeier. 2016. "CoVaR." *American Economic Review* 106 (7): 1705–41.

Aiyar, S., C. Calomiris, and T. Wieladek. 2014. "Does Macro-Prudential Regulation Leak? Evidence from a UK Policy Experiment." *Journal of Money, Credit and Banking* 46 (s1): 181–214.

———. 2016. "How Does Credit Supply Respond to Monetary Policy and Bank Minimum Capital Requirements?" *European Economic Review* 82 (February): 142–65.

Andrieş, A. M., S. Nistor, S. Ongena, and N. Sprincean. 2020. "On Becoming an O-SII ('Other Systemically Important Institution')." *Journal of Banking and Finance* 111 (February): Article 105723.

Bank for International Settlements (BIS). 2012. "Operationalising the Selection and Application of Macroprudential Instruments." Committee on the Global Financial System (CGFS) Paper No. 48, December.

Basel Committee on Banking Supervision (BCBS). 2011. "Global Systemically Important Banks: Assessment Methodology and the Additional Loss Absorbency Requirement." Bank for International Settlements.

———. 2012. "A Framework for Dealing with Domestic Systemically Important Banks." Bank for International Settlements.

———. 2013. "Global Systemically Important Banks: Updated Assessment Methodology and the Higher Loss Absorbency Requirement." Bank for International Settlements.

Bats, J. V., and A. C. Houben. 2020. "Bank-based versus Market-based Financing: Implications for Systemic Risk." *Journal of Banking and Finance* 114 (May): Article 105776.

Bekaert, G., and J. Breckenfelder. 2019. "The (re)allocation of Bank Risk." Columbia Business School Research Paper, forthcoming.

Benoit, S., J. E. Colliard, C. Hurlin, and C. Pérignon. 2017. "Where the Risks Lie: A Survey on Systemic Risk." *Review of Finance* 21 (1): 109–52.

Bisias, D., M. Flood, A. W. Lo, and S. Valavanis. 2012. "A Survey of Systemic Risk Analytics." Working Paper No. 12-01, Office of Financial Research, U.S. Department of the Treasury.

Bostandzic, D., F. Irresberger, R. Juelsrud, and G. Weiß. 2022. "Do Capital Requirements Make Banks Safer? Evidence from a Quasi-natural Experiment." *Journal of Financial and Quantitative Analysis* 57 (5): 1805–33.

Braouezec, Y., and K. Kiani. 2021. "Target Capital Ratio and Optimal Channel(s) of Adjustment: A Simple Model with Empirical Applications to European Banks." *Applied Economics* 53 (13): 1435–62.

Bridges, J., D. Gregory, M. Nielsen, S. Pezzini, A. Radia, and M. Spaltro. 2014. "The Impact of Capital Requirements on Bank Lending." Bank of England Working Paper No. 486.

Brownlees, C., B. Chabot, E. Ghysels, and C. Kurz. 2020. "Back to the Future: Backtesting Systemic Risk Measures during Historical Bank Runs and the Great Depression." *Journal of Banking and Finance* 113 (April): Article 105736.

Brownlees, C., and R. Engle. 2017. "SRISK: A Conditional Capital Shortfall Measure of Systemic Risk." *Review of Financial Studies* 30 (1): 48–79.

Cappelletti, G., A. Ponte Marques, P. Varraso, Ž. Budrys, and J. Peeters. 2019. "Impact of Higher Capital Buffers on Banks' Lending and Risk-taking: Evidence from the Euro Area Experiments." ECB Working Paper No. 2292.

Cappelletti, G., A. Reghezza, C. R. d'Acri, and M. Spaggiari. 2020. "Compositional Effects of O-SII Capital Buffers and the Role of Monetary Policy." ECB Working Paper No. 2440.

Cohen, B. H., and M. Scatigna. 2016. "Banks and Capital Requirements: Channels of Adjustment." *Journal of Banking and Finance* 69 (Supplement 1): S56–S69.

Coleman, T. F., A. LaPlante, and A. Rubtsov. 2018. "Analysis of the SRISK Measure and Its Application to the Canadian Banking and Insurance Industries." *Annals of Finance* 14 (4): 547–70.

Dautović, E. 2020. "Has Regulatory Capital Made Banks Safer? Skin in the Game vs. Moral Hazard." ECB Working Paper No. 2449.

Drehmann, M., and L. Gambacorta. 2012. "The Effects of Countercyclical Capital Buffers on Bank Lending." *Applied Economics Letters* 19 (7): 603–8.

Engle, R. 2002. "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroscedasticity Models." *Journal of Business and Economic Statistics* 20 (3): 339–50.

———. 2009. *Anticipating Correlations: A New Paradigm for Risk Management.* Princeton, NJ: Princeton University Press.

Engle, R., and S. Manganelli. 2004. "CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles." *Journal of Business and Economic Statistics* 22 (4): 367–81.

Engle, R., and T. Ruan. 2019. "Measuring the Probability of a Financial Crisis." *Proceedings of the National Academy of Sciences* 116 (37): 18341–46.

European Banking Authority (EBA). 2014. "Guidelines on Criteria for the Assessment of O-SIIs." EBA Guidelines No. 2014/10.

European Systemic Risk Board (ESRB). 2013. "Recommendation of the European Systemic Risk Board of 4 April 2013 on Intermediate Objectives and Instruments of Macro-prudential Policy (ESRB/2013/1)."

———. 2021. "A Review of Macroprudential Policy in the EU in 2020." July.

Glosten, L. R., R. Jagannathan, and D. E. Runkle. 1993. "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks." *Journal of Finance* 48 (5): 1779–1801.

Grinderslev, O. J., and K. L. Kristiansen. 2016. "Systemic Risk in Danish Banks: Implementing SRISK in a Danish Context." Danmarks Nationalbank Working Paper No. 105.

Gropp, R., T. Mosk, S. Ongena, and C. Wix. 2019. "Bank Response to Higher Capital Requirements: Evidence from a Quasi-natural Experiment." *Review of Financial Studies* 32 (1): 266–99.

Gündüz, Y. 2020. "The Market Impact of Systemic Risk Capital Surcharges." Deutsche Bundesbank Discussion Paper No. 09/2020.

Hansen, L. 2014. "Challenges in Identifying and Measuring Systemic Risk." In *Risk Topography: Systemic Risk and Macro Modeling,* ed. M. Brunnermeier and A. Krishnamurthy, 15–30. Chicago: University of Chicago Press.

International Monetary Fund (IMF). 2011. "Macroprudential Policy: An Organizing Framework."

Jiménez, G., S. Ongena, J. L. Peydró, and J. Saurina. 2017. "Macroprudential Policy, Countercyclical Bank Capital Buffers and Credit Supply: Evidence from the Spanish Dynamic Provisioning Experiments." *Journal of Political Economy* 125 (6): 2126–77.

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

Mayordomo, S., and M. Rodríguez-Moreno. 2020. "How Do European Banks Cope with Macroprudential Capital Requirements?" *Finance Research Letters* 38 (January): Article 101459.

Stern, G. H., and R. J. Feldman. 2004. *Too Big to Fail: The Hazards of Bank Bailouts.* Brookings Institution Press.

Tavolaro, S., and F. Visnovsky. 2014. "What is the Information Content of the SRISK Measure as a Supervisory Tool?" Banque de France Débats Économiques et Financiers (Economic and Financial Debates) No. 10.

Vogel, U. 2020. "O-SII Designation and Deposit Funding Costs." *Economics Letters* 192 (July): Article 109261.

Heterogeneous Expectations and the Business Cycle*

Tolga Özden
Bank of Canada

We analyze the empirical relevance of heterogeneous expectations and central bank credibility in a canonical New Keynesian model subject to the effective lower bound (ELB). Agents switch between an anchored rational expectations (RE) and an adaptive learning forecast rule, where the latter may result in a de-anchoring of inflation expectations. We estimate the model for the U.S. economy using aggregate macrodata and survey data on inflation expectations. We use the estimated model to examine the interaction between the risk of deflationary spirals and central bank credibility at the ELB. A loss of central bank credibility increases the probability of deflationary spirals, highlighting the importance of keeping inflation expectations anchored during periods of uncertainty.

JEL Codes: E37, E65, C11, C32.

1. Introduction

Following the Global Financial Crisis (GFC) of 2007–08, leading central banks around the globe cut their nominal interest rates to near-zero levels and encountered the effective lower bound (ELB) constraint on their rates. This has led to an increased volume of research about the relevance and impact of the constraint on the economy. In the aftermath of the Great Recession, central banks have increasingly relied on communication policies in the form of forward guidance and signaling, which have become an important

*Stimulating comments and suggestions from Borağan Aruoba, Cars Hommes, Sergey Slobodyan, Jacek Suda, Florian Wicknig, Raf Wouters, Yang Zhang, and anonymous referees are gratefully acknowledged. The views expressed in this paper are solely those of the author and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank. Author e-mail: tozden@bank-banque-canada.ca.

pillar of many central banks' unconventional policy toolkit. Given the rise in the frequency and intensity of communication-based tools by central banks, an important and closely related issue is central bank credibility and its interaction with the business cycle dynamics at the ELB.

The macroeconomic literature often examines the impact of ELB on business cycles using the rational expectations (RE) approach. In a standard model, this approach assumes that agents in the economy have complete knowledge of the central bank's objective function and trust that future policy actions will align with that objective. Consequently, there is little room to study the relevance of endogenous central bank credibility in a model with rational expectations. In this paper we relax the full rationality assumption and propose a heterogeneous expectation model with limited information, where agents are allowed to switch between two types of forecasting rules.

As our starting point, we use the canonical three-equation hybrid New Keynesian model, subject to the ELB constraint on nominal interest rates. We introduce heterogeneous expectations to this framework, where agents are allowed to switch between an anchored pseudo-rational expectation model and an adaptive learning model where expectations may become de-anchored if certain conditions are met. The switching mechanism between these two types of expectations is endogenous in the model, where the relative agent shares using each type of forecasting rule depend on their past predictive performance.

A key novelty of our model is that when a high proportion of adaptive learners is combined with the ELB constraint, the economy loses its stability. In these cases, a rising share of adaptive learners corresponds to a loss of trust in the central bank's ability to circumvent the ELB constraint through unconventional monetary policy measures. Consequently, more agents abandon the rational expectations rule and switch to adaptive learning. The presence of more adaptive learners weakens the feedback channel from the central bank's desired interest rate path (shadow rate) to inflation and output gap, which intensifies deflationary pressures. Combined with the ELB, this leads to a higher real interest rate and depresses aggregate demand. Adverse shocks can trigger deflationary spirals under such circumstances if the share of adaptive learners exceeds a critical threshold, where expectations become de-anchored on the downside and the central bank is unable to combat ever-falling inflation and output gap due to the ELB constraint.

We estimate the model for the United States using historical data on consumer price index (CPI) inflation, federal funds rate, and gross domestic product (GDP) as well as short-term (one-quarter) and long-term (10-year) inflation expectations. For inflation expectations, we utilize a novel index of the term structure of inflation expectations, ATSIX, proposed in Aruoba (2020) and regularly published by the Federal Reserve Bank of Philadelphia.[1] To account for the ELB constraint in our sample, we use a regime-switching approximation in the estimation procedure. With the introduction of adaptive learning and time-varying shares of agents, the model is characterized by a conditionally linear structure. We combine this with the standard filtering algorithm in regime-switching literature à la Kim and Nelson (1999) and reformulate the model in state space form with time-varying parameters, which allows us to estimate the structural parameters of the model with standard Bayesian Markov chain Monte Carlo (MCMC) methods. We use the estimated model to pin down the conditions that are needed for deflationary spirals to occur, as well as to assess the likelihood of encountering such scenarios.

The paper is closely related to Özden and Wouters (2021), where a medium-scale dynamic stochastic general equilibrium (DSGE) model is estimated with various representative agent learning rules to examine their fitness before and after the Great Recession. In this study we make use of the same estimation methodology for adaptive learning models developed in that paper. There are two key features that distinguish our current analysis: First, we allow for heterogeneity of expectations, which is a crucial channel both for fitting the data and to study endogenous central bank credibility. Second, we use inflation expectations to estimate the model, which allows for a more robust identification of the parameters related to learning and heterogeneity.

It is important to distinguish the deflationary spiral channel studied in this paper from those that emerge in RE models. As shown in Bianchi, Melosi, and Rottner (2021), deflationary spirals can also

[1]The details of the index can be found at https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/atsix.

arise in a fully rational setup. When low long-run interest rates are combined with agents' expectations about future ELB regimes, a deflationary bias can occur even when the ELB is not binding. If the deflationary bias becomes excessive, the RE equilibrium loses its determinacy and deflationary spirals occur. This highlights the *monetary policy rule* channel of deflationary spirals, where the possibility of hitting the ELB in a future period renders symmetric monetary policy rules sub-optimal. The central bank can mitigate the risk of deflationary spirals by implementing an asymmetric rule instead, whereby its response to inflation above target is slower than its response to inflation below target. This reduces the risk of encountering the ELB regime in the future, which in turn reduces the risk of deflationary spirals. In contrast, our analysis focuses on the *central bank credibility* channel, where deflationary spirals can be mitigated by anchoring expectations at the desired equilibrium. This underscores the importance of effective central bank communication policies to minimize the associated risk.

The key results of the paper are as follows: (i) The heterogeneous expectation model fits the data better than a pure RE or pure adaptive learning model. The model performs particularly well in terms of generating realistic inflation expectation dynamics. (ii) The estimated shares of rational and adaptive agents during the ELB regime 2009–15 are close to 50 percent for the United States, suggesting that expectations have partially reacted to the shadow rate over this episode. (iii) The presence of adaptive learners contributes to a de-anchoring of inflation expectations both on the upside when inflation is high, such as during the Great Inflation period, and on the downside when the ELB constraint is binding, as observed during the Great Recession period. (iv) A high share of adaptive learners and a loss of central bank credibility increase the risk of deflationary spirals. When agents start to extrapolate recent data more, the risk of deflationary spirals increases further.

At the time of writing this paper, many advanced economies have been experiencing rising and persistent inflationary pressures. This has brought many central banks' focus back to inflation expectations, with worries over potential de-anchoring risks.[2] While the

[2]See, e.g., Blanchard (2022) and the recent speech by Carstens (2022).

main focus of this paper is on business cycle dynamics at the ELB regime, the estimation results over the high-inflation period of the '60s and '70s also shed light on today's issues. The paper is organized as follows: Section 2 introduces the key model features, as well as assumptions on heterogeneous expectations and ELB, together with some theoretical results to illustrate the model properties. Section 3 discusses the estimation methodology, key results, and the empirical properties of the model. Section 4 discusses a number of counterfactual exercises at the ELB to study the interaction between heterogeneous expectations and the risk of deflationary spirals. Section 5 concludes.

1.1 Literature Review

Our paper relates to several strands of literature on adaptive learning, heterogeneous expectation, and regime switching in DSGE models. Earlier work on heterogeneous expectations in New Keynesian models considers a variety of topics; e.g., Branch (2004) studies the empirical properties of heterogeneous expectations with survey data on inflation expectations; Branch and McGough (2009) analyze the microfoundations of New Keynesian models with heterogeneous expectations; Anufriev et al. (2013) consider different interest rate rules and macroeconomic stability under heterogeneous expectations; Di Bartolomeo, Di Pietro, and Giannini (2016) study how heterogeneous expectations affect the design of optimal monetary policy in a New Keynesian model; Cornea-Madeira, Hommes, and Massaro (2019) estimate the New Keynesian Phillips curve with heterogeneous expectations; and Hommes, Massaro, and Weber (2019) test a number of heterogeneous and bounded rationality models in a learning-to-forecast experiment.

More recently, there have been a number of papers that study the interactions between the ELB, unconventional monetary policy, and heterogeneous expectations. A closely related study is Busetti et al. (2017), where the authors study how prolonged periods of weak inflation in the euro zone may induce a de-anchoring of expectations. Other closely related papers include Andrade et al. (2019), who consider forward guidance in a heterogeneous expectations framework with optimistic and pessimistic agents; Hommes and Lustenhouwer (2019), who study the theoretical properties of a New Keynesian

(NK) model with an ELB under heterogeneous expectations; Goy, Hommes, and Mavromatis (2020), who analyze the effects of different types of forward guidance in a New Keynesian model with heterogeneous expectations and the ELB constraint; Lansing (2021), where a representative agent contemplates between a targeted equilibrium and a deflationary equilibrium, and where a non-trivial probability on the deflationary equilibrium becomes partially self-fulfilling by lowering the averages of observed variables; Arifovic et al. (2020), who study heterogeneous expectations through a novel mechanism called social learning; and Carvalho et al. (2021), who estimate a model where agents are allowed to switch between decreasing and constant gain algorithms to form their expectations. The marginal contribution of our paper to this literature is to estimate the model under heterogeneous expectations together with survey data and endogenous regime switching in monetary policy (MP).

When it comes to regime-switching models, the RE framework plays a central role in the DSGE literature. These models focus on the theoretical properties and solution methods within the RE framework.[3] More recently, a number of papers also study DSGE models with endogenous regime switching under RE.[4] While there is ample research in regime-switching models with rational agents, research in this class of models with imperfect information/learning agents has been scarce. Examples include Branch, Davig, and McGough (2007), who establish theoretical properties of learning about both regime switches and structural relations, and Gust,

[3]Examples include Farmer, Waggoner, and Zha (2009, 2011) and Cho (2016), who study the theoretical properties and determinacy conditions associated with RE equilibria in Markov-switching models; Bianchi (2016), who proposes new methods for measuring expectations and uncertainty in Markov-switching models; and Kulish and Pagan (2017), who propose solution and estimation methods for forward-looking models with structural changes under a variety of assumptions for agents' beliefs about those structural changes. Other empirical applications in regime-switching DSGE models include, among others, Sims and Zha (2006), Liu and Mumtaz (2011), Bianchi (2016), and Bianchi and Ilut (2017).

[4]See, e.g., Barthélemy and Marx (2017), who use perturbation methods to solve and estimate endogenous regime-switching models; Chang, Maih, and Fei (2018), who propose an efficient filtering method to handle the estimation of state space models with endogenous switching parameters depending on latent autoregressive factors; and Benigno et al. (2020), who consider an endogenous regime-switching framework to study financial crises.

Herbst, and Lopez-Salido (2018), who study the effectiveness of forward guidance in a model where agents are aware of regime switches but do not know the transition probabilities and instead infer about them using a form of Bayesian learning.

The paper also relates to models studying the effects of ELB and unconventional monetary policy under imperfect information and adaptive learning. Examples include Evans, Guse, and Honkapohja (2008), who study the global dynamics of liquidity traps under adaptive learning; Haberis, Harrison, and Waldron (2014), who analyze macroeconomic effects of transient interest rate pegs in an imperfect information model; Eusepi and Preston (2010), who consider central bank communication in a model where agents' expectations are not consistent with the central bank policy; Cole (2018), who studies the effectiveness of learning on forward guidance, where forward guidance is introduced into monetary policy with a sequence of shocks; and similarly Cole and Martínez-García (2019), who study the effectiveness of forward guidance in a New Keynesian model with imperfect central bank credibility. The present paper contributes to this literature by allowing a fraction of agents to use adaptive learning rules through an evolutionary selection mechanism, and by estimating the model using survey data on inflation expectations.

2. Model Setup

2.1 Structural Equations and Rational Expectations

We consider the simple canonical version of the three-equation New Keynesian model as in Clarida, Gali, and Gertler (1999).[5] We first present the skeleton form of the model without any regime switching, given by the following structural equations:

$$\begin{cases} y_t = (1-\iota_y)E_t y_{t+1} + \iota_y y_{t-1} - \frac{1}{\tau}(r_t - E_t\pi_{t+1}) + u_{y,t}, \\ \pi_t = \beta((1-\iota_p)E_t\pi_{t+1} + \iota_p\pi_{t-1}) + \kappa y_t + u_{\pi,t}, \\ r_t = \rho_r r_{t-1} + (1-\rho_r)(\phi_\pi\pi_t + \phi_y y_t) + \phi_{\Delta y}(y_t - y_{t-1}) + \varepsilon_{r,t}, \end{cases} \tag{1}$$

[5] Similar setups have been considered in closely related papers of Busetti et al. (2017), Goy, Hommes, and Mavromatis (2020), and Lansing (2021), among others.

where y_t, π_t, and r_t denote the output gap, inflation, and nominal interest rate, respectively. The first equation represents the IS curve, where ι_y is the intrinsic level of inertia (or indexation) in output gap, and τ is the intertemporal elasticity of substitution for households. The second equation is the Phillips curve, with ι_p the price indexation and κ denoting the slope of the Phillips curve. The last equation is the monetary policy reaction function, with ρ_r the interest smoothing rate, ϕ_π inflation reaction, ϕ_y output gap reaction, and $\phi_{\Delta y}$ output gap growth reaction. The model is supplemented with three shocks, where the demand shock $u_{y,t}$ and cost-push shock $u_{\pi,t}$ follow AR(1) processes given by

$$\begin{cases} u_{y,t} = \rho_y u_{y,t-1} + \varepsilon_{y,t}, \\ u_{\pi,t} = \rho_\pi u_{\pi,t-1} + \varepsilon_{\pi,t}. \end{cases} \tag{2}$$

The monetary policy shock $\varepsilon_{r,t}$ is assumed to be an i.i.d. process. Before introducing the ELB constraint on the nominal rates and the regime-switching setup, it is useful to start with the rational expectations (RE) equilibrium of the model, associated with the minimum state variable (MSV) solution. The model can be written in the standard matrix form:

$$\begin{cases} \boldsymbol{A} X_t = \boldsymbol{B} X_{t-1} + \boldsymbol{C} E_t X_{t+1} + \boldsymbol{D} u_t, \\ u_t = \boldsymbol{\rho} u_{t-1} + \varepsilon_t, \end{cases} \tag{3}$$

for conformable matrices $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$, $\boldsymbol{D}$, and $\boldsymbol{\rho}$, with $X_t = [y_t, \pi_t, r_t]'$, $u_t = [u_{y,t}, u_{\pi,t}, 0]'$, and $\varepsilon_t = [\varepsilon_{y,t}, \varepsilon_{\pi,t}, \varepsilon_{r,t}]'$. The standard deviations of the i.i.d. shocks are denoted by $\eta_t = [\eta_y, \eta_\pi, \eta_r]'$. Under RE, the equilibrium solution takes the following form, along with the implied one-step-ahead expectations:

$$\begin{cases} X_t = \boldsymbol{b} X_{t-1} + \boldsymbol{d} u_t, \\ E_t X_{t+1} = \boldsymbol{b} X_t + \boldsymbol{d}\boldsymbol{\rho} u_t. \end{cases} \tag{4}$$

Plugging the expectations back into the law of motion (3) yields

$$(\boldsymbol{A} - \boldsymbol{C}\boldsymbol{b}) X_t = \boldsymbol{B} X_{t-1} + (\boldsymbol{C}\boldsymbol{d}\boldsymbol{\rho} + \boldsymbol{D}) u_t. \tag{5}$$

The RE solution is then pinned down by the following fixed-point conditions:[6]

$$\begin{cases} \boldsymbol{b} = (\boldsymbol{A} - \boldsymbol{C}\boldsymbol{b})^{-1}\boldsymbol{B}, \\ \boldsymbol{d} = (\boldsymbol{A} - \boldsymbol{C}\boldsymbol{b})^{-1}(\boldsymbol{C}\boldsymbol{d}\boldsymbol{\rho} + \boldsymbol{D}). \end{cases} \tag{6}$$

2.2 ELB and Regime Switching

In this paper our main objective is to evaluate the effects of the ELB constraint on macroeconomic outcomes. Introducing the constraint on the interest rate rule leads to the following form:

$$r_t = max\{\bar{r}, \rho_r r_{t-1} + (1-\rho_r)(\phi_\pi \pi_t + \phi_y y_t) + \phi_{\Delta y}(y_t - y_{t-1}) + \varepsilon_{r,t}\}, \tag{7}$$

which is an occasionally binding constraint (OBC) on the nominal rates, with $\bar{r}$ corresponding to the ELB value. In the literature, a popular method for approximating this OBC-induced non-linearity is to consider a regime-switching approach, used in, e.g., Binning and Maih (2016), Chen (2017), and Lindé, Maih, and Wouters (2017). In this setup, monetary policy is subject to two different regimes: a Taylor-rule regime where interest rates follow the intended reaction function when the ELB constraint does not bind, and an ELB regime where monetary policy becomes inactive when the reaction function becomes constrained by the lower bound. If we denote by s_t the regime-switching process, which can take on values $s_t = E$ (ELB regime) and $s_t = T$ (Taylor-rule regime), the monetary policy rule evolves according to

$$\begin{cases} r_t(s_t = T) = \rho_r r_{t-1} + (1-\rho_r)(\phi_\pi \pi_t + \phi_y y_t) + \\ \qquad\qquad\qquad \phi_{\Delta y}(y_t - y_{t-1}) + \varepsilon^T_{r,t}, \\ r_t(s_t = E) = \bar{r} + \varepsilon^E_{r,t}. \end{cases} \tag{8}$$

The transition matrix is given by time-varying probabilities as follows:[7]

---

[6]We make use of the methods introduced in Uhlig (1995) to solve for the fixed-point conditions.

[7]For standard deviations of monetary policy shocks, we use the notation $\eta_{r,T}$ and $\eta_{r,E}$ at Taylor and ELB regimes, respectively.

$$\boldsymbol{Q}_t = \begin{bmatrix} q_t^T & 1 - q_t^T \\ 1 - q_t^E & q_t^E \end{bmatrix},$$

where the probabilities q_t^T and q_t^E depend on the central bank's desired policy rate at every period, which is defined as the shadow rate henceforth. More formally, we assume that the shadow rate r_t^* follows:

$$\begin{cases} r_t^*(s_t = T) = \rho_r r_{t-1} + (1 - \rho_r)(\phi_\pi \pi_t + \phi_y y_t) + \phi_{\Delta y}(y_t - y_{t-1}), \\ r_t^*(s_t = E) = \rho_r r_{t-1}^* + (1 - \rho_r)(\phi_\pi \pi_t + \phi_y y_t) + \phi_{\Delta y}(y_t - y_{t-1}). \end{cases} \tag{9}$$

This structure makes use of the following assumptions: The shadow rate r_t^* is the central bank's desired level of nominal interest rate *in the absence of monetary policy shocks and the ELB constraint.* During normal times with the Taylor rule, the shadow rate is smoothed over the observed nominal interest rate. Therefore during normal times, the only difference between these two rates is the presence of i.i.d. monetary policy shocks. During ELB periods when nominal rates are constrained, the shadow rate is smoothed over itself, which allows for persistent deviations from the nominal rate beyond the i.i.d. monetary policy shocks. This captures the idea of keeping the interest rates *lower for longer*, where the central bank wants to keep the policy rate at near-ELB levels until the shadow rate recovers back to a level above the ELB.

Given the shadow rate r_t^*, the transition probabilities are determined according to

$$\begin{aligned} q_t^T &= \frac{\theta_1}{\theta_1 + exp(-\Phi_1(r_t^* + (\bar{r_T} - \bar{r_E})))}, \\ q_t^E &= \frac{\theta_2}{\theta_2 + exp(\Phi_2(r_t^* + (\bar{r_T} - \bar{r_E})))}, \end{aligned} \tag{10}$$

where $\bar{r_T}$ and $\bar{r_E}$ are the constant trend values of the nominal interest rate during Taylor and ELB regimes, respectively.[8] These parameters are introduced into the measurement equations as constants and are estimated jointly with the structural parameters of the model, which is discussed further in Section 3.

[8]Given the trend values $\bar{r_T}$ and $\bar{r_E}$, the identity for ELB constraint in (7) is given by $\bar{r} = -\bar{r_T} + \bar{r_E}$.

In a regime-switching world, the RE solution makes use of two key assumptions: Agents are aware of the current underlying regime s_t, and they know the transition matrix $\boldsymbol{Q_t}$ associated with the regimes. In other words, RE models equate agents' subjective expectations about regime switches to the objective model expectations, leading to regime-dependent expectations in the following form:

$$\begin{aligned} E_t[X_{t+1}|s_t = T] &= q_t^T(\boldsymbol{b}(s_{t+1} = T)X_t + \boldsymbol{d}(s_{t+1} = T)\boldsymbol{\rho} u_t) + \\ &(1 - q_t^T)(\boldsymbol{b}(s_{t+1} = E)X_t + \boldsymbol{d}(s_{t+1} = E)\boldsymbol{\rho} u_t), \\ E_t[X_{t+1}|s_t = E] &= q_t^E(\boldsymbol{b}(s_{t+1} = E)X_t + \boldsymbol{d}(s_{t+1} = E)\boldsymbol{\rho} u_t) + \\ &(1 - q_t^E)(\boldsymbol{b}(s_{t+1} = T)X_t + \boldsymbol{d}(s_{t+1} = T)\boldsymbol{\rho} u_t). \end{aligned} \tag{11}$$

The RE solution in the baseline version of the model in (1) is unique and determinate when the Taylor principle of $\phi_\pi > 1$ is satisfied. The equilibrium becomes indeterminate at the ELB when monetary policy is not active. Davig and Leeper (2007) establish that in a regime-switching environment with RE, the equilibrium determinacy can continue to hold even if one of the underlying regimes is indeterminate. They define this property as the *long-run Taylor principle* (LRTP). The implications of this for the canonical New Keynesian model with Taylor and ELB regimes is that, as long as the passive (indeterminate) periods are sufficiently short lived relative to the active (determinate) periods, the model dynamics can still be characterized by a determinate equilibrium. This property allows for the estimation and simulation of RE models with regime switching in the presence of indeterminate regimes.

2.3 Heterogeneous Expectations

In this paper we deviate from the standard full rationality assumption by breaking the tight link between subjective expectations and objective model-implied expectations. In particular, we relax the assumption that agents are aware of the underlying regime s_t and the transition probability matrix $\boldsymbol{Q_t}$. We further relax the assumption that agents are rational; instead we introduce a heterogeneous expectation mechanism with anchored and de-anchored expectation rules, explained in further detail below.

2.3.1 Anchored Rational Expectations

The first type of agents form their expectations using the rational solution in (6) associated with an active Taylor rule $\phi_\pi > 1$. In other words, this type of agent always follows expectations based on a determinate RE solution. During normal times with the Taylor regime ($s_t = T$), this assumption boils down to the standard model solution associated with RE. During ELB periods ($s_t = E$), expectations associated with this type take on a different interpretation: Nominal rates are constrained by the ELB, but expectations evolve *as if* the central bank's desired interest rate path, i.e., the shadow rate r_t^*, is what matters for the economy.

The assumption that agents always use the RE solution associated with active policy rule implicitly means that they know the shadow rate at any given period, even though the shadow rate is not directly observable during ELB periods. Therefore this assumption can be interpreted as a successful central bank communication and correctly anchored expectations on the desired interest rate, which proxies for the impact of a central bank's unconventional policy tools on expectations. We assume that forward guidance communications and quantitative easing measures allow the central bank to correctly signal the desired interest rate and anchor this class of agents' expectations on the targeted equilibrium. Put differently, these agents believe that unconventional monetary policy measures perfectly substitute for the slack on the nominal rates introduced by the ELB constraint.

It is important to note that this expectation formation rule ignores not only the presence of the ELB constraint but also the presence of other agents in the economy that form their expectations differently. Therefore, these expectations correspond to a form of pseudo-rationality only, i.e., what would happen if all expectations were rational and if the monetary policy was not constrained by the ELB. Such behavior is usually referred to as a *fundamentalist* rule in heterogeneous expectations studies.[9] In this paper, we refer to this type as rational agents with *anchored* expectations.

[9]See, e.g., Hommes and Lustenhouwer (2019) and Goy, Hommes, and Mavromatis (2020), where *fundamentalist* agents use the steady-state values or long-run averages of the relevant endogenous variables when forming their expectations.

2.3.2 Adaptive Learning

The second class of agents use a constant gain recursive least squares (RLS) learning rule based on the observable variables of output gap, inflation, and nominal interest rates. Specifically, we assume that agents have the following regression model, along with the implied one-step-ahead expectations:

$$\begin{cases} X_t = \boldsymbol{\alpha}_{t-1} + \boldsymbol{\beta}_{t-1} X_{t-1} + \delta_t, \\ E_t X_{t+1}^L = \boldsymbol{\alpha}_{t-1} + \boldsymbol{\beta}_{t-1} X_t, \end{cases} \tag{12}$$

where $\boldsymbol{\alpha}_{t-1}$ is a vector of perceived means, $\boldsymbol{\beta}_{t-1}$ is the perceived first-order correlation matrix, and δ_t is a vector of i.i.d. shocks. The first equation in (12) is referred to as the agents' *perceived law of motion* (PLM) henceforth. This particular VAR(1) form of learning has been frequently used in the learning literature; see, e.g., Jääskelä and McKibbin (2010), Milani (2011), and Chung and Xiao (2013). It has the advantage of being close to the beliefs consistent with the MSV solution of the model.[10]

We use a *t-timing* assumption on expectations, which means that agents are able to use period-t information when forming their expectations. This corresponds to a joint determination of expectations and period-t variables.[11] Agents update the perceived parameters in their PLM after the endogenous variables are determined, hence these parameters appear with a lag in (12) in the form of $\boldsymbol{\alpha}_{t-1}$ and $\boldsymbol{\beta}_{t-1}$. Under constant gain RLS, the parameters evolve according to

$$\begin{cases} R_t = R_{t-1} + \gamma(\tilde{X}_{t-1}\tilde{X}'_{t-1} - R_{t-1}), \\ \Phi_t = \Phi_{t-1} + \gamma R_t^{-1} \tilde{X}_{t-1}(X_t - \Phi_{t-1}\tilde{X}_{t-1})', \end{cases} \tag{13}$$

[10] The only difference between the MSV solution and VAR(1) expectations is that in the latter, the exogenous AR(1) cost-push and demand shocks are not included in the regression. This keeps the state space of the PLM small and more tractable.

[11] The alternative is to use the assumption of $t-1$ *dating* for both types of agents, which takes on a sequential structure where first expectations are formed using information from period $t-1$ and then period-t variables are determined given the expectations. We abstract away from this approach in this paper.

where $\tilde{X}_{t-1} = [1, X'_{t-1}]'$, $\Phi_t = [\boldsymbol{\alpha}_t, \boldsymbol{\beta}_t]$, and R_t is the second moments matrix of perceived autocovariances. γ denotes the constant gain value, which determines the weight that agents place on the latest available observations. When nominal rates are constrained by the ELB, the learning rule in (13) loses its stability. During ELB regimes, we interpret the share of these agents as a measure of central bank credibility: More agents that use the anchored rational expectations rule with shadow rate reflect more trust in the central bank's ability to circumvent the ELB constraint with unconventional monetary policy tools. A lower share weakens the transmission channel from shadow rate to inflation and output gap, thereby reflecting a lower central bank credibility. A sufficiently high share of adaptive learners at the ELB creates the risk of deflationary spirals, which is illustrated in further detail in Section 2.4.

2.3.3 Aggregate Dynamics

Given the RE-based (anchored) and learning-based (de-anchored) expectation formation rules, the fraction of agents using each rule evolves according to a fitness measure based on their one-step-ahead forecasting performance as in Busetti et al. (2017), Hommes and Lustenhouwer (2019), Goy, Hommes, and Mavromatis (2020), and Lansing (2021). In particular, we assume the following fitness measures ζ_t^{RE} and ζ_t^{L} associated with each rule:[12]

$$\begin{cases} \zeta_t^{RE} = -(1-\omega) FE_t^{RE} + \omega \zeta_{t-1}^{RE}, \\ \zeta_t^{L} = -(1-\omega) FE_t^{L} + \omega \zeta_{t-1}^{L}, \end{cases} \tag{14}$$

where FE_t^{RE} and FE_t^{L} denote the sum of squared forecast errors for inflation and output gap under for the RE- and learning-based PLMs, respectively. Given the fitness measures, agents' fractions are determined by

$$n_t^{RE} = \frac{exp(\chi \zeta_t^{RE})}{exp(\chi \zeta_t^{RE}) + exp(\chi \zeta_t^{L})}, \quad n_t^{L} = \frac{exp(\chi \zeta_t^{L})}{exp(\chi \zeta_t^{RE}) + exp(\chi \zeta_t^{L})}, \tag{15}$$

[12]The fitness measures follow the standard assumption in the heterogeneous expectations literature as in the aforementioned studies.

where n_t^{RE} (rational) and n_t^{L} (learning) denote the fractions of agents associated with each type. χ is an *intensity of choice* measure, common across both types, which determines the frequency of switching between the rules. Finally, the implied one-step-ahead and N-step-ahead inflation expectations are given by[13]

$$\begin{cases} E_t X_{t+1} = n_{t-1}^{RE} E_t X_{t+1}^{RE} + n_{t-1}^{L} E_t X_{t+1}^{L}, \\ E_t X_{t+N} = n_{t-1}^{RE} E_t X_{t+N}^{RE} + n_{t-1}^{L} E_t X_{t+N}^{L}. \end{cases} \tag{16}$$

The model dynamics evolve according to the aggregate law of motion in (3); rational and adaptive expectations in (4) and (12); monetary policy and shadow rate rules in (8) and (9); the learning rule in (13); the rule for updating agent fractions in (14)–(15); and finally the rule to determine aggregate expectations in (16).

2.4 Adaptive Learning and Instability at the ELB: Illustration

A well-known result in the adaptive learning literature is that, akin to the determinacy condition in RE models, the learning dynamics are expectationally stable (E-stable) when the Taylor principle $\phi_\pi > 1$ is satisfied (Bullard and Mitra 2002). During ELB regimes where monetary policy is constrained, the E-stability principle breaks down for standard model parameterizations, and learning dynamics become unstable.[14]

In our heterogeneous expectation setup, the presence of adaptive learners serves as a source of potential instability at the ELB. If the share of adaptive learners becomes sufficiently high, aggregate dynamics of the model become unstable. In such an environment, adverse shocks can push the economy into self-fulfilling deflationary spirals with ever-falling inflation and output gap.

To understand the intuition behind the instability, we illustrate the key mechanism at the ELB regime in a simplified setting in order

[13] Shares of agents n_t^{RE} and n_t^{L} enter into aggregate expectations with a one-period lag to obtain a sequential timing structure of expectations in the model. This is discussed in further detail in Appendix D.

[14] E-stability refers to the stability of constant gain learning algorithms. When the E-stability condition is not satisfied, learning dynamics are characterized by divergent behavior; see Evans and Honkapohja (2001) for further details.

to obtain analytical stability conditions. Consider the three-equation model in (1) without shocks and lagged state variables:

$$\boldsymbol{A}X_t = \boldsymbol{C}E_tX_{t+1}, \tag{17}$$

with $\boldsymbol{A} = \begin{bmatrix} 1+\frac{\phi_y}{\tau} & \frac{\phi}{\tau} \\ -\kappa & 1 \end{bmatrix}$ and $\boldsymbol{C} = \begin{bmatrix} 1 & \frac{1}{\tau} \\ 0 & \beta \end{bmatrix}$. Under RE, the agents' PLM takes the form of $E_tX_{t+1} = \boldsymbol{a}$. Plugging back into the law of motion and solving for the equilibium yields $\boldsymbol{a} = \boldsymbol{0}$ as the unique RE solution if the Taylor principle $\phi_\pi > 1$ is satisfied. Under adaptive learning, agents' PLM is time varying:

$$E_tX_{t+1} = \boldsymbol{\alpha}_{t-1}, \tag{18}$$

where the vector $\boldsymbol{\alpha}_{t-1}$ is updated every period as new observations become available. Assuming shares of adaptive learners $\boldsymbol{n^L_{t-1}} = [n^L_{\pi,t-1}, n^L_{y,t-1}]'$ for inflation and output gap, respectively, the implied actual law of motion (ALM) is given by[15]

$$\boldsymbol{A}X_t = \boldsymbol{C}[\boldsymbol{n^L_{t-1}}\boldsymbol{\alpha}_{t-1} + (\boldsymbol{1} - \boldsymbol{n^L_{t-1}})\boldsymbol{a}]. \tag{19}$$

The T-map associated with the adaptive agents' PLM is given by $\boldsymbol{\alpha}_{t-1} \Rightarrow T(\boldsymbol{\alpha}_{t-1}) = \boldsymbol{\Gamma_1 n^L_t}\boldsymbol{\alpha}_{t-1}$, with $\boldsymbol{\Gamma_1} = \boldsymbol{A}^{-1}\boldsymbol{C}$.[16] The law of motion is E-stable under learning if all eigenvalues of $\frac{\partial T(\boldsymbol{\alpha}_{t-1})}{\partial \boldsymbol{\alpha}_{t-1}} = \boldsymbol{\Gamma_1 n^L_{t-1}}$ have real parts less than 1.

In this simple setting, when the state variables X_t deviate from the deterministic equilibrium $X_t = \boldsymbol{0}$, the forecasting performance of adaptive learners outpaces those of rational agents. This is due to the adaptive learners adjusting their beliefs based on their previous errors, while rational agents' forecasts remain fixed at the equilibrium. Consequently, the share of adaptive learners increases and the central bank loses credibility whenever the state variables move away from the equilibrium. In turn, the lower credibility and the high

[15] In our empirical section, we use the restriction $n^L_{\pi,t} = n^L_{y,t}$. In this section, for illustrative purposes, we allow for different shares of adaptive learning agents for inflation and output gap $[n^L_{\pi,t}, n^L_{y,t}]$.

[16] T-map refers to the mapping from agents' PLM to the implied ALM of the model. See Evans and Honkapohja (2001) for a detailed treatment.

Figure 1. E-unstability Region in the Skeleton New Keynesian Model as a Function of the Share of Adaptive Learners

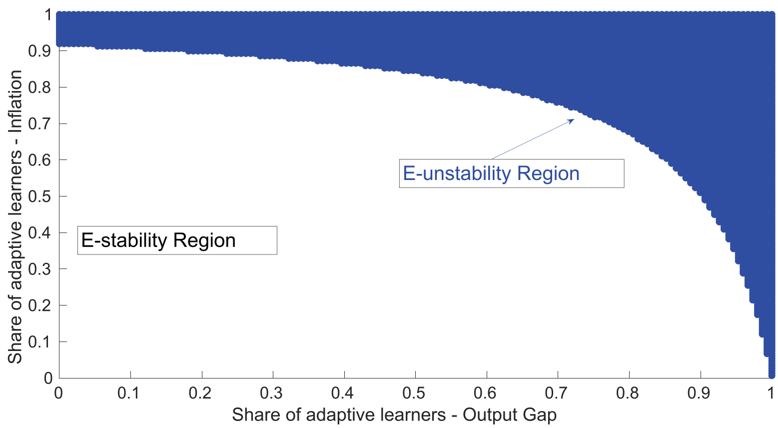

Note: The blue area shows the region where learning dynamics become unstable in the model.

share of adaptive learners slow down the economy's return to equilibrium.[17] As we will show in Section 3, the same principle applies to the empirical estimates of inflation expectations under RE and learning. While expectations under RE tend to be centered around the equilibrium, those of adaptive learners follow the data more closely and become de-anchored during periods when inflation persistently deviates from its trend.

The instability in the model arises when persistent deviations from the equilibrium coincide with the ELB constraint on nominal interest rates. In these scenarios, the rising share of adaptive learners is combined with the central bank's inability to combat the falling inflation and output gap. As a result, the economy is stuck in a self-fulfilling deflationary spiral at the ELB, which generates a de-anchoring of inflation expectations and a loss of central bank credibility.

Figure 1 shows the instability region in the model as a function of the share of adaptive learners on inflation and output

[17]In Appendix A, we derive the analytical relationship between agent shares and forecast errors for the simple law of motion (19).

gap.[18] As the proportion of adaptive learners increases significantly for either inflation or output gap expectations, the system becomes E-unstable. In Section 4, we delve into a more detailed discussion of the potential occurrence of these scenarios using our full-fledged model estimated for the United States.

3. Estimation

3.1 Methodology and Data

This section discusses the estimation methodology, along with the data set used in estimations and prior distributions for estimated parameters. The regime-switching model described in the previous section can be summarized as a state space system with time-varying matrices as follows:

$$S_t = \boldsymbol{\gamma}^{s_t}_{1,\Phi_{t-1}} + \boldsymbol{\gamma}^{s_t}_{2,\Phi_{t-1}} S_{t-1} + \boldsymbol{\gamma}^{s_t}_{3,\Phi_{t-1}} \varepsilon_t, \tag{20}$$

with $S_t = [X_t, \varepsilon_t]'$ and conformable matrices $\boldsymbol{\gamma}^{s_t}_{1,\Phi_{t-1}}$, $\boldsymbol{\gamma}^{s_t}_{2,\Phi_{t-1}}$, and $\boldsymbol{\gamma}^{s_t}_{3,\Phi_{t-1}}$ with two layers of time variation in the system matrices. The time-varying adaptive learning parameters $\boldsymbol{\alpha}_{t-1}$, $\boldsymbol{\beta}_{t-1}$ and shares of agents n^L_{t-1}, n^{RE}_{t-1} are captured by Φ_{t-1}. Monetary policy regime switches (ELB regime or Taylor rule) are captured by s_t. The timing assumptions of the expectations in the model admit a conditionally linear structure, where the likelihood is evaluated using the Kim and Nelson (1999, henceforth KN) filter. The parameters are estimated using standard Bayesian methods; see Appendix D for further details of the implementation.[19]

To estimate the model, we use historical U.S. data on output gap, inflation, and nominal interest rates over the period 1960:Q1–2019:Q4.[20] The output gap series is obtained by detrending GDP

[18]We use a standard parameterization with $\beta = 0.99$, $\tau = 1$, and $\kappa = 0.05$ for this illustration. The instability boundary in Figure 1 depends on the parameter values, but the main intuition is robust to alternative parameterizations.

[19]The working paper version of this study (Özden 2021) provides an alternative estimation method, where the heterogeneity of expectations is approximated as an endogenous regime-switching mechanism and the model is rewritten in a four-regime setup.

[20]We also use three years of data over 1957:Q1–1959:Q4 as a burn-in sample to initialize the likelihood.

using the method proposed in Hamilton (2018).[21] We further use short-term (one-quarter) and long-term (10-year) inflation expectations as observables in our estimation.

There is a wide variety of survey data on inflation expectations, with their availability ranging over different sample periods. In this paper, we utilize the ATSIX index introduced in Aruoba (2020). This is a composite index combining data from the Survey of Professional Forecasters (SPF) (Croushore 1993) and Blue Chip forecasts to obtain a reliable term structure of inflation expectations. The index has the advantage of avoiding the fixed-horizon versus fixed-event issues that are prevalent in many surveys, and also yields better forecasts of realized inflation than its alternatives as outlined in Aruoba (2020).[22] ATSIX series are available from 1992:Q1 for long-run expectations, and from 1998:Q1 for short-run expectations. For the earlier sample over 1960:Q1–1991:Q4, we treat inflation expectations as latent variables when estimating the model in order to test the model's predictions about these series during the Great Inflation period. To discuss model-implied expectation dynamics, we splice the ATSIX index with data from the SPF, which allows us to extend the expectation series back to 1979:Q1. We use the combined series to qualitatively examine model-implied inflation expectations over the early part of our sample. We use the following measurement equations in the estimation:

$$\begin{cases} y_t = y_t^{obs}, \\ \pi_t = \bar{\pi} + \pi_t^{obs}, \\ r_t = \bar{r}(s_t) + r_t^{obs}, \\ E_t \pi_{t+1} = \bar{\pi} + E_t \pi_{t+1}^{obs,ATSIX}, \\ E_t \pi_{t+40} = \bar{\pi} + E_t \pi_{t+40}^{obs,ATSIX}, \end{cases} \tag{21}$$

[21] Appendix F provides a sensitivity check around alternative measures of output gap, where we reestimate the model with a quadratically detrended output gap as in Cornea-Madeira, Hommes, and Massaro (2019), and output gap based on the Congressional Budget Office's (CBO's) measure of potential output (Shackleton 2018).

[22] In many surveys the forecasters are not consistently asked about their forecasts over a fixed horizon but rather over a fixed event, which can lead to an inconsistency about the timing assumptions. The ATSIX index does not suffer from this drawback; see Aruoba (2020) for further details.

where the right-hand-side variables are the historical data (observables), and the left-hand-side variables are the model variables. To include inflation expectations data in the estimation, we introduce two measurement error shocks. The law of motion for one-quarter and 10-year inflation expectations becomes[23]

$$\begin{cases} E_t\pi_{t+1} = n^{RE}_{t-1}E_t\pi^{RE}_{t+1} + n^{L}_{t-1}E_t\pi^{L}_{t+1} + \varepsilon^{exp,1}_{\pi,t}, \\ E_t\pi_{t+40} = n^{RE}_{t-1}E_t\pi^{RE}_{t+40} + n^{L}_{t-1}E_t\pi^{L}_{t+40} + \varepsilon^{exp,40}_{\pi,t}. \end{cases} \tag{22}$$

We assume a constant trend inflation $\bar{\pi}$ and a regime-switching constant trend interest rate $\bar{r}(s_t)$, which takes on values $\bar{r}_T$ and $\bar{r}_E$ as shown in (10). This approach closely follows that of Gust, Herbst, and Lopez-Salido (2018), who assume a shift in the intercept of interest rate $\bar{r}(s_t)$, which switches to a lower value during the ELB period.[24] We further impose the inflation trend $\bar{\pi}$ on measurement equations for inflation expectations. The constants are included in the measurement equations and are estimated along with the structural parameters, rather than detrending the data prior to estimation.

We estimate three additional models together with the heterogeneous expectation model: (i) the RE benchmark, without adaptive learners and with no regime switching in monetary policy; (ii) the RE model with regime switching in monetary policy; and (iii) a pure adaptive learning model with regime switching in monetary policy. Together, these models help us disentangle the marginal impact of adaptive learning, heterogeneous expectations, and monetary policy switching on model fitness.[25]

[23]Standard deviations of the measurement errors on inflation expectations are denoted by $\eta^{SR}_{\pi,exp}$ (short term) and $\eta^{LR}_{\pi,exp}$ (long term), respectively.

[24]The same intercept shift is also assumed for the shadow rate over the same period.

[25]The regime-switching RE model is approximated with a constant transition matrix $\boldsymbol{Q}$ when we estimate the model, to avoid the non-linearity induced by the expectational equations in (11). The heterogeneous expectation and adaptive learning models are instead estimated with the time-varying matrix $\boldsymbol{Q_t}$. Since expectations do not directly interact with the transition matrix in these models, their estimation still admits a conditionally linear structure that can be handled by the standard Kim and Nelson (1999) filter. Özden and Wouters (2021) show that the impact of a time-varying transition matrix has a negligible impact on

All structural, learning, and switching parameters are assigned prior distributions consistent with previous values used in the literature. This is discussed in detail in Appendix C, and Table C.1 provides a summary of all distributions used. The initial beliefs for heterogeneous expectations and adaptive learning models are derived from the estimated RE model, where we first estimate the baseline model in (1) under RE without regime switching. Using the estimated RE model, we retrieve the implied VAR(1) beliefs consistent with the estimated equilibrium, which are used to initialize the beliefs of adaptive learners. We use Sims's (1999) *csminwel* algorithm to obtain the posterior mode, which is used to initialize the MCMC algorithm with random-walk Metropolis-Hastings. We use 500,000 parameter draws for all models under consideration. The first 50 percent of the draws are discarded as a burn-in sample, and highest posterior density (HPD) intervals are computed using the remaining 50 percent of the sample.

3.2 Posterior Estimation Results

In this section we discuss the posterior estimation results for the heterogeneous expectation (HE) model along with the three accompanying models described in the previous section, i.e., (i) baseline rational expectations (RE), (ii) rational expectations with regime switching (RE-RS), and (iii) adaptive learning (AL).

The posterior moments of parameter distributions, together with the marginal likelihoods of all models, are reported in Table 1.[26] Based on the marginal likelihoods and Bayes factors, three key results emerge: First, all three models with regime switching in monetary policy fit the data better than the RE benchmark, regardless of the underlying expectation mechanism (i.e., rational, learning, or heterogeneous expectations). Second, both the HE and the AL model perform better than the RE-RS model. This suggests that the presence of adaptive learners improves the model fitness. Third, the HE model performs better than the AL model, which shows

estimation results, therefore the results with constant matrix $\boldsymbol{Q}$ and time-varying matrix $\boldsymbol{Q_t}$ are comparable.

[26]The (log-) marginal likelihood values reported in the table are based on the modified harmonic mean estimator. The Bayes factors are calculated using a log base 10, following Jeffreys Guidelines (Greenberg 2012).

Table 1. Posterior Distribution Moments for Baseline RE (No Switching), RE with Switching in MP, Heterogeneous Expectations, and Adaptive Learning Models

	RE			RE—Switching in MP			Hetero Expectations			Learning		
Parameter	**LB**	**Mean**	**UB**	**LB**	**Mean**	**UB**	**LB**	**Mean**	**UB**	**LB**	**Mean**	**UB**
$\overline{\pi}$ (Inflation Trend)	0.63	0.66	0.69	0.62	0.64	0.65	0.55	0.58	0.6	0.03	0.15	0.27
$\bar{r_T}$ (Int. Rate Trend—Taylor)	0.67	1.02	1.43	0.78	1	1.24	0.25	0.26	0.27	0.32	0.33	0.34
$\bar{r_E}$ (Int. Rate Trend—ELB)				0.01	0.01	0.01	0.03	0.03	0.04	0.029	0.032	0.035
κ (NKPC Slope)	0.0018	0.0033	0.0071	0.0011	0.0018	0.0028	0.0017	0.0032	0.0052	0.009	0.014	0.017
τ (Risk Aversion)	0.43	0.68	1.16	1.55	2.19	3.05	1.03	1.13	1.3	1.22	1.28	1.43
ι_y (Indexation—IS Curve)	0.15	0.22	0.3	0.09	0.15	0.21	0.03	0.06	0.12	0.63	0.69	0.74
ι_p (Indexation—NKPC)	0.03	0.07	0.15	0.02	0.05	0.08	0.09	0.15	0.2	0.36	0.44	0.52
ϕ_π (MP Inflation Reaction)	1.18	1.45	1.77	1.16	1.41	1.72	1.24	1.53	1.79	1.08	1.33	1.55
ϕ_y (MP Output Gap Reaction)	0.19	0.28	0.39	0.17	0.26	0.38	0.34	0.42	0.51	0.23	0.31	0.44
$\phi_{\Delta y}$ (MP Output Gap Growth Reaction)	0.11	0.15	0.19	0.06	0.09	0.11	0.09	0.12	0.14	0.03	0.05	0.06
ρ_r (MP Smoothing)	0.87	0.9	0.92	0.87	0.9	0.93	0.9	0.92	0.94	0.91	0.93	0.95
ρ_y (Persistence—Demand Shock)	0.96	0.97	0.98	0.88	0.92	0.95	0.93	0.95	0.96	0.22	0.24	0.28
ρ_π (Persistence—Supply Shock)	0.05	0.2	0.43	0.01	0.04	0.08	0.01	0.04	0.08	0.02	0.05	0.12
η_y (St. Dev.—Demand Shock)	0.2	0.31	0.48	0.13	0.17	0.22	0.28	0.32	0.36	0.91	1.01	1.11
η_π (St. Dev.—Supply Shock)	0.3	0.44	0.58	0.65	0.71	0.77	0.35	0.38	0.41	0.34	0.37	0.41
$\eta_{r,T}$ (St. Dev.—MP Shock Taylor)	0.22	0.26	0.31	0.2	0.22	0.24	0.22	0.24	0.27	0.19	0.2	0.22
$\eta_{r,E}$ (St. Dev.—MP Shock at ELB)				0.03	0.03	0.04	0.01	0.01	0.01	0.008	0.01	0.013
$\eta^{SR}_{\pi,exp}$ (St. Dev.—SR Inflation Exp.)	0.09	0.16	0.24	0.09	0.1	0.11	0.09	0.11	0.12	0.13	0.15	0.17
$\eta^{LR}_{\pi,exp}$ (St. Dev.—LR Inflation Exp.)	0.08	0.09	0.10	0.05	0.06	0.07	0.07	0.08	0.09	0.11	0.12	0.13
Φ_1 (MP Switching—Taylor to ELB)							0.06	0.1	0.15	0.11	0.2	0.39
Φ_2 (MP Switching—ELB to Taylor)							0.07	0.16	0.31	0.09	0.15	0.26
γ (Constant Gain)							0.0584	0.0585	0.0587	0.0575	0.0579	0.0581
ω (Memory)							0.49	0.62	0.74			
χ (Intensity of Choice)							0.38	0.51	0.66			
$1-q^T$ (Exog. Exit Probability—Taylor)			0.13	0.19	0.27							
$1-q^E$ (Exog. Exit Probability—ELB)			0.03	0.04	0.06							
Marg. (log-) Likl.		−479.63			−435.86			−410.06			−432.66	
Bayes Factor		1			19			29.86			20.39	

Note: The estimation period is from 1960:Q1 to 2019:Q4 using historical U.S. data.

that the expectational heterogeneity mechanism also improves the model fit. Taken together, these results suggest that both monetary policy switching and heterogeneity of expectations are important mechanisms to fit the data.[27]

Before analyzing the model-implied dynamics and inflation expectations, we discuss the differences in estimated parameter values. First, comparing the baseline and regime-switching RE models, it is readily seen that most parameters have similar posterior HPD intervals. There are two exceptions: First, the estimated slope of the Phillips curve κ is lower in the regime-switching model, implying a higher degree of price stickiness when the ELB constraint is accounted for. This is in line with the findings in Del Negro, Giannoni, and Schorfheide (2015), Lindé, Smets, and Wouters (2016), and Lindé, Maih, and Wouters (2017). Second, the risk-aversion parameter τ is considerably higher in the regime-switching model than in the baseline. This higher value is explained by the expectational feedback channel in the IS curve: When monetary policy is constrained by ELB, agents' expectations take into account the constraint in the regime-switching model. Therefore the ex ante real interest rate $r_t - E_t[\pi_{t+1}]$ has a larger feedback effect on output gap y_t in the IS equation once the ELB constrained is accounted for. The higher risk-aversion parameter in the regime-switching model has the effect of dampening this feedback channel.

Next we compare the HE and AL models with the regime-switching RE. The differences are more pronounced in this comparison: NKPC is steeper in the HE model, and it becomes even more steep in the AL model. The estimated NKPC slope is in line with previous findings in adaptive learning literature; e.g., Milani (2007), Jääskelä and McKibbon (2010), and Slobodyan and Wouters (2012b) all report lower Calvo parameters or steeper NKPC slope in their estimation results under learning compared with the RE benchmark. This result suggests that learning dynamics can partially substitute for nominal price stickiness. The risk-aversion parameter τ is lower in both the HE and the AL model compared with the RE-RS model, which relates to the expectational feedback channel discussed above:

[27] As a robustness check, in Appendix G we provide estimations of all models without inflation expectations. The relative ranking of the models remains the same.

Agents in the RE-RS model switch their expectations immediately once the economy becomes constrained by the ELB, which strengthens the feedback channel from ex ante real interest rate to output gap. In the HE and AL models, expectations adapt gradually over time as agents learn about the consequences of the ELB. Therefore the resulting risk aversion τ is lower than the RE-RS model but still higher than the baseline RE model.

The constant trend parameters for inflation and interest rate in measurement equations, $\bar{\pi}$ and $\bar{r_T}$, are lower in HE and AL models compared with RE and RE-MS. This is due to the time variation about the perceived mean $\boldsymbol{\alpha}_t$ in the HE and AL models.[28] While agents' expectations about the mean of inflation and output gap are zero in the RE and RE-RS models, the time-varying intercepts in the HE and AL models introduce a non-zero mean in their expectations. In other words, the HE and AL models have an endogenous inflation and interest rate trend induced by time-varying beliefs. This results in a level shift and lower estimates for the intercepts in measurement equations. The results are in line with, e.g., Carvalho et al. (2021), who interpret time-varying learning dynamics as a source of endogenous inflation trend.

The remaining structural parameter estimates are similar under all models, with HPD bands well within the range of each other. The posterior means for ϕ_π range over the interval [1.33, 1.53], whereas the output gap reaction ϕ_y and output gap growth reaction $\phi_{\Delta y}$ range over the intervals [0.26, 0.42] and [0.05, 0.15], respectively. The same argument also applies to interest rate smoothing ρ_r, which fluctuates between 0.90 and 0.95. All models except AL are characterized by a highly persistent demand shock (ρ_y ranging between [0.92, 0.96]) and a near-white-noise supply shock (ρ_π ranging between [0.04, 0.2]). This is accompanied by low indexation parameters in these models, with ι_y ranging over $[0.06, 0.22]$ and ι_π over $[0.05, 0.15]$. The AL model is instead characterized by a less persistent demand shock with $\rho_y = 0.24$, which is substituted by higher indexation parameters $\iota_y = 0.69$ and $\iota_\pi = 0.44$. Some studies in the past have suggested that learning dynamics in DSGE models

[28] The time variation in all PLM coefficients, $\boldsymbol{\alpha}_t$ and $\boldsymbol{\beta}_t$, is reported in Figure E.1, Appendix E.

can substitute for mechanical sources of persistence such as indexation, habits, capital adjustment costs, and persistence of structural shocks. Other studies have found learning dynamics have a negligible impact on these parameter estimates.[29] Hence the evidence in the literature on the impact of learning dynamics on mechanical sources of persistence is mixed and depends on the particular model setup. In our setting, learning and heterogeneity do not substitute for mechanical sources of persistence.

In the HE and AL models, the estimated constant gain values have similar posterior means with 0.0585 and 0.579, respectively. This implies that approximately 50 percent of adaptive learners' expectations are determined by three years of most recent data.[30] For the HE model the estimated memory parameter ω in expectational switching is 0.6, whereas the intensity of choice χ is 0.51. Our estimated constant gain value is somewhat higher than other studies in the literature that have only used aggregate macrodata in their estimation. Furthermore, our intensity of choice χ is significantly lower than other studies that have estimated similar mechanisms in the absence of inflation expectations, e.g., Cornea-Madeira, Hommes, and Massaro (2019). Therefore our findings suggest that using inflation expectations in the estimation is crucial for correctly identifying the parameters that determine the expectation formation process.[31]

For the remainder of this section, we discuss model-implied dynamics of the HE model to better understand whether and how it generates more realistic expectation formation dynamics. Figure 2 shows the 90 percent HPD band of the estimated shadow rate under

[29] For example, Milani (2007) documents that learning dynamics result in substantially lower degrees of habit and indexation. Examples of papers that do not find important differences in estimated RE and learning models include Jääskelä and McKibbon (2010) and Slobodyan and Wouters (2012a).

[30] This suggests a geometric discount rate of $(1-\gamma)^T$ for T periods in the past.

[31] As a robustness check, in Appendix G we report estimation results without using any inflation expectations data. This yields a lower gain coefficient and a significantly higher intensity of choice with a larger uncertainty band. This provides further support for the argument that having inflation expectations data as observables plays an important role in identification of learning- and heterogeneity-related parameters.

Figure 2. Estimated Shadow Rate Together with the Nominal Interest Rates for United States over the Period 1960:Q1–2019:Q4

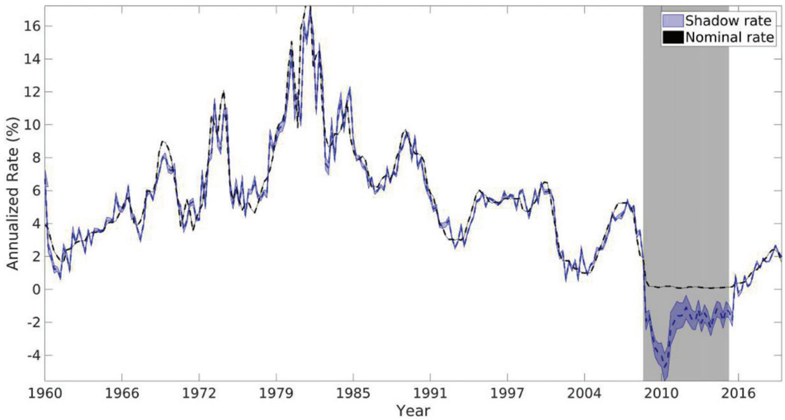

Note: The blue area depicts the 90 percent HPD band of the shadow rate estimate. The gray area depicts the ELB regime following the GFC.

the HE model, together with nominal interest rates over our estimation sample.[32] The estimated shadow rate is crucial in determining the expectations of rational agents in the model. It is readily seen that during the Taylor-rule regime before the Great Recession, the shadow rate closely follows the nominal interest rate path. As discussed in the previous section, this close relationship is by construction since the shadow rate is smoothed over the observed nominal rate during Taylor regime. The rates start diverging when the economy enters the ELB regime, and the shadow rate reaches a trough in 2010:Q2 with a range of [3.9, 5.56]. This is consistent with other studies in the literature, e.g., Kulish, Morley, and Robinson (2014), where the authors report an annual rate of –4 percent as the trough of their shadow rate estimate. The rate starts to gradually pick up after the initial crisis period, and the two rates converge again by the end of 2015 as nominal rates starts rising and the economy

[32]To obtain the HPD bands for the shadow rate and other latent variables in the model, we simulate the model 1,000 times using parameter values from the MCMC chain.

switches back to the Taylor-rule regime. The observed pattern in the shadow rate is also consistent with other empirical studies, e.g., Aruoba et al. (2022), who use a structural VAR with occasionally binding constraints to estimate the shadow rate.

Figure 3 shows survey expectations data together with model-implied short- and long-run inflation expectations.[33] The top two panels show the results with the heterogeneous expectations model: The blue areas depict the 90 percent HPD interval of short-run (panel A) and long-run (panel B) inflation expectations for the HE model. The figures include two types of survey data: ATSIX index (blue line) is used in the estimation of the model; it is available from 1992:Q1 for long-run (10-year) and 1998:Q1 for short-run (one-quarter) inflation expectations. Data from the SPF (red line) are not included in the estimation. We use this to splice the ATSIX index and examine the model-implied inflation expectations over the earlier part of the sample. It is readily seen that model-implied inflation expectations match the survey data fairly well. In particular, over the Great Inflation period, the model captures the de-anchoring of inflation expectations very well. Over the period where SPF data is available (i.e., 1979:Q1 onwards), model-implied series match the survey data closely for both short- and long-run inflation expectations, despite the fact that no inflation expectations data are included in the estimation over this period.

To see how well the HE model performs in terms of model-implied expectations, panels C–D in Figure 3 compare the median model-implied short- and long-run inflation expectations under the RE, AL, and HE models against survey data. This comparison helps us understand whether the improvement in model fitness for the HE model is accompanied by a better fit on inflation expectations. Two results become evident from these figures: First, the RE model falls short of explaining the survey data during the Great Inflation period, when inflation was high and inflation expectations were de-anchored. In particular, long-run inflation expectations under RE remain stable throughout the entire sample, regardless of the realized inflation. The AL and HE models are both more successful along

[33] By model-implied expectations, we refer to expectation series generated by the models in absence of measurement errors.

Figure 3. Short-Term and Long-Term Inflation Expectations over the Period 1960:Q1–2019:Q4

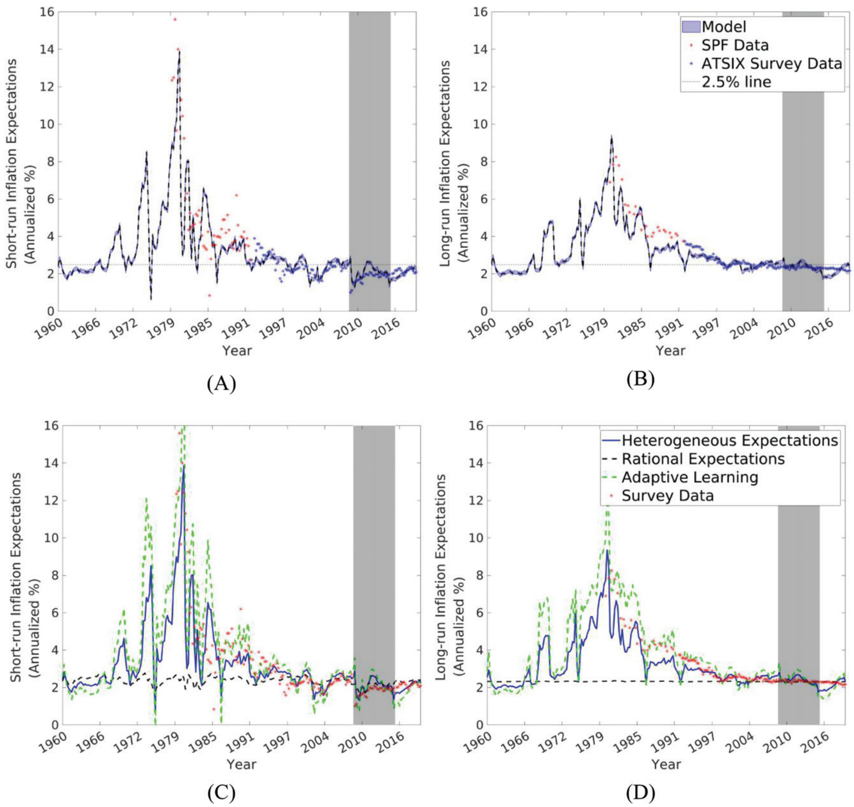

Note: Panels A and B show the 90 percent HPD bands of model-implied expectations (short- and long-run expectations, respectively) from the HE model, together with the ATSIX index, SPF data, and a constant 2.5 percent line. Panels C and D compare the posterior medians of model-implied expectations under the RE, HE, and AL models against survey data (ATSIX combined with SPF).

this dimension, and they both match periods with de-anchored inflation expectations fairly well. Second, the AL model typically has more trouble during periods with relatively stable inflation over the post–Great Moderation period. Model-implied data from the '80s onwards are typically too volatile, particularly for long-run inflation expectations. The HE model overcomes these two shortcomings by

Table 2. In-Sample RMSEs and Biases for Inflation Expectations in RE, HE, and AL Models

	Hetero. Exp.	Rational Exp.	AL
Long-Run Inflation Expectations			
RMSE	0.59	1.4	0.73
Bias	–0.21	–0.72	0.22
Short-Run Inflation Expectations			
RMSE	1.43	2.54	1.68
Bias	–0.19	–0.94	0.31

Note: The sample period for expectations covers 1979:Q1–2019:Q4, which consists of SPF data between 1979:Q1 and 1991:Q4, and ATSIX index between 1992:Q2 and 2019:Q4.

allowing the agents to endogenously switch between learning and rational expectations.

To make this point more clear, in Table 2 we report the in-sample root mean square errors (RMSEs) and biases of inflation expectation forecast errors for the RE, HE, and AL models. Not surprisingly, the RE model yields the worst statistics both for short-run and for long-run inflation expectations. On average, RE-implied expectations are negatively biased due to the models' inability to produce de-anchored expectations over the Great Inflation period. The AL model yields better RMSEs and biases than the RE benchmark, but it is still outperformed by the HE model. In particular, the AL model suffers from positive biases on average, suggesting that it tends to overshoot the degree of de-anchoring. These results confirm that having both types of expectations with endogenous shares is vital for explaining periods of both de-anchored and anchored inflation expectations.

Figure 4 shows the estimated share of adaptive learners over history together with inflation (panel A), and their correlation over the sample period (panel B).[34] To understand how the estimate of these shares are pinned down, recall from Section 2.4 that there is

[34] To calculate the correlation series, we use a rolling sample of 10 years, starting with the sample over 1960:Q1–1969:Q4.

Figure 4. Estimated Share of Adaptive Learners over History

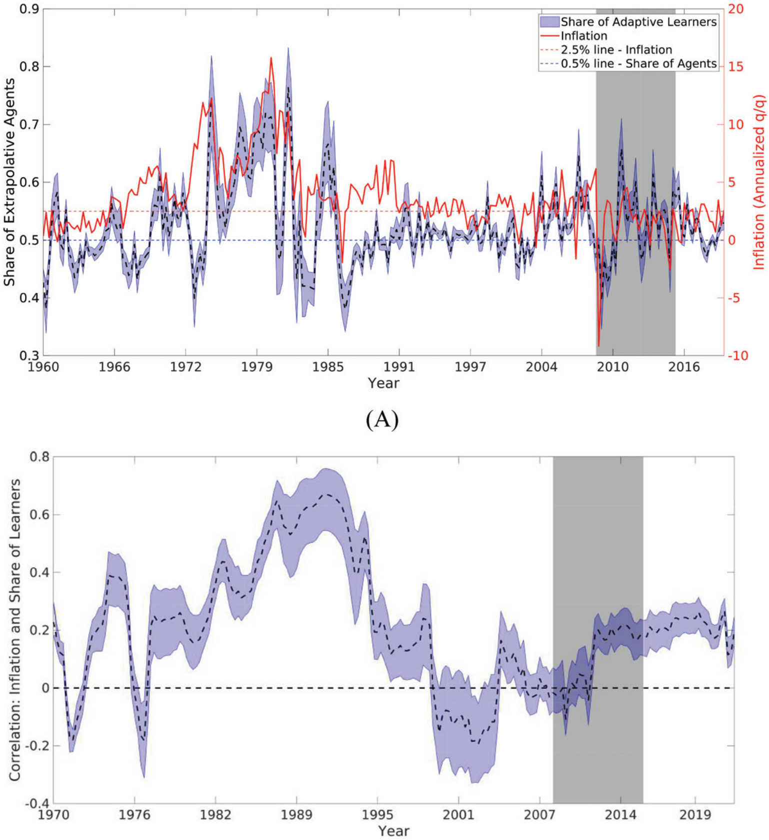

(A)

(B)

Note: Panel A shows the estimated 90 percent HPD band of the share of adaptive learners (left y-axis) and CPI inflation (right y-axis) over the estimation period 1960:Q1–2019:Q4. Panel B shows the estimated 90 percent HPD band of rolling-window correlation between the share of adaptive learners and CPI inflation. The correlations are based on a sample size of 10-year rolling window.

a tight relationship between agents' forecast errors and their shares. Inflation expectations under the RE benchmark tend to gravitate toward the inflation target, whereas those of adaptive learners follow realized inflation more closely, as is readily seen in Figure 3.

Consequently, periods where inflation is persistently above the trend, such as during the Great Inflation period, are typically dominated by a high share of adaptive learners. By incorporating data on both inflation and inflation expectations in the estimation, the model is able to identify the parameters governing the speed of learning γ, the intensity of choice χ, and the degree of memory in switching ω. These parameters collectively determine persistence and volatility of the estimated share of adaptive learners based on agents' realized forecast errors.

During the initial part of the sample, the co-movement between the share of adaptive learners and inflation is remarkably high, with a correlation of up to 0.78 during the '80s. The model explains the high inflation and de-anchored inflation expectations over this period with a high share of adaptive learners. As adaptive learners extrapolate recent data to form their expectations, this has a tendency to put further upward pressure on inflation. When inflation starts to stabilize after the 1990s, the tight positive correlation between inflation and share of adaptive learners breaks down. During the early 2000s, we observe a reversal in the correlation, which becomes weakly negative until 2008. A higher share of adaptive learners creates a weak deflationary effect over this period with stable inflation, before changing signs again following the GFC. At the beginning of GFC the share of adaptive learners temporarily falls down, which is partly how the model explains the missing deflation puzzle. During this period adaptive learners expect a stronger deflation that is not observed in the data, and the model explains this as a temporary fall in their share. Throughout the rest of the Great Recession the shares remain balanced around 50 percent, which suggests that expectations have at least partially remained anchored and responded to the shadow rate over this period. This is in line with other empirical studies in the literature, e.g., Mavroeidis (2021), who suggest that inflation and output gap have partially responded to shadow rate over the post-GFC period. This result can be interpreted as a successful central bank communication by the Federal Reserve Board over this period.

From a narrative perspective, model dynamics under heterogeneous expectations suggest that endogenous central bank credibility plays an important role in driving inflation. During the Great Inflation period, the model shows that the share of adaptive learners is high and central bank credibility is low. As the central bank brings inflation under control, the share of adaptive learners stabilizes around 50 percent from the '90s onwards. These are are in line with previous studies on the subject. For example, Carvalho et al. (2021) analyze a model where agents switch between a constant gain and a decreasing gain learning rule. They find that constant gain learning was dominant during the early '70s and '80s, whereas decreasing gain learning has become more prevalent from the '90s onwards, which is consistent with our results.

Along similar lines, Malmendier and Nagel (2016) provide a demographic interpretation in a model where agents overweight inflation experienced during their lifetimes. In this context, the authors document a divergence in expectations between younger and older cohorts during the late '70s and '80s. Younger individuals' experience with high inflation over this period contributes to a high perceived inflation persistence, which in turn creates more persistence and sluggishness in inflation expectations. This demographic dispersion in inflation expectation only goes away in the '90s. In our framework, this is reflected as a declining share of adaptive learners. The authors show that learning from experience can be seen as a microfoundation of constant gain learning models, since aggregate dynamics from the model can be approximated quite closely with a constant gain mechanism. As such, the success of our heterogeneous expectation model in explaining survey data, as well as the dominance of constant gain learners over the Great Inflation period, can be interpreted as a validation of their learning-from-experience framework in a DSGE setup. Lower average inflation over the Great Recession period, combined with the ELB constraint, creates a risk of de-anchored inflation expectations in the negative direction. We study this channel in further detail in the next section.

4. Model Dynamics at the ELB

The estimation results in the previous section highlight that the HE model fits the data better, and the heterogeneity mechanism is crucial in explaining the historical inflation dynamics. As shown in Section 2.4, the expectational switching mechanism creates the possibility of observing deflationary spirals at the ELB when the share of adaptive learners becomes too high. With this in mind, in this section we focus on the ELB regime over the post-GFC period and investigate the properties of ELB and deflationary spiral episodes in the heterogeneous expectation model. The discussion is focused on two key questions: (i) What is the risk of a deflationary spiral occurring in the model? (ii) How do ELB regimes and deflationary spirals interact with the heterogeneity mechanism and endogenous central bank credibility?

We use two exercises to analyze these issues. In the first exercise, we use U.S. data in 2008:Q4 as a starting point and generate density forecasts between 2009:Q1 and 2016:Q4 at the estimated posterior mean values. This is helpful to understand the estimated risk of de-anchoring and deflationary spirals, and how this risk interacts with key parameters that determine the learning and switching mechanism. In the second exercise, we use standard stochastic simulations of the model to discuss key moments and statistics at ELB episodes. This is useful to discuss the model-implied unconditional distributions of ELB and deflationary spiral episodes. In both exercises, we formally define a deflationary spiral as an episode where quarter-on-quarter inflation falls below 10 percent.[35]

Starting with the density forecasts of the model, Figure 5 shows an example of a deflationary spiral in the heterogeneous expectation model. Following the GFC period, the share of adaptive learners remains above the estimated baseline for an extended period from 2010:Q1 onwards and deflationary pressures keep building up. The shadow rate becomes increasingly more accommodative. Due to falling inflation and the ELB constraint on nominal rates, real interest rates rise and depress aggregate demand. As agents lose their trust in the central bank's ability to make up for this increasingly large slack in nominal interest rates, more agents switch to the adaptive learning rule. When the share of adaptive learners becomes critically high, the economy enters into a deflationary spiral episode with ever-falling inflation and output gap. This is an illustration

[35]We use 1,000 simulations in both exercises to compute the HPD bands. The shocks are drawn from normal distributions using the estimated standard deviations at the posterior mean in Table 1. For stochastic simulations we use a maximum simulation length of 5,000 periods, and simulations are terminated in both exercises when a deflationary spiral is detected. Further note that the parameters are fixed at the posterior mean in both exercises, therefore the confidence bands reported in this section do not reflect any parameter uncertainty.

Figure 5. Example of Deflationary Spiral Occurring over the Period 2009:Q1–2016:Q4

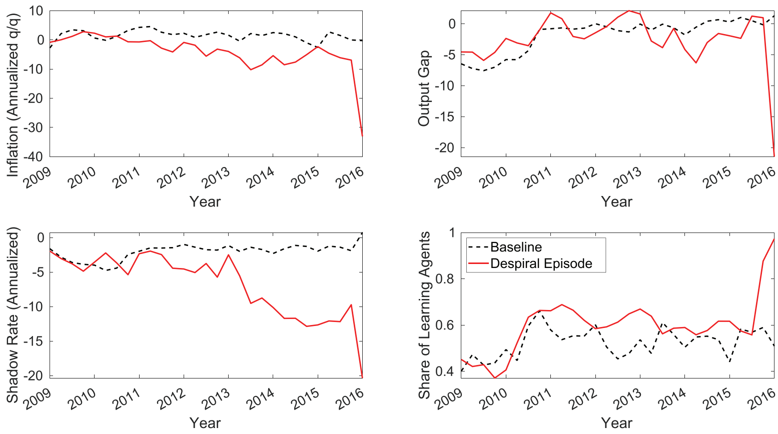

of the analytical results discussed in Section 2.4. While the stability conditions are no longer tractable in the full model, the density forecasts and stochastic simulations of the model show that the main intuition continues to hold in a more empirically relevant setup.

To see how often and under what conditions these deflationary spirals occur in the model, Figure 6 shows the 90 percent HPD interval of density forecasts from 2009:Q1 onwards. We divide the simulations into two categories when reporting the confidence bands: episodes that result in a deflationary spiral (gray area), and all other ELB episodes that do not result in deflationary spirals (blue area). It is readily seen that despiral episodes are characterized by a large downside risk on not only inflation but also output gap and shadow rate. More importantly, despiral episodes are characterized by a higher average share of adaptive learners, i.e., lower central bank credibility. It is also worth noting that the median forecasts under non-spiral ELB episodes are close to realized inflation, output gap, and the estimated values of shadow rate and share of adaptive learners. Baseline results for all variables fall within the range of 90 percent HPD interval over the forecast horizon. This suggests that unconventional monetary policy actions over this period have kept the share of adaptive learners low enough to make a switch to a deflationary spiral episode unlikely.

Figure 6. Density Forecasts with Heterogeneous Expectations Model over the Period 2009:Q1–2016:Q4

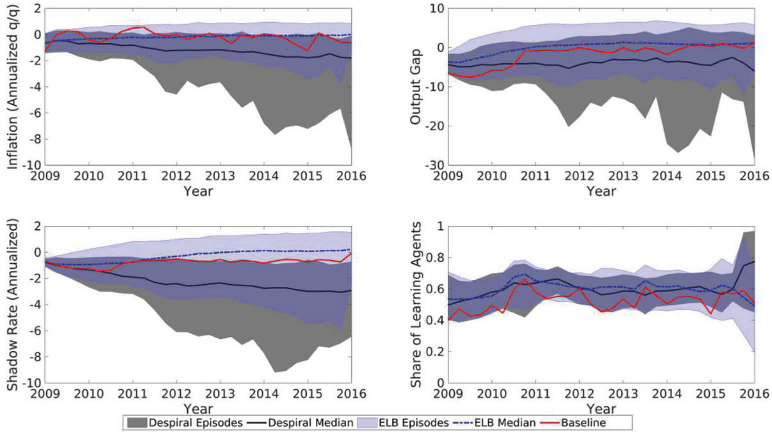

Note: The 90 percent HPD bands are reported for simulations that result in a deflationary spiral (gray area) and standard ELB episodes that do not result in deflationary spirals (blue area). The purple area corresponds to the region where the HPD bands overlap.

To make the connection between the share of adaptive learners and deflationary outcomes more concrete, Figure 7 shows the distribution of inflation against the share of adaptive learners over the counterfactual period. The unconditional distributions (i.e., both despiral and non-spiral episodes), depicted by gray dots, are characterized by a weakly negative correlation between inflation and share of adaptive learners (−0.17). In these simulations, inflation does not fall strongly into the negative territory and therefore an increase in the share of adaptive learners only creates a weak deflationary effect. On the contrary, deflationary spiral episodes, depicted by red dots, are associated with not only a higher share of adaptive learners and lower central bank credibility but also a stronger negative correlation between inflation and the share of adaptive learners (−0.44). When credibility is low to begin with (i.e., share of learners is high), a further decline in credibility tends to create a stronger deflationary effect than a starting point with high credibility and low share of adaptive learners.

Figure 7. Distribution of Inflation (x-axis) Plotted Against the Distribution of the Share of Adaptive Learners (y-axis)

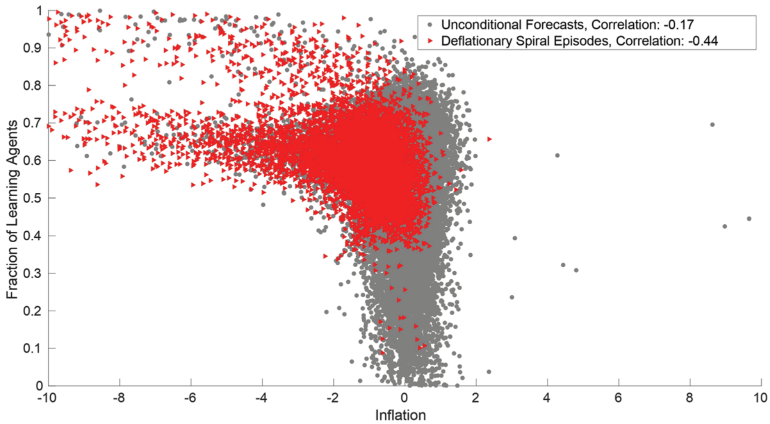

Note: We report the distributions for all ELB episodes (gray area) and deflationary spirals (red area) separately.

Table 3. Probability of Despiral Episodes under Alternative Parameterizations of the Model

Scenario	Deflationary Spiral Probability
Baseline	27.6%
Large Gain ($\gamma = 0.1$)	38.8%
High Intensity of Choice ($\chi = 2$)	46.6%

Table 3 shows the estimated probability of a despiral episode underlying the density forecasts in Figures 6 and 7. At the estimated parameter values, 27.6 percent of the simulations result in deflationary spirals. This is accompanied by two additional parameterizations of the model: If we increase the constant gain parameter γ to 0.1 from its estimate of 0.0585, i.e., when adaptive learners pay more attention to recent data, then the probability of deflationary spirals increases to 38.8 percent. If we change the intensity of choice parameter χ to 2 from its estimate of 0.51, i.e., agents switch more frequently between rational expectations and adaptive learning, then the probability increases further to 46.6 percent. Both counterfactuals represent scenarios where expectations can become de-anchored more easily, and as a result they both result in more frequent deflationary spirals.

Our second exercise is based on unconditional stochastic simulations of the model as discussed above. This helps us examine model properties and key statistics associated with ELB regimes, and it serves as a robustness check to see if the results based on U.S. data continue to hold in an unconditional environment. The simulations are mainly characterized by short-lived ELB episodes, with occasional long-lived ELB episodes: The probability that an ELB episode lasts for at least one, two, and five years are 27.5 percent, 12 percent, and 1.5 percent, respectively. The corresponding distributions of ELB episodes, together with other related summary statistics such as the frequency of ELB episodes and how long it takes to encounter deflationary spirals, are discussed in further detail in Appendix E. Here we instead focus on the averages of key variables at the ELB and despiral episodes, reported in Table 4: the top two columns in the table show the averages across all ELB and despiral episodes. Similar to density forecasts, despiral episodes are on average characterized by lower inflation, output gap, shadow rate, and both short-run and long-run expectations, together with a substantially higher average share of adaptive learners (0.8) compared with non-spiral ELB regimes (0.54). The bottom two columns in the table show the average *entry* values into ELB and despiral episodes. When the economy enters into an ELB regime with an already high share of adaptive learners, i.e., low central bank credibility, then the regime is more likely to turn into a deflationary spiral. Taken together, these results confirm the takeaways from U.S.-based density forecasts in an unconditional setting.

Our results in this section show that the share of adaptive learners and initial beliefs in ELB regimes play an important role in driving deflationary spirals in the model. It is important to highlight the difference between deflationary spirals in our endogenous central bank credibility setup and those that have been studied in a fully rational setup. Most notably, Bianchi, Melosi, and Rottner (2021) study deflationary spirals in a rational expectations model,

Table 4. Averages and Average Entry Values into ELB and Despiral Regimes

	Average—ELB	**Average—Defl. Spiral**
Inflation	–0.21	–1.26
Inf. Exp.—SR	–0.15	–1.16
Inf. Exp.—LR	–0.03	–0.29
Shadow Rate	–1.56	–7.16
Output Gap	–5.04	–29.52
Output Gap Exp.	–4.46	–27.68
Fraction of Learners	0.54	0.8
	Average Entry—ELB	**Average Entry—Defl. Spiral**
Inflation	–0.12	–0.44
Inf. Exp.—SR	–0.1	–0.36
Inf. Exp.—LR	–0.02	–0.07
Shadow Rate	–1.31	–2.43
Output Gap	–3.84	–9.19
Output Gap Exp.	–3.33	–8.01
Fraction of Learners	0.54	0.63

Note: The results are based on 1,000 stochastic simulations of the model.

where agents' expectations about future ELB regimes may lead to a deflationary bias. When the possibility of hitting the ELB regime becomes too large, the deflationary bias increases. For sufficiently large values of the bias, the equilibrium loses its determinacy and deflationary spirals occur. In this fully rational environment, the central bank can mitigate the risk of deflationary spirals by implementing an asymmetric monetary policy rule, whereby its response to inflation above target is slower than its response to inflation below target. This emphasizes the channel of monetary policy rule in mitigating the risk of deflationary spirals, where an asymmetric rule reduces the risk of encountering ELB episodes. Our model instead emphasizes the central bank credibility channel: Deflationary spirals occur when agents lose their trust in the central bank's ability to circumvent the ELB constraint through unconventional monetary policy measures. Therefore the risk of deflationary spirals can be mitigated by managing expectations at the ELB through central bank communication channels.

5. Conclusions

In this paper we estimate a heterogeneous expectation model based on the canonical New Keynesian model, with monetary policy subject to the ELB constraint on nominal interest rates. We use aggregate macrodata as well as survey data on inflation expectations to identify the learning and heterogeneous switching mechanisms in the model. Several results stand out. The heterogeneous expectation model fits the data better than models with fully rational agents or with agents using only adaptive learning. The results suggest that private-sector inflation expectations in the United States over the sample period 1960:Q1–2019:Q4 can be described as a mixture of anchored, rational expectations and de-anchored expectations based on adaptive learning. The latter plays a particularly important role during high inflation periods with de-anchored expectations, such as the Great Inflation period. The model also shows that during the U.S. experience with ELB after the GFC, expectations have remained partially anchored and responded to the shadow rate. Third and most importantly, our counterfactual experiments show that a high degree of de-anchoring and a loss of central bank credibility are associated with an increased likelihood of deflationary spirals and prolonged recessions. This emphasizes the importance of central bank communication channels in managing expectations and mitigating deflationary spiral risk. The paper also opens potential avenues of future research. The current framework only incorporates unconventional monetary policy through its expectational channel. Future studies should also account for the direct effects of unconventional tools, in particular quantitative easing measures. Moreover, the heterogeneous expectation and endogenous central bank credibility framework laid out in this paper is likely to have important insights into the liftoff from the ELB, and for the post-pandemic inflationary environment that many central banks in advanced economies have been experiencing.

Appendix A. Forecast Errors and Shares of Agents

In this section we use the simple deterministic version of the three-equation model described in Section 2.4 to derive an analytical relationship between agents' forecast errors and their shares when the model deviates from equilibrium. Recall that the economy's law of motion is given by

$$\boldsymbol{A}X_t = \boldsymbol{C}[\boldsymbol{n^L_{t-1}}\boldsymbol{\alpha}_{t-1} + (1 - \boldsymbol{n^L_{t-1}})\boldsymbol{a}],$$

with $\boldsymbol{a} = \mathbf{0}$ in equilibrium. We first rewrite the adaptive learning rule given in (13), which can be simplified in the absence of stochastic shocks and lagged state variables. Given that agents are only learning about the intercepts in this case, we have $\tilde{X}_{t-1} = \boldsymbol{c}$ and $\Phi_t = \boldsymbol{\alpha}_t$, where $\boldsymbol{c}$ is a vector of constants. Without loss of generality, set $\boldsymbol{c} = 1$ and rewrite the perceived volatility term R_t as follows:

$$R_t = R_{t-1} + \gamma(1 - R_{t-1}) = \gamma \sum_{j=0}^{t-1} (1 - \gamma)^j + R_0,$$

for some initial value R_0. As $t \to \infty$, we get $R_t = \gamma \sum_{j=0}^{\infty}(1-\gamma)^j = 1$. Using this, the equation for $\boldsymbol{\alpha}_t$ can be simplified as $t \to \infty$:

$$\boldsymbol{\alpha}_t = \boldsymbol{\alpha}_{t-1} + \gamma R_t^{-1}(X_t - \boldsymbol{\alpha}_{t-1}) = \gamma \sum_{j=0}^{\infty} (1 - \gamma)^j X_{t-j}.$$

Then it follows that $\frac{\partial \boldsymbol{\alpha}_t}{\partial X_{t-j}} = \gamma(1 - \gamma)^j > 0$ for any constant gain value $\gamma > 0$. In other words, whenever X_t deviates from the equilibrium, agents revise their beliefs about $\boldsymbol{\alpha}_t$ in the same direction as X_t.

Given the forecasting rules, we rewrite agents' shares in terms of their forecast errors. Note that rational and adaptive agents' squared forecast error vectors are given by X_t^2 and $(X_t - \boldsymbol{\alpha}_{t-1})^2$, respectively. Setting $\chi = 1$ and $\omega = 0$ in the switching function without loss of generality, Equations (14) and (15) together reduce to

$$\boldsymbol{n_t^{RE}} = \frac{exp(-X_t^2)}{exp(-X_t^2) + exp(-(X_t - \boldsymbol{\alpha}_{t-1})^2)},$$
$$\boldsymbol{n_t^{L}} = \frac{exp(-(X_t - \boldsymbol{\alpha}_{t-1})^2)}{exp(-X_t^2) + exp(-(X_t - \boldsymbol{\alpha}_{t-1})^2)},$$

with $\boldsymbol{n_t^{RE}} = [n_{\pi,t}^{RE}, n_{y,t}^{RE}]'$ and $\boldsymbol{n_t^{L}} = [n_{\pi,t}^{L}, n_{y,t}^{L}]'$. Given the fact that $\frac{\partial \boldsymbol{\alpha}_t}{\partial X_{t-j}} > 0$, we have $exp(-X_t^2) < exp(-(X_t - \boldsymbol{\alpha}_{t-1})^2)$. This implies $\boldsymbol{n_t^{L}} > \boldsymbol{n_t^{RE}}$ whenever X_t deviates from the equilibrium ($X_t^2 > 0$). Therefore any deviations from the equilibrium are met with a rising share of adaptive learners until the economy converges back to the equilibrium.

Appendix B. Data Descriptions

This section describes the quarterly time series used in the estimations. The data set spans from 1960:Q1 to 2019:Q4 and all time series except inflation expectations are retrieved from the Federal Reserve Economic Data (FRED) database.

- Real Gross Domestic Product (FRED mnemonic: GDPC1), denoted as GDP_t and available at https://fred.stlouisfed.org/series/GDPC1.
- Consumer Price Index for All Urban Consumers (FRED mnemonic: CPIAUCSL), denoted as P_t and available at https://fred.stlouisfed.org/series/CPIAUCSL.
- Effective Federal Funds Rate (FRED mnemonic: FEDFUNDS), denoted as R_t and available at https://fred.stlouisfed.org/series/FEDFUNDS.
- CBO's Measure of Real Potential GDP (FRED mnemonic: GDPPOT), denoted as GDP_t^{pot} and available at https://fred.stlouisfed.org/series/GDPPOT.
- Aruoba Term Structure of Inflation Expectations, denoted as $ATSIX_t$ and available at https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/atsix. We use the notation $ATSIX_t^{t+j}$ to refer to the measure of j-quarter-ahead forecasts made at period t.

The following variables are used in the measurement equations:

- Output gap y_t^{obs} is based on the cycle component of the Hamilton filter, applied to $log(GDP_t)$ over the estimation sample. For the CBO-based measure of output gap, which is used as a robustness check in Appendix F, output gap is computed as $y_t^{obs} = log(GDP_t) - log(GDPPOT_t)$.
- Inflation $\pi_t^{obs} = \frac{P_t}{P_{t-1}}$.
- Nominal interest rate $r_t^{obs} = R_t$.
- Short-term (one-quarter-ahead) inflation expectations $E_t\pi_{t+1}^{obs,ATSIX} = ATSIX_t^{j+1}$.
- Long-term (10-year-ahead) inflation expectations $E_t\pi_{t+40}^{obs,ATSIX} = ATSIX_t^{j+40}$.

Appendix C. Prior Distributions

This section discusses the prior distributions of all structural parameters used in the estimation. Table C.1 provides a summary all parameter distributions.

The risk-aversion parameter τ has a gamma distribution with a mean 2 and standard deviation 0.5 as in An and Schorfheide (2007). The monetary policy reaction coefficients are all based on the Smets-Wouters (2007) model. Accordingly, inflation reaction ϕ_π is assigned a gamma distribution with mean 1.5 and standard deviation 0.25; output gap reaction coefficients ϕ_y and $\phi_{\Delta y}$ are assigned gamma distributions with mean 0.25 and standard deviation 0.1. The interest rate smoothing parameter ρ_r is assigned a beta distribution with mean 0.75 and standard deviation 0.1. Similarly, shock parameters are based on the same model, where shock persistence parameters ρ_y and ρ_π are assigned a beta distribution with mean 0.5 and standard deviation 0.2, and shock standard deviations are assigned inverted gamma distributions with mean 0.1 and standard deviation 2. The standard deviation of the monetary policy shock over the ELB regime is an exception, which is instead assigned a uniform distribution over the unit interval. For the slope of the Phillips curve κ, we use a relatively tight prior of a beta distribution with mean 0.05 and standard deviation 0.025.

Table C.1. Prior Distributions for the Estimated Parameters in the New Keynesian Model

Parameter	Dist.	Prior Mean	Prior St. Dev.	Lower B.	Upper B.
$\overline{\pi}$ (Inflation Trend)	Uniform	0.5	0.29	0	∞
$\bar{r_T}$ (Int. Rate Trend—Taylor)	Uniform	0.5	0.29	0	∞
$\bar{r_E}$ (Int. Rate Trend—ELB)	Normal	0.1	0.25	0	∞
κ (NKPC Slope)	Beta	0.05	0.025	0	1
τ (Risk Aversion)	Gamma	2	0.5	0	∞
ι_y (Indexation—IS Curve)	Beta	0.25	0.1	0	1
ι_π (Indexation—NKPC)	Beta	0.25	0.1	0	1
ϕ_π (MP Inflation Reaction)	Gamma	1.5	0.25	1	∞
ϕ_y (MP Output Gap Reaction)	Gamma	0.25	0.1	0	∞
$\phi_{\Delta y}$ (MP Output Gap Growth Reaction)	Gamma	0.25	0.1	0	∞
ρ_y (Persistence—Demand Shock)	Beta	0.5	0.2	0	1
ρ_π (Persistence—Supply Shock)	Beta	0.5	0.2	0	1
ρ_r (MP Smoothing)	Beta	0.5	0.2	0	1
η_y (St. Dev.—Demand Shock)	Inv. Gamma	0.1	2	0	∞
η_π (St. Dev.—Supply Shock)	Inv. Gamma	0.1	2	0	∞
$\eta_{r,T}$ (St. Dev.—MP Shock at Taylor)	Inv. Gamma	0.1	2	0	∞
$\eta_{r,E}$ (St. Dev.—MP Shock at ELB)	Uniform	0.5	0.29	0	∞
$\eta^{SR}_{\pi,exp}$ (St. Dev.—SR Inflation Exp.)	Inv. Gamma	0.1	2	0	∞
$\eta^{LR}_{\pi,exp}$ (St. Dev.—LR Inflation Exp.)	Inv. Gamma	0.1	2	0	∞
Φ_1 (MP Switching—Taylor to ELB)	Gamma	0.2	0.1	0	∞
Φ_2 (MP Switching—ELB to Taylor)	Gamma	0.2	0.1	0	∞
γ (Constant Gain)	Gamma	0.035	0.015	0	1
ω (Memory)	Beta	0.5	0.2	0	1
χ (Intensity of Choice)	Gamma	5	2	0	∞
$1-q^T$ (Exog. Exit Probability—Taylor)	Uniform	0.5	0.29	0	1
$1-q^E$ (Exog. Exit Probability—ELB)	Uniform	0.5	0.29	0	1

This corresponds to a lower mean and standard deviation compared with previous studies; e.g., An and Schorfheide (2007) use a wider beta distribution with mean 0.3 and standard deviation 0.15. Nevertheless, the prior used here encompasses parameter values consistent with most empirical studies as its credible interval. The indexation parameters ι_y and ι_π are assigned beta distributions with mean 0.25 and standard deviation 0.1. The constant trend parameters in the measurement equations are assigned uniform distributions over the interval [0.2], except for the output gap trend, which is fixed at 0 and is not included in the estimation. The constant trend for interest rates during the ELB period, $\bar{r_E}$, is assigned a more informative normal prior with a mean of 0.1 and standard deviation 0.25 in order to restrict the range of parameter values over this period. For the constant transition probabilities in the RE model, $1 - q^T$ and $1 - q^E$, we assign uniform priors over the unit interval.[36] These parameters correspond to the exit probabilities from Taylor and ELB regimes, respectively. For the endogenous switching models, the parameters θ_1 and θ_2 in the monetary policy switching functions are fixed at 1. For the other two parameters on monetary policy switching, we assign gamma distributions with mean 0.2 and standard deviation 0.1 on $\frac{\Phi_1}{1000}$ and $\frac{\Phi_2}{1000}$, which covers both gradual and abrupt transitions for monetary policy regime switching. The persistence of expectational switching, ω, is assigned the same distribution as the shock persistence parameters, i.e., a beta distribution with mean 0.5 and standard deviation 0.2. The intensity of choice χ is assigned a gamma distribution with mean 5 and standard deviation 2, which is based on the findings of Cornea-Madeira, Hommes, and Massaro (2019) on inflation expectations. Finally, the constant gain parameter γ is assigned a gamma distribution with mean 0.035 and standard deviation 0.015, which is based on Slobodyan and Wouters (2012b).

[36]This differs from previous studies that assume tighter beta distributions, e.g., Chen (2017) and Lindé, Maih, and Wouters (2017).

Appendix D. System of Equations and Timing Assumptions

The full system of equations characterizing the heterogeneous expectation model is as follows:

$$\begin{cases} \underline{\text{Law of motion:}} \\ \boldsymbol{A}(s_t)X_t = \boldsymbol{B}(s_t)X_{t-1} + \boldsymbol{C}(s_t)E_tX_{t+1} + \boldsymbol{D}(s_t)u_t, \\ u_t = \boldsymbol{\rho}u_{t-1} + \varepsilon_t, \\ \underline{\text{Expectations:}} \\ E_tX_{t+1} = n_{t-1}^{RE}E_tX_{t+1}^{RE} + n_{t-1}^{L}E_tX_{t+1}^{L}, \\ E_tX_{t+1}^{RE} = \boldsymbol{b}X_t + \boldsymbol{d\rho}u_t, \\ E_tX_{t+1}^{L} = \boldsymbol{\alpha}_{t-1} + \boldsymbol{\beta}_{t-1}X_t, \\ \underline{\text{Agent shares, fitness and forecast errrors:}} \\ n_t^{RE} = \frac{exp(\chi\zeta_t^{RE})}{exp(\chi\zeta_t^{RE})+exp(\chi\zeta_t^{L})}, \\ n_t^{L} = \frac{exp(\chi\zeta_t^{L})}{exp(\chi\zeta_t^{RE})+exp(\chi\zeta_t^{L})}, \\ \zeta_t^{RE} = -(1-\omega)FE_t^{RE} + \omega\zeta_{t-1}^{RE}, \\ \zeta_t^{L} = -(1-\omega)FE_t^{L} + \omega\zeta_{t-1}^{L}, \\ FE_t^{RE} = (X_t - E_{t-1}X_t^{RE})^2, \\ FE_t^{L} = (X_t - E_{t-1}X_t^{L})^2, \\ \underline{\text{Adaptive Learning:}} \\ R_t = R_{t-1} + \gamma(\tilde{X}_{t-1}\tilde{X}_{t-1}' - R_{t-1}), \\ \Phi_t = \Phi_{t-1} + \gamma R_t^{-1}\tilde{X}_{t-1}(X_t - \Phi_{t-1}\tilde{X}_{t-1})'. \end{cases} \tag{D.1}$$

(D.2)

The intraperiod timing structure of the model at period t is as follows:

- Given state variables X_{t-1} and regime transition matrix $\boldsymbol{Q_{t-1}}$ from period $t-1$, the shadow rate r_t^*, new transition matrix $\boldsymbol{Q_t}$, and new regime probabilities $(s_t = T)$ and $(s_t = E)$ are realized.
- State variables X_t and expectations E_tX_{t+1} are jointly determined, given beliefs $\boldsymbol{\alpha}_{t-1}$, $\boldsymbol{\beta}_{t-1}$; share of agents n_{t-1}^{RE}, n_{t-1}^{RE} from period $t-1$; and regime probabilities $(s_t = T)$ and $(s_t = E)$.

- Given new state variables X_t, forecast errors FE_t^{RE}, FE_t^L; fitness measures ζ_t^{RE}, ζ_t^L; and new shares of agents n_t^{RE}, n_t^L are realized.
- Given new state variables X_t, beliefs of adaptive learners $\boldsymbol{\alpha}_t$, $\boldsymbol{\beta}_t$ are updated.

We use a modified version of the Kim-Nelson filter (KN filter) to estimate the latent variables, regime probabilities, and the likelihood function. Given the sequential timing of events in the model, the filter admits a conditionally linear structure and consists of the following main blocks: (i) a standard Kalman filter to estimate the latent variables for given beliefs and agent shares, (ii) a Hamilton filter to estimate the latent regime probabilities (Taylor regime or ELB), (iii) a collapsing step to average out the state variables and state covariance matrix, and (iv) updating agent fractions and beliefs conditional on the collapsed state variables. Then the Kalman-filter steps of the next period are applied conditional on the updated fractions and beliefs. Further details of the filtering approach can be found in the appendix of Özden and Wouters (2021).

Appendix E. Additional Model Statistics

This appendix reports additional results related to the HE model. Recall that adaptive learners' PLM is assumed to take the following VAR(1) form:

$$X_t = \boldsymbol{\alpha}_{t-1} + \boldsymbol{\beta}_{t-1} X_{t-1} + \delta_t, \tag{E.1}$$

where $X_t = [y_t, \pi_t, r_t]'$, $\boldsymbol{\alpha}_{t-1} = [\alpha_\pi^{t-1}, \alpha_y^{t-1}, \alpha_r^{t-1}]'$, and

$$\boldsymbol{\beta}_{t-1} = \begin{bmatrix} \beta_{y,y}^{t-1} & \beta_{y,\pi}^{t-1} & \beta_{y,r}^{t-1} \\ \beta_{\pi,y}^{t-1} & \beta_{\pi,\pi}^{t-1} & \beta_{\pi,r}^{t-1} \\ \beta_{r,y}^{t-1} & \beta_{r,\pi}^{t-1} & \beta_{r,r}^{t-1} \end{bmatrix}.$$

Figure E.1 shows the estimated time variation in the PLM coefficients of adaptive learners throughout the sample period 1960:Q1–2019:Q4. Figure E.2 shows some additional summary statistics from the stochastic simulations of the HE model discussed in Section 4.

Figure E.1. Estimated Time Variation in PLM Coefficients of Adaptive Learners in the Heterogeneous Expectation Model over the Period 1960:Q1–2019:Q4

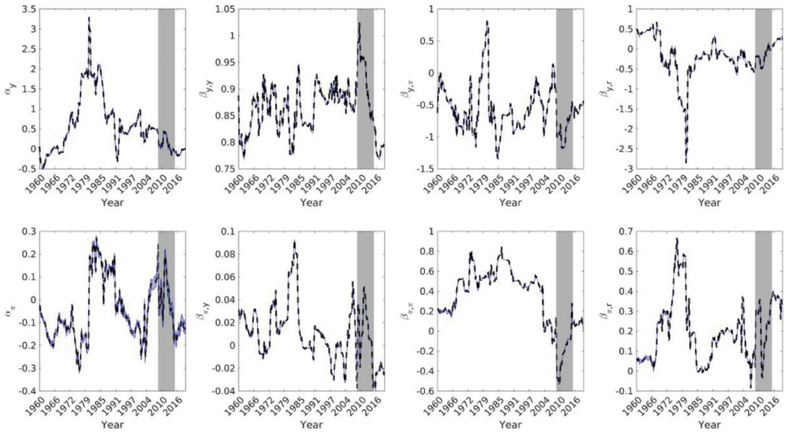

Figure E.2. Summary Statistics from Stochastic Simulations of the Model at the ELB

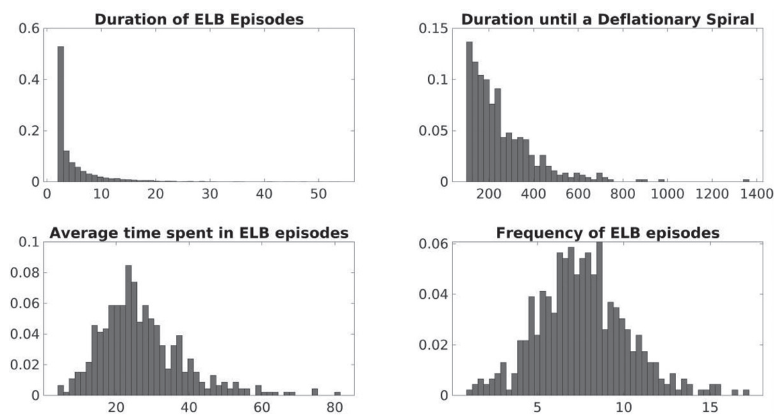

Note: The results are based on 1,000 simulations of the model.

In the simulations, the average duration of an ELB regime (top-left panel) is 3.5 quarters, with a probability of 1.5 percent for durations exceeding five years. In other words, most ELB regimes are short lived, mixed in with the occasional long-lived ELB regimes. The top-right panel shows the distribution of durations until a deflationary spiral is observed in the model. While not all ELB regimes result in deflationary spirals, as discussed in Section 4, all simulations eventually result in a deflationary spiral when the economy is hit with a large enough shock to push the share of adaptive learners into a critically high value. On average, it takes 257 quarters in the model for a deflationary spiral to occur. The bottom-left panel shows the average time spent in ELB regimes in the model: A simulation spends 25 percent of its duration in ELB regimes on average. The bottom-right panel shows the frequency of ELB regimes in the simulations: On average, the model encounters six ELB regimes once every 100 quarters. The frequency and duration of ELB regimes in the model are generally in the range of numbers reported in the literature. For example, Hills, Nakata, and Schmidt (2016) report a range of 10–27 percent as the time spent in the ELB regime in their model calibrated for the U.S. economy. Similarly, Chu and Zhang (2022) report a range of 16–29 percent of ELB regimes in Bank of Canada's main DSGE model ToTEM under a variety of monetary policy rules.

Appendix F. Robustness: Alternative Measures of Output Gap

In this appendix, as a robustness check, we discuss the estimation results of the heterogeneous expectation model under alternative measures of output gap. As discussed in Section 3, our baseline measure of output gap utilizes the Hamilton filter. This is constructed by computing the cyclical component based on the two-year-ahead forecast error of the series using a random-walk model; see Hamilton (2018) for further details. We provide two alternative measures to this output gap. The first one is based on a simple quadratic detrending of real GDP series, as in, e.g., Cornea-Madeira, Hommes, and Massaro (2019). The second one is based on the CBO's estimate of potential output, where output gap is computed as the difference between output and its potential. The resulting measures of output gap are shown in Figure F.1, whereas the parameter estimates for the model under alternative measures are reported in Table F.1. All three measures are qualitatively similar and generally agree over

Figure F.1. Alternative Measures of Output Gap Based on Hamilton Filter, CBO's Measure of Output Gap, and Quadratically Detrended Output over the Period 1960:Q1–2019:Q4

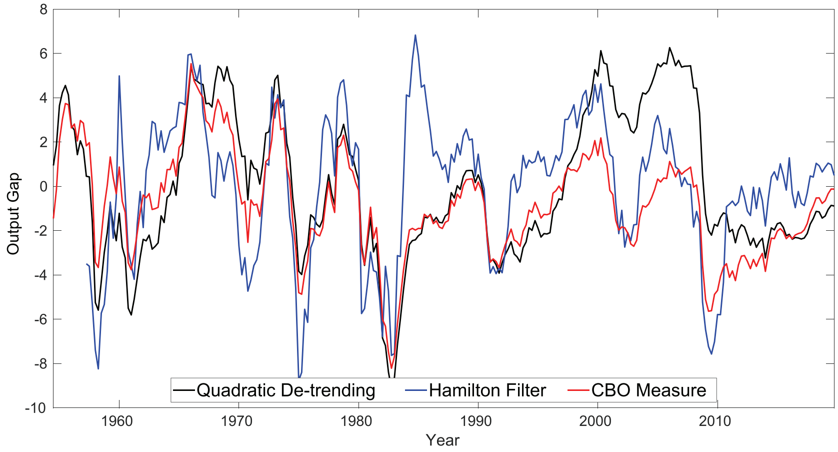

periods with excesss demand and excess supply. The measure based on the Hamilton filter is more volatile than its alternatives, suggesting that the estimated trend (i.e., potential output) under this filter is smoother. The results in Table F.1 suggest that the parameter estimates are generally robust to alternative measures of output gap. There are a few exceptions: The NKPC slope κ and risk aversion τ are both higher under the Hamilton filter, whereas indexation in IS curve ι_y is lower compared with its alternatives. All of these are consequences of the more volatile and less persistent output gap measure under the Hamilton filter. The remaining parameter estimates are very similar across different measures, with parameter bands well within the range of each other.

Table F.1. Posterior Distribution Moments for the Heterogeneous Expectation Model with Alternative Measures of Output Gap, Estimated with Data over Period 1960:Q1–2019:Q4

	Hamilton Filter			CBO Measure of Potential			Quadratic Detrending		
Parameter	**LB**	**Mean**	**UB**	**LB**	**Mean**	**UB**	**LB**	**Mean**	**UB**
$\overline{\pi}$ (Inflation Trend)	0.55	0.58	0.6	0.55	0.58	0.61	0.53	0.55	0.59
$\bar{r_T}$ (Int. Rate Trend—Taylor)	0.26	0.26	0.26	0	0.13	0.3	–0.15	0	0.2
$\bar{r_E}$ (Int. Rate Trend—ELB)	0.03	0.03	0.04	0.03	0.03	0.04	0.03	0.03	0.04
κ (NKPC Slope)	0.0017	0.0032	0.0052	0.0006	0.0014	0.0028	0.0001	0.0003	0.0007
τ (Risk Aversion)	1.03	1.13	1.3	0.53	0.61	0.71	0.82	0.99	1.12
ι_y (Indexation—IS Curve)	0.03	0.06	0.12	0.4	0.45	0.51	0.27	0.36	0.46
ι_p (Indexation—NKPC)	0.09	0.15	0.2	0.05	0.09	0.13	0.07	0.12	0.17
ϕ_π (MP Inflation Reaction)	1.24	1.53	1.79	1.34	1.71	2.15	1.4	1.65	1.83
ϕ_y (MP Output Gap Reaction)	0.34	0.42	0.51	0.23	0.34	0.52	0.22	0.3	0.39
$\phi_{\Delta y}$ (MP Output Gap Growth Reaction)	0.09	0.12	0.14	0.22	0.25	0.28	0.2	0.26	0.3
ρ_y (Persistence—Demand Shock)	0.9	0.92	0.94	0.93	0.95	0.97	0.93	0.94	0.96
ρ_π (Persistence—Supply Shock)	0.93	0.95	0.96	0.92	0.95	0.97	0.94	0.96	0.98
ρ_r (MP Smoothing)	0.01	0.04	0.08	0.01	0.02	0.06	0.01	0.05	0.1
η_y (St. Dev.—Demand Shock)	0.28	0.32	0.36	0.38	0.42	0.5	0.26	0.32	0.39
η_π (St. Dev.—Supply Shock)	0.35	0.38	0.41	0.34	0.37	0.4	0.37	0.4	0.43
$\eta_{r,T}$ (St. Dev.—MP Shock)	0.22	0.24	0.27	0.23	0.25	0.27	0.23	0.25	0.28
$\eta_{r,E}$ (St. Dev.—MP Shock at ELB)	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
$\eta^{SR}_{\pi,exp}$ (St. Dev.—SR Inflation Exp.)	0.09	0.11	0.12	0.08	0.1	0.11	0.09	0.1	0.12
$\eta^{LR}_{\pi,exp}$ (St. Dev.—LR Inflation Exp.)	0.07	0.08	0.09	0.09	0.11	0.12	0.11	0.13	0.14
Φ_1 (MP Switching—Taylor to ELB)	0.06	0.1	0.15	0.07	0.11	0.16	0.06	0.1	0.16
Φ_2 (MP Switching—ELB to Taylor)	0.07	0.16	0.31	0.06	0.12	0.29	0.07	0.16	0.35
γ (Constant Gain)	0.0584	0.0585	0.0587	0.059	0.0593	0.0602	0.0588	0.0589	0.0589
ω (Memory)	0.49	0.62	0.74	0.58	0.68	0.77	0.45	0.66	0.88
χ (Intensity of Choice)	0.38	0.51	0.66	1.12	1.47	1.96	0.24	0.36	0.46
Marg. (log-) Likl.		–410.05			–280.03			–333.92	

Appendix G. Robustness: Estimations without Inflation Expectations

This appendix presents the estimation results of the models in Section 3 without using inflation expectations data. The goal is to check the sensitivity of parameter estimates to inflation expectations data, and to examine whether the relative fitness of the models change when data on inflation expectations is excluded. The results are presented in Table G.1. In terms of model fitness, the relative ranking of the models remains the same as in Section 3. The HE model provides the best fit, followed by the AL model, the RE model with switching in MP, and finally the baseline RE model. This shows that the HE model improves model fitness not only along the margin of inflation expectations data but also on aggregate macrovariables. Parameters related to learning and heterogeneous expectations are all sensitive to expectations data: The estimated constant gain γ is higher in both the HE and the AL model when inflation expectations are included. The estimated memory in heterogeneous switching ω, as well as the intensity of choice χ, are both lower when estimated with expectations. These results show that including expectations data in the data set plays an important role in identifying parameters related to the learning process. It is also important to note that the results in this section are consistent with previous studies in the literature that have estimated learning and heterogeneous expectation models without using any survey data, e.g., Milani (2007), Slobodyan and Wouters (2012a, 2012b), and Cornea-Madeira, Hommes, and Massaro (2019), to name a few.

Table G.1. Posterior Distribution Moments for All Models, Estimated with Data on Only Inflation, Output Gap, and Nominal Interest Rate over the Period 1960:Q1–2019:Q4

Parameter	RE			RE—Switching in MP			Hetero			Learning		
	LB	Mean	UB	LB	Mean	UB	LB	Mean	UB	LB	Mean	UB
$\overline{\pi}$ (Inflation Trend)	0.71	0.92	1.12	0.74	1.11	1.31	0.29	0.44	0.73	0.28	0.6	1.06
$\bar{r_T}$ (Int. Rate Trend—Taylor)	0.84	1.21	1.57	0.74	1.3	1.6	0.2	0.39	1.06	0.4	0.89	1.48
$\bar{r_E}$ (Int. Rate Trend—ELB)				0.03	0.03	0.03	0.03	0.03	0.04	0.03	0.03	0.04
κ (NKPC Slope)	0.0066	0.0135	0.0238	0.0037	0.0086	0.0158	0.0125	0.0214	0.0276	0.0472	0.0571	0.0704
τ (Risk Aversion)	1.49	2.04	2.81	2.12	2.59	3.31	1.72	2.29	3.05	1.7	2.28	3.08
ι_y (Indexation—IS Curve)	0.09	0.15	0.21	0.08	0.14	0.19	0.05	0.07	0.13	0.11	0.16	0.22
ι_p (Indexation—NKPC)	0.03	0.07	0.13	0.03	0.07	0.11	0.02	0.04	0.09	0.02	0.04	0.09
ϕ_π (MP Inflation Reaction)	1.37	1.66	2	1.3	1.56	1.83	1.28	1.64	1.97	1.08	1.35	1.64
ϕ_y (MP Output Gap Reaction)	0.17	0.25	0.36	0.18	0.26	0.35	0.25	0.36	0.43	0.18	0.28	0.42
$\phi_{\Delta y}$ (MP Output Gap Growth Reaction)	0.08	0.1	0.13	0.07	0.08	0.1	0.07	0.1	0.13	0.04	0.06	0.08
ρ_r (MP Smoothing)	0.88	0.9	0.93	0.88	0.9	0.92	0.89	0.92	0.94	0.9	0.93	0.95
ρ_y (Persistence—Demand Shock)	0.86	0.91	0.95	0.83	0.88	0.92	0.88	0.93	0.95	0.53	0.61	0.68
ρ_π (Persistence—Supply Shock)	0.66	0.75	0.82	0.68	0.76	0.82	0.68	0.8	0.86	0.63	0.72	0.8
η_y (St. Dev.—Demand Shock)	0.15	0.2	0.26	0.16	0.21	0.27	0.18	0.21	0.27	0.33	0.38	0.46
η_π (St. Dev.—Supply Shock)	0.1	0.14	0.17	0.1	0.13	0.17	0.13	0.19	0.23	0.22	0.25	0.27
$\eta_{r,T}$ (St. Dev.—MP Shock)	0.2	0.22	0.24	0.2	0.22	0.24	0.2	0.22	0.25	0.19	0.21	0.22
$\eta_{r,E}$ (St. Dev.—MP Shock at ELB)				0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Φ_1 (MP Switching—Taylor to ELB)							0.07	0.11	0.19	0.1	0.18	0.33
Φ_2 (MP Switching—ELB to Taylor)							0.4	0.11	0.27	0.03	0.05	0.13
γ (Constant Gain)							0.0017	0.0036	0.0137	0.0062	0.0081	0.0105
ω (Memory)							0.67	0.78	0.94	0.57	0.76	0.9
χ (Intensity of Choice)							1.36	4.63	7.45	2.3	4.79	8.39
$1-q^T$ (Exog. Exit Probability—Taylor)				0.16	0.23	0.34						
$1-q^E$ (Exog. Exit Probability—ELB)				0.0028	0.012	0.03						
Marg. (log-) Likl.		–551.22			–527.88			–513.48			–516.69	
Bayes Factor		1			10.13			16.39			14.99	

References

An, S., and F. Schorfheide. 2007. “Bayesian Analysis of DSGE Models.” *Econometric Reviews* 26 (2–4): 113–72.

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

Anufriev, M., T. Assenza, C. Hommes, and D. Massaro. 2013. “Interest Rate Rules and Macroeconomic Stability under Heterogeneous Expectations.” *Macroeconomic Dynamics* 17 (8): 1574–1604.

Arifovic, J., A. Grimaud, I. Salle, and G. Vermandel. 2020. “Social Learning and Monetary Policy at the Effective Lower Bound.” Bank of Canada Working Paper 2020-2.

Aruoba, S. B. 2020. “Term Structures of Inflation Expectations and Real Interest Rates.” *Journal of Business and Economic Statistics* 38 (3): 542–53.

Aruoba, S. B., M. Mlikota, F. Schorfheide, and S. Villalvazo. 2022. “SVARs with Occasionally-Binding Constraints.” *Journal of Econometrics* 231 (2): 477–99.

Barthélemy, J., and M. Marx. 2017. “Solving Endogenous Regime Switching Models.” *Journal of Economic Dynamics and Control* 77 (April): 1–25.

Benigno, G., A. Foerster, C. Otrok, and A. Rebucci. 2020. “Estimating Macroeconomic Models of Financial Crises: An Endogenous Regime-Switching Approach.” NBER Working Paper No. 26935.

Bianchi, F. 2016. “Methods for Measuring Expectations and Uncertainty in Markov-Switching Models.” *Journal of Econometrics* 190 (1): 79–99.

Bianchi, F., and C. Ilut. 2017. “Monetary/Fiscal Policy Mix and Agents’ Beliefs.” *Review of Economic Dynamics* 26 (October): 113–39.

Bianchi, F., L. Melosi, and M. Rottner. 2021. “Hitting the Elusive Inflation Target.” *Journal of Monetary Economics* 124 (November): 107–22.

Binning, A., and J. Maih. 2016. “Implementing the Zero Lower Bound in an Estimated Regime-Switching DSGE Model.” Norges Bank Working Paper No. 3/2016.

Blanchard, O. 2022. "Why I Worry about Inflation, Interest Rates, and Unemployment." RealTime Economics Blog, Peterson Institute for International Economics, March 14.

Branch, W., T. Davig, and B. McGough. 2007. "Expectational Stability in Regime-Switching Rational Expectations Models." Federal Reserve Bank of Kansas City Research Working Paper No. 08-09. Available at https://ssrn.com/abstract=1073150.

Branch, W. A. 2004. "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations." *Economic Journal* 114 (497): 592–621.

Branch, W. A., and B. McGough. 2009. "A New Keynesian Model with Heterogeneous Expectations." *Journal of Economic Dynamics and Control* 33 (5): 1036–51.

Bullard, J., and K. Mitra. 2002. "Learning about Monetary Policy Rules." *Journal of Monetary Economics* 49 (6): 1105–29.

Busetti, F., D. Delle Monache, A. Gerali, and A. Locarno. 2017. "Trust, but Verify. De-anchoring of Inflation Expectations under Learning and Heterogeneity." ECB Working Paper No. 1994.

Carstens, A. 2022. "The Return of Inflation." Speech at the International Center for Monetary and Banking Studies, Geneva, April 5.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston. 2021. "Anchored Inflation Expectations." CAMA Working Paper No. 25/2020 (Updated August 2021). Available at https://ssrn.com/abstract=3555532.

Chang, Y., J. Maih, and T. Fei. 2018. "State Space Models with Endogenous Regime Switching." Mimeo.

Chen, H. 2017. "The Effects of the Near-Zero Interest Rate Policy in a Regime-Switching Dynamic Stochastic General Equilibrium Model." *Journal of Monetary Economics* 90 (October): 176–92.

Cho, S. 2016. "Sufficient Conditions for Determinacy in a Class of Markov-Switching Rational Expectations Models." *Review of Economic Dynamics* 21 (July): 182–200.

Chu, V., and Y. Zhang. 2022. "Harnessing the Benefit of State-Contingent Forward Guidance." Bank of Canada Staff Analytical Note No. 2022-13.

Chung, H., and W. Xiao. 2013. "Cognitive Consistency, Signal Extraction, and Macroeconomic Persistence." Mimeo, SUNY Binghamton.

Clarida, R., J. Gali, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707.

Cole, S. J. 2018. "The Effectiveness of Central Bank Forward Guidance under Inflation and Price-Level Targeting." *Journal of Macroeconomics* 55 (March): 146–61.

Cole, S. J., and E. Martínez-García. 2019. "The Effect of Central Bank Credibility on Forward Guidance in an Estimated New Keynesian Model." Globalization and Monetary Policy Institute Working Paper No. 375.

Cornea-Madeira, A., C. Hommes, and D. Massaro. 2019. "Behavioral Heterogeneity in US Inflation Dynamics." *Journal of Business and Economic Statistics* 37 (2): 288–300.

Croushore, D. D. 1993. "Introducing the Survey of Professional Forecasters." *Business Review* (Federal Reserve Bank of Philadelphia) 6 (November/December): 3–15.

Davig, T., and E. M. Leeper. 2007. "Generalizing the Taylor Principle." *American Economic Review* 97 (3): 607–35.

Del Negro, M., M. P. Giannoni, and F. Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics* 7 (1): 168–96.

Di Bartolomeo, G., M. Di Pietro, and B. Giannini. 2016. "Optimal Monetary Policy in a New Keynesian Model with Heterogeneous Expectations." *Journal of Economic Dynamics and Control* 73 (December): 373–87.

Eusepi, S., and B. Preston. 2010. "Central Bank Communication and Expectations Stabilization." *American Economic Journal: Macroeconomics* 2 (3): 235–71.

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

Evans, G. W., and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics.* Princeton University Press.

Farmer, R. E., D. F. Waggoner, and T. Zha. 2009. "Understanding Markov-Switching Rational Expectations Models." *Journal of Economic Theory* 144 (5): 1849–67.

———. 2011. "Minimal State Variable Solutions to Markov-Switching Rational Expectations Models." *Journal of Economic Dynamics and Control* 35 (12): 2150–66.

Goy, G., C. Hommes, and K. Mavromatis. 2020. "Forward Guidance and the Role of Central Bank Credibility under Heterogeneous Beliefs." *Journal of Economic Behavior and Organization* 200 (August): 1240–74.

Greenberg, E. 2012. *Introduction to Bayesian Econometrics*. Cambridge University Press.

Gust, C., E. Herbst, and J. D. Lopez-Salido. 2018. "Forward Guidance with Bayesian Learning and Estimation." FEDS Working Paper No. 2018-072.

Haberis, A., R. Harrison, and M. Waldron. 2014. "Transitory Interestrate Pegs under Imperfect Credibility." Centre for Macroeconomics Discussion Paper No. 2014-22.

Hamilton, J. D. 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *Review of Economics and Statistics* 100 (5): 831–43.

Hills, T., T. Nakata, and S. Schmidt. 2016. "The Risky Steady State and the Interest Rate Lower Bound." ECB Working Paper No. 1913.

Hommes, C., and J. Lustenhouwer. 2019. "Inflation Targeting and Liquidity Traps under Endogenous Credibility." *Journal of Monetary Economics* 107 (November): 48–62.

Hommes, C., D. Massaro, and M. Weber. 2019. "Monetary Policy under Behavioral Expectations: Theory and Experiment." *European Economic Review* 118 (September): 193–212.

Jääskelä, J., and R. McKibbin. 2010. "Learning in an Estimated Small Open Economy Model." Reserve Bank of Australia Research Discussion Paper No. 2010-02.

Kim, C.-J., and C. R. Nelson. 1999. "Has the US Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle." *Review of Economics and Statistics* 81 (4): 608–16.

Kulish, M., J. Morley, and T. Robinson. 2014. "Estimating DSGE Models with Forward Guidance." Discussion Paper No. 2014-32A, School of Economics, University of New South Wales.

Kulish, M., and A. Pagan. 2017. "Estimation and Solution of Models with Expectations and Structural Changes." *Journal of Applied Econometrics* 32 (2): 255–74.

Lansing, K. J. 2021. "Endogenous Forecast Switching Near the Zero Lower Bound." *Journal of Monetary Economics* 117 (January): 153–69.

Lindé, J., J. Maih, and R. Wouters. 2017. "Estimation of Operational Macromodels at the Zero Lower Bound." Mimeo.

Lindé, J., F. Smets, and R. Wouters. 2016. "Challenges for Central Banks' Macro Models." In *Handbook of Macroeconomics*, Vol. 2, ed. J. B. Taylor and H. Uhlig, 2185–2262. Elsevier.

Liu, P., and H. Mumtaz. 2011. "Evolving Macroeconomic Dynamics in a Small Open Economy: An Estimated Markov-Switching DSGE Model for the UK." *Journal of Money, Credit and Banking* 43 (7): 1443–74.

Malmendier, U., and S. Nagel. 2016. "Learning from Inflation Experiences." *Quarterly Journal of Economics* 131 (1): 53–87.

Mavroeidis, S. 2021. "Identification at the Zero Lower Bound." *Econometrica* 89 (6): 2855–85.

Milani, F. 2007. "Expectations, Learning and Macroeconomic Persistence." *Journal of Monetary Economics* 54 (7): 2065–82.

———. 2011. "Expectation Shocks and Learning as Drivers of the Business Cycle." *Economic Journal* 121 (552): 379–401.

Özden, T. 2021. "Heterogeneous Expectations and the Business Cycle at the Effective Lower Bound." De Nederlandsche Bank Working Paper No. 714.

Özden, T., and R. Wouters. 2021. "Restricted Perceptions, Regime Switches and the Zero Lower Bound." Mimeo, University of Amsterdam.

Shackleton, R. 2018. "Estimating and Projecting Potential Output Using CBO's Forecasting Growth Model." Congressional Budget Office Working Paper No. 2018-03.

Sims, C. 1999. "Matlab Optimization Software." Quantitative Macroeconomics and Real Business Cycles (QM&RBC) Codes 13.

Sims, C. A., and T. Zha. 2006. "Were There Regime Switches in US Monetary Policy?" *American Economic Review* 96 (1): 54–81.

Slobodyan, S., and R. Wouters. 2012a. "Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models." *American Economic Journal: Macroeconomics* 4 (2): 65–101.

———. 2012b. “Learning in an Estimated Medium-Scale DSGE Model.” *Journal of Economic Dynamics and Control* 36 (1): 26–46.

Smets, F., and R. Wouters. 2007. “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review* 97 (3): 586–606.

Uhlig, H. F. 1995. “A Toolkit for Analyzing Nonlinear Dynamic Stochastic Models Easily.” Discussion Paper No. 101, Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis.

Synchronization vs. Transmission: The Effect of the German Slowdown on the Italian Business Cycle*

Alessandro Mistretta
Banca d'Italia

This work studies the transmission of the business cycle across countries by analyzing the effects of the 2018 German slowdown on Italian activity. We apply a difference-in-differences strategy to expectations data from Banca d'Italia's Survey of Inflation and Growth Expectations (SIGE). Firms exporting to Germany had lower expectations for the Italian economy (sentiment) and for their own demand, investment, and employment (assessment) than firms exporting to other countries or not exporting at all. We quantify the response of key Italian macroeconomic aggregates to worsening sentiment and assessment of Italian firms using a forecasting model. A significant contemporaneous impact on Italian GDP highlights the role of the expectations of firms exposed to foreign markets in transmitting foreign business cycle.

JEL Codes: E2, E32, F15, F44, L6.

1. Introduction

The existence of business cycle synchronization, especially in a currency union, is extensively discussed in the economic literature. Despite the clear evidence that business cycle synchronization plays an essential role in the European economy, the relative degree of this

*I am indebted to Federico Cingano, Simone Emiliozzi, Marco Flaccadoro, Elisa Guglielminetti, Alberto Locarno, Claudia Pacella, Concetta Rondinelli, Stefano Siviero, Martin Uribe, Stefania Villa, Giordano Zevi, Roberta Zizza, Francesco Zollino, and five anonymous referees for their valuable comments. An earlier version of this paper circulated under the title "Business Cycle Synchronization or Business Cycle Transmission? The Effect of the German Slowdown on the Italian Economy." The views expressed do not necessarily reflect those of the Banca d'Italia. Author e-mail: alessandro.mistretta@bancaditalia.it.

varies over time. According to recent literature, there is no clear consensus about the degree of synchronization in the most recent period, especially after the double-dip recession: de Lucas Santos and Delgado Rodríguez (2016) and Gomez et al. (2017) report an increase in business cycle co-movement, while other authors show evidence of business cycle divergence (see, among others, Ferroni and Klaus 2015; Grigoraş and Stanciu 2016; Beck 2021). In particular, Beck (2021) suggests that the declining share of manufacturing in the European Union (EU) explains the increased divergence. However, there is no consensus on the determinants of business cycle co-movement that distinguish between the possibility of a common (namely determined by a common economic shock) and a transmitted business cycle (Garnier 2004; di Giovanni, Levchenko, and Mejean 2018).

In this paper, given Germany's economic importance for the whole euro-area economy, we study the relationship between the German and Italian business cycles.

The relationship between the German business cycle and the Italian economic performance is likely to be significant, as the two countries are closely interconnected through trade. Germany and Italy have open economies, with exports representing a significant portion of their gross domestic product (GDP). In 2019, Germany was the EU's top exporter and Italy was the third, with exports accounting for 45 and 32 percent of their respective GDPs. Germany is the top sales market for Italian firms, accounting for 13 percent of Italian goods exports in 2019. Additionally, 17 percent of Italy's imported goods come from Germany.[1] These close ties are due to both countries being part of the euro area and having significant manufacturing sectors, which account for 23 and 17 percent of their respective GDPs.

The contemporaneous correlations between the key economic activity indicators (GDP and industrial production) of these economies were exceptionally high during the double-dip recession. The correlation for industrial production (IP) has remained relatively high. On the contrary, the correlation for GDP declined since 2014 and, after reaching a historical minimum in 2018:Q1,

[1]The share of goods originating in Germany is double those originating in France, which is Italy's second biggest trading partner.

Figure 1. Correlation between Italian and German Economic Indicators

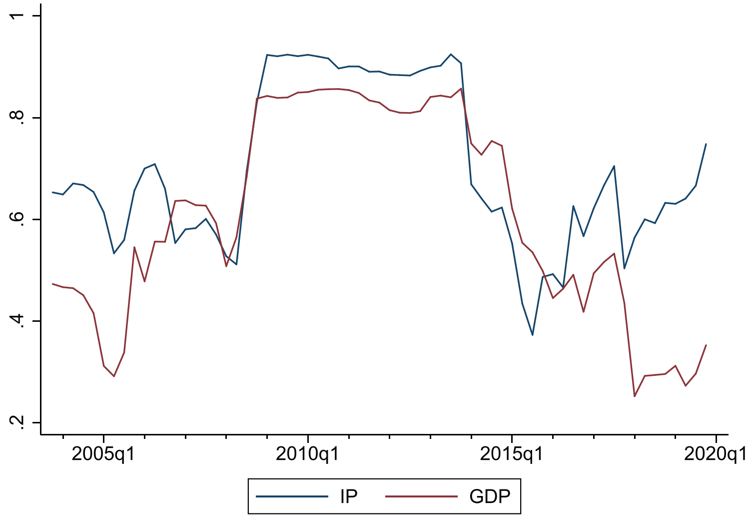

Note: Rolling correlation (five year) on *q-o-q* growth rates; Eurostat data.

returned to growth, reaching a peak during the COVID-19 recession (Figure 1). On the whole, the German and Italian business cycles are closely synchronized.

In this paper, starting from an important economic shock that hit the German economy in 2018, we analyze whether this negative shock affected the Italian economy.

The German economic cycle started slowing down in 2018:Q1; the weakening was particularly marked from 2018:Q3 in the manufacturing sector: the growth rate of manufacturing value-added has been subdued since then, while services have proved to be more resilient (Figure 2).

This slowdown has been caused by some country-specific shocks rather than common euro-area shocks. Differently from before, the German IP dynamic has been significantly worse since 2018, with respect to those recorded in Italy and the other euro-area countries (Figure 3).

Several temporary factors have hampered German growth since the beginning of 2018, such as the high levels of sick leave due

Figure 2. Germany, Main Economic Indicators

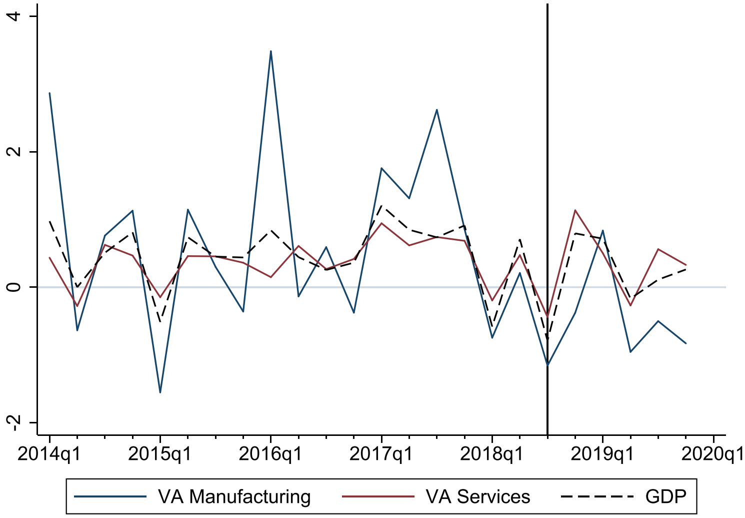

Note: q-o-q growth rates on Eurostat data.

Figure 3. Industrial Production, Main Euro-Area Economies

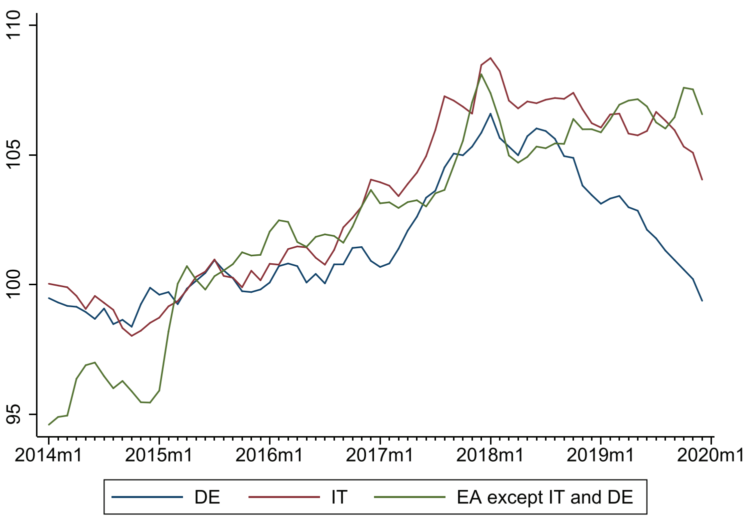

Note: MA(3), Indices 2015=100; Eurostat data.

Figure 4. German Industrial Production

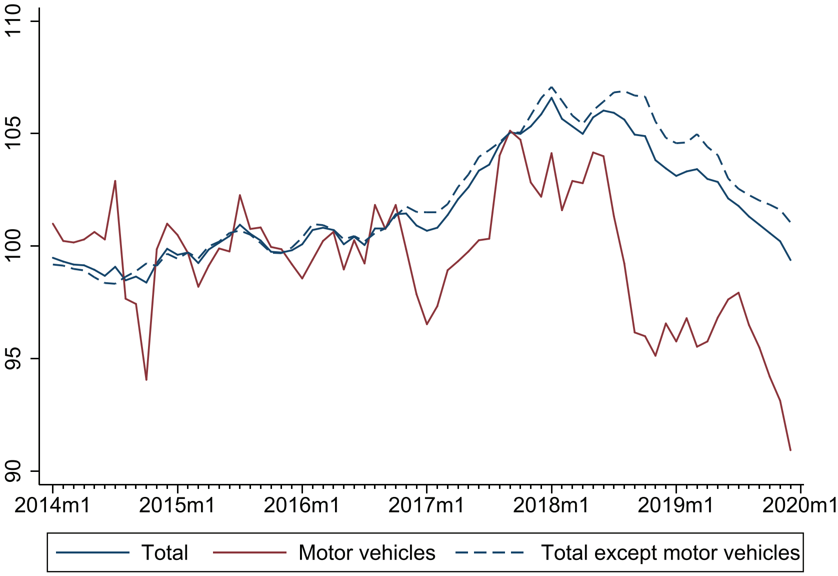

Note: MA(3), Indices 2015=100; Eurostat data.

to the unusually virulent influenza, the cold winter weather conditions, and industrial strikes; additionally, there was already growing evidence that the automotive sector may have reached its peak (Camba-Méndez and Forsells 2018).

During 2018, German growth was curbed by bottlenecks in the automotive sector: due to difficulties in the introduction of a new emissions testing procedure (the Worldwide Harmonised Light Vehicle Test Procedure, or WLTP), the production of motor vehicles fell sharply (see Figure 4); delays in obtaining certificates of compliance with these new standards led German manufacturers to suspend the production of many car models,[2] causing severe disruption to both delivery and sales (European Commission 2019).

[2]Some producers even waited to request WLTP approval for selected models at the end of their life cycle, thus effectively ceasing production until new models were introduced.

Figure 5. Economic Shocks in German Industry

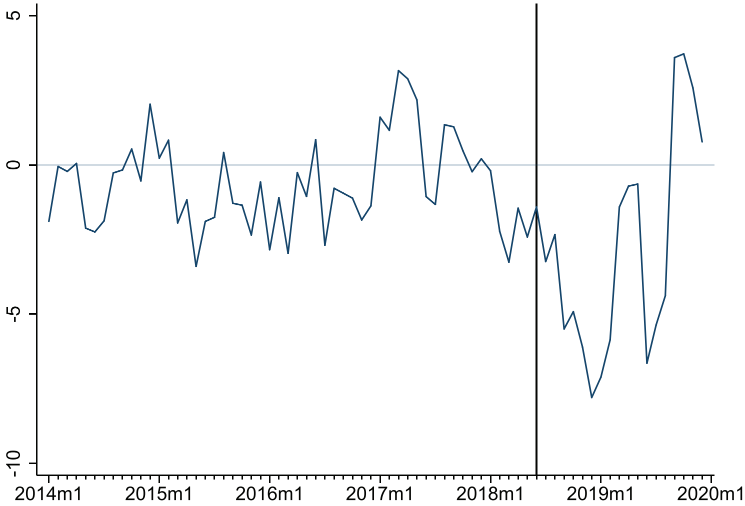

Note: Difference between expected and effective production levels; deviation from historical mean. Business and consumer surveys—European Commission.

As a result, the decline in industrial production was not confined to the automotive sector but widespread across manufacturing and more persistent than previously expected.[3]

Finally, the difference between the actual and expected production levels became significant compared with the historical mean;[4] this suggests that the economic slowdown in industrial activity was unexpected (Figure 5).

Considering the nature of the German slowdown, which was exogenous to the Italian economy until 2020:Q1, in this paper we analyze whether there was a transmission of the economic shock to the Italian economy.

---

[3]According to the European Commission (2019), German GDP in 2018 would have been 0.6 percent higher without such a fall in the automotive sector. According to the national accounts, between 2014–17, manufacturing contributed, on average, to total German growth by about 0.8 percent per year; this contribution became modest in 2018 (0.2 percent) and negative in 2019 (–0.8 percent).

[4]The unexpected assessment error was more than twice its historical standard deviation.

We apply a difference-in-differences (diff-in-diff) strategy to expectations data from Banca d'Italia's Survey of Inflation and Growth Expectations (SIGE) to investigate if and how the slowdown in Germany is hitting the Italian economy. We focus on the "direct effect"—namely, the effect on the activity of firms exporting to the German market—as this approach does not enable us to identify "indirect effects" that may transit through other channels, such as global value chains or domestic demand. Therefore, this evaluation probably underestimates the effect of the decline in German manufacturing on the Italian economy.

Although many works have exploited this data set to study different issues relating to inflation expectations (see, among others, Bartiloro, Bottone, and Rosolia 2019; Coibion, Gorodnichenko, and Ropele 2020; Conflitti and Zizza 2021), to the best of our knowledge, only one paper uses this data set to analyze issues relating to the business cycle (Cesaroni and Iezzi 2017).

In 2019, the sentiment indicators were worse for Italian companies exposed to the German market. Expectations for demand and plans for investment and employment were significantly worse for these firms. The effects on investment and employment were also observed, but with a delay compared with the effects on demand. Additionally, the disagreement about the economic forecast increased for exporters to Germany, representing the main contribution to the increase in total uncertainty. After discussing how well the SIGE series mimics the national economic aggregate, we quantify the effect of the German slowdown on Italian GDP using a forecasting model. According to the estimates, the effect on GDP was about 1 percentage point, mainly concentrated in 2019; the negative effect is equal to 2.5 percentage points on firms' investment; conversely, we do not find any effect on employment.

This work's contributions are twofold: firstly, we address the macroeconomic issue using a microeconometric approach (and policy evaluation techniques in particular) to survey microdata; and secondly, we investigate the relationship between the German and the Italian business cycles from the standpoint of transmission rather than of "simple" synchronization.

The rest of the paper is structured as follows. Section 2 reviews the literature, and Section 3 describes the data set used. Section 4 proposes a microeconometric exercise to estimate the effect of the German slowdown on Italian firms' economic activity, while Section 5 quantifies this effect from a macroeconomic point of view. Finally, Section 6 concludes.

2. Literature

Since Dellas (1986), a common business cycle across countries has been extensively studied from both a theoretical and an empirical point of view. Dellas (1986) proposed a model that predicts a positive and persistent co-movement in trade and gross national products (GNPs) across countries; he showed empirically that the primary source of this positive covariance is the existence of common shocks rather than trade interdependence. Canova and Dellas (1993) confirmed this view by finding a positive (moderate) effect of trade interdependence on the common business cycle, though it is not statistically significant.

The determinants of business cycle co-movements between countries were investigated by Baxter and Kouparitsas (2005), who found controversial results. Using a large data set with more than 100 countries, they showed empirically that (i) the correlation between business cycles is increasing in trade relationships; (ii) the industrial structure does not affect the business cycle's synchronization; and (iii) the existence of a currency union does not have a significant impact on the correlated business cycle.

The importance of a currency union for business cycle synchronization has been analyzed extensively since the late 1990s. Frankel and Rose (1998) studied the effects of a common currency area on the business cycle in their seminal paper. They argued that these effects are ambiguous: (i) on the supply side, by reducing trade barriers, a common currency union can lead to more industry specialization by a country and then to more asynchronous business cycles resulting from industry-specific shocks; and (ii) on the other hand, increased integration may result in more highly correlated business cycles because of demand shocks or intra-industry trade. However, this ambiguity was more theoretical than empirical since they found empirically that greater integration involves a more highly integrated cycle.

Many papers have analyzed the impact of adopting the euro on business cycle synchronization. Gonçalves, Rodrigues, and Soares

(2009) found that the euro increased the correlation among the economic cycles of euro-area members. Other studies have classified countries by their importance to the euro-area business cycle, distinguishing between *core* and *peripheral* countries (e.g., Ahlborn and Wortmann 2018). Enders, Jung, and Müller (2013) found that domestic shocks generate more significant cross-country spillovers under the European Monetary Union (EMU) than before the EMU was created. Campos, Fidrmuc, and Korhonen (2019) found that the correlation between business cycles across European countries has significantly increased since the introduction of the euro in 1999 (from an average of 0.4 to 0.6), confirming the view previously expressed by Frankel and Rose (1997). However, the business cycle correlation is lower than in the United States due to the existence of European national borders (Clark and van Wincoop 2001). Despite increased synchronization after the euro's adoption, recent papers have shown evidence of business cycle divergence in the EU, particularly after the double-dip recession (e.g., Ferroni and Klaus 2015; Grigoraş and Stanciu 2016; Beck 2021).

To summarize, the empirical literature explains the existence of business cycle synchronization because of (i) the presence of common shocks that hit different economies at the same time (Dellas 1986; Canova and Dellas 1993; Imbs 2004); and (ii) the possibility that shocks are transmitted through trade and multinational linkages (Frankel and Rose 1998; Eickmeier 2007; Burstein, Kurz, and Tesar 2008; Kleinert, Martin, and Toubal 2015; di Giovanni, Levchenko, and Mejean 2018).

From the theoretical point of view, the interconnection of the business cycle in a two-country model is extensively studied. A significant strand of literature explains the channels for and the persistence of business cycle synchronization (see, among others, Chiarella, Flasher, and Hung 2006). In this vein of literature, the model proposed by Charpe et al. (2016) is particularly relevant to the present work, in which the role of business confidence is exploited as an independent transmission channel for the business cycle in a two-country model. In particular, the state of confidence, which depends on the current state of the business cycle in the countries considered, would play a reinforcing effect through the expected profit and aggregate investment.

Due to the importance of the interconnection within the euro area in line with the transmission view, this paper investigates how the German business cycle affects the Italian one. As stated previously, we study the effects on the Italian economy of some country-specific shocks that occurred in Germany; this is particularly suited to investigating whether a negative German economic shock is transmitted to the Italian business cycle.

3. Data

In this paper, we use the Survey of Inflation and Growth Expectations (henceforth SIGE) carried out quarterly by Banca d'Italia, on a sample of about 1,000 industrial and service firms with more than 50 employees.[5] The survey collects, among other things, data regarding firms' expectations for consumer price inflation, developments in their own selling prices, and views on the broad macroeconomic outlook, as well on their own business.[6]

Questions regarding economic activity included in the SIGE can be broadly classified into two different groups: those aimed at assessing a firm's sentiment, both on the general economic situation and its own economic situation (henceforth *sentiment indicators*); and those that elicit firms' projections/assessments about their own decisions such as investment or employment plans or their economic total or external demand (henceforth *assessment indicators*).

In this paper, we measure the impact of the German economy's slowdown on the following SIGE indicators:

- The *sentiment indicators* include firms' sentiment on the general Italian economic situation; opinions on the current conditions for investing; the probability of observing an

[5]The survey has been conducted since 1999; from 2019:Q4, the sample has been extended to 1,200 firms. The sample represents about 4 percent of the entire reference population (about 5 percent from 2019:Q4); however, the results refer to the reference population thanks to sampling weights (Banca d'Italia 2019).

[6]Like the typical diffusion indices, the question allows you to choose between three options that indicate an improvement, a worsening, or a stabilization in a specific aspect of a firm's activity. To derive a macroeconomic message, these responses are aggregated using the balances between the share of those companies that indicate an improvement and those that signal a worsening.

Figure 6. SIGE Balances and Corresponding Aggregates in the National Accounts

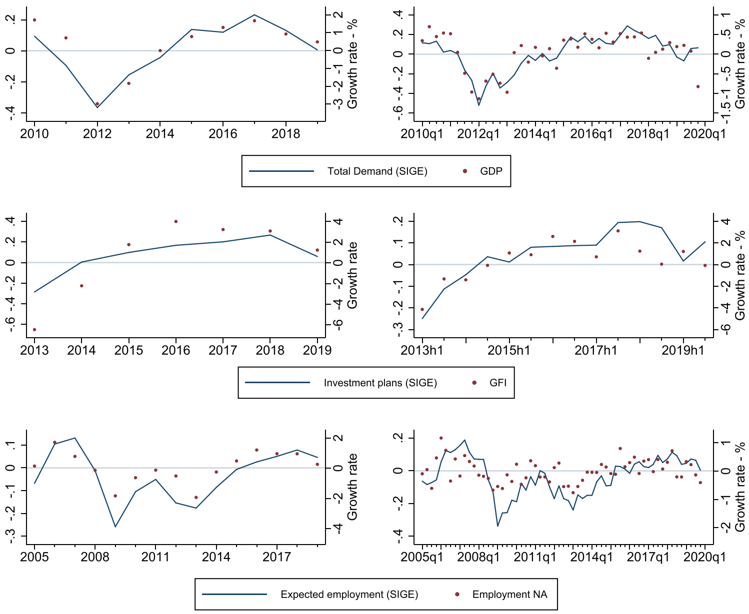

Note: Banca d'Italia SIGE and Istat National Accounts.

improvement in the Italian economy in the following three months; and sentiment indicators about companies' own expected business conditions in the following three months and over a three-year horizon;

- The *assessment indicators* include opinions on firms' current and expected demand for their products (both total and external); investment plans at different time horizons; and the number of employees in the next three months.

The information in the SIGE is very helpful for analyzing the business cycle, as it tracks the corresponding aggregates from the national accounts quite reliably (similar results hold for other business surveys; see, among others, Bachmann and Zorn 2020).

Figure 6 illustrates the close alignment between the SIGE's balances (blue lines) and national account aggregates (red dots). The

Table 1. Regressions

	(1) ΔGDP	(2) ΔIFL	(3) ΔEMPL
	y-o-y		
SIGE	7.924*** (0.00)	19.67*** (0.00)	8.498*** (0.00)
N r2	10 0.835	7 0.909	15 0.831
	q-o-q	***hy-o-hy***	***q-o-q***
SIGE	2.650*** (0.00)	12.32*** (0.00)	1.906*** (0.00)
N r2	40 0.713	14 0.624	60 0.300

Note: p-values in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

SIGE's total demand dynamic corresponds closely to GDP growth. Additionally, the SIGE's investment plans align well with the gross fixed investment (GFI) growth rate and the question on employment with employment growth.[7]

These graphical findings are corroborated by simple regression models where the national account series are regressed on the corresponding SIGE balances. As shown in Table 1, the SIGE balances seem to account for more than 80 percent of the variation in the response variable; this percentage appears to be higher when yearly data are considered (row 1).

Furthermore, we examine the impact on Italian firms' uncertainty by utilizing a simplified version of the measures proposed by Giordani and Soderlind (2003). These include individual uncertainty $(E(\sigma_i^2))$, aggregate uncertainty (V_A), and disagreement among firms' expectations $(V(\mu_{it}))$.[8] Following a large body of literature

[7]In this work, employment growth is based on the number of *employees (domestic concept)* released by Eurostat.

[8]For more information, refer to Appendix A.

Figure 7. Uncertainty Measures

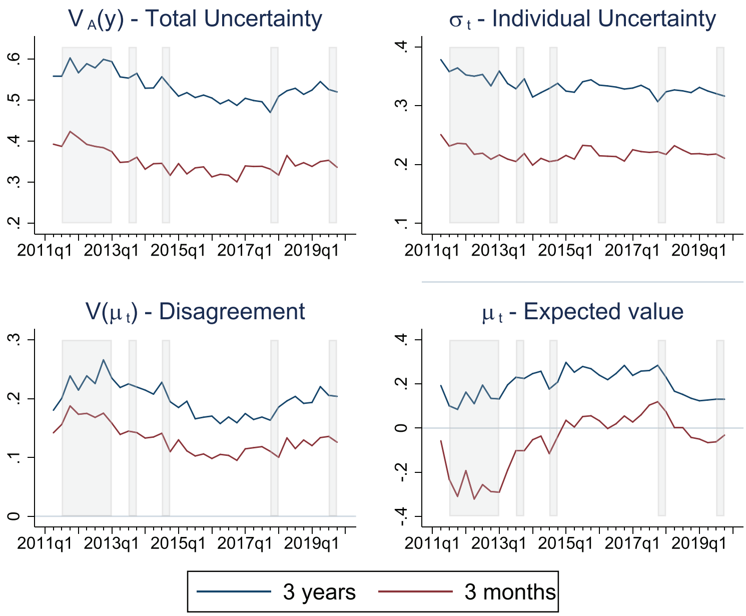

Note: Our calculations based on Banca d'Italia's SIGE survey.

that has investigated the effect of uncertainty on firms' activity, finding that (i) there is a negative relationship between demand uncertainty and firms' decisions (see, among others, Guiso and Parigi 1999; Bloom 2009) and that (ii) uncertainty itself rises sharply during recessions (Bloom et al. 2018), we investigate whether demand uncertainty could be an important channel to explain business decisions (such as investment end employment) for exporters to Germany.

According to our estimates, firms seem to have more optimistic expectations about their economic conditions in the medium run (three years) compared with the short run (three months); however, higher expectations are associated with higher uncertainty (see Figure 7). Disagreement is higher during recession periods, and this is the primary source of uncertainty at aggregate level, confirming the main findings of Giordani and Soderlind (2003).

The SIGE contains some additional structural information, such as a firm's export propensity, which is used to classify firms into four different classes.

3.1 Firms' Exposure to the German Market

The SIGE questionnaire (Appendix D) occasionally includes specific questions to address important issues from a policy perspective when the survey is conducted. In 2019:Q1 and 2019:Q3–2020:Q1, the survey included the following questions aimed at gauging firms' expectations on current and future external demand from Germany:

Comparedwith three months ago, is the foreign demand for your products . . . ?	*Higher*	*Unchanged*	*Lower*	*I do not export to this market*
In Germany	☐	☐	☐	☐
How will the foreign demand for your products vary **in the next three months?**	*Increase*	*No change*	*Decrease*	*I do not export to this market*
In Germany	☐	☐	☐	☐

Using these replies, firms are divided into three groups: exporters to Germany,[9] exporters to other markets, and non-exporters.[10] This division is key to implementing the empirical strategy.

Due to the data set's lack of information, we assume that exporters to Germany both in 2019:Q1 and in 2019:Q3 have been exporting to that country since 2014:Q1. This assumption is justified because decisions concerning destination markets are strategic, as entering a new market entails non-negligible initial costs.[11]

[9] Firms that declare that they export to Germany in at least two of three of the quarters in which they were interviewed are classified as exporters to Germany. Conversely, we exclude from our analysis firms that rarely declared that they export to Germany.

[10] The questionnaire includes a specific question to distinguish between exporters and non-exporters (see question A.2 in Appendix D).

[11] Indeed, according to official statistics (Istat and ICE 2019), the number of firms exporting to Germany remained roughly stable during the period considered: there were 25,024 in 2014 and 24,408 in 2018.

In our sample, about 49 percent of firms only sell in the domestic market; about 70 percent of the remaining firms export to Germany.

Additionally, in 2019:Q4, we asked for information about the propensity to export to the German market.

	Zero	*Up to 1/3 but more than zero*	*Between 1/3 and 2/3*	*Over 2/3 of export*
Considering your firm's total exports in 2019, please indicate the share of exports to the German market.	☐	☐	☐	☐

This information is crucial since it allows a proxy to be computed for the degree of the *German shock* that hits a specific firm according to its exposure to the German market.

We define the exposure as

$$Exposure_{it} = PropensityExport_{it} * ProportionExportGermany_i. \tag{1}$$

This represents the share of total sales from exports to the German market. Due to data limitations, we cannot obtain a continuous variable.[12] Additionally, we assume that the proportion of exports to the German market remains in the same range during the whole period.[13]

Using this strategy, we can define $Exposure_{it}$ for about 5,000 observations throughout the period (see Table 2). Those who export to Germany sell about 10 percent of their total sales in Germany on average; less than 1 percent of the observations are related to firms that export more than 60 percent of their sales to Germany (see Figure 8).

[12] For export propensity and proportion of exports to the German market, firms indicate a range instead of a precise number. To compute $Exposure_{it}$, we use the median value within the provided range.

[13] We know that this assumption, namely a constant share of exports to Germany in a specific range during the time considered, is stronger than those about the decision to export to the German market. However, this is the best information we have, and we only use it in a robust exercise.

Table 2. Classification According to the Exposure to the German Market

Share of Exports to Germany Compared with Total Exports	Median	Share of Sales Exported				Total
		0	0–1/3	1/3–2/3	2/3–1	
		0	.165	.495	.83	
0	0	8,391	986	385	355	10,117
0–1/3	.165	0	1,222	1,692	1,271	4,185
1/3–2/3	.495	0	98	299	186	583
2/3–1	.83	0	2	47	122	171
Not Classified		0	445	403	377	1,225
Total		8,391	2,753	2,826	2,311	16,281

Note: Banca d'Italia SIGE.

Figure 8. Exposure to the German Market

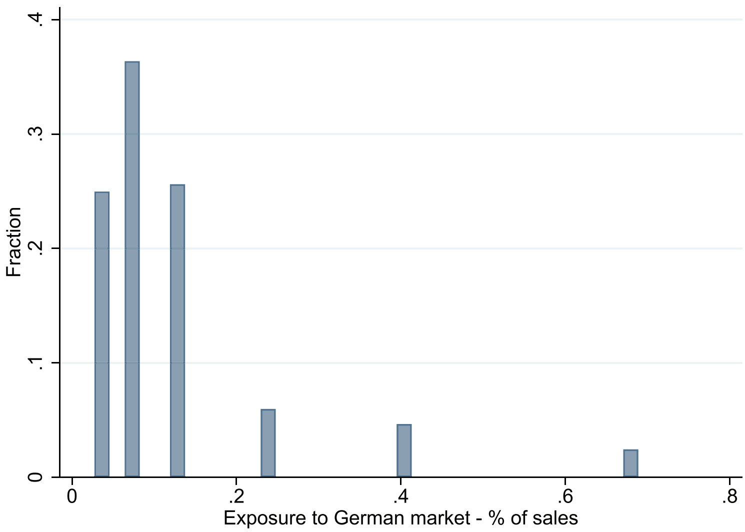

Note: In this graph, firms with zero exposure to the German market are not considered. Banca d'Italia SIGE.

In the empirical analysis, we consider the period between 2014:Q1 and 2019:Q4[14] and exclude firms in the construction sector whose questionnaire does not include questions relating to the German market; non-respondents to those questions belonging to the remaining sectors were dropped.

Additionally, we excluded export-oriented firms that exited the sample before 2019:Q1, since we cannot identify those selling to the German market. At the same time, we keep the firms that are no longer in the sample but declared that they only sell to the domestic market, since they can be univocally classified as part of the control group.

These criteria exclude about 5 percent of the firms from the sample in recent waves (30 percent at the beginning of the sample period; see Table 3). We end up with a sample of about 16,300 observations.

4. The Effect of the German Slowdown: A Microeconometric Approach

4.1 Empirical Strategy

We use a diff-in-diff strategy to analyze the causal link between the German economic slowdown and Italian firms' sentiment and economic behavior.

Following the literature on diff-in-diff estimators (see, among others, Angrist and Pischke 2009; Imbens and Wooldridge 2009), we define the German slowdown as the *treatment*, which can be interpreted as an external shock to exporters to that market. Exporters to Germany thus comprise the treated group (henceforth *treated*), while the control group includes the rest of the sample (non-exporters and exporters to markets different from Germany; henceforth *control*).

As mentioned in the previous section, firms selling to Germany in 2019:Q1 and 2019:Q3 are assumed to have been exporting to that country throughout the whole sample period. According to this definition, the sample is classified as shown in Table 3. The treatment period is set to begin in 2018:Q3, the first quarter after the

[14]Since we observed exports to Germany in 2019 alone, using previous data might be less reasonable. Additionally, we decided to exclude data from 2020:Q1, since the common economic shock of COVID-19 could affect the results.

Table 3. Sample Composition

Quarters	Control		Treated	Not Classified	Total
	Non-exporter	Exporter to Other Countries	Exporter to Germany		
2014:Q1	363	33	193	240	829
2014:Q2	336	35	213	234	818
2014:Q3	327	43	207	226	803
2014:Q4	338	36	204	218	796
2015:Q1	349	43	215	218	825
2015:Q2	348	44	213	199	804
2015:Q3	314	44	223	205	786
2015:Q4	319	35	224	198	776
2016:Q1	331	41	227	192	791
2016:Q2	348	36	225	194	803
2016:Q3	344	43	249	191	827
2016:Q4	340	47	247	176	810
2017:Q1	328	46	252	162	788
2017:Q2	332	52	259	159	802
2017:Q3	348	62	277	134	821
2017:Q4	354	65	294	108	821
2018:Q1	375	80	320	126	901
2018:Q2	345	80	322	104	851
2018:Q3	366	78	326	89	859
2018:Q4	339	77	310	95	821
2019:Q1	365	84	351	42	842
2019:Q2	360	88	331	67	846
2019:Q3	377	92	352	45	866
2019:Q4	445	132	440	50	1,067
Total	8,391	1,416	6,474	3,672	19,953

Note: In this table, the observations used are classified according to their exposure to the external market. The construction sector is excluded from this paper. Banca d'Italia SIGE.

growth of German manufacturing value-added turned negative.[15] We know that different economic aggregates may have a different delay in responding to a similar shock. However, to avoid an arbitrary treatment period for the evaluated series, we chose to initiate

[15] As discussed in Section 1, since the beginning of 2018, some temporary factors have hampered the German economy; however, only after 2018:Q2 did the slowdown in manufacturing become evident and persistent.

the treatment period in the first quarter in which German manufacturing displayed consecutive negative quarter-on-quarter (q-o-q) fluctuations.

The estimated equation is the following:

$$\begin{aligned} y_{it} = {} & \beta GER_i Treat_{t>2018Q2} + \alpha_1 GER_i + \alpha_2 Treat_{t>2018Q2} \\ & + \varphi_t + q_t + \varphi_i + \epsilon_{it}. \end{aligned} \tag{2}$$

In this equation, y_{it} represents the outcome variable that may be affected by the German slowdown, GER_i is a dummy variable identifying the *treated* group (those who export to Germany), $Treat_{t>2018Q2}$ is the post-treatment dummy equal to one during the period of the German slowdown (from 2018:Q3 to 2019:Q4), and φ_i are (vectors of) fixed effects that may vary across specifications.

Since we are using quarterly data, seasonality must be taken into account. For this reason, we control for at least four seasonal dummies (q_t) in each regression.[16] Finally, to control for different cycles at the industry/area level, we interact time dummies with the area/industry ones.

The parameter of interest is β, representing the causal effect of the German slowdown shock on the different outcomes considered. This parameter assumes a particular relevance for the *assessment indicators* since they can be used as proxies for the national account aggregates.

This parameter represents the average causal effect over the period 2018:Q3–2020:Q1. However, depending on the length of exposure to the *treatment* (i.e., the German slowdown), the causal effect may change over time. For this reason, using a dynamic treatment effects model (Callaway and Sant'Anna 2021; Goodman-Bacon 2021), we explore time-varying diff-in-diff effects for a group of variables,[17] in which we estimate the dynamic effects of the treatment for each semester (Jacobson, LaLonde, and Sullivan 1993). The estimated equation changes as follows:

[16] Alternatively, 24 different dummies are used (φ_t, one for each quarter), which bundle trend and seasonal effects together.

[17] Notably total demand, investment plans, and the number of employees.

$$y_{it} = \beta GER_i Treat_{t>2018Q2} + \sum_{h=2018H2}^{2020H1} \beta_h \mathbb{1}_h GER_i$$
$$+ \alpha_1 GER_i + \alpha_2 Treat_{t>2018Q2} + \mu_t + q_t + \mu_i + \epsilon_{it}, \qquad (3)$$

where the causal effect for a given semester h is equal to $\beta + \beta_h$.[18]

Finally, in a robustness exercise, we use the heterogeneity in treatment intensity, namely the exposure to the German market. Using the dose-response function (DRF) approach proposed by Cerulli (2015) based on Hirano, Imbens, and Ridder (2003), we can check whether the firms more exposed to German demand are those that recorded the worse effect.

4.2 Results

For the sake of robustness, we estimate several specifications for each variable of interest, differing as regards time and firm fixed effects. In column 1 (Table 4), we only control for seasonal effects using quarterly dummies, while in the second specification (column 2), we control for the sectors (at the two-digit NACE Rev. 2 level), geographical area (Northwest, Northeast, Center, South), and firm size ("50–200 employees," "200–1,000 employees," and "more than 1,000 employees"); this specification also includes a set of time dummies. In the third specification (column 3), in addition to firm size and geographical area, we control for sector-specific cycles, using ad hoc time-trend-seasonal dummies, while in column 4, we control for different time effects at the geographical level in addition to industry and size fixed effects. In the last two specifications, we use firm fixed effects, only considering quarterly seasonal effects (column 5) or both sector- and area-specific trends (column 6).

In all regressions, standard errors are clustered at the firm level, and the sample weights provided in the data set are used to obtain results referring to the underlying population as a whole. Econometric estimates are supplemented with graphical representations, with a twofold goal: first, to give an intuitive representation of the

[18]Data on 2020:H1 constitute the projection collected for some variables in 2019:Q4, before the COVID-19 disruption.

Table 4. Diff-in-Diff Exercise

	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment Indicators						
SITGEN	–0.123***	–0.131***	–0.175***	–0.118***	–0.112***	–0.174***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
PROMIG	–1.842*	–2.270**	–3.793***	–2.017*	–1.767*	–3.603***
	(0.08)	(0.03)	(0.01)	(0.06)	(0.06)	(0.01)
SITINV	–0.148***	–0.163***	–0.198***	–0.154***	–0.146***	–0.190***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
SITIMP5	–0.0916***	–0.0920***	–0.0938**	–0.0814***	–0.0659**	–0.0703*
	(0.00)	(0.00)	(0.01)	(0.01)	(0.03)	(0.06)
SIMP36M	–0.0657	–0.0645	–0.137**	–0.0595	–0.0445	–0.0955**
	(0.14)	(0.15)	(0.01)	(0.19)	(0.28)	(0.04)
Assessment Indicators						
DOMTOT	–0.314***	–0.301***	–0.282***	–0.292***	–0.314***	–0.239***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
PRETOT	–0.193**	–0.200***	–0.214***	–0.194***	–0.144*	–0.189***
	(0.02)	(0.00)	(0.00)	(0.00)	(0.05)	(0.00)
DOMEST	–0.133***	–0.125**	–0.110*	–0.116**	–0.135***	–0.123**
	(0.01)	(0.01)	(0.06)	(0.02)	(0.01)	(0.03)
PREEST	–0.126	–0.0709	–0.0440	–0.0726	–0.100	–0.0719
	(0.14)	(0.30)	(0.51)	(0.28)	(0.18)	(0.21)
INVPRE	–0.150***	–0.180***	–0.166**	–0.166***	–0.118**	–0.113*
	(0.00)	(0.00)	(0.01)	(0.00)	(0.03)	(0.10)
INVSEM	–0.194***	–0.217***	–0.160***	–0.204***	–0.178***	–0.140**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.02)
OCCTOT	–0.0805***	–0.0899***	–0.0805**	–0.0901***	–0.0592**	–0.0486
	(0.00)	(0.00)	(0.02)	(0.00)	(0.05)	(0.21)
Uncertainty Measures: Three Months Ahead						
V_A	–0.00489	–0.00649	–0.0133	–0.00951	–0.00381	–0.00755
	(0.75)	(0.68)	(0.49)	(0.56)	(0.81)	(0.70)
$E(\sigma_i^2)$	–0.00980	–0.0159	–0.0191	–0.0169	–0.00604	–0.00880
	(0.34)	(0.13)	(0.18)	(0.11)	(0.51)	(0.47)
$V(\mu_{it})$	0.00491	0.00941	0.00580	0.00744	0.00224	0.00125
	(0.71)	(0.49)	(0.73)	(0.59)	(0.87)	(0.94)
μ_{it}	–0.0537***	–0.0559***	–0.0677***	–0.0507**	–0.0350*	–0.0499**
	(0.01)	(0.00)	(0.00)	(0.01)	(0.08)	(0.03)

(*continued*)

Table 4. (Continued)

	(1)	(2)	(3)	(4)	(5)	(6)
Uncertainty Measures: Three Years Ahead						
V_A	0.0232	0.0265	0.0342*	0.0224	0.0295*	0.0419**
	(0.16)	(0.12)	(0.07)	(0.19)	(0.07)	(0.02)
$E(\sigma_i^2)$	–0.0213*	–0.0269**	–0.0331**	–0.0303**	–0.0176*	–0.0181
	(0.06)	(0.02)	(0.03)	(0.01)	(0.08)	(0.12)
$V(\mu_{it})$	0.0445***	0.0533***	0.0673***	0.0527***	0.0471***	0.0601***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
μ_{it}	–0.0419	–0.0364	–0.0921***	–0.0332	–0.0336	–0.0712***
	(0.11)	(0.17)	(0.00)	(0.22)	(0.18)	(0.01)
FE						
Quarter	X				X	
Time		X				
Time × Industry			X			X
Time × Area				X		X
Firm					X	X
Industry		X		X		
Area		X	X			
Size		X	X	X		
N	15,891	14,517	14,322	14,517	15,507	13,948

Note: p-values in parentheses. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$. Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

impact, and second, to show that the common trend assumption is fulfilled.[19]

The results show that the German slowdown adversely affected Italian firms' sentiment and economic choices. The worsening is considerable for firms that export to Germany. The effects are, in most cases, statistically and economically significant. The results are shown in Table 4, where each parameter is estimated in a different diff-in-diff regression. (Table 5 provides more details on the variables used in the diff-in-diff exercise.)

[19]The graphs represent the averages for the seasonally adjusted variables belonging to the three groups. These variables usually have a range of responses between -1 and 1, where zero represents a neutral response. For some questions, to guarantee the possibility of distinguishing both the direction and the magnitude of the variation, the range is set between -2 and 2.

Table 5. Variables Used in the Diff-in-Diff Exercise

Acronyms		Question
		Sentiment Indicators
SITGEN	B2	Compared with three months ago, do you think Italy's general economic situation is ...?
PROMIG	B3	What do you think the probability is of Italy's general economic situation improving in the next three months?
SITINV	C7	Compared with three months ago, do you think conditions for investment are ...?
SITIMP5	C1	What do you think business conditions for your company will be like in the next three months?
SIMP36M	C2	What do you think business conditions for your company will be like in the next three years?
		Assessment Indicators
DOMTOT	C9	Change in demand for residential buildings compared with three months ago...?
PRETOT	C10	How will the total demand for your products vary in the next three months?
DOMEST	C11	Compared with three months ago, is the foreign demand for your products...?
PREEST	C12	How will the foreign demand for your products vary in the next three months?
INVPRE	F1	What do you expect nominal expenditure will be on (tangible and intangible) fixed investment in YEAR compared with that in YEAR?
INVSEM	F2	What do you expect nominal expenditure will be in the first half of YEAR compared with that in the second half of YEAR?
OCCTOT	E1	Your firm's total number of employees in the next three months will be...
		Uncertainty Measures
V_A		Aggregate Uncertainty
$V(\mu_{it})$		Disagreement
$E(\sigma_i^2)$		Average Individual Uncertainty
μ_{it}		Point Forecast

Figure 9. Sentiment Indicators

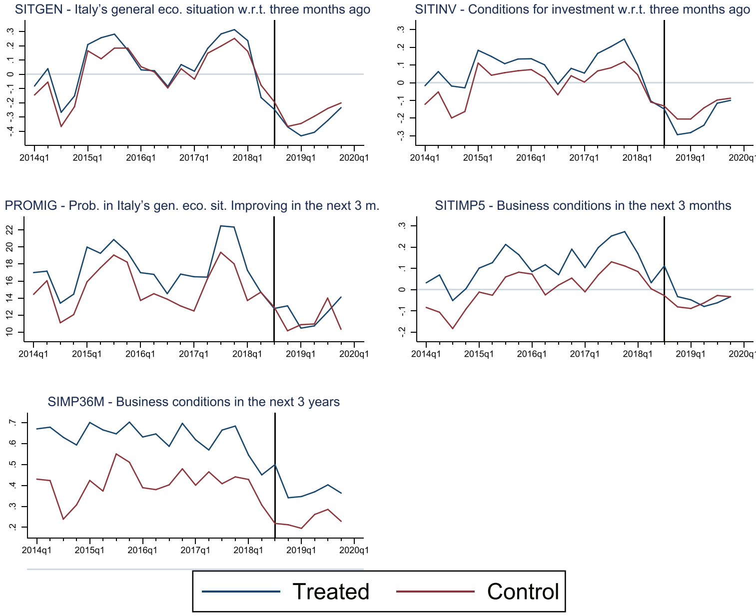

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

4.2.1 Sentiment Indicators

In the pre-treatment period, exporters to Germany had a very similar perception of Italy's current situation compared with that of the other firms (Figure 9). After the treatment, the former group's opinions became markedly worse, with the balance between expectations of improvement and worsening being lower by about 12 percentage points (see SITGEN in Table 4).

Concerning the probability of an improvement in Italy's general economic situation in the following three months,[20] the average for the replies of firms exporting to Germany before the treatment was

[20] For this question, firms can choose between different ranges of probability; we assign each firm the median value of the range chosen. Unlike the other questions, in this case, the results are in terms of probability points instead of balance.

higher than that of the other firms by about 3 points. This difference declined by about 2 points after the treatment (see PROMIG in Table 4 and Figure 9). Finally, when focusing on the opinions about the conditions for investing, while treated firms had a better assessment than the control group before 2019, the roles were reversed after the German slowdown (see SITINV in Table 4 and Figure 9). In this case, the negative effect is highly significant from both a statistical and an economic viewpoint: the balance between expectations of an improvement and a deterioration in conditions for investing is 14 points worse for treated firms with respect to the pre-treatment period; namely, the share of firms in favor of deterioration was greater than those in favor of improvement by about 14 percentage points. Focusing on firms' sentiments about their business situation, exporters to Germany are relatively more optimistic about their medium-run outlook than the short-term one, historically speaking. The German slowdown had a negative impact, particularly on the short-run opinions. Among treated firms, the (weighted balance of the) sentiment regarding their expected situation in the following three months is lower by about 8 percentage points (see SITIMP5 in Table 4 and Figure 9). Instead, no effect is found for the sentiment regarding the medium run (see SITIMP36M in Table 4 and Figure 9).

4.2.2 Assessment Indicators

The impact of the German slowdown is evident and significant for the variables included in the *assessment indicators*, namely those that track national accounts measures well.

With regard to firms' total current demand for their products, after the treatment, the opinions of the affected firms (those exporting to Germany) worsened significantly more than those of the firms in the control group, with a negative impact of approximately 30 points (as shown by DOMTOT in Table 4 and Figure 10). Weaker results were found for the expected demand in the next three months (with an average decrease of 20 points; PRETOT).

The German slowdown has hit total demand significantly since 2018:H2; the effect became greater in 2019 (see Table 6 and Figure 11).

Figure 10. Assessment Indicators

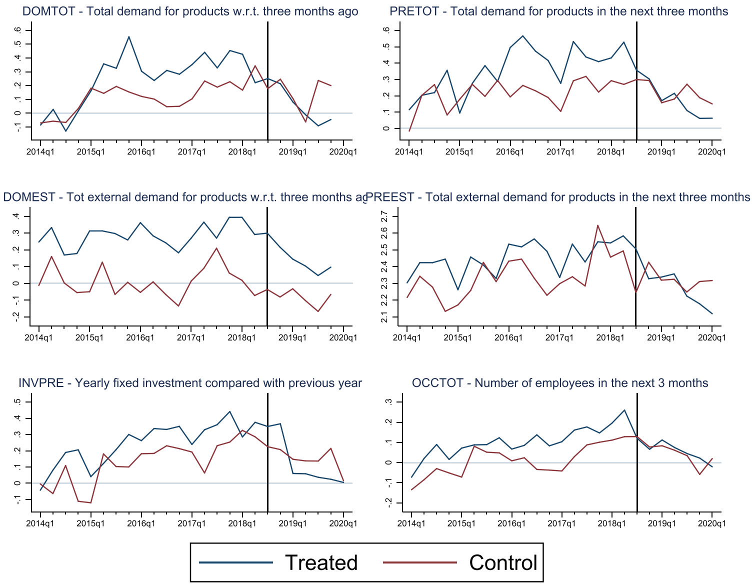

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D. When the questions refer to projections, the balances are plotted over the forecast period; for this reason, in some graphs there is one more observation than in the others.

The impact is also significant for the opinions relating to external demand: the negative effect of the German slowdown on external demand is negative and significant, amounting to approximately 13 points (DOMEST). However, there is no statistically significant evidence of an impact on expected external demand (PREEST).

The effect on firms' investment plans for the current year is also sizable. Before 2019 the balance for exporters to Germany was higher on average by about 14 points; this gap turned negative after the treatment (–15 points on average) across all specifications (see INVPRE in Table 4 and Figure 10). Similar results are found for the capital accumulation planned for the current semester (INVSEM). In this particular case, the effect seems to be significant from 2019:H1

Table 6. Total Demand, Dynamic Effects

	(1)	(2)	(3)	(4)	(5)	(6)
2018:H2–2019:H2	–0.314***	–0.301***	–0.282***	–0.292***	–0.314***	–0.239***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
2018:H2	–0.211***	–0.241***	–0.289***	–0.230***	–0.214***	–0.253***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
2019:H1	–0.327***	–0.339***	–0.274***	–0.324***	–0.323***	–0.212***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
2019:H2	–0.390***	–0.318***	–0.283***	–0.319***	–0.399***	–0.252***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
FE						
Quarter	X				X	
Time		X				
Time × Industry			X			X
Time × Area				X		X
Firm					X	X
Industry		X		X		
Area		X	X			
Size		X	X	X		
N	16,053	14,681	14,478	14,681	15,665	14,099

Note: p-values in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 11. Total Demand, Dynamic Effects

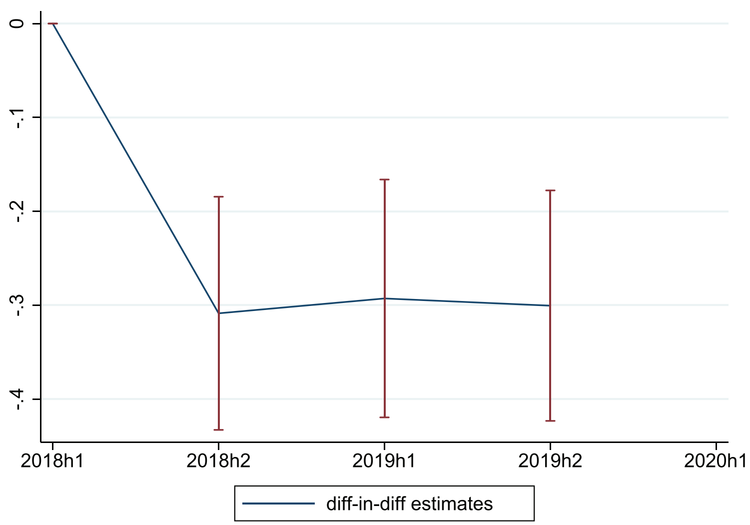

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

Table 7. Investment Plans, Dynamic Effects

	(1)	(2)	(3)	(4)	(5)	(6)
2018:H2–2019:H2	–0.150***	–0.180***	–0.166**	–0.166***	–0.118**	–0.113*
	(0.00)	(0.00)	(0.01)	(0.00)	(0.03)	(0.10)
2018:H2	0.061	–0.010	–0.105	0.001	0.105	–0.0721
	(0.33)	(0.88)	(0.20)	(0.99)	(0.11)	(0.41)
2019:H1	–0.200***	–0.250***	–0.178*	–0.243***	–0.173**	–0.150
	(0.00)	(0.00)	(0.06)	(0.00)	(0.01)	(0.12)
2019:H2	–0.275***	–0.285***	–0.232**	–0.264***	–0.256***	–0.179*
	(0.00)	(0.00)	(0.02)	(0.00)	(0.00)	(0.08)
2020:H1	–0.197***	–0.150*	–0.127	–0.131	–0.174***	0.0372
	(0.00)	(0.06)	(0.22)	(0.11)	(0.02)	(0.76)
FE						
Quarter	X				X	
Time		X				
Time × Industry			X			X
Time × Area				X		X
Firm					X	X
Industry		X		X		
Area		X	X			
Size		X	X	X		
N	16,616	15,196	14,989	15,196	16,235	14,616

Note: p-values in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

onwards and should be weakly significant from 2020:H1 (see Table 7 and Figure 12).[21]

The intention to hire new workers in the next three months also decreased more for treated firms by about 8 points (see OCCTOT in Table 4 and Figure 10). The causal effect on the intention to hire seems negative from 2018:H2 onwards; however, it became significant from 2019:H2 and, according to firms' expectations, it should be greater in 2020:H1.

Taking into account that these variables are reliable proxies for the corresponding national account aggregates (henceforth target variables), these results appear particularly important, suggesting that the German slowdown had a (contemporaneous) impact

[21] In addition, in this case, the effects on 2020:H1 are those relating to plans declared in 2019:Q4.

Figure 12. Investment Plans, Dynamic Effects

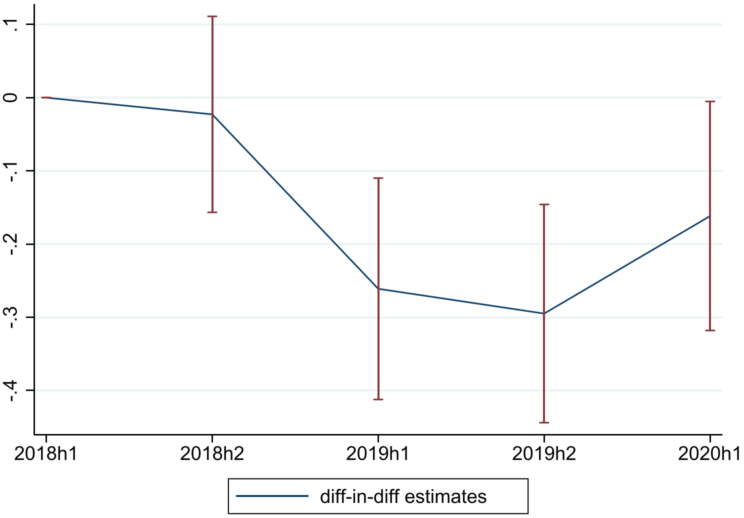

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

on total demand and (lagged) for investment plans and intention to hire.

4.2.3 Uncertainty Measures

According to the measures proposed in Section 3, firms' points forecast are historically higher for companies that export to Germany (Figure 13). At the same time, exporters to Germany are characterized by a higher level of individual uncertainty since, on average, they have a forecast distribution with fatter tails.

According to our model, the treated group reduced their short-term point forecast by about 0.05 points (μ_{it}; see Table 4 and Figure 13). The treatment seems to have no effect on individual uncertainty ($E(\sigma_i^2)$) in the short run, probably because the treated group had slightly more conservative expectations in favor of economic stability during the treatment period. We do not find any effect on disagreement on total uncertainty (V_A).

Figure 13. Uncertainty Measures, Short Term (Three Months Ahead)

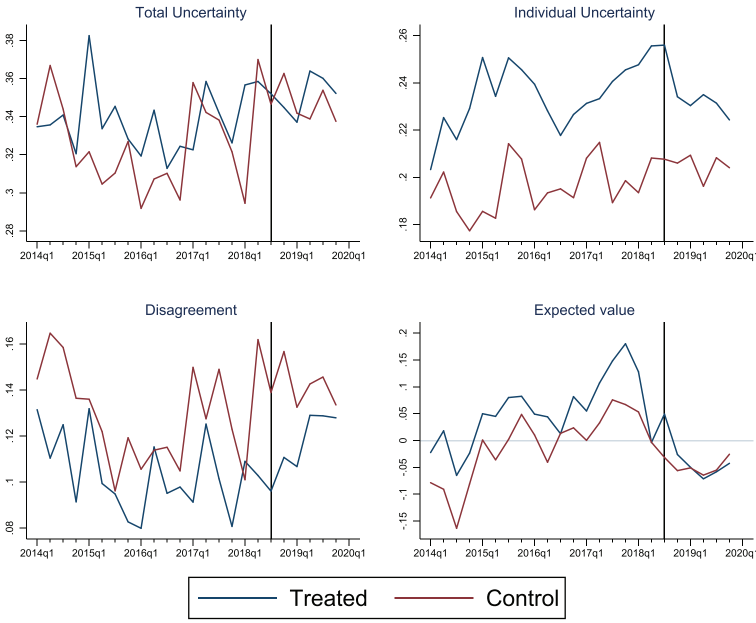

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

In contrast, looking at the medium run (three years ahead), the causal effect on total uncertainty is weakly positive (see Table 4 and Figure 14). Although individual uncertainty seems to be negatively affected by the German slowdown (namely, the treated group becomes less uncertain with respect to the control one), disagreement ($V(\mu_{it})$) within the treated group increased after the German slowdown, representing the main contribution to the increment of total uncertainty. The effect on individual forecasts seems to be very weak.

4.3 Robustness

Since the share of the sample excluded by the analysis is greater for quarters further back in the past (see Section 3 and Table 3),

Figure 14. Uncertainty Measures, Medium Term (Three Years Ahead)

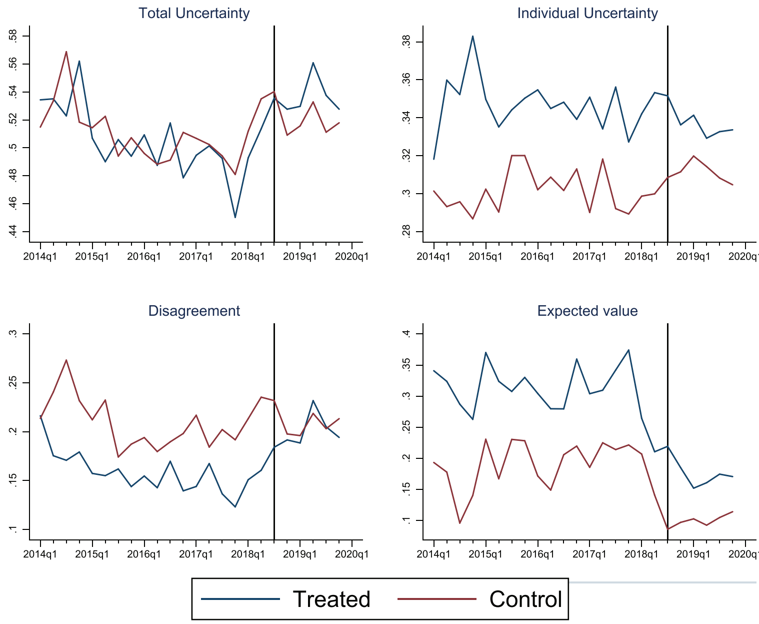

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

the results could be affected by a selection bias problem. To address this issue, we propose two different robust regressions: (i) we use a symmetric pre- and post-treatment period considering only the last 12 quarters (2017:Q1–2019:Q4; see column 2 of Table 8); and (ii) we only consider one balanced panel since 2016:Q1 (see column 3 of Table 8). Finally, we propose an additional specification considering both the symmetric period and the balanced panel (column 4). The results are confirmed in all three cases, suggesting that they are also robust for the selection process in the data.[22]

[22]In these specifications, several observations are dropped, suggesting that there is a trade-off between robustness and representativeness.

Table 8. Robustness Exercise

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sentiment Indicators							
SITGEN	–0.131***	–0.122***	–0.149**	–0.137*	–0.131***	–0.138***	–0.033
	(0.00)	(0.00)	(0.04)	(0.08)	(0.00)	(0.00)	(0.61)
PROMIG	–2.270**	–2.281**	–4.852**	–4.227**	–2.27**	–2.293**	–0.174
	(0.03)	(0.02)	(0.02)	(0.04)	(0.02)	(0.03)	(0.92)
SITINV	–0.163***	–0.152***	–0.162***	–0.164***	–0.163***	–0.182***	–0.074
	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.12)
SITIMP5	–0.092***	–0.093***	–0.094*	–0.105*	–0.092***	–0.100***	–0.046
	(0.00)	(0.00)	(0.09)	(0.09)	(0.00)	(0.00)	(0.28)
SIMP36M	–0.065	–0.016	–0.025	–0.037	–0.065	–0.068	–0.061
	(0.15)	(0.71)	(0.67)	(0.60)	(0.15)	(0.15)	(0.41)
Assessment Indicators							
DOMTOT	–0.301***	–0.320***	–0.300***	–0.230***	–0.304***	–0.309***	–0.067
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.45)
PRETOT	–0.200***	–0.208***	–0.136***	–0.137***	–0.207***	–0.208***	–0.055
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.42)
DOMEST	–0.125**	–0.164***	–0.098	–0.128	–0.135***		
	(0.01)	(0.00)	(0.34)	(0.23)	(0.01)		
PREEST	–0.071	–0.057	–0.044	–0.058	–0.073		
	(0.30)	(0.41)	(0.49)	(0.36)	(0.31)		
INVPRE	–0.180***	–0.184***	–0.252**	–0.264**	–0.189***	–0.187***	–0.065
	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.39)
INVSEM	–0.217***	–0.236***	–0.264***	–0.218**	–0.222***	–0.229***	–0.083
	(0.00)	(0.00)	(0.01)	(0.04)	(0.00)	(0.00)	(0.22)
OCCTOT	–0.090***	–0.092***	–0.055	–0.052	–0.094***	–0.099***	–0.065
	(0.00)	(0.00)	(0.26)	(0.32)	(0.00)	(0.00)	(0.13)
Uncertainty Measures: Three Months Ahead							
V_A	–0.006	0.003	–0.045	–0.050*	–0.006	–0.007	–0.006
	(0.68)	(0.84)	(0.14)	(0.10)	(0.73)	(0.68)	(0.79)
$E(\sigma_i^2)$	–0.016	–0.014	–0.005	–0.008	–0.016	–0.015	–0.015
	(0.13)	(0.18)	(0.72)	(0.57)	(0.13)	(0.17)	(0.38)
$V(\mu_{it})$	0.009	0.017	–0.040	–0.042	0.010	0.008	0.009
	(0.49)	(0.25)	(0.12)	(0.12)	(0.45)	(0.56)	(0.64)
μ_{it}	–0.056***	–0.059***	–0.071**	–0.082**	–0.056***	–0.061***	–0.024
	(0.00)	(0.00)	(0.03)	(0.03)	(0.00)	(0.00)	(0.38)
Uncertainty Measures: Three Years Ahead							
V_A	0.026	0.040**	0.032	0.035	0.030*	0.033*	0.034
	(0.12)	(0.01)	(0.23)	(0.19)	(0.09)	(0.07)	(0.17)
$E(\sigma_i^2)$	–0.027**	–0.022*	–0.013	–0.022	–0.027**	–0.022*	0.001
	(0.02)	(0.06)	(0.38)	(0.14)	(0.02)	(0.07)	(0.94)
$V(\mu_{it})$	0.053***	0.062***	0.045	0.056**	0.057***	0.055***	0.032
	(0.00)	(0.00)	(0.11)	(0.03)	(0.00)	(0.00)	(0.17)
μ_{it}	–0.036	–0.022	–0.046	–0.053	–0.041	–0.035	–0.000
	(0.17)	(0.40)	(0.23)	(0.20)	(0.13)	(0.22)	(1.00)

(*continued*)

Table 8. (Continued)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
N	14,517	8,099	3,280	2,456	14,224	13,222	8,733
	Baseline: Specification (2) in Table 4	Symmetric around 2018:Q2 '17:Q1–'19:Q4	Balanced Panel Since 2016:Q1	Symmetric and Balanced '17:Q1–'19:Q4	Excluding Automotive Sector	Excluding Exporter to Other Markets	Falsification Test

Note: p-values in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. As in the baseline specification, all regressions consider industry, area, and size fixed effects; additionally, we include a set of time dummies. Errors are clustered at firm level. Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

As argued in Section 1, a primary cause of the German slowdown was the bottlenecks in the German automotive sector that were significant and probably had spillover effects on the Italian one.[23] To prevent results from being driven by a specific economic issue relating to a particular sector, we exclude the automotive industry from the sample considered.[24] Results confirm previous estimations (see column 5 of Table 8), suggesting that the effect of the German slowdown was not confined to the Italian automotive sector alone but was widespread in the economy as a whole.

As an additional check, we exclude exporters to markets different from Germany from the control group (and then from the entire analysis; see column 6 of Table 8). This should reduce the possibility of the "second-order effect," relating to indirect global value chains, resulting in downward biases. The results are confirmed in this case too. Finally, to address the same issue, we propose a falsification test excluding exporters to Germany from the analysis. In this case, we designate the exporters to a country other than Germany as a treated group, while the control group is composed of firms that do not export. In this specification, we test the presence of a "secondary effect." Results suggest the irrelevance of this effect:

[23]The automotive sector in Italy accounts for about 4.3 percent of the IP index (of which 2.5 percent is component production). A considerable amount of (automotive component) producers export to Germany. Unlike before, during 2018, the German automotive cycle returned to leading the Italian one, supporting this hypothesis.

[24]We exclude firms belonging to the NACE two-digit 29 and 30 classifications.

Figure 15. HICP

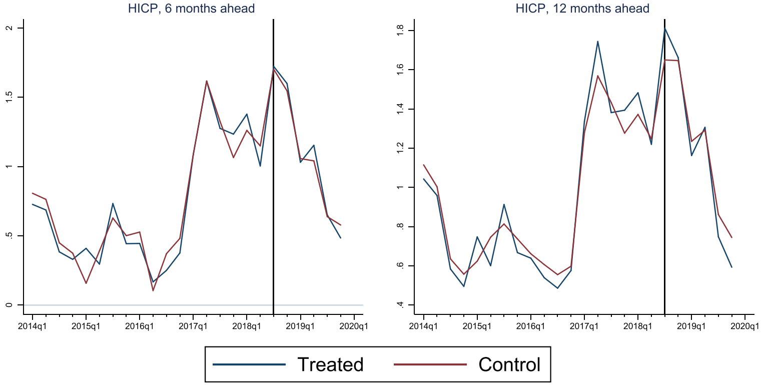

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

the magnitude of the estimates is negligible, compared with the reference estimation, and not statistically significant (see column 7 of Table 8).

To ensure that both groups are comparable in their exposure to the German market, we test the effect of the treatment by applying the same modeling strategy to the firms' 6- and 12-month-ahead expectations for the year-on-year growth in the Italian Harmonised Index of Consumer Prices (HICP). These variables should not be affected by the German slowdown, as there is no reason for exporters to Germany to have different expectations for the Italian HICP due to their nominal nature. This implies that both groups should have similar expectations in both pre- and post-treatment periods.

The findings support this hypothesis: in neither case does the treatment have an effect (as shown in Figure 15 and Table 9). The expectations are roughly the same for both groups, both pre- and post-treatment, indicating that the two groups are comparable, except for their exposure to the German economic outlook.

Finally, the last robustness exercise tests how heterogeneity in treatment among exporters affects a firm's performance: we hypothesize that the firms most exposed to the German market should record a greater negative effect.

Table 9. Effects on HICP

	(1)	(2)	(3)	(4)	(5)	(6)
HICP 6 Months	0.0173	–0.0215	–0.0286	–0.0216	–0.0361	0.0288
	(0.81)	(0.66)	(0.47)	(0.65)	(0.65)	(0.53)
HICP 12 Months	–0.0301	–0.0392	–0.0782	–0.0250	–0.0191	0.0268
	(0.72)	(0.50)	(0.11)	(0.66)	(0.82)	(0.61)
FE						
Quarter	X				X	
Time		X				
Time × Industry			X			X
Time × Area				X		X
Firm					X	X
Industry		X		X		
Area		X	X			
Size		X	X	X		
N	10,184	9,262	9,000	9,262	9,852	8,665

Note: p-values in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We test this hypothesis using the dose-response function approach with a third-order polynomial approximation.[25] The results support a negative relationship between the level of (treatment) exposure and the decline of sentiment indicators during the slowdown period. This means that firms with higher exposure experienced a more significant drop in demand and a stronger negative impact on investment decisions and future employment (as shown in Figure 16).[26]

In our view, this negative relationship between treatment intensity and causal effect is an additional finding that confirms our main argument, meaning that the German cycle is relevant and affects the Italian one.

[25] We use the Stata command *ctreatreg* proposed by Cerulli (2015), which estimates the causal effect according to treatment dose, namely the presence of heterogeneity treatment among the affected firms.

[26] Unfortunately, in our sample, only about 2.5 percent of firms export more than 40 percent of their total production to Germany; for this reason, the confidence interval becomes larger when there is a high degree of treatment.

Figure 16. Dose-Response Function Approach

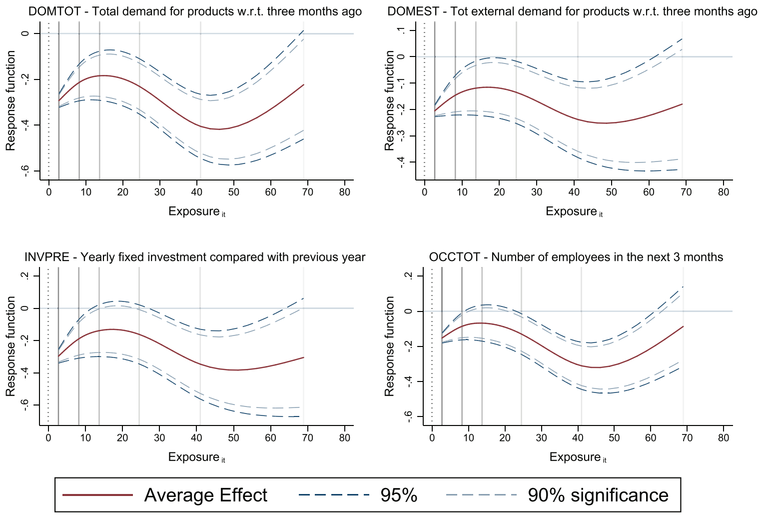

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

5. A Proposal for the Macroeconometric Quantification

The SIGE *assessment indicators* (henceforth proxies) track some national account economic aggregates very well (GDP, GFI, and employment growth rates; henceforth target variables). Additionally, these proxies seem to have good out-of-sample forecasting accuracy for the corresponding target variables (for more details, see Appendix C; on the same argument, see, among others, Lahiri and Monokroussos 2013; Milani 2017).

Economic theory justifies these properties by using two different arguments: (i) the *"animal spirits" view* posits autonomous fluctuations in beliefs that, in turn, have causal effects on economic activity (Blanchard 1993; Hall 1993) and (ii) the *information view* points out that confidence measures contain essential information about the current and future states of the economy (Beaudry and Portier 2004, 2014; Barsky and Sims 2012). This paper focuses on the ability of these variables to mimic economic activity rather than analyzing the relevance of one point of view to the other.

Let's define y_t as the growth rate of the target variables and $\hat{Y}_t$ as the value predicted by the forecasting model, using the corresponding SIGE balances (proxy), as regressors.

$$\hat{Y}_t = \hat{\gamma} B_{Tot,t} + \hat{\alpha}_1 y_{t-1} \tag{4}$$

To quantify the economic loss (in terms of GDP, GFI, and employment) relating to the German slowdown, we calculate the unobserved counterfactual dynamics of proxy variables ($B^{UC}_{Tot,t}$; see Figure 17) and remove the effect of the German slowdown on the Italian economy (Angrist and Pischke 2009), as explained in Appendix C. The resulting counterfactual proxy variables show a more positive trend than the actual data (B_{Tot}), indicating that the economic shock in Germany had an impact on the Italian economy.

By incorporating these counterfactual balances into our forecasting model, we estimate the target aggregates ($Y_t^{\hat{U}C}$) without the effect of the German slowdown.

$$\hat{Y_t}^{UC} = \hat{\gamma} B^{UC}_{Tot,t}(\beta) + \hat{\alpha}_1 y_{t-1}, \tag{5}$$

where $\hat{\gamma}$ and $\hat{\alpha}_1$ are the parameters estimated according to the model selected in Appendix C that maximize the one-step-ahead out-of-sample accuracy.[27]

Then, we estimate the effect (E_t) of the German slowdown on the Italian economy as the difference between the growth rate, predicted by the model ($\hat{Y}_t$) using the real balances, and that ($\hat{Y_t}^{UC}$) obtained using the counterfactual proxies ($B^{UC}_{Tot,t}$) as a regressor.

$$E_t = \hat{Y}_t - \hat{Y_t}^{UC} \tag{6}$$

In Table 10, for each target variable, we show the real growth rate according to the national accounts data (y_t), the growth rate from our forecasting model ($\hat{Y}_t$), and, finally, the one estimated as the counterfactual measure ($\hat{Y_t}^{UC}$).

[27]In Equation (4), we use an ARX(1) model since the linear model is a particular case with $\alpha_1 = 0$. However, to quantify the effect, we use the best model chosen according to Appendix C.

Figure 17. Actual vs. Counterfactual Balances and National Account Aggregates

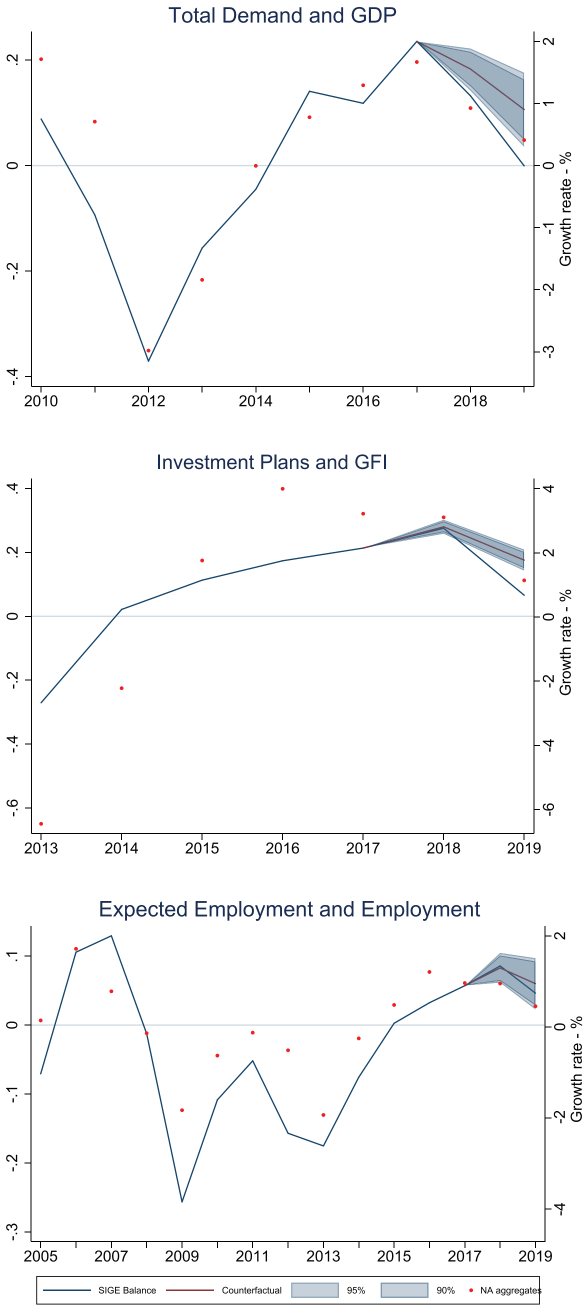

Note: Our calculations based on Banca d'Italia's SIGE survey.

Table 10. Estimated Effect

Variable	**Year**	y_t	$\widehat{Y}_t$	$\widehat{Y}_t^{UC}$				$E_t = \widehat{Y}_t - \widehat{Y}_t^{UC}$			
				2.5%	**Mean**	**Median**	**97.5%**	**2.5%**	**Mean**	**Median**	**97.5%**
GDP	2018	0.94	0.95	0.98	1.11	1.11	1.24	–0.03	–0.16	–0.17	–0.30
	2019	0.29	0.34	0.47	1.18	1.17	1.90	–0.13	–0.84	–0.83	–1.56
Investments	2018	3.11	4.17	4.07	4.55	4.55	5.03	0.11	–0.37	–0.37	–0.86
	2019	1.14	1.79	2.57	4.31	4.31	6.1	–0.78	–2.52	–2.52	–4.31
Employment	2018	0.95	1.4	1.33	1.4	1.4	1.47	0.07	0.00	0.00	–0.07
	2019	0.45	1.04	0.74	1.09	1.09	1.45	0.30	–0.06	–0.05	–0.41

As the latter itself depends on an estimation procedure, we propose a confidence interval.[28] Finally, we compute the effect E_t for the Italian economy deriving from the German slowdown based on Equation (6). We propose an average effect and the relative confidence interval in this case.

According to our estimates, the impact of the German slowdown may have been negative for Italian GDP growth by about 0.2 and 0.8 percentage point in 2018 and 2019, respectively, signaling that the German slowdown was immediate and significant for Italian GDP.

The effect on investment decisions may have been delayed: we do not find any significant effect on investment in 2018,[29] while the impact may have been about 2.5 percentage points in 2019.

Finally, we do not find any statistically significant effect on employment decisions, in line with the results shown in Table 11 and Figure 18, which only predict a significant effect for 2020:H1.[30]

6. Conclusions

The novelty of this work is twofold: (i) we study a macroeconomic issue using both micro- and macrotechniques, specifically by combining policy evaluation techniques with forecasting methods; and (ii) we show a transmission channel from the German cycle to the Italian one.

We investigate to what extent the German economic slowdown that occurred in 2018:Q2–2019:Q4 affected Italian firms using a diff-in-diff strategy, based on microdata from the Survey of Inflation and Growth Expectations, collected quarterly by Banca d'Italia. In particular, we study whether that external shock affected firms' opinions about the general Italian economic situation, their business situation, and their expectations for accumulation, hiring, and demand, which are good predictors of the corresponding national account aggregates. We find that since late 2018, the developments in the sentiment and assessment indicators, particularly for the short term,

[28] To obtain this measure, we use the confidence interval at 5 percent used in Figure 17 as input for our model.

[29] The related confidence interval includes 0.

[30] As explained before, the results for 2020 are based on the assessment collected in 2019:Q4. We do not quantify the effect for 2020.

Table 11. Intention to Hire, Dynamic Effects

	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
2018:H2–2019:H2	–0.0899***	–0.0809***	–0.0805**	–0.0901***	–0.0596**	–0.0486
	(0.00)	(0.00)	(0.02)	(0.00)	(0.04)	(0.21)
2018:H2	–0.0458	–0.0580	–0.0911*	–0.0514	–0.0142	–0.0495
	(0.17)	(0.12)	(0.06)	(0.17)	(0.67)	(0.31)
2019:H1	–0.048	–0.065	–0.050	–0.073	–0.032	–0.034
	(0.18)	(0.11)	(0.30)	(0.07)	(0.38)	(0.49)
2019:H2	–0.116***	–0.110***	–0.0829	–0.114***	–0.0978***	–0.0553
	(0.00)	(0.00)	(0.11)	(0.00)	(0.00)	(0.23)
2020:H1	–0.137***	–0.153***	–0.114*	–0.144***	–0.138***	–0.067
	(0.00)	(0.00)	(0.07)	(0.00)	(0.00)	(0.36)
FE						
Quarter	X				X	
Time		X				
Time × Industry			X			X
Time × Area				X		X
Firm					X	X
Industry		X		X		
Area		X	X			
Size		X	X	X		
N	16,749	15,317	15,108	15,317	16,367	14,734

Note: p-values in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 18. Intention to Hire, Dynamic Effects

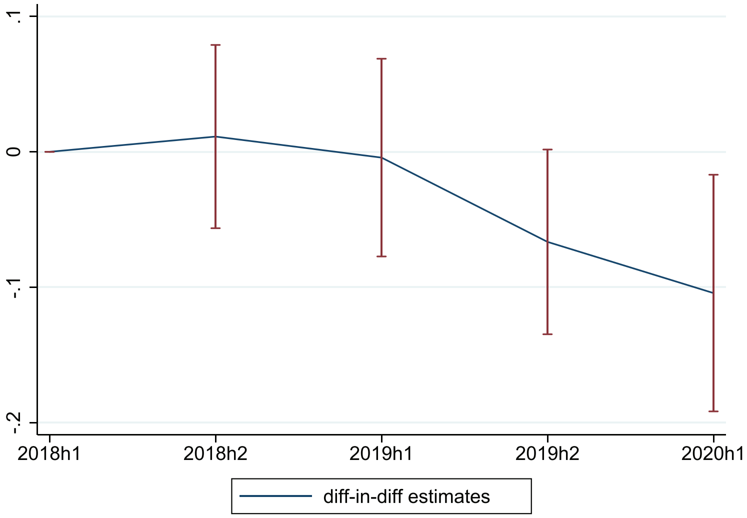

Note: Our calculations based on Banca d'Italia's SIGE survey. For more details on the variables, see Table 5 and Appendix D.

were worse for the Italian companies exposed to the German market; firms' assessments for demand and plans regarding investment and employment were significantly worse as well.

Firms exposed to the German market declared the worst expectations for their activity in the short term (three months ahead); moreover, for the medium term (three years ahead), the German slowdown only slightly affected total uncertainty for the Italian economy because those exporters to Germany disagreed more with each other.

Our results demonstrate that the SIGE series can accurately predict the corresponding national account aggregates (GDP, total investment, and employment). By utilizing the diff-in-diff method, we remove the effect of the German slowdown from the SIGE assessments and obtain the unobserved counterfactual series. By using these series in a forecasting model in a partial equilibrium context, we can estimate the impact of the German slowdown on Italian GDP, investment, and employment growth. By comparing these counterfactual figures with those derived from the actual SIGE balances, we quantify the negative effect of the German slowdown on the Italian economy.

Our findings suggest that the German slowdown had a negative and contemporaneous impact on Italian GDP, estimated at about 1 percentage point over two years (2018–19). The effect appears to have been considerable and delayed for investment but negligible for employment, whose effects are not statistically different from zero for both years (2018–19). These results suggest transmission channels in these two economies where the commercial trade relationship plays an important role. The influence on the business climate could play an additional role by reinforcing the effect through its impact on the expected profit and, by extension, on aggregate investment (Charpe et al. 2016).

Appendix A. Uncertainty Measures

Following Giordani and Soderlind (2003), let's define μ_i as the point forecast of firm i about its future economic condition, namely the firm's expected value based on three possible states. Assuming that its subjective forecast distribution is known, we define a measure of individual uncertainty, which is informative about the distribution probability attached to the different states, as the standard deviation (σ_i) of this forecast distribution.

We compute a simple version of these measures thanks to the SIGE information.

In particular, in each quarter t, the SIGE questionnaire asks about the probability assigned by the firm i to better (p_b), worse (p_w), and unchanged (p_u) business conditions for the next three months and three years.

We assume a payoff scheme (π_j) for each of these three (j) states, in particular

$$\pi_j = \begin{cases} -1 \; with \; probability \; p_w; \\ \;\; 0 \; with \; probability \; p_u; \\ \;\; 1 \; with \; probability \; p_b; \end{cases}$$

Using this information, we define the individual point forecast as

$$\mu_{it} = \sum_{j=w,u,b} p_{ijt}\pi_{ijt} = -1 \cdot p_{iwt} + 0 \cdot p_{iut} + 1 \cdot p_{ibt} = -p_{iwt} + p_{ibt} \tag{A.1}$$

and the individual (forecast) uncertainty as

$$\sigma_{it}^2 = \sum_{j=w,u,p} p_{itj}(\pi_{itj} - \mu_{it})^2. \tag{A.2}$$

The average individual uncertainty $(E(\sigma_t^2))$ across firms contributes to determining a measure of aggregate uncertainty.

According to Giordani and Soderlind (2003), an additional source of uncertainty comes from differences between firms' expectations. In particular, they define disagreement with the variance of the point estimates across firms $(V(\mu_t))$.

Finally, aggregate uncertainty $(V_A(y))$ is equal to the sum of disagreement and the average individual uncertainty:

$$V_A(y) = V(\mu_t) + E(\sigma_t^2). \tag{A.3}$$

Appendix B. Counterfactual Balances

Let's define an aggregated balance B_{Tot} as the weighted average of the balances referring to the three different groups: treated firms (B_{tr}), those in the control group (B_{co}), and those excluded by our analysis (B_{NC}).

$$B_{Tot,t} = w_{tr,t}B_{tr,t} + w_{co,t}B_{co,t} + w_{NC,t}B_{NC,t} \tag{B.1}$$

Let's rewrite Equation (B.1) as

$$B_{Tot,t} = w_{tr,t}\underbrace{(B_{tr,t} - B_{co,t})}_{\alpha_1} + (w_{tr,t} + w_{co,t})B_{co,t} + w_{NC,t}B_{NC,t}. \tag{B.2}$$

Then define the unobserved counterfactual balance $B_{Tot,t}^{UC}$ as the weighted average of balances for the three groups where, for the treated firms, we subtract the time-varying effects as estimated in Section 4.2 from the actual balance.[31]

[31] The effects are estimated for each different semester h. To be conservative, we decided to correct the actual balances using the smallest causal effect estimated in the previous section; namely specification (3) of Table 5 and specification (2) for both investment (Table 6) and employment (Table 10).

$$B^{UC}_{Tot,t}(\beta) = w_{tr,t}(\alpha_1 - \overbrace{(\beta + \beta_h)}^{\text{Causal Effect}}) + (w_{tr,t} + w_{co,t})B_{co,t} + w_{NC,t}B_{NC,t} \tag{B.3}$$

Appendix C. A Simple Forecasting Model

In this appendix, we test the predictive properties of the SIGE balance for the corresponding variables in the national accounts.

We implement this test using two simple models: (i) a simple linear regression, where the SIGE series are regressors; (ii) an ARX(1) model that considers an autoregressive component.

Linear Model	*ARX(1)*
$y_t = \gamma B_{Tot,t} + \epsilon_t$	$y_t = \gamma B_{Tot,t} + \alpha_1 y_{t-1} + \epsilon_t$

These two models are estimated using quarterly[32] and annual data; however, since we focus on the effect over 2018 and 2019, when the quarterly model is used, we aggregate quarterly figures to obtain the annual frequency.

To analyze the forecasting performance, we split the sample into two subperiods and, starting from 2016:Q1, we estimate one-step-ahead (out-of-sample) forecasts. We obtain the relative forecasting performance using both average bias and the mean absolute forecast error (MAFE).

Let's define

$$Bias = \sum_{t=t_0}^{T} \frac{1}{T - t_0} \hat{e}_t = \sum_{t=t_0}^{T} \frac{1}{T - t_0} (y_t - \hat{y}_{(t|t-1)}) \tag{C.1}$$

and

$$MAFE = \sum_{t=t_0}^{T} \frac{1}{T - t_0} \mid \hat{e}_t \mid, \tag{C.2}$$

where y_t is the growth rate of the target variable considered in the forecast exercise and $\hat{y}_{(t|t-1)}$ is the one-step-ahead forecast for time t

[32] Half-yearly data in the case of investment.

Table C.1. Observations Used in the Forecast Exercise

	GDP	Employment	Investment
Period	2010:Q1–2019:Q4	2005:Q1–2019:Q4	2013:H1–2019:H2
Quarterly* Obs. (n)	40	60	14
Annual Obs. (n)	10	15	7

*For investments, we consider half-yearly instead of quarterly data.

computed using the information at time $t-1$; finally, t_0 and T are the first and the last quarters involved in the out-of-sample prediction (2016:Q1 and 2019:Q4, respectively).

Due to the different data availability, the information considered in each model differs for different variables. See Table C.1.

Table C.2 shows statistics on forecast performance for both quarterly and annual growth rates. Since annual models are based on just a few observations, in order to guarantee more robust results, we also aggregate quarterly figures with two different procedures to obtain the annual frequency.

There is no particular advantage to using models with an autoregressive component: the relative coefficient is only statistically different from zero when regressions consider recent quarters, probably due to the procedure used to estimate provisional data.

Additionally, models based on annual data perform better with respect to quarterly models because they display lower volatility for dependent and regressor variables.

In general, models only based on SIGE (proxy) variables perform similarly to those usually used for the short-term forecast.

According to our results, models that minimize both bias and MAFE criteria are linear models based on annual data; however, they only consider a few observations. For this reason, Section 5 uses linear models[33] based on quarterly data and particularly those that quantify the annual figure, and aggregate quarterly data in the standard way.

[33]For reasons of consistency, we chose the same model for all target variables.

Table C.2. Forecast Performance

Growth Rate	Models		Statistics	Linear Model			ARX(1)		
				GDP	***INV****	***EMPL***	***GDP***	***INV****	***EMPL***
q-o-q	Quarterly		MAFE	0.231	1.421	0.338	0.219	1.421	0.344
			Bias	0.162	0.536	0.01	0.156	0.584	0.016
y-o-y	Quarterly	Simple Mean	MAFE	0.464	4.626	0.299	0.425	4.816	0.314
			Bias	0.43	4.626	–0.14	0.407	4.816	–0.115
		Standard	MAFE	0.545	1.626	0.285	0.526	1.649	0.296
			Bias	0.545	0.874	–0.078	0.526	0.965	–0.056
	Annual		MAFE	0.299	1.24	0.322	0.375	1.502	0.335
			Bias	0.282	–0.002	–0.085	0.346	1.502	–0.131

*For investments, we use half-yearly instead of quarterly data.

Note: To obtain annual figures, we aggregate quarterly data by using two different methods: (i) the standard one that uses a different weight for each *q-o-q* growth rate according to their realization during the year; and (ii) the simple average of the *q-o-q* growth rates, to prevent forecast bias from being amplified by the position of the quarters in which it is verified. The red box shows the best models for each NA aggregate based on annual data; the orange one shows those based on quarterly data.

Appendix D. Questionnaire

SURVEY ON INFLATION AND GROWTH EXPECTATIONS
BANCA D'ITALIA

December 2019

Company Name ______________________________

A0. Which is your firm's main sector? |___| SETTONS

- **(1) Manufacturing**
- **(2) Other Industry**
  - Mineral extraction from mines
  - Elettrical. gas. vapour. air conditioning supply
  - Water supply
  - Sewerage, waste management, and redevelopment
- **(3) Trading**
- **(4) Other Servicies**

$\}$ Fill in GREEN questionnaire

- **(5) Construction**
  - Buildings
  - Engineering
  - Special construction works (demolition and preparation of building sites, plant installation, completion and finishing. etc.)

$\}$ Fill in LIGHT BLUE questionnaire

INDUSTRY EXCLUDING CONSTRUCTION AND SERVICES

Firm
Instructions: For percentage changes, indicate the sign in the first box on the left (+ :for increases; —: for decreases).

SECTION A – General Information

A1. Number of employees : |__|__|__|__| ADD

A2. Share of sales revenues coming from exports: |__|
(1= more than 2/3; 2= Between 1/3 and 2/3; 3= Up to 1/3 and more than zero; 4=Zero) EXPORT4

SECTION B – General economic situation of the country

	…in June 2020? IT6	…in December 2020? IT12	…in December 2021? IT24	… on average between December 2022 and December 2024? IT48
B1a. (about 3/5 of the sample) In October consumer price inflation, measured by the 12-month change in the harmonized index of consumer prices was +0.2 per cent in Italy and +0.7 per cent in the euro area. What do you think it will be in Italy...	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%
B1b. (about 1/5 of the sample) What do you think consumer price inflation in Italy, measured by the 12-month change in the harmonized index of consumer prices, will be…	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%
B1c. (about 1/5 of the sample) The European Central Bank has as an objective the maintenance of the 12-month change in the harmonized index of consumer prices in the euro area close but below 2 per cent in the medium term. What do you think consumer price inflation in Italy, measured by the 12-month change in the harmonized index of consumer prices, will be…	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%

B2. Compared with 3 months ago, do you consider Italy's general economic situation is …? SITGEN ☐ Better ☐ The same ☐ Worse

B3. What do you think is the probability of an improvement in Italy's general economic situation **in the next 3 months**? PROMIG
☐ Zero ☐ 1-25 per cent ☐ 26-50 per cent ☐ 51-75 per cent ☐ 76-99 per cent ☐ 100 per cent

SECTION C – Your firm's business conditions

How do you think business conditions for your company will be:
C1. in the **next 3 months**? ☐ Much better ☐ Better ☐ The same ☐ Worse ☐ Much worse SITIMP5

C2. in the **next 3 years**? ☐ Much better ☐ Better ☐ The same ☐ Worse ☐ Much worse SIMP36C5

For each of the above forecasts imagine there are 100 points available; distribute them among the possible forecasts according to the probability assigned to each one. How do you think business conditions for your company will be:

	Better SITM3M SITM3A			The same SITU3M SITU3A			Worse SITP3M SITP3A			Total		
C3. In the next **3 months**										1	0	0
C4. In the next **3 years**										1	0	0

Please indicate whether and with what intensity the following FACTORS will affect your firm's business **in the next 3 months.**

Factors affecting your firm's business In the **next 3 months**	*Effect on business*			*Intensity (if not nil)*		
	Negative	*Nil*	*Negative*	*Nil*	*Negative*	*Nil*
C5. Changes in demand DISIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C6. Changes in your prices PRSIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7. Availability and the cost of credit CRSIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.1 Uncertainty due to econ. and political factors POLIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.2 Exchange rate dynamics TACAM	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.3 Oil price dynamics PRPET	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.4 Tensions on liberalization policies of international trade POLIB	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|

C8. Compared with 3 month ago, do you think conditions for investment are ... ? SITINV ☐ Better ☐ The same ☐ Worse

C9. What do you think your liquidity situation will be in the next 3 months. given the expected change in the conditions of access to credit?
☐ Insufficient ☐ Sufficient ☐ More than sufficient **LIQUID**

C10. Compared with three months ago, is the total demand for your products ... ? **DOMTOT** ☐ Higher ☐ Unchanged ☐ Lower

C11. How will the total demand for your products vary in the next 3 months? **PRETOT** ☐ Increase ☐ No change ☐ Decrease

(Answer to questions C12-C14.1 only if the share of sales revenues coming from exports is positive. otherwise go to C15)

Compared with three months ago, is the foreign demand for your products…?	*Higher*	*Unchanged*	*Lower*	*I do not sell in this market*
C.12 Total **DOMEST**	|__|	|__|	|__|	
C.12.1 In Germany **RTEU_GE**	|__|	|__|	|__|	|__|

(Please answer question C13 only if your answer to question C.12.1 was not 'I do not sell in this market')
C.13 Considering your firm's total exports in 2019, please indicate the share of exports to the German market. |__|
(1= Over 2/3 of turnover; 2= Between 1/3 and 2/3; 3= Up to 1/3 but more than zero; 4=Zero) **EXPGE**

How will the foreign demand for your products vary **in the next 3 months**?	*Increase*	*No change*	*Decrease*	*I do not sell in this market*
C.14 Total **PREEST**	|__|	|__|	|__|	
C.14.1 In Germany **ETEU_GE**	|__|	|__|	|__|	|__|

C15. Compared with three months ago, are credit conditions for your company ...? **SITCRE** ☐ Better ☐ Unchanged ☐ Worse

SECTION D – Changes in your firm's selling prices

D1. In the last 12 months, what has been the average change in your firm's prices? **DPRE** |__| |__|__|.|__|%

D2. For the next 12 months, what do you expect will be the average change in your firm's prices? **DPREZ** |__| |__|__|.|__|%

Please indicate direction and intensity of the following factors as they will affect your firm's selling prices **in the next 12 months**:

	Effect on firm's selling prices			***Intensity*** *(if not nil)*		
Factors affecting your firm's prices in the next 12 months	*Downward*	*Neutral*	*Upward*	*Low*	*Average*	*High*
D2.1. Total demand **DPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D2.2. Raw materials prices **MPPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D2.3. Intermediate Input **IICT**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D2.4. Labour costs **CLPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D2.5. Pricing policies of your firm's main competitors **PRPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D2.6 Exchange rate dynamics **TCPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D2.7 Inflation expectations dynamics **AINF**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D2.8 Financial conditions **CFIN**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|

D3. In the last 12 months, what has been the average change in your firm's prices of goods and services bought in Italy and abroad ?
|__| |__|__|.|__|% **DPRE_INT**

D4. In the next 12 months, what has been the average change in your firm's prices of goods and services bought in Italy and abroad?
|__| |__|__|.|__|% **DPREZ_INT**

SECTION E – Workforce

	Lower	*Unchanged*	*Higher*
E1. Your firm's total number of employees in the next 3 months will be: **OCCTOT**	1|__|	2|__|	3|__|

SEZIONE F – Investment

F1. What do you expect will be the nominal expenditure on (tangible and intangible) fixed investment in 2020 compared with that in 2019?
☐ Much higher ☐ A little higher ☐ About the same ☐ A little lower ☐ Much lower **INVPRE**

F2. And what do you expect will be the nominal expenditure in the first half of 2020 compared with that in the second half of 2019:
☐ Much higher ☐ A little higher ☐ About the same ☐ A little lower ☐ Much lower **INVSEM**

NOTE: The responses "much higher" and "much lower" also apply when. in the two periods compared. investments are zero.

F3. Please rank in order of importance the following SOURCES OF INFORMATION choosing from those that you use the most to obtain financial information to support your business decisions (e.g. production, investments or entry into new markets).

(*Please indicate no more than 3*)

1. Newspapers (paper or online). 2. TV news. 3. Publications by public institutions (e.g. Bank of Italy, Istat or Ministry of the Economy and Finance) and business associations (e.g. Confindustria or Confartigianato). 4. Market consultancy and analysis services provided by private firms. 5. Direct contact with clients and/or suppliers. 6. Social media (e.g. Twitter or Facebook) ☐ **FON1** ☐ **FON2** ☐ **FON3**

CONSTRUCTION

Firm
Instructions: For percentage changes, indicate the sign in the first box on the left (+ :for increases; —: for decreases).

SECTION A – General Information

A1. Number of employees : |__|__|__|__| ADD

A2. Share of sales revenues coming from exports: |__|
(1= more than 2/3; 2= Between 1/3 and 2/3; 3= Up to 1/3 and more than zero; 4=Zero) EXPORT4

A3. Share of revenue from residential building: |__|
(1= more than 2/3; 2= Between 1/3 and 2/3; 3= Up to 1/3 and more than zero; 4=Zero) COMPRES4

SECTION B – General economic situation of the country

	…in June 2020? IT6	…in December 2020? IT12	…in December 2021? IT24	… on average between December 2022 and December 2024? IT48
B1a. (about 3/5 of the sample) In October consumer price inflation, measured by the 12-month change in the harmonized index of consumer prices was +0.2 per cent in Italy and +0.7 per cent in the euro area. What do you think it will be in Italy...	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%
B1b. (about 1/5 of the sample) What do you think consumer price inflation in Italy, measured by the 12-month change in the harmonized index of consumer prices, will be…	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%
B1c. (about 1/5 of the sample) The European Central Bank has as an objective the maintenance of the 12-month change in the harmonized index of consumer prices in the euro area close but below 2 per cent in the medium term. What do you think consumer price inflation in Italy, measured by the 12-month change in the harmonized index of consumer prices, will be…	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%	|__| |__|__|,|__|%

B2. Compared with 3 months ago, do you consider Italy's general economic situation is …? SITGEN ☐ Better ☐ The same ☐ Worse

B3. What do you think is the probability of an improvement in Italy's general economic situation **in the next 3 months**? PROMIG
☐ Zero ☐ 1-25 per cent ☐ 26-50 per cent ☐ 51-75 per cent ☐ 76-99 per cent ☐ 100 per cent

SECTION C – Your firm's business conditions

How do you think business conditions for your company will be:

C1. in the **next 3 months**? ☐ Much better ☐ Better ☐ The same ☐ Worse ☐ Much worse SITIMP5

C2. in the **next 3 years**? ☐ Much better ☐ Better ☐ The same ☐ Worse ☐ Much worse SIMP36C5

For each of the above forecasts imagine there are 100 points available; distribute them among the possible forecasts according to the probability assigned to each one. How do you think business conditions for your company will be:

	Better SITM3M SITM3A	The same SITU3M SITU3A	Worse SITP3M SITP3A	Total
C3. in the next **3 months**				1 0 0
C4. in the next **3 years**				1 0 0

Please indicate whether and with what intensity the following FACTORS will affect your firm's business **in the next 3 months.**

Factors affecting your firm's business	*Effect on business*			*Intensity (if not nil)*		
In the **next 3 months**	*Negative*	*Nil*	*Positive*	*Low*	*Average*	*High*
C5a. Trend in new sites CNSIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C5b. Trend in existing sites CASIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C6. Changes in your prices PRSIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7. Availability and the cost of credit CRSIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.1 Uncertainty due to economic and political factors POLIT	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.2 Exchange rate dynamics TACAM	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.3 Oil prices dynamics PRPET	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
C7.4 Tensions on liberalization policies of international trade POLIB	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|

C8. Compared with 3 month ago, do you think conditions for investment are ... ? SITINV ☐ Better ☐ The same ☐ Worse

C9. What do you think your liquidity situation will be in the next 3 months. given the expected change in the conditions of access to credit?
☐ Insufficient ☐ Sufficient ☐ More than sufficient **LIQUID**

C10. Change in demand for residential building compared with 3 months ago... ? **DOMTOT** ☐ Higher ☐ Unchanged ☐ Lower

C11. How will the total demand for your products vary in the next 3 months? **PRETOT** ☐ Increase ☐ No change ☐ Decrease

(Answer to questions C12-C13 only if the share of sales revenues coming from residential building is positive. otherwise go to C14)

C12. Compared with three months ago, is the demand for residential building... ? **DOMRES** ☐ Higher ☐ Unchanged ☐ Lower

C13. How will the demand for residential building vary **in the next 3 months**? **PRERES** ☐ Increase ☐ No change ☐ Decrease

C14. Compared with three months ago, are credit conditions for your company ...? **SITCRE** ☐ Better ☐ Unchanged ☐ Worse

SECTION D – Changes in your firm's selling prices

D1. In the last 12 months, what has been the average change in your firm's prices? **DPRE** |__| |__|__|.|__|%

D2. For the next 12 months, what do you expect will be the average change in your firm's prices? **DPREZ** |__| |__|__|.|__|%

Please indicate direction and intensity of the following factors as they will affect your firm's selling prices **in the next 12 months**:

Factors affecting your firm's prices in the next 12 months	***Effect on firm's selling prices***			***Intensity*** *(if not nil)*		
	Downward	*Neutral*	Upward	*Low*	*Average*	High
D3. Total demand **DPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D4. Raw materials prices **MPPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D5. Intermediate input **IITC**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D6. Labour costs **CLPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D7. Pricing policies of your firm's main competitors **PRPR**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D8. Inflation expectations dynamics **AINF**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|
D9. Financial conditions **CFIN**	1|__|	2|__|	3|__|	1|__|	2|__|	3|__|

D10. In the last 12 months, what has been the average change in your firm's prices of goods and services bought in Italy and abroad? |_|_|_|,|_|%
DPRE_INT

D11. In the last 12 months, what has been the average change in your firm's prices of goods and services bought in Italy and abroad? |_|_|_|,|_|%
DPREZ_INT

SECTION E – Workforce

	Lower	*Unchanged*	*Higher*
E1. Your firm's total number of employees in the next 3 months will be: **OCCTOT**	1|__|	2|__|	3|__|

SEZIONE F – Investment

F1. What do you expect will be the nominal expenditure on (tangible and intangible) fixed investment in 2019 compared with that in 2018?
☐ Much higher ☐ A little higher ☐ About the same ☐ A little lower ☐ Much lower **INVPRE**

F2. And what do you expect will be the nominal expenditure in the second half of 2019 compared with that in the first half of 2019:
☐ Much higher ☐ A little higher ☐ About the same ☐ A little lower ☐ Much lower **INVSEM**

NOTE: The responses "much higher" and "much lower" also apply when in the two periods compared investments are zero.

F3. Please rank in order of importance the following SOURCES OF INFORMATION choosing from those that you use the most to obtain financial information to support your business decisions (e.g. production, investments or entry into new markets).

(*Please indicate no more than 3*)

1. Newspapers (paper or online). 2. TV news. 3. Publications by public institutions (e.g. Bank of Italy, Istat or Ministry of the Economy and Finance) and business associations (e.g. Confindustria or Confartigianato). 4. Market consultancy and analysis services provided by private firms. 5. Direct contact with clients and/or suppliers. 6. Social media (e.g. Twitter or Facebook) ☐ **FON1** ☐ **FON2** ☐ **FON3**

References

Ahlborn, M., and M. Wortmann. 2018. "The Core Periphery Pattern of European Business Cycles: A Fuzzy Clustering Approach." *Journal of Macroeconomics* 55 (March): 12–27.

Angrist, J., and J.-S. Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Bachmann, R., and P. Zorn. 2020. "What Drives Aggregate Investment? Evidence from German Survey Data." *Journal of Economic Dynamics and Control* 115 (June): Article 103873.

Banca d'Italia. 2019. "Survey on Inflation and Growth Expectations." Methods and Sources: Methodological Notes, January 14.

Barsky, R. B., and E. R. Sims. 2012. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." *American Economic Review* 102 (4): 1343–77.

Bartiloro, L., M. Bottone, and A. Rosolia. 2019. "The Heterogeneity of the Inflation Expectations of Italian Firms along the Business Cycle." *International Journal of Central Banking* 15 (5): 175–205.

Baxter, M., and M. A. Kouparitsas. 2005. "Determinants of Business Cycle Comovement: A Robust Analysis." *Journal of Monetary Economics* 52 (1): 113–57.

Beaudry, P., and F. Portier. 2004. "An Exploration into Pigou's Theory of Cycles." *Journal of Monetary Economics* 51 (6): 1183–1216.

———. 2014. "News-Driven Business Cycles: Insights and Challenges." *Journal of Economic Literature* 52 (4): 993–1074.

Beck, K. 2021. "Why Business Cycles Diverge? Structural Evidence from the European Union." *Journal of Economic Dynamics and Control* 133 (December): Article 104263.

Blanchard, O. J. 1993. "Consumption and the Recession of 1990-1991." *American Economic Review* 83 (2): 270–74.

Bloom, N. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–85.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* 86 (3): 1031–65.

Burstein, A., C. Kurz, and L. Tesar. 2008. "Trade, Production Sharing, and the International Transmission of Business Cycles." *Journal of Monetary Economics* 55 (4): 775–95.

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

Camba-Méndez, G., and M. Forsells. 2018. "The Recent Slowdown in Euro Area Output Growth Reflects Both Cyclical and Temporary Factors." ECB Economic Bulletin No. 4/2018.

Campos, N. F., J. Fidrmuc, and I. Korhonen. 2019. "Business Cycle Synchronisation and Currency Unions: A Review of the Econometric Evidence Using Meta-analysis." *International Review of Financial Analysis* 61 (January): 274–83.

Canova, F., and H. Dellas. 1993. "Trade Interdependence and the International Business Cycle." *Journal of International Economics* 34 (1): 23–47.

Cerulli, G. 2015. "ctreatreg: Command for Fitting Dose-response Models under Exogenous and Endogenous Treatment." *Stata Journal* 15 (4): 1019–45.

Cesaroni, T., and S. Iezzi. 2017. "The Predictive Content of Business Survey Indicators: Evidence from SIGE." *Journal of Business Cycle Research* 13 (1): 75–104.

Charpe, M., C. Chiarella, P. Flaschel, and C. R. Proaño. 2016. "Business Confidence and Macroeconomic Dynamics in a Nonlinear Two-Country Framework with Aggregate Opinion Dynamics." In *Dynamic Modeling, Empirical Macroeconomics, and Finance: Essays in Honor of Willi Semmler,* ed. L. Bernard and U. Nyambuu, 289–310. Cham: Springer International Publishing.

Chiarella, C., P. Flasher, and H. Hung. 2006. "Interacting Business Cycle Fluctuation: A Two-Country Model." *Singapore Economic Review* 51 (03): 365–94.

Clark, T. E., and E. van Wincoop. 2001. "Borders and Business Cycles." *Journal of International Economics* 55 (1): 59–85.

Coibion, O., Y. Gorodnichenko, and T. Ropele. 2020. "Inflation Expectations and Firm Decisions: New Causal Evidence." *Quarterly Journal of Economics* 135 (1): 165–219.

Conflitti, C., and R. Zizza. 2021. "What's Behind Firms' Inflation Forecasts?" *Empirical Economics* 61 (5): 2449–75.

de Lucas Santos, S., and M. J. Delgado Rodríguez. 2016. "Core-Periphery Business Cycle Synchronization in Europe and the Great Recession." *Eastern European Economics* 54 (6): 521–46.

Dellas, H. 1986. "A Real Model of the World Business Cycle." *Journal of International Money and Finance* 5 (3): 381–94.

di Giovanni, J., A. A. Levchenko, and I. Mejean. 2018. "The Micro Origins of International Business-Cycle Comovement." *American Economic Review* 108 (1): 82–108.

Eickmeier, S. 2007. "Business Cycle Transmission from the US to Germany—A Structural Factor Approach." *European Economic Review* 51 (3): 521–51.

Enders, Z., P. Jung, and G. J. Müller. 2013. "Has the Euro Changed the Business Cycle?" *European Economic Review* 59 (April): 189–211.

European Commission. 2019. “European Economic Forecast. Spring 2019.” Institutional Paper No. 102, Economic and Financial Affairs, European Commission.

Ferroni, F., and B. Klaus. 2015. “Euro Area Business Cycles in Turbulent Times: Convergence or Decoupling?” *Applied Economics* 47 (34–35): 3791–3815.

Frankel, J. A., and A. K. Rose. 1997. “Is EMU More Justifiable Ex Post Than Ex Ante?” *European Economic Review* 41 (3): 753–60.

———. 1998. “The Endogenity of the Optimum Currency Area Criteria.” *Economic Journal* 108 (449): 1009–25.

Garnier, J. 2004. “To What Extent Does Intra-industry Trade Matter in Business Cycles Comovements? Distinguishing Common and Transmitted Cycles.” Mimeo.

Giordani, P., and P. Soderlind. 2003. “Inflation Forecast Uncertainty.” *European Economic Review* 47 (6): 1037–59.

Gomez, D. M., H. J. Ferrari, B. Torgler, and G. J. Ortega. 2017. “Synchronization and Diversity in Business Cycles: A Network Analysis of the European Union.” *Applied Economics* 49 (10): 972–86.

Gonçalves, C. E. S., M. Rodrigues, and T. Soares. 2009. “Correlation of Business Cycles in the Euro Zone.” *Economics Letters* 102 (1): 56–58.

Goodman-Bacon, A. 2021. “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics* 225 (2): 254–77.

Grigoraş, V., and I. E. Stanciu. 2016. “New Evidence on the (de)synchronisation of Business Cycles: Reshaping the European Business Cycle.” *International Economics* 147 (October): 27–52.

Guiso, L., and G. Parigi. 1999. “Investment and Demand Uncertainty.” *Quarterly Journal of Economics* 114 (1): 185–227.

Hall, R. E. 1993. “Macro Theory and the Recession of 1990-1991.” *American Economic Review* 83 (2): 275–79.

Hirano, K., G. W. Imbens, and G. Ridder. 2003. “Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score.” *Econometrica* 71 (4): 1161–89.

Imbens, G. W., and J. M. Wooldridge. 2009. “Recent Developments in the Econometrics of Program Evaluation.” *Journal of Economic Literature* 47 (1): 5–86.

Imbs, J. 2004. "Trade, Finance, Specialization, and Synchronization." *Review of Economics and Statistics* 86 (3): 723–34.

Istat and ICE. 2019. "Annuario Statistico — Commercio Estero e Attivit Internazionali delle Imprese." ("Statistical Yearbook—Foreign Trade and International Business Activities.")

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83 (4): 685–709.

Kleinert, J., J. Martin, and F. Toubal. 2015. "The Few Leading the Many: Foreign Affiliates and Business Cycle Comovement." *American Economic Journal: Macroeconomics* 7 (4): 134–59.

Lahiri, K., and G. Monokroussos. 2013. "Nowcasting US GDP: The Role of ISM Business Surveys." *International Journal of Forecasting* 29 (4): 644–58.

Milani, F. 2017. "Sentiment and the U.S. Business Cycle." *Journal of Economic Dynamics and Control* 82 (September): 289–311.

Disagreement and Discretionary Monetary Policy*

Marvin Goodfriend, Pierre Jinghong Liang,
and Gaoqing Zhang
Carnegie Mellon University

This paper identifies a new coordination motive endogenously induced by a central bank's lack of commitment in the presence of information imperfection. We show that when differentially informed economic agents disagree about the central bank's inflation incentives, discretion in monetary policymaking induces agents to coordinate by "forecasting the forecasts of others" in order to forecast the central bank's policy actions. In particular, the induced coordination mechanism compels the central bank to choose monetary policy that responds to fluctuations in the average belief about its incentive. As a result, discretion has the potential to vastly increase fluctuations in employment and inflation, especially when the disagreement among agents is low. More broadly, our paper makes an argument for the inclusion of information diversity among agents in monetary policy discussions and in the characterization of the inflation dynamics.

JEL Codes: E52, E61, D82.

*Marvin Goodfriend passed away in December 2019. He was our colleague, mentor, and friend. We dedicate this paper to his memory and his contributions to the field of economics. Goodfriend was a faculty member of the Tepper School of Business at the Carnegie Mellon University and also had a research affiliation with NBER. Pierre Jinghong Liang and Gaoqing Zhang are from the Tepper School of Business at the Carnegie Mellon University. Pierre Liang gratefully acknowledges the Dean's Summer Research funding of the Tepper School. We thank Boragan Aruoba (editor), John McDermott, Pingyang Gao, Ariel Zetlin-Jones, Xiaodong Zhu, participants at Chicago-Minnesota Theory Conference and HKU Economics seminar, and especially two anonymous IJCB referees, for valuable comments. Author e-mails: liangj@andrew.cmu.edu or gaoqingz@andrew.cmu.edu.

1. Introduction

This paper identifies a new coordination motive endogenously induced by a central bank's lack of commitment in the presence of information imperfection. We show that when differentially informed economic agents disagree about the central bank's inflation incentive, discretion in monetary policymaking induces agents to coordinate by "forecasting the forecasts of others" in order to forecast the central bank's policy actions. We abstract from inherent interdependencies that have been studied in the past to isolate the cause and the effect of the newly identified coordination motive.[1] In particular, our discretion-disagreement coordination mechanism compels the central bank to choose monetary policy that responds to fluctuations in the average belief about its inflation incentive which, in turn, is what forces agents to coordinate by forecasting the forecasts of others. As a result, discretion has the potential to vastly increase fluctuations in employment and inflation, especially when the disagreement among agents is low. More broadly, our paper makes an argument for the inclusion of information diversity among agents in monetary policy discussions and in the characterization of the inflation dynamics.

We adopt two information imperfections involving the central bank's inflation incentive: (i) disagreement among individual agents, and (ii) the average forecast error of all agents. Specifically, agents forecast future inflation incentives imperfectly and asymmetrically with a private signal that contains a common noise and an idiosyncratic noise. The common noise yields a stochastic average forecast error, and the volatility of the idiosyncratic noise governs dispersion of the individual forecasts around the average forecast. Surprisingly, we find that in equilibrium under discretion holding fixed average forecast accuracy, *more agreement* among agents destabilizes employment and inflation. Equally surprising, we find that *more accurate* average forecasts also destabilize employment and

[1] Coordination motives may also arise from inherent interdependencies among actors, either through technology linkage (Angeletos and Pavan 2004, 2007), information extraction (Townsend 1983), monopolistic competition (Woodford 2001), trading (Angeletos and La'O 2013), or beauty-contest preference (Morris and Shin 2002).

inflation in equilibrium under discretion when agents are sufficiently in agreement with each other. In effect, more agreement among agents coordinates forecasts more tightly and magnifies the effect of the common noise on employment and inflation via the central bank's more aggressive reaction to the average forecast. The magnified common noise constitutes an information-based source of macroeconomic instability which has not been identified previously.

The coordination problem we identify exists only under discretionary monetary policy. The problem goes away when the central bank follows a credible rule, even in the presence of imperfect information. Specifically, with commitment, the central bank can unilaterally and uniformly anchor each individual firm's expectation by credibly specifying both current and future policies. As a result, the central bank's equilibrium policy actions would be based upon the predetermined decision rule known to all firms. Although this rule may depend on future shocks to the inflation incentive that are imperfectly known to firms, firms can simply estimate these shocks by constructing first-order beliefs, without necessarily constructing (higher-order) beliefs about others' beliefs.

Without commitment, the central bank's current action loses control of the average current expectation of its future policy actions and, worse, it must react to its assessment of the average expectation. Therefore, when an individual agent forecasts future policy actions, the agent must forecast the average expectation which depends on other agents' forecasts. As a result, the average (first-order) belief of future policy actions now depends on the average forecast of other agents' forecasts, or the average second-order belief. Similarly, the average second-order belief would, in turn, depend on third-order beliefs, and so on.

In equilibrium under discretion, the aggregate variables such as output and inflation are functions of the average forecast of future monetary policy actions; the average forecast is determined by a hierarchy of higher-order beliefs which, in turn, depends on the properties of the average forecast and the degree of disagreement. Recognizing the complexity of the problem, we characterize the effects of higher-order beliefs on aggregate variables in a linear manner within a New Keynesian (New Synthesis) macroeconomic model, solved in closed form with a class of normally distributed signals.

With this specification, equilibrium inflation and aggregate output respond to the average forecast error linearly. The degree of disagreement, among other model parameters, affects the equilibrium sensitivity of inflation and output to the average forecast error because the degree of disagreement affects the aggressiveness with which the central bank, under discretion, is forced to respond to the average forecast error.

Given the equilibrium monetary policy, inflation and output become more volatile due to the addition of the shock on the average forecast. The induced coordination—the cause of heightened macrofluctuations—makes the problem especially pronounced due to a "multiplier" effect of the average forecast error. This is because the same private signal is used for each level (ladder) in the individual higher-order belief hierarchy; averaging across all individuals eliminates the idiosyncratic noise but not the common noise. Thus, the common noise is retained at each higher-order level of the average belief, magnifying the noise contained in the average forecast in equilibrium. When a discretionary central bank reacts to the average forecast, the magnified average forecast error enters into aggregate inflation and output, generating volatility due to the information imperfection beyond those "real" shocks commonly studied such as cost-push or demand shocks.

Facing such a pronounced problem caused by information imperfection, the conventional wisdom would suggest that reducing the volatilities of the information shocks would be desirable. We find that this intuition does not hold generally. Holding fixed the average forecast-error volatility, a higher degree of disagreement among agents makes them less coordinated, as each firm relies less on its signal when forming expectations about future policy actions. The central bank, in turn, becomes less responsive to the average forecast. Thus, fluctuations in employment and inflation due to information shocks are lower with a higher degree of disagreement. So narrowing the degree of disagreement would introduce more economic fluctuations and destabilize inflation and output.

On the other hand, holding the degree of disagreement fixed, an increase in the precision of the average forecast creates a trade-off between a direct reduction in the size of common noise and an indirect increase in sensitivity to the noise. Specifically, an increase in the precision directly reduces the common noise in the central bank's equilibrium monetary policy, leading to less volatility in the output and inflation series holding fixed the central bank's reaction sensitivity. However, when the precision of average forecast increases, all agents are better informed about future inflation and adjust their inflation expectation; the discretionary central bank's reaction to the aggregate expected inflation, which includes the error therein, adds more volatility to equilibrium inflation and output. When the degree of disagreement is high (low), the trade-off favors the direct (indirect) effect. Consequently, reducing the average forecast-error volatility stabilizes employment and inflation in equilibrium only when the degree of disagreement is high enough.

Our model specification borrows key elements from two distinct literatures: (i) macroeconomic research focusing on monetary policy and (ii) information economics research focusing on information structure. Based on the large literature on monetary policy research, we deploy the structural equations summarizing key insights from the New Keynesian model as described in the survey paper by Clarida, Gali, and Gertler (1999). Our model shares with Cukierman and Meltzer (1986), an early work on central bank opacity, a key feature that the central bank inflation incentive is stochastic and perpetually obscured. From the information economic research, we draw from research on information structure by economists as well as accounting researchers. The key element of our information structure—correlated private signals—has been used in the studies of financial markets (Holthausen and Verrecchia 1990) and recently has been studied in coordination settings with inherent interdependencies (Myatt and Wallace 2012; Liang and Zhang 2019).

To appreciate the connection this paper makes, consider the two debates in monetary policy that have received much academic, practical, and policy attention. The rules-versus-discretion debate has a long and varied standing. According to McCallum (1999, p. 1485), a "major reorientation" dates back to Barro and Gordon (1983a, 1983b) "built upon the insights of Kydland and Prescott (1977)." As is well known, the main insight identified by this literature is that discretionary policies suffer from the time-inconsistency problem: the market participants' rational expectation renders discretionary policies, designed by a benevolent central bank, ineffective and, worse, generates unnecessary inflation and economy fluctuations. The transparency–opacity debate in monetary policy can be traced back to Cukierman and Meltzer (1986) and Goodfriend (1986). This debate explicitly considers the potential information asymmetry between the central bank and the market participants (see, e.g., the survey by Geraats 2002). For example, collectively the public may perceive a lack of access into the workings of the central bank in promulgating monetary policy, leading to a perceived opacity of the central bank (e.g., Winkler 2002). Our paper bridges the two debates by identifying the link in between. In this light, our paper is related to that of Morris and Shin (2005), who also point to the connection between central bank discretion and transparency.[2] Interestingly, Morris and Shin (2005) also stress the preeminent role of managing expectations in linking the debate of central bank transparency and monetary policy to the extent that the central bank manipulates market expectations via communication and, at the same time, extracts information from market prices to guide monetary policy. In a sense, our paper complements the insight of Morris and Shin (2005) by outlining an alternative mechanism through which market expectations about the central bank's policy target interact with the monetary policy the central bank sets at its discretion.

Students of central banks have long noted the importance of the disagreement among individuals. For instance, Brunner (1981) studies the disagreement among individual agents' subjective perceptions of the monetary policy. King (1982, 1983) and Dotsey and King (1986) study the informational implication to monetary policy when differentially informed agents extract endogenous information from prices. Outside the two debates on central bank discretion and transparency, the pioneering idea by Phelps (1983) has stressed the lack of common knowledge in explaining the aggregate economic dynamics. The initial work by Townsend (1983) analytically formulated the idea of forecasting the forecasts of others. Woodford (2001), among other recent works, relies on finite information-processing capacity

[2] Morris and Shin (2005, p. 1) articulate a general point about this link from a political economy perspective: "In light of the considerable discretion enjoyed by independent central banks, the standards of accountability that they must meet are perhaps even higher than for most other public institutions. Transparency allows for democratic scrutiny of the central bank and hence is an important precondition for central bank accountability."

(Sims 2003) to show that informational disagreement among individuals leads them to construct beliefs about others' beliefs, or higher-order beliefs, within the contemporary framework of macromodels.

More recently, Benhabib, Wang, and Wen (2015) study an information friction caused by a timing friction: firms must make a production decision before demand is realized, while consumers must make labor supply and consumption plans before production is realized. This timing friction, coupled with aggregate sentiment in consumer demand, gives rise to endogenous aggregate fluctuations without any assumptions on technological externalities, non-convexities, etc. Absent in Benhabib, Wang, and Wen (2015) is the role of a central bank; thus the information friction we study is not the focus of their paper. The monetary policy is present in Paciello and Wiederholt (2014) in which information friction comes from the (in)attention to aggregate conditions paid by the firms, either exogenously or endogenously, when making individual production decisions in an economy described by standard New Keyesian model, the same framework we used. However, Paciello and Wiederholt (2014) consider a monetary policymaker with full commitment capabilities, an assumption we do not make in our paper. In Hellwig and Veldkamp (2009), the focus is on how exogenous coordination incentives affect agents' individual endogenous information acquisition, leading to the idea of "knowing what others know." The agents in our paper do indeed desire to know what others know (i.e., forecast of others' forecasts), but their coordination incentives are induced by a central bank unable to commit, as opposed to being exogenously given in Hellwig and Veldkamp (2009). In addition, the roles of diverse information and higher-order beliefs in large economies are also studied by Angeletos and Lian (2018), Angeletos and La'O (2020), and Angeletos and Huo (2021) in an emerging literature on incomplete information in macroeconomics (see Section 8 of the excellent survey by Angeletos and Lian 2016). However, they focus on dispersed information about the state of the economy as opposed to our focus on the dispersed information about the discretionary policy target of the central bank. Our paper complements this literature by adding a new coordination motive driven by central bank discretion and transparency.

The rest of the paper proceeds as follows. Section 2 lays out the basic macroeconomic framework and the key elements of our information assumptions. Section 3 analyzes the resulting model and constructs the central higher-order-belief arguments. Section 4 analyzes a parameterized version of the model. Section 5 concludes.

2. Model Setup

2.1 A Simple Macroeconomic Framework

The economy is populated with a central bank that takes the nominal interest rate as the instrument of monetary policy, a representative household, and a continuum of firms, indexed by $[0, 1]$. Rather than deriving the optimal conditions for the household and firms, we describe the operation of the economy by a set of structural equations that can be derived from log-linearizing optimal consuming and profit-maximizing conditions (as in Galí 2008). Let y_t and y_t^p denote the logs of the aggregate economy output and the potential level of the output. The potential output is the level of output that would arise if wages and prices were fully flexible, but it may be lower than the efficient level due to existing frictions such as monopolistic competition, taxes, and subsidies. Define the output gap x_t as the difference between y_t and y_t^p:

$$x_t \equiv y_t - y_t^p. \tag{1}$$

In addition, let π_t be the inflation rate from period $t-1$ to t.

First, there is a New Keynesian Phillips curve that links inflation π_t to output gap x_t, generated by the firms in the economy:

$$\pi_t = \lambda x_t + \beta \bar{E}_t^F \pi_{t+1} + u_t, \tag{2}$$

where $\beta \in (0, 1)$ denotes the discounting factor and $\bar{E}_t^F[\cdot]$ denotes the average belief of the firms, i.e., $\bar{E}_t^F[\cdot] = \int_0^1 E_t\left[\cdot | I_t^i\right] di$, with firm i's information set I_t^i.[3] The shock u_t follows

$$u_t = \rho_u u_{t-1} + \hat{u}_t, \tag{3}$$

[3]For simplicity, we assume in our main analysis a reduced form of Phillips curve (2) in line with the standard New Keynesian Phillips curve when information is complete. As noted in Angeletos and Lian (2018) and Angeletos and Huo (2021), with incomplete information, the Phillips curve varies from the standard one. Accordingly, to assess the robustness of our analysis, we analyze in Section 5 a variant of our model with a microfounded Phillips curve.

where $\rho_u \in [0, 1)$ and $\hat{u}_t$ are i.i.d. random variables with mean 0 and variances σ_u^2. The Phillips curve can be derived by profit-maximizing conditions by firms that compete with each other monopolistically and face nominal price rigidities (Calvo 1983; Yun 1996; Woodford 2008). The key feature of the Phillips curve is that the average expected inflation $\bar{E}_t^F [\cdot]$ enters, which creates a role for the beliefs of the firms in affecting equilibrium inflation and output levels. This role, in turn, influences the central bank's monetary policy in equilibrium, making it partially self-fulling.

Second, there is a dynamic "IS" equation that describes the relation between real interest rate and output gap generated by the representative household in the economy:

$$x_t = -\phi r_t + E_t^H x_{t+1} + g_t, \tag{4}$$

where r_t is the real interest rate (from period t to period $t+1$) and $E_t^H [\cdot]$ denotes the expectation by the representative household. g_t is a shock that follows,

$$g_t = \rho_g g_{t-1} + \hat{g}_t, \tag{5}$$

where $\rho_g \in [0, 1)$ and $\hat{g}_t$ are i.i.d. random variables with zero mean and variances σ_g^2. The IS equation can be derived from log-linearizing the Euler equation of the representative household.

Third, a Fisher equation links the nominal interest rate to the real interest rate and the representative household's expected inflation. Let i_t be the nominal interest rate from period t to $t+1$:

$$i_t = r_t + E_t^H \pi_{t+1}. \tag{6}$$

Replacing r_t in the IS equation with $r_t = i_t - E_t \pi_{t+1}$ in the Fisher equation gives a modified IS equation,

$$x_t = -\phi \left(i_t - E_t^H \pi_{t+1} \right) + E_t^H x_{t+1} + g_t. \tag{7}$$

The central bank in period t minimizes deviations of aggregate output gap and inflation from their respective targets:

$$\frac{1}{2} E_t \left\{ \sum_{\tau=0}^{\infty} \beta^\tau \left[\alpha \left(x_{t+\tau} - k_{t+\tau} \right)^2 + \pi_{t+\tau}^2 \right] \right\}, \tag{8}$$

subject to the Phillips curve (2) and the IS curve (7), where α is the relative weight on output deviations. We interpret the loss function as endowing the central bank a dual mandate: a zero-inflation target and output gap target. We assume that the target for the output gap is k_t. Adapting Barro and Gordon (1983a), k_t represents the extent the central bank intends to raise actual output above potential (toward efficient output). For example, $k_t = 0$ implies that the central bank is satisfied with aggregate output at the potential output level (but below the efficient level).[4] When $k_t > 0$, the central bank has an incentive to target actual output above the potential level, generating an incentive to inflate. As is typically done, we will use k_t to represent both higher output target than potential and inflation incentives interchangeably.

Critically, the inflation incentive is thus time varying in this paper, unlike in standard models such as that of Clarida, Gali, and Gertler (1999). In essence, this assumption implies that the central bank knows something more about its own preferred output gap target than the market collectively. Empirically, the assumption is consistent with the facts that financial markets respond to U.S. Federal Reserve (Fed) actions and market participants spend significant effort in Fed watching, a point made by a recent paper by Stein and Sunderam (2018). Further, there is a long literature on central bank secrecy/transparency dating back to the 1981 Supreme Court case favoring the Federal Open Market Committee's (FOMC's) position of delaying the release of meeting minutes, a practice criticized by academics such as Goodfriend (1986). While Fed transparency and accountability has increased after many years, some ambiguity and flexibility (about its own internal policy targets) remain in the process of Fed policymaking according to long-time observer Lars Svensson as recently as 2022 (see King and Wolman 2022). We believe these observations and past academic work support our assumption on the time-varying inflation incentive k_t.

[4]One may interpret a zero or low k_t as either the central bank truly believes that potential output is very close to efficient output or that potential output is far below efficient output but a discretionary central bank recognizes its own limitation due to the time-inconsistence problem and chooses to tolerate the inefficiencies and thus lower inflation incentive.

2.2 Information Environment

We first describe in detail the information environment, followed by a description of the resulting updating mechanism used by each firm when forecasting central bank's future inflation incentive. Every period, two standard macroshocks $\{u_t, g_t\}$ and the inflation incentive shock k_t are contemporaneously observable to all firms, the representative household, and the central bank. The incentive shock k_t follows,

$$k_t - \bar{k} = \rho_k \left(k_{t-1} - \bar{k}\right) + \nu_t, \tag{9}$$

where $\rho_k \in [0, 1)$ and $\nu_t \sim N(0, \frac{1}{q})$. As a result, $k_t \sim N\left(\bar{k}, \frac{1}{q_k}\right)$, where $\bar{k} > 0$ and $q_k = q\left(1 - \rho_k^2\right)$.

In addition, each individual firm receives a *private* foreknowledge about future inflation incentive. Specifically, at time t firm i receives a signal s^i_{t+j} informative about the j-period-ahead inflation incentive shock k_{t+j}. The signal is modeled as

$$s^i_{t+j} = k_{t+j} + \eta_{t+j} + \varepsilon^i_{t+j}, \tag{10}$$

where $\eta_{t+j} \sim N\left(0, \frac{1}{m}\right)$ is common across firms and $\varepsilon^i_{t+j} \sim N\left(0, \frac{1}{n}\right)$ is idiosyncratic among firms.[5] Each private signal contains two shocks representing the two information imperfections. First, the average signal is a forecast of the future inflation incentive but with error, measured by η_{t+j}. Denoting $\bar{s}_{t+j} = \int_0^1 s^i_{t+j} di$ the average signal of all firms, we have

$$\text{Average Forecast Error: } \bar{s}_{t+j} - k_{t+j} = \eta_{t+j} \quad \text{and} \quad Var(\eta_{t+j}) = \frac{1}{m}. \tag{11}$$

The volatility of average forecast error is measured by its variance $\frac{1}{m}$. The larger m is, the more precise $\bar{s}_t$ is about the central bank's

[5] s^i_{t+j} can be interpreted as a sufficient signal summarizing any new information regarding k_{t+j} that arrives in period t. It can be interpreted as from (unmodeled) private information acquisition, central bank disclosure, or learning from observing noisy signals of past inflation and output gap, etc. (Cukierman and Meltzer 1986; Stein and Sunderam 2018). Section 5 includes an extension that firms learn k_{t+j} based on their observation of an endogenous aggregate variable.

incentive k_t. Second, the idiosyncratic shock in each signal generates disagreement among agents:

$$\text{Disagreement: } s^i_{t+j} - \bar{s}_{t+j} = \varepsilon^i_{t+j} \qquad \text{and} \quad Var(\varepsilon^i_{t+j}) = \frac{1}{n}, \quad (12)$$

at any time t and firm i. The degree of disagreement among firms is measured by $\frac{1}{n}$, the variance of ε^i_t. The larger n is, the smaller the disagreement across firms. Notice that our specification of the information structure allows us to capture the precision of average forecast error independently from the disagreement among firms. Adopting an information structure capturing disagreement independently from average forecast error is critical for our model. If each private signal only contains idiosyncratic noise, the average forecast would be perfect by assumption, no matter what other imperfect public information is available.[6] In this regard, our modeling choice is motivated by insights generated by the decades of theoretical research on accounting information structure.[7]

[6]In effect, making this seemingly common and innocuous information assumption would inadvertently build in a collective rationality that precludes analysis of the kind of coordination mechanism that we study here. To see this more explicitly, consider an alternative two-signal structure as in Morris and Shin (2002). To fix ideas, suppose that each agent observes a purely public signal $z_{t+j} = k_{t+j} + \phi_{t+j}$ and a purely private signal $s^i_{t+j} = k_{t+j} + \chi^i_{t+j}$, where the noise terms are all independent of each other. Note that in this structure, we no longer have separate parameters capturing the degrees of disagreement and collective knowledge. To elaborate, note first that in the two-signal structure, collective knowledge is perfect through aggregating all the private signals s^i_{t+j}, i.e., the average of all private signals $\bar{s}_{t+j} = \int s^i_{t+j} di = k_{t+j}$. Second, the disagreement is jointly determined by the precision of the public and the private signals. To see this, note that each agent's posterior belief about k_{t+j} is a weighted average of the public and the private signals. Intuitively, when the public signal becomes more precise, each agent places more weight on the public signal, resulting in less disagreement in their posteriors about k_{t+j}. Conversely, when the private signal becomes more precise, each agent places more weight on the private signal, contributing to more disagreement in their posteriors about k_{t+j}. Accordingly, given the different informational properties of the two-signal structure from those of our information structure, switching to the two-signal structure would alter the implications of our analysis for the roles of disagreement and collective knowledge considerably. We thank an anonymous reviewer for suggesting that we discuss the two-signal structure.

[7]Starting in the late 1960s and early 1970s, accounting researchers began linking accounting concepts to information economics concepts (see American

Every period, firm i uses information set I_t^i to forecast relevant future shocks in order to form beliefs about future inflation. We assume firm i's relevant information set is

$$I_t^i = \left\{ \{u_\tau\}_{\tau=0}^{t}, \{g_\tau\}_{\tau=0}^{t}, \{k_\tau\}_{\tau=0}^{t}, \left\{s_\tau^i\right\}_{\tau=0}^{t+j} \right\}, \tag{13}$$

which includes all the past observations of k_τ up to period t and all the past acquired signals s_τ^i up to period $t+j$.[8] Using I_t^i to update beliefs about future k follows Bayes's rule.[9]

Every period t, the central bank's information set is $I_t^{CB} = \{\{u_\tau\}_{\tau=0}^{t}, \{g_\tau\}_{\tau=0}^{t}, \{k_\tau\}_{\tau=0}^{t}, \{\bar{s}_\tau\}_{\tau=0}^{t+j}\}$, and it chooses policy instrument i_t to achieve its objective.[10] This assumption is supported by the observation that a central bank is typically endowed with

Accounting Association monographs by Feltham 1973 and Mock 1976). The agenda is to build on the traditional approach under a purely measurement perspective and to tie the accounting measurement concepts to economic trade-off in decisionmaking under uncertainty. A seminal contribution is by Ijiri and Jaedicke (1966), who framed objectivity within statistical sampling setting as interpersonal agreement and related it to reliability. Ijiri and Jaedicke introduced two properties of accounting measurement structure. One is the distance between the true state and the average measurements, which we define as average forecast error in our paper. The other one is the distance between the average measurements and measurements by different measurers, which we define as disagreement.

[8]The sources of information available for each firm are exogenously given. We view I_t^i as sufficient statistics for firm i to forecast future inflation at time t. Endogenous information sources may include potentially noisy observations of output and prices such as nominal interest rate and inflation series. We abstract away from these endogenous sources to focus on the role of disagreement, however it is generated, on macrovariables.

[9]The computations of first-, second-, or higher-order expectations can be very simple or quite complex depending on parameters. Consider a simple case of $j = 2$ and $\rho_k = 0$; in order to form a first-order belief about next period's inflation incentive k_{t+1}, the firm i would only use s_{t+1}^i to compute its individual conditional expectation of k_{t+1} (as all other signals are useless due to the independence assumptions). For a computation of the (higher-order) beliefs, see the discussion of Proposition 2 in Section 4. When ρ_k is not zero, these expectation computations involve more terms, as more signals are now informative about future central bank incentives. For example, with a non-zero ρ_k, $\{k_t, s_{t+1}^i, s_{t+2}^i, \ldots s_{t+j}^i\}$ are all informative about k_{t+1}. See Proposition 2 in Section 4 for a detailed account of such first-, second-, and higher-order expectations when the inflation incentives k_t's are serially correlated.

[10]Technically in a simultaneous-move game, a Nash equilibrium only requires the central bank to choose a best response to average expectations, not necessarily to observe the actual average expectation. Therefore, the observability of average signals by the central bank is inconsequential.

more information than an individual firm. The representative household's information set is $I_t^H = \left\{ \{u_\tau\}_{\tau=0}^t, \{g_\tau\}_{\tau=0}^t, \{k_\tau\}_{\tau=0}^t \right\}$, and it chooses intertemporal consumption with rational expectation. As we will show later, the Phillips curve that effectively determines the equilibrium inflation and output gap does not include the expectation of the representative household $E_t^H[\cdot]$. As a result, $E_t^H[\cdot]$ (thus the household's information set) affects nominal interest rate through the dynamic IS curve but does not affect the equilibrium inflation and output.

3. Preliminary Policy Analysis with Disagreement and Discretion

We assume that the central bank conducts a discretionary monetary policy each period. In a typical period t, the firms and the central bank simultaneously decide their actions. Specifically, the central bank chooses the nominal interest i_t given its information set I_t^{CB}, while each firm forms an expectation (forecasts) about the inflation rate in the next period, given its information set I_t^i and its conjecture about the central bank's future actions. In short, the players play a simultaneous-move game according to their best response given their own information set. This section provides the preliminary analysis needed to construct the closed-form equilibrium outcome in Section 4.

3.1 First-Order Condition for the Central Bank

Since the central bank cannot commit, it only chooses the current nominal interest rate i_t (but not future rates) that solves the following optimization program:

$$\begin{aligned} \min_{i_t} \quad & \frac{1}{2} E_t \left\{ \sum_{\tau=0}^{\infty} \beta^\tau \left[\alpha \left(x_{t+\tau} - k_{t+\tau} \right)^2 + \pi_{t+\tau}^2 \right] \right\}, \\ s.t. \quad & x_t = -\phi \left[i_t - E_t^H \pi_{t+1} \right] + E_t^H x_{t+1} + g_t, \\ & \pi_t = \lambda x_t + \beta \bar{E}_t^F \pi_{t+1} + u_t. \end{aligned} \tag{14}$$

Following Clarida, Gali, and Gertler (1999), we solve the optimization program in two stages: first, we solve for the pair of (x_t, π_t) that maximizes the objective given the Phillips curve (2); second, we use the IS curve (7) to determine the nominal interest rate i_t that supports the optimal pair of (x_t, π_t). Throughout the paper, since we are mostly interested in the equilibrium properties of inflation and output, we will focus on analyzing the first stage. Accordingly, we will omit the superscript F in expectation notation $\bar{E}_t^F$ and use $\bar{E}_t$ to represent average expectation by the firms in the rest of the paper to simplify exposition. In the first stage, notice that since the central bank, under discretion, cannot credibly change the firms' beliefs about its future actions, it takes the firms' expectations as given. As a result, the optimization problem for the central bank can be simplified into

$$\begin{aligned} &\min_{\{x_t, \pi_t\}} \frac{1}{2}\left[\alpha\left(x_t - k_t\right)^2 + \pi_t^2\right] + F_t, \\ s.t. \quad &\pi_t = \lambda x_t + f_t, \end{aligned} \tag{15}$$

where $f_t = \beta \bar{E}_t \pi_{t+1} + u_t$, and $F_t = \frac{1}{2} E_t \left\{\sum_{\tau=1}^{\infty} \beta^i \left[\alpha\left(x_{t+\tau} - k_{t+\tau}\right)^2 + \pi_{t+\tau}^2\right]\right\}$.[11] The first-order condition on x_t gives

$$x_t = -\frac{\lambda}{\alpha}\pi_t + k_t. \tag{16}$$

The first-stage solution reveals that the central bank must choose its policy instrument (in the second stage) to respect Equation (16). Holding k_t constant, a central bank seeing a positive (cost-push) shock u_t that pushes current inflation π_t higher via the Phillips curve would choose a policy to reduce current output, thus lowering the output gap x_t. Equation (16) also shows that the central bank is tempted to raise the output gap by k_t, holding the (cost-push) shock u_t constant. The higher the k_t, the higher the central bank's temptation to push up the output gap.

[11]To focus our attention on the role of higher-order beliefs on monetary policy, we ignore alternative equilibria involving reputation (using, e.g., grim-trigger strategies) which could support a more efficient outcome (see text by Mailath and Samuelson 2006). See footnote 26 on page 1671 of Clarida, Gali, and Gertler (1999) for background and explanations.

Substituting the first-order condition (16) into the Phillips curve (2) reveals inflation expectation dynamics generated by the central bank's best response:

$$\pi_t = \frac{\alpha\lambda}{\alpha+\lambda^2}k_t + \frac{\alpha\beta}{\alpha+\lambda^2}\bar{E}_t\left[\pi_{t+1}\right] + \frac{\alpha}{\alpha+\lambda^2}u_t. \tag{17}$$

Equation (17) suggests that the central bank must respond to changes in average expectations $\bar{E}_t\left[\pi_{t+1}\right]$ in determining inflation. The higher the expected future inflation $\bar{E}_t\left[\pi_{t+1}\right]$, the higher the actual current inflation π_t. In this sense, (17) captures the self-fulfilling nature in monetary policymaking. The coefficient before $\bar{E}_t\left[\pi_{t+1}\right]$, $\frac{\alpha\beta}{\alpha+\lambda^2} \in (0,1)$, thus measures how responsive the actual current inflation is to the expected future inflation.

3.2 Forward-Recursive Solutions of Phillips Curve under Disagreement and Discretion

We solve for π_t through forward-looking iteration. Iterating (17) once gives

$$\begin{aligned}\pi_t = {} & \frac{\alpha}{\alpha+\lambda^2}u_t + \frac{\alpha\beta}{\alpha+\lambda^2}\frac{\alpha}{\alpha+\lambda^2}\bar{E}_t\left[u_{t+1}\right] + \frac{\alpha\lambda}{\alpha+\lambda^2}k_t \\ & + \frac{\alpha\beta}{\alpha+\lambda^2}\frac{\alpha\lambda}{\alpha+\lambda^2}\bar{E}_t\left[k_{t+1}\right] + \left(\frac{\alpha\beta}{\alpha+\lambda^2}\right)^2 \bar{E}_t\bar{E}_{t+1}\left[\pi_{t+2}\right]. \end{aligned} \tag{18}$$

The key observation is that, in contrast to symmetric information case (i.e., no disagreement), the law of iterated expectation does not hold for average beliefs by differentially informed firms (Morris and Shin 2002).[12] That is,

$$\bar{E}_t\bar{E}_{t+1}\left[\cdot\right] \neq \bar{E}_t\left[\cdot\right]. \tag{19}$$

In fact, $\bar{E}_t\bar{E}_{t+1}\left[\cdot\right]$ corresponds to the second-order average beliefs of the firms, i.e., the firms' beliefs about the others' beliefs, which may differ substantially from the first-order average beliefs $\bar{E}_t\left[\cdot\right]$ when

[12] We will verify this point once we specify information structure for firms in the next section.

firms are differentially informed about central bank incentives. Similarly, a third-order belief term would show up when Equation (17) is iterated twice, and so on. Because of the failure of the law of iterated expectation, we must characterize the entire hierarchy of higher-order beliefs, all of which depend on the firms' current information set (I_t^i) and affect the equilibrium monetary policies. To simplify notations, we denote the l-th order beliefs as $\bar{E}_t^l[\cdot]$, where

$$\bar{E}_t^l[\cdot] \equiv \bar{E}_t \bar{E}_{t+1} ... \bar{E}_{t+l-1}[\cdot]. \tag{20}$$

We find that the iteration of (17) converges and gives π_t as a function of the higher-order beliefs, as summarized in the proposition below.

Proposition 1. *In equilibrium, the inflation rate π_t depends on the sum of the higher-order beliefs about $\{k_{t+l}\}_{l=0}^{l=\infty}$, i.e.,*

$$\begin{aligned}\pi_t = {} & \frac{\alpha\lambda}{\alpha+\lambda^2} k_t + \frac{\alpha u_t}{\alpha(1-\beta\rho_u)+\lambda^2} \\ & + \left\{ \sum_{l=1}^{\infty} \left(\frac{\alpha\beta}{\alpha+\lambda^2} \right)^l \frac{\alpha\lambda}{\alpha+\lambda^2} \bar{E}_t^l [k_{t+l}] \right\}. \end{aligned} \tag{21}$$

Before we consider the specific linear-normal information structure laid out earlier, we note that the coordination problem we identify exists only under discretionary monetary policy. The problem goes away when the central bank follows a credible rule, even in the presence of imperfect information. Specifically, with commitment, the central bank can unilaterally and uniformly *anchor* each individual firm's expectation by credibly specifying both current and future policies $\{\pi_{t+\tau}\}_{\tau=0}^{\infty}$. As a result, the central bank's equilibrium policy actions would be based upon the predetermined decision rule that is known to all firms. In equilibrium, firms' forecasting problem would then reduce to simply estimating the unobservable shocks in the predetermined rule—for instance, the central bank's inflation incentives $\{k_{t+\tau}\}_{\tau=0}^{\infty}$—by constructing first-order beliefs about them, without necessarily constructing (higher-order) beliefs about others' beliefs.

Without commitment, however, the central bank's current action loses control of the average current expectation of its future policy actions and, worse, it must react to its assessment of the average expectation. All firms know that the central bank will adjust its actions in response to the firms' aggregate expectations, $\bar{E}_t\left[\pi_{t+1}\right]$, every period. In mechanical terms, the central bank will choose the pair of $\{x_t, \pi_t\}$ for a given $\bar{E}_t\left[\pi_{t+1}\right]$, the aggregate expectation of its own future action, based on the Phillips curve relation. From an individual firm's perspective, since others' forecasts collectively affect the central bank's monetary actions—which, in turn, affect the very inflation rate it wants to forecast to begin with—it must also forecast the forecasts of others. In this process, rationality dictates that it must form beliefs about others' beliefs about $\{k_{t+\tau}\}_{\tau=0}^{\infty}$, others' beliefs about others' beliefs, and even higher-order beliefs. These beliefs in turn determine individual forecasts of all firms, which collectively influence the equilibrium inflation through the self-fulfilling feature embedded in the modified Phillips curve (17). Notice that from Equation (21), the relative importance of higher-order beliefs is determined by $\frac{\alpha\beta}{\alpha+\lambda^2}$, the responsiveness of the actual inflation to the expected future inflation. If $\frac{\alpha\beta}{\alpha+\lambda^2}=0$, the equilibrium inflation becomes independent of the aggregate expectation, making it unnecessary for each firm to forecast others' forecasts. As a result, all the higher-order-belief terms vanish.

3.3 A Closed-Form Forward-Recursive Solution

As a matter of exposition and practice, we believe allowing two-period-ahead foreknowledge (i.e., setting $j=2$) is sufficient, in part, because it allows a closed-form solution to the full equilibrium (the derivations for the cases with $j>2$ are similar but less analytically tractable).[13] Specifically, at any period t, a firm's information set is $I_t^i=\left\{\{u_\tau\}_{\tau=0}^{t},\{g_\tau\}_{\tau=0}^{t},\{k_\tau\}_{\tau=0}^{t},\left\{s_\tau^i\right\}_{\tau=0}^{t+2}\right\}$. To proceed, we first remove redundant elements in the firm's information set. First, at

[13] As it turns out, if firms only have one-period-ahead foreknowledge about the inflation incentive (i.e., $j=1$), the higher-order beliefs would become degenerate such that all higher-order beliefs would coincide with the first-order beliefs. Accordingly, absent non-degenerate higher-order beliefs, firms' signals would be used to forecast the inflation incentive but not the forecasts of other firms. Hence the implication of our model for central bank transparency would also change. Detailed analysis is available upon request.

each period t, observe that $\{k_\tau\}_{\tau=0}^{t}$ are commonly known and are sufficient statistics for signals, $\left\{s_\tau^i\right\}_{\tau=0}^{t}$, so the only useful signals are $\left\{s_{t+1}^i, s_{t+2}^i\right\}$ when forecasting future k's. Second, since k_t follows an AR(1) process, k_t is a sufficient statistics for all the past $\{k_\tau\}_{\tau=0}^{t-1}$. To sum, a firm's information set can be simplified into $I_t^i = \left\{k_t, s_{t+1}^i, s_{t+2}^i\right\}$ for the purpose of forecasting future k's.

The following proposition provides the closed-form solutions to the higher-order beliefs terms:

PROPOSITION 2. *When $j = 2$, the l-th order average beliefs become*

$$\bar{E}_t^l\left[k_{t+l}\right] = \bar{k} + \rho_k^{l-1}\left\{\left[1 - w\left(l\right)\right]\bar{E}_t\left[k_{t+1} - \bar{k}|s_{t+1}^i, k_t\right] + w\left(l\right)\frac{\bar{s}_{t+2} - \bar{k}}{\rho_k}\right\}, \tag{22}$$

where $\bar{E}_t\left[k_{t+1} - \bar{k}|s_{t+1}^i, k_t\right] = \frac{q}{q+\frac{mn}{m+n}}\rho_k\left(k_t - \bar{k}\right) + \frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\left(\bar{s}_{t+1} - \bar{k}\right)$ *and $w\left(l\right)$ is a constant given in the appendix.*

Proposition 2 suggests that $\bar{E}_t\left[k_{t+1}|s_{t+1}^i, k_t\right]$ and $\bar{s}_{t+2}$ are the two sufficient statistics for period-t firms to forecast the average higher-order beliefs about the central bank's future inflation incentive. To further illustrate the construction of the higher-order-belief hierarchy, consider first a special case in which the central bank's inflation incentive k_t is serially uncorrelated ($\rho_k = 0$). In this case, firms share a common prior on $k_t \sim N\left(\bar{k}, \frac{1}{q}\right)$. In addition, when forecasting k_{t+l}, the only useful signals are the prior $\bar{k}$ and s_{t+l}^i, and all the other signals, $\left\{s_\tau^i\right\}_{\tau \neq t+l}^{t}$, are not useful, since k_t is serially uncorrelated. In this case, the first-order belief $\bar{E}_t\left[k_{t+1}\right]$ is a weighted average of the prior and the average signal $\bar{s}_{t+1}$, with the weights simply the ones under Bayesian updating and similarly for the first-order belief $\bar{E}_t\left[k_{t+2}\right]$, i.e.,

$$\bar{E}_t\left[k_{t+1}\right] = \bar{k} + \frac{\frac{1}{q}}{\frac{1}{q} + \frac{1}{m} + \frac{1}{n}}\left(\bar{s}_{t+1} - \bar{k}\right), \tag{23}$$

$$\bar{E}_t\left[k_{t+2}\right] = \bar{k} + \frac{\frac{1}{q}}{\frac{1}{q} + \frac{1}{m} + \frac{1}{n}}\left(\bar{s}_{t+2} - \bar{k}\right). \tag{24}$$

To form the average second-order belief, $\bar{E}_t^2\left[k_{t+2}\right]$, first consider an individual firm i's expectation of next period's average belief:

$$\begin{aligned} E_t^i\left[\bar{E}_{t+1}\left[k_{t+2}\right]\right] &= \bar{k} + \frac{\frac{1}{q}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}}\left(E_t^i\left[\bar{s}_{t+2}\right]-\bar{k}\right) \\ &= \bar{k} + \frac{\frac{1}{q}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}}\left(\bar{k} + \frac{\frac{1}{q}+\frac{1}{m}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}}\left(s_{t+2}^i-\bar{k}\right)-\bar{k}\right) \\ &= \bar{k} + \frac{\frac{1}{q}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}} \frac{\frac{1}{q}+\frac{1}{m}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}}\left(s_{t+2}^i-\bar{k}\right), \qquad (25) \end{aligned}$$

and aggregating all firms' expectations, the average second-order belief becomes[14]

$$\begin{aligned} \bar{E}_t^2\left[k_{t+2}\right] &\equiv \bar{E}_t\left[\bar{E}_{t+1}\left[k_{t+2}\right]\right] \\ &= \bar{k} + \frac{\frac{1}{q}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}} \frac{\frac{1}{q}+\frac{1}{m}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}}\left(\bar{s}_{t+2}-\bar{k}\right). \qquad (26) \end{aligned}$$

Notice that in forming the average second-order belief $\bar{E}_t^2\left[k_{t+2}\right]$, the average signal $\bar{s}_{t+2}$ is assigned a lower weight relative to the typical Bayesian weight, i.e., $\frac{\frac{1}{q}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}} \frac{\frac{1}{q}+\frac{1}{m}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}} < \frac{\frac{1}{q}}{\frac{1}{q}+\frac{1}{m}+\frac{1}{n}}$, while the prior is assigned a higher weight. As a result, $\bar{E}_t\left[k_{t+2}\right] \neq \bar{E}_t^2\left[k_{t+2}\right]$, consistent with literature on the role of public information in coordination settings (Morris and Shin 2002). Notice in a standard model without disagreements ($n = \infty$), each firm's information set contains only public information; no "overweighting" takes place, making the higher-order-beliefs degenerate. In the special case of $\rho_k = 0$, for beliefs higher than the second order, the higher-order expectations become degenerate and equal to the prior $\bar{k}$, i.e., $\bar{E}_t^l\left[k_{t+l}\right] \equiv \bar{k}$ for $l > 2$. This is because period-t firms only receive private signals about, and thereby disagree on, the central bank's future inflation incentive up to period $t+2$. For any other future k_{t+l}, period-t firms share the same common prior $\bar{k}$ and agree with each other. Such

[14]One can also verify that taking a limit of the expression of the higher-order beliefs, i.e., expression (22), in Proposition 2 at $\rho_k = 0$ produces the same expressions of $\bar{E}_t\left[k_{t+1}\right]$ and $\bar{E}_t\left[\bar{E}_{t+1}\left[k_{t+2}\right]\right]$ given in the text.

perfect agreement among firms makes the higher-order beliefs that are higher than the second-order degenerate.

In the general case of serially correlated inflation incentive k_t ($\rho_k \neq 0$), the entire hierarchy of higher-order beliefs, including the ones that are higher than the second-order belief, remain non-degenerate. This is because, since k_t is serially correlated, period-t firms can utilize their private signals about k_{t+1} and k_{t+2} to forecast future $\{k_{t+l}\}_{l=3}^{\infty}$, thus disagreeing with each other in their beliefs about all of the central bank's future inflation incentive. Such disagreement in turn makes the higher-order beliefs $\left\{\bar{E}_t^l\left[k_{t+l}\right]\right\}_{l=3}^{\infty}$ non-degenerate. The proof of Proposition 2 contains the derivation of these expectations explicitly.

3.4 Symmetric Information Benchmark

Before we characterize the equilibrium in the model with informational imperfection, for comparison purposes, consider an identical model except no firm receives any signal about future inflation incentives (see Clarida, Gali, and Gertler 1999, Sections 3 and 4.1, for similar results.). It is well known in this setting that under discretionary monetary policy, the equilibrium output and inflation contain an inflation bias driven by k_t and $\bar{k}$:

$$\begin{aligned}
\pi_t^* &= \frac{\alpha}{\alpha\left(1-\beta\rho_u\right)+\lambda^2}u_t + \frac{\alpha\lambda}{\alpha\left(1-\beta\right)+\lambda^2}\bar{k} \\
&\quad + \frac{\alpha\lambda}{\alpha\left(1-\beta\rho_k\right)+\lambda^2}\left(k_t-\bar{k}\right), \\
x_t^* &= -\frac{\lambda}{\alpha\left(1-\beta\rho_u\right)+\lambda^2}u_t + \frac{\alpha\left(1-\beta\right)}{\alpha\left(1-\beta\right)+\lambda^2}\bar{k} \\
&\quad + \frac{\alpha\left(1-\beta\rho_k\right)}{\alpha\left(1-\beta\rho_k\right)+\lambda^2}\left(k_t-\bar{k}\right), \\
i_t^* &= \frac{g_t}{\phi} + \frac{\alpha\lambda}{\alpha\left(1-\beta\right)+\lambda^2}\bar{k} + \frac{\alpha\rho_u+\lambda\left(1-\rho_u\right)/\phi}{\alpha\left(1-\beta\rho_u\right)+\lambda^2}u_t \\
&\quad + \frac{\alpha\lambda\rho_k-\alpha\left(1-\beta\rho_k\right)\left(1-\rho_k\right)/\phi}{\alpha\left(1-\beta\rho_k\right)+\lambda^2}\left(k_t-\bar{k}\right).
\end{aligned} \tag{27}$$

The sources of aggregate output and inflation fluctuations are shocks to the Phillips curve and the central bank's inflation incentive. Equilibrium nominal interest rate also reacts to shocks to the dynamic IS curve.

4. Equilibrium in Closed Form

In this section, we first derive, in closed form, the equilibrium inflation, output gap, and nominal interest rate under the imperfect information environment for the special case of $j = 2$ (i.e., firms receive only two-period-ahead k_{t+2}). Then, we conduct comparative stochastic dynamic analysis of how information imperfections affect the volatilities of equilibrium inflation and output.

4.1 The Stochastic Stationary Equilibrium

Substituting the expressions for the higher-order expectations $\bar{E}_t^l\left[k_{t+l}\right]$ given by Equation (22) into the solution for π_t (in Equation (21)) and then x_t (in Equation (16)) gives the equilibrium inflation π_t^{**} and output gap x_t^{**}. The equilibrium nominal interest rate can be derived by substituting the pair (π_t^{**}, x_t^{**}) into the IS curve (7), giving the complete equilibrium characterization.

PROPOSITION 3. *Assuming $j = 2$, the equilibrium $\{\pi_t^{**}, x_t^{**}, i_t^{**}\}$is given by*

$$\pi_t^{**} = \frac{\alpha u_t}{\alpha\left(1-\beta\rho_u\right)+\lambda^2} + \frac{\alpha\lambda}{\alpha\left(1-\beta\right)+\lambda^2}\bar{k}$$

$$+ \frac{\alpha\lambda}{\alpha\left(1-\beta\rho_k\right)+\lambda^2}\left(k_t-\bar{k}\right)$$

$$+ \frac{\alpha\lambda}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1}$$

$$\times\left(\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}} - \rho_k\left(k_t-\bar{k}\right)\right),$$

$$
\begin{aligned}
x_t^{**} = & -\frac{\lambda u_t}{\alpha(1-\beta\rho_u)+\lambda^2} + \frac{\alpha(1-\beta)}{\alpha(1-\beta)+\lambda^2}\bar{k} \\
& + \frac{\alpha(1-\beta\rho_k)}{\alpha(1-\beta\rho_k)+\lambda^2}\left(k_t - \bar{k}\right) \\
& - \frac{\lambda^2}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1} \\
& \quad \times \left(\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}} - \rho_k\left(k_t-\bar{k}\right)\right), \\
i_t^{**} = & \frac{g_t}{\phi} + \frac{\alpha\lambda}{\alpha(1-\beta)+\lambda^2}\bar{k} + \frac{\alpha\rho_u + \lambda(1-\rho_u)/\phi}{\alpha(1-\beta\rho_u)+\lambda^2}u_t \\
& + \frac{\alpha\lambda\rho_k - \alpha(1-\beta\rho_k)(1-\rho_k)/\phi}{\alpha(1-\beta\rho_k)+\lambda^2}\left(k_t-\bar{k}\right) \\
& + \frac{1}{\phi}\left[\frac{\lambda^2}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1}\right. \\
& \quad \left. \times \left(\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}} - \rho_k\left(k_t-\bar{k}\right)\right)\right],
\end{aligned} \tag{28}
$$

where the "demeaned" higher-order beliefs are

$$
\begin{aligned}
\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}} - \rho_k\left(k_t-\bar{k}\right) = & \left[1-w(l)\right]\frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\left(\nu_{t+1}+\eta_{t+1}\right) \\
& + w(l)\left(\frac{\nu_{t+2}+\rho_k\nu_{t+1}+\eta_{t+2}}{\rho_k}\right).
\end{aligned} \tag{29}
$$

Proposition 3 shows that the equilibrium inflation π_t^{**} and hence the output gap x_t^{**} are determined by three factors, the contemporaneous (cost-push) shock u_t, the inflation bias $\frac{\alpha\lambda}{\alpha(1-\beta)+\lambda^2}\bar{k} + \frac{\alpha(1-\beta\rho_k)}{\alpha(1-\beta\rho_k)+\lambda^2}\left(k_t-\bar{k}\right)$, and the higher-order expectations $\bar{E}_t^l\left[k_{t+l}\right]$. The first two factors have been extensively examined in the literature and appear even in the benchmark model without informational imperfection (see Equation (27)). Specifically, consistent with standard results (Clarida, Gali, and Gertler 1999), we verify that the cost-push shock u_t is inflationary. In addition, consistent with Barro and Gordon (1983a), we find that the discretion in monetary policy can lead to a persistent inflation bias $\frac{\alpha\lambda}{\alpha(1-\beta)+\lambda^2}\bar{k}$.

In addition to the two well-known effects in the literature, the proposition above shows that the combination of the discretion in the monetary policy and the disagreement among firms can lead to another potentially detrimental effect, as captured in the third terms of the equilibrium inflation and output gap. Through the channel of the induced-coordination problem we identify, the discretionary monetary policy causes the equilibrium inflation and output to react to firms' higher-order beliefs about the central bank's inflation incentive, leading to heightened fluctuations in output and inflation. Our findings thereby suggest that the combination of lack of commitment by the central bank and the imperfect information known to firms makes the central bank less capable to stabilize output and inflation, compared with a central bank in an alternate economy without such commitment and information frictions.

The source of the heightened output and inflation fluctuations comes from the volatilities of the primitive variables in our model. Specifically, the equilibrium inflation and output will respond not only to the central bank's current inflation incentive k_t but also to the noises, $\{\eta_{t+1}, \eta_{t+2}\}$, contained in firms' average signals $\{\bar{s}_{t+1}, \bar{s}_{t+2}\}$, as well as innovations in the central bank's future inflation incentive, $\{\nu_{t+1}, \nu_{t+2}\}$. Furthermore, the coordination problem induced by the discretionary monetary policy makes the destabilizing effect of the monetary policy more prominent, due to a "multiplier" effect. When forecasting the forecasts of others, firms' average forecast is determined by a hierarchy of higher-order beliefs, each of which depends on the noises in firms' current information. As the monetary policy reacts to firms' forecast, the entire hierarchy of higher-order beliefs enters into the equilibrium inflation and output and the noise contained in these beliefs leads to heightened volatility. In particular, Equation (28) shows precisely that since the same private signals $\left\{s^i_{t+1}, s^i_{t+2}\right\}$ are used for each level in the individual higher-order belief hierarchy, the common information noise $\{\eta_{t+1}, \eta_{t+2}\}$ in these private signals is retained at every term of the higher-order beliefs, magnifying the noises contained in the average inflation forecast. When the central bank responds to the average forecast, the magnified information noises enter into aggregate inflation and output, generating heightened macrofluctuations.

4.2 Comparative Stochastic Dynamic Analysis

In the face of the heightened volatility caused by firms' imperfect information, the conventional wisdom would suggest that reducing the volatilities of the two informational shocks, η_{t+j} (average forecast error) and ε^i_{t+j} (degree of disagreement), would be desirable. We find that this intuition does not hold generally. Importantly, we show that reducing the volatilities of the two informational shocks can increase the macrofluctuations by inducing a more aggressive monetary policy response. To see the effect of informational properties on volatilities, from Proposition 3, the volatility of inflation is computed as

$$\begin{aligned} Var\left(\pi^{**}_t\right) = &\left(\frac{\alpha}{\alpha\left(1-\beta\rho_u\right)+\lambda^2}\right)^2 \frac{\sigma^2_u}{1-\rho^2_u} \\ &+ \left(\frac{\alpha\lambda}{\alpha\left(1-\beta\rho_k\right)+\lambda^2}\right)^2 Var\left(k_t\right) + \left(\frac{\alpha\lambda}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\right)^2 \\ &\times Var\left(\textstyle\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1}\left(\frac{\bar{E}^l_t[k_{t+l}]-\bar{k}}{\rho^{l-1}_k} - \rho_k\left(k_t-\bar{k}\right)\right)\right), \end{aligned} \tag{30}$$

where the first term represents the volatility stemming from the shock u_t, the second term represents the volatility stemming from the central bank's *current* inflation incentive k_t, and the third term represents the volatility stemming from higher-order expectations about the central bank's *future* inflation incentive. By the first-order condition, $x^{**}_t = -\frac{\lambda}{\alpha}\pi^{**}_t + k_t$, the volatility of x^{**}_t is proportional to the volatility of π^{**}_t, and the two share similar properties. Thus we will focus on analyzing the volatility of π^{**}_t. For notational convenience, we define the sensitivities of the equilibrium inflation to the future signals $\bar{s}_{t+1}$ and $\bar{s}_{t+2}$ as

$$\begin{aligned} W_{\bar{s}_{t+1}}\left(m,n\right) = &\frac{\alpha\lambda}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1} \\ &\times\left\{\left[1-w\left(l\right)\right]\frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\right\}, \end{aligned} \tag{31}$$

$$W_{\bar{s}_{t+2}}(m,n) = \frac{\alpha\lambda}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1}\frac{w(l)}{\rho_k}, \tag{32}$$

which depend on the informational properties, the average forecast error m, and the degree of disagreement n. Using notations $W_{\bar{s}_{t+1}}(m,n)$ and $W_{\bar{s}_{t+2}}(m,n)$, $Var(\pi_t^{**})$ can be rewritten as

$$\begin{aligned}&\left(\frac{\alpha}{\alpha(1-\beta\rho_u)+\lambda^2}\right)^2\frac{\sigma_u^2}{1-\rho_u^2}+\left(\frac{\alpha\lambda}{\alpha(1-\beta\rho_k)+\lambda^2}\right)^2 Var(k_t)\\&+\left[W_{\bar{s}_{t+1}}(m,n)+\rho_k W_{\bar{s}_{t+2}}(m,n)\right]^2 Var(\nu_{t+1})\\&+\left[W_{\bar{s}_{t+1}}(m,n)\right]^2 Var(\eta_{t+1})\\&+\left[W_{s_{t+2}}(m,n)\right]^2\left[Var(\nu_{t+2})+Var(\eta_{t+2})\right].\end{aligned}\tag{33}$$

$Var(\nu_{t+1}) = Var(\nu_{t+2}) = \frac{1}{q}$ and $Var(\eta_{t+1}) = Var(\eta_{t+2}) = \frac{1}{m}$. Therefore, $Var(\pi_t^{**})$ becomes

$$\begin{aligned}&\left(\frac{\alpha}{\alpha(1-\beta\rho_u)+\lambda^2}\right)^2\frac{\sigma_u^2}{1-\rho_u^2}+\left(\frac{\alpha\lambda}{\alpha(1-\beta\rho_k)+\lambda^2}\right)^2 Var(k_t)\\&\quad+\frac{\left[W_{\bar{s}_{t+1}}(m,n)+\rho_k W_{\bar{s}_{t+2}}(m,n)\right]^2+\left[W_{s_{t+2}}(m,n)\right]^2}{q}\\&\quad+\frac{\left[W_{\bar{s}_{t+1}}(m,n)\right]^2+\left[W_{s_{t+2}}(m,n)\right]^2}{m}.\end{aligned}\tag{34}$$

Equation (34) suggests that in addition to the volatility driven by the shocks u_t and k_t, inflation volatility is also driven by two other shocks. The third term in (34) represents the fundamental volatility stemming from the innovations in the central bank's future inflation incentive $\{\nu_{t+1},\nu_{t+2}\}$, and the fourth term is the non-fundamental volatility stemming from the noises in firms' signals, i.e., $\{\eta_{t+1},\eta_{t+2}\}$. Equation (34) shows that the informational properties can influence the macrofluctuations in two ways. First, improving the precision of the average forecast error (increasing m)

directly reduces the size of the noises $\{\eta_{t+1}, \eta_{t+2}\}$ in the equilibrium monetary policy, leading to less volatility in the output and inflation. We capture this effect in the $\frac{1}{m}$ term in (34) and call this effect a *noise-diminishing* effect. Second, when either the precision of average forecast error or the agreement among firms increases, the average forecast becomes more sensitive to the firms' imperfect information, thus the average forecast error. The central bank, in turn, reacts more aggressively to the average expectation; unfortunately, this reaction adds more volatility to equilibrium inflation and output. In other words, increasing m or n can increase the sensitivity of the monetary policy to firms' signals and noises (i.e., $W_{\bar{s}_{t+1}}(m, n)$ and $W_{\bar{s}_{t+2}}(m, n)$). We capture this effect in $W_{\bar{s}_{t+1}}(m, n)$ and $W_{\bar{s}_{t+2}}(m, n)$ and call it a *sensitivity* effect. Whether improving firms' information (increasing m and n) reduces the volatilities thereby depends on the trade-off between the sensitivity effect and the noise-diminishing effect. We summarize the effect of the informational properties on the volatilities of inflation and output in the proposition below.

PROPOSITION 4. *Information properties* (m, n) *influence the volatilities of inflation and output as follows:*

(i) Volatilities increase strictly in n*, i.e., more agreement always increases volatilities.*

(ii) There exists a unique $\hat{n}$*, such that volatilities decrease strictly in* m *if and only if* $n < \hat{n}$*, i.e., more accurate average forecast decreases volatilities when disagreement is sufficiently high.*

Proposition 4 suggests that holding fixed the average forecast-error volatility (m), a higher degree of agreement among agents (increasing n) leads to higher fluctuations in output and inflation. On the other hand, holding the degree of disagreement fixed, reducing the size of the average forecast error has a non-monotonic effect on the volatility. We find that increasing m helps to stabilize inflation and output if and only if the disagreement among the firms is sufficiently high.

The intuition for these results is due to a trade-off between the *noise-diminishing* and the *sensitivity* effects. Specifically, as we explained earlier, the lack of commitment by the central bank induces an implicit coordination motive among the firms, making it necessary for an individual firm to forecast the forecasts of others. That is, in forming its best forecast of the future inflation, a firm uses its information to estimate not only the central bank's inflation incentive but also others' beliefs about the incentive. We call the first use of information the *fundamental* value of information and the second use the *strategic* value of information. Under the information structure specified in our model, improvements in the precision of the average forecast m and the agreement n play different roles in affecting the two uses of information (see Liang and Zhang 2019). First, increasing either m or n diminishes the size of (common or idiosyncratic) noises and moves the firms' signals closer to the central bank's true target, which enhances the fundamental value of information. Second, increasing n increases the strategic value of information, while increasing m decreases the strategic value of information. This is because the strategic value of information is determined by the correlation between firms' private signals, $corr\left(s_\tau^i, s_\tau^{i'}\right) = \frac{\frac{1}{m}+\frac{1}{q}}{\frac{1}{m}+\frac{1}{n}+\frac{1}{q}}$ for $i \neq i'$, which is strictly increasing in n but decreasing in m. Intuitively, increasing m reduces the size of common noises and hence the common variation among the firms' signals, reducing the correlation between the signals, while increasing n decreases the size of idiosyncratic noises and hence the idiosyncratic variation, increasing the correlation.

The different role of m and n in influencing the value of information determines their effects on the volatilities. We first explain the effect of higher agreement. Since increasing n (higher agreement) increases both the fundamental and the strategic value of the information, all firms respond more sensitively to their signals $\left\{s_{t+1}^i, s_{t+2}^i\right\}$ in forming their forecasts ($W_{\bar{s}_{t+1}}(m,n)$ and $W_{\bar{s}_{t+2}}(m,n)$ both increase). After the idiosyncratic noises $\left\{\varepsilon_{t+1}^i, \varepsilon_{t+2}^i\right\}$ are diversified away in the aggregation, the average expectation of the firms becomes more responsive to the average signals $\{\bar{s}_{t+1}, \bar{s}_{t+2}\}$. This is the sensitivity effect of increasing n. When the central bank cannot commit, it is tempted to respond more to the aggregate expectation, making its monetary policy more sensitive to the errors in firms' average expectation as well. As a result, the equilibrium inflation induced by the monetary policy is driven by the errors in the aggregate expectation to a larger extent and becomes more volatile.

The effect of increasing m differs from that of increasing n in two ways. First, increasing m increases the fundamental value of the information but decreases the strategic value. Overall, increasing m still increases the firms' sensitivity to their signals (increases $W_{\bar{s}_{t+1}}(m,n)$ and $W_{\bar{s}_{t+2}}(m,n)$). The higher sensitivity leads to higher volatilities through the transmission mechanism illustrated above; however, this sensitivity effect of m is weaker than that of n because the decrease in the strategic value of the information led by higher m dampens the increase in the sensitivity. Second, in aggregating the firms' forecasts, the common noise $\{\eta_{t+1}, \eta_{t+2}\}$ is not diversified away as the idiosyncratic noises $\{\varepsilon^i_{t+1}, \varepsilon^i_{t+2}\}$. This captures the noise-diminishing effect of increasing m: a higher m directly diminishes the size of the common noise and makes the average expectation and hence the inflation rate less volatile. The net effect of m on the volatilities thus depends on the trade-off between the sensitivity effect and the noise-diminishing effect. When the disagreement is sufficiently high, the strategic value of information in forecasting the forecasts of others becomes important. Due to the adverse effect of m on the strategic value, the firms are more reluctant to respond to their information, despite the fact that the increase in m improves the fundamental value. As a result, the sensitivity effect becomes weak and dominated by the noise-diminishing effect. Accordingly, increasing m leads to lower volatilities. Otherwise, when the disagreement is low, the strategic value of information becomes less important, making the sensitivity effect strong and dominate the noise-diminishing effect. In these cases, increasing m amplifies the volatilities.

5. Additional Analysis

In this section, we derive some additional results to enrich the implications of our paper.[15]

[15]We thank two annoymous reviewers for suggesting this additional analysis, which helps to enrich the implications of our paper.

5.1 Fixing the Total Precision of Firms' Signals

In our main analysis, we have analyzed the effects of separately varying the accuracy of firms' average forecast m and the agreement among firms' forecasts n. It is also interesting to examine how the volatilities change with either m or n, fixing the total precision of firms' signals $\frac{1}{m} + \frac{1}{n}$. In this case, the precision of the first-order expectation stays constant, so any variations in the volatilities are triggered by changes in the higher-order expectations. We summarize our results in the following proposition.

PROPOSITION 5. *Fixing the total precision of firms' signals* $\frac{1}{m} + \frac{1}{n}$, *information properties* (m, n) *influence the volatilities of inflation and output as follows:*

(i) Volatilities increase strictly in n, *i.e., more agreement always increases volatilities.*

(ii) Volatilities decrease strictly in m, *i.e., more accurate average forecast always decreases volatilities.*

Proposition 5 suggests that, holding the total precision of firms' signals constant, a higher degree of agreements among firms still magnifies the volatilities in output and inflation, similar to the result in Proposition 4. However, while Proposition 4 points to a non-monotonic effect of changing the accuracy of the average forecast on the volatilities, Proposition 5 shows that, when the total precision is fixed, reducing the average forecast error always helps to stabilize inflation and output. The intuition for this result, again, lies in how varying the degree of agreement n and the average forecast accuracy m affects the fundamental and the strategic value of firms' signals. First, since the total precision of firms' signals is fixed, changing either n or m will not alter the signals' fundamental value. Second, recall that the strategic value of the signals is determined by the correlation between firms' private signals, $corr\left(s_\tau^i, s_\tau^{i'}\right) = \frac{\frac{1}{m}+\frac{1}{q}}{\frac{1}{m}+\frac{1}{n}+\frac{1}{q}}$ for $i \neq i'$. Fixing the total precision $\frac{1}{m} + \frac{1}{n}$, the correlation is strictly increasing in n and decreasing in m. Accordingly, a higher agreement n improves the strategic value of firms' signals and makes all firms respond more sensitively to their signals, which, in turn, contributes to higher fluctuations in output and inflation through the transmission mechanism discussed in our main analysis. Conversely, a higher average forecast precision m impairs the strategic value of firms' signals and makes all firms respond less sensitively to their signals and the noises in the signals. In addition, a higher m also diminishes the size of the common noise. Combining the two effects, improving the accuracy of firms' average forecasts helps to mitigate volatilities in output and inflation.

To illustrate the implications of Proposition 5, consider a shock to the information environment such that after the shock, the disagreement among firms vanishes (i.e., the idiosyncratic shock in the signal ε^i_{t+j} is muted, $n = \infty$) but the total precision of firms' signals is unaffected. Note that from Proposition 5, since increasing the agreement always amplifies volatilities, imposing full agreement would in fact result in maximal fluctuations in the economy among all scenarios with the same level of total precision. Stated differently, our analysis cautions against efforts to shrink disagreement/dispersion in market participants' understanding of central banks' operations, even if these efforts do not reduce market participants' total knowledge about central banks.

5.2 Fundamental and Non-fundamental Volatilities

In our main analysis, we have examined how varying the informational properties $\{m, n\}$ affects the volatilities in output and inflation. Examining Equation (34) suggests that the properties m and n can influence the volatilities through two components: (i) the fundamental volatility stemming from shocks to the central bank's inflation incentives (i.e., the third term in (34)) and (ii) the non-fundamental volatility from the common noises in firms' signals (i.e., the fourth term in (34)). To illustrate the underlying economic forces of our paper, it is helpful to decompose the effects of $\{m, n\}$ on the fundamental and non-fundamental volatilities. The following proposition summarizes our analysis of such decomposition.

Proposition 6. *Information properties (m, n) influence the fundamental and non-fundamental volatilities of inflation and output as follows:*

(i) *Both the fundamental and non-fundamental volatilities increase strictly in n, i.e., more agreement always increases both volatilities.*

(ii) *The fundamental volatilities increase strictly in m, i.e., more accurate average forecast always increases the fundamental volatilities.*

(iii) *There exists a threshold $\hat{n}'$, such that the non-fundamental volatilities decrease strictly in m if and only if $n < \hat{n}'$, i.e., more accurate average forecast decreases the non-fundamental volatilities when disagreement is sufficiently high.*

The message of Proposition 6 echoes our main result in Proposition 4. In particular, recall from our discussion of Proposition 4 that increasing the agreement always makes firms more responsive to their signals. Since the signals commingle the fundamental shocks to the central bank's inflation incentives with some noises, firms also react more sensitively to both the fundamental shocks and the non-fundamental noises. Through the transmission mechanism discussed previously, firms' higher sensitivities lead to heightened fundamental and non-fundamental volatilities. This explains part (i) of Proposition 6.

Similarly, recall that although increasing the average forecast precision m has conflicting effects on the fundamental and the strategic value of firms' information, its overall effect on firms' sensitivity to their signals is still positive. This, in turn, contributes to higher fundamental volatilities. This explains part (ii) of Proposition 6.

The last part of Proposition 6 is also in line with Proposition 4. Note that a reduction of the average forecast error (a higher m) affects the non-fundamental volatilities in two ways: it not only increases firms' sensitivity to their signals but also diminishes the size of the common noises. From the discussion of Proposition 4, the noise-diminishing effect dominates the sensitivity effect when the disagreement is high. Accordingly, in these cases, increasing m helps to reduce the non-fundamental volatilities.

5.3 Adjusting the Central Bank's Objective Function

A main takeaway from our analysis is that the informational frictions faced by firms can potentially lead to amplified volatilities in inflation and output. These volatilities are driven by the central bank's equilibrium choice of the discretionary monetary rule, and hence depend on the central bank's objective function (8). Accordingly, one may argue that, to mitigate these aggregate volatilities, the central bank ex ante may have incentives to adjust its objective function. While a full characterization of the central bank's optimal objective function is beyond the scope of our paper, we now explore a specific way of adjusting the objective function, that is, the central bank adjusting the relative weight α placed on output deviations $(x_t - k_t)^2$. In practice, a lower α can be interpreted as appointing a more conservative central banker (Rogoff 1985; Clarida, Gali, and Gertler 1999). We summarize the effect of the weight α on the inflation volatility in the proposition below.

PROPOSITION 7. *The inflation volatility increases strictly in the weight α on output deviations.*

Proposition 7 suggests that the informational frictions faced by firms induce greater inflation volatilities if the central bank is more concerned about the output gap target. Stated differently, to stabilize aggregate fluctuations, the central bank should reduce the weight placed on output deviations. To see the intuition, recall from Equation (21) that the equilibrium inflation is more sensitive to the terms of higher-order beliefs if the central bank is more responsive to changes in average expectations of inflation in choosing its discretionary monetary policy, as captured by the response coefficient $\frac{\alpha\beta}{\alpha+\lambda^2}$. Note that this coefficient is strictly increasing in the weight α. A central bank focusing more on output gaps is more tempted to push inflation π_t higher in order to lower the output gap. This, in turn, sets the trap for the central bank to respond to firms' inflation expectations, thus amplifying the inflation volatility through the coordination channel identified in our main analysis.

5.4 A Microfoundation of the Common Noise in Firms' Signals

Our analysis suggests that the common noise in firms' signals η_t plays a crucial role, as it contributes to the non-fundamental uncertainties in the output and inflation. In our main model, we do not specify how the common noise may arise in firms' information environment. We study an extension below to offer a microfoundation in which firms learn the central bank's policy objective from an endogenous variable and such learning is subject to a common noise. For simplicity, in this extension, we adopt the macroeconomic framework in Barro and Gordon (1983a) and capture endogenous learning following Stein and Sunderam (2018). To illustrate the main idea, we only consider the static version of the learning model in Stein and Sunderam (2018, Section II); Stein and Sunderam (2018, Section IV) also extend the learning model to a fully dynamic model.

Specifically, the operation of the economy is described by a "Phillips curve" (Equation 1 on pp. 592, Barro and Gordon 1983a):

$$U_t = U_t^n - a\left(\pi_t - \pi_t^e\right), \tag{35}$$

where U_t denotes the unemployment rate and is a proxy for the overall state of real activity, U_t^n denotes the natural unemployment rate, π_t denotes the inflation rate, and π_t^e denotes firms' expected inflation. The coefficient $a > 0$ represents the "Phillips curve slope." The central bank's objective is to minimize a social loss function (Barro and Gordon 1983a, Equation (3) on p. 593):

$$Z_t = E\left[\left(U_t - k_t\right)^2\right] + bE\left[\pi_t^2\right], \tag{36}$$

where $k_t \sim N\left(U_t^n, \frac{1}{\tau_k}\right)$ represents the central bank's preferred unemployment rate target, as discussed in Barro and Gordon (1983a). k_t can be either higher or lower than the natural unemployment rate, depending on the central bank's target. The parameter $b > 0$ captures the central bank's relative weight on minimizing inflation in its objective function. The same as in Stein and Sunderam (2018), we assume that k_t is the central bank's private information and unknown to firms at the beginning of period t. Specifically, we assume that the central bank sets the inflation rate π_t by following a partial adjustment rule of the form

$$\pi_t = \mu\left(k_t - U_t^n\right) + \varepsilon_t, \tag{37}$$

where $\varepsilon_t \sim N\left(0, \frac{1}{\tau_\varepsilon}\right)$ represents a noise "that is overlaid onto the rate-setting process" by Stein and Sunderam (2018, p. 1024), who interpret the noise as either a "tremble" in the central bank's optimal choice of π_t or "coming from the Fed's use of round numbers (typically in 25 bps) for the funds rate settings that it communicates to the market," whereas its private information about k_t is presumably continuous. As we will show soon, the noise ε_t generates a common noise in the endogenous signal learned by firms and prevents firms from fully recovering the central bank's private information k_t. The parameter μ is the central bank's response coefficient to its unemployment target, and the central bank will set μ optimally in equilibrium to minimize the social loss function Z_t.

Firms try to infer the central bank's private information k_t based on their observation of the aggregate state of the economy, proxied by the unemployment rate U_t. To do so, firms conjecture that in equilibrium, the unemployment rate and the inflation rate take the following forms:

$$U_t = U_t^n + \hat{\delta}\left(k_t - U_t^n\right) - \hat{\lambda}\varepsilon_t, \tag{38}$$

$$\pi_t = \hat{\mu}\left(k_t - U_t^n\right) + \varepsilon_t. \tag{39}$$

That is, firms correctly conjecture the equilibrium forms of the unemployment rate and the inflation rate but do not observe the sets of actual response coefficients. Instead, firms conjecture the values of these coefficients as $\left\{\hat{\delta}, \hat{\lambda}, \hat{\mu}\right\}$, where $\hat{x}$ denotes firms' conjecture of variable x. Rational expectations require that in equilibrium, firms' conjectures are correct and coincide with the equilibrium values (yet unable to fully invert, similar to Stein and Sunderam 2018). Importantly, note that the unemployment rate U_t is an endogenous signal about the central bank's private information subject to a common noise. The noise arises because of the noise ε_t in the central bank's rate-setting process in Equation (37), which, in turn, enters the unemployment rate through the Phillips curve in Equation (35).

We will verify that $\lambda > 0$ so there is indeed a common noise term in the endogenous signal U_t.

Using their conjectured form of the unemployment rate in Equation (38), firms infer imperfectly the central bank's private information k_t from their observation of U_t. Rewriting (38) yields

$$\hat{U}_t \equiv U_t^n + \frac{U_t - U_t^n}{\hat{\delta}} = k_t - \frac{\hat{\lambda}}{\hat{\delta}}\varepsilon_t. \tag{40}$$

$\hat{U}_t$ can be viewed as an adjusted unemployment rate that constitutes an endogenous signal of k_t, where the precision of $\hat{U}_t$ is $\frac{\hat{\delta}^2}{\hat{\lambda}^2}\tau_\varepsilon$. Standard Bayesian updating gives

$$E\left[k_t|\hat{U}_t\right] = \frac{\frac{\hat{\delta}^2}{\hat{\lambda}^2}\tau_\varepsilon}{\frac{\hat{\delta}^2}{\hat{\lambda}^2}\tau_\varepsilon + \tau_k}\hat{U}_t + \frac{\tau_k}{\frac{\hat{\delta}^2}{\hat{\lambda}^2}\tau_\varepsilon + \tau_k}U_t^n = U_t^n + \chi\left(U_t - U_t^n\right), \tag{41}$$

where $\chi \equiv \frac{\hat{\delta}\tau_\varepsilon}{\hat{\delta}^2\tau_\varepsilon + \hat{\lambda}^2\tau_k}$. Intuitively, when there is less noise in the unemployment rate U_t (i.e., a smaller $\hat{\lambda}$), firms learn more information about k_t from U_t and hence react more to a change in U_t (i.e., a higher χ).

Given firms' inference of k_t and the conjectured form of the equilibrium inflation rate in Equation (39), firms also form an expectation of inflation:

$$\pi^e = E\left[\pi_t|U_t\right] = \hat{\mu}\left(E\left[k_t|U_t\right] - U_t^n\right) = \hat{\mu}\chi\left(U_t - U_t^n\right). \tag{42}$$

Substituting Equation (42) into Equation (35) solves the equilibrium unemployment rate:

$$U_t = U_t^n - \frac{a\pi_t}{1 - a\hat{\mu}\chi}. \tag{43}$$

Given the central bank's choice of inflation rate in Equation (37) and the equilibrium unemployment rate in Equation (43), the central bank's objective function Z_t in Equation (36), taking expectations over the noise ε_t, can be written as

$$Z_t = \left(1 + \frac{a\mu}{1 - a\hat{\mu}\chi}\right)^2 \frac{1}{\tau_k} + \left(\frac{a}{1 - a\hat{\mu}\chi}\right)\frac{1}{\tau_\varepsilon} + b\left(\frac{\mu^2}{\tau_k} + \frac{1}{\tau_\varepsilon}\right). \tag{44}$$

The central bank minimizes Z_t by choosing the optimal response coefficient μ in Equation (37). The central bank does so taking as given firms' conjectures $\left\{\hat{\delta}, \hat{\lambda}, \hat{\mu}\right\}$. Taking the first-order condition with respect to μ gives

$$\left(1+\frac{a\mu}{1-a\hat{\mu}\chi}\right)\frac{a}{1-a\hat{\mu}\chi}+b\mu=0. \tag{45}$$

In the rational expectations equilibrium, firms' conjectures are correct, i.e., $\hat{\mu}=\mu$, $\hat{\delta}=\delta$, and $\hat{\lambda}=\lambda$. Imposing these requirements in the first-order condition (45) and matching the coefficients in the equilibrium unemployment rate in Equation (43) with firms' conjecture in Equation (38) determine the sets of coefficients, $\{\delta, \lambda, \mu\}$, in equilibrium. We summarize the equilibrium in the following proposition.

PROPOSITION 8. *Consider a model in which firms learn about the central bank's private information from an endogenous variable. The equilibrium unemployment rate and inflation take the following form:*

$$U_t = U_t^n + \delta\left(k_t - U_t^n\right) - \lambda\varepsilon_t, \tag{46}$$

$$\pi_t = \mu\left(k_t - U_t^n\right) + \varepsilon_t, \tag{47}$$

where the coefficients are given by

$$\delta = \frac{a\sqrt{b\left(1-\delta\right)\delta\tau_k}}{\left(1-\delta\right)\delta\tau_\varepsilon + b\tau_k} \in \left(0,1\right), \tag{48}$$

$$\mu = -\sqrt{\frac{1}{b}\left(1-\delta\right)\delta} < 0, \tag{49}$$

$$\lambda = \sqrt{\frac{b\delta}{1-\delta}} > 0. \tag{50}$$

Proposition 8 confirms that in equilibrium, since $\lambda > 0$, the unemployment rate U_t indeed constitutes an endogenous signal about the central bank's unemployment rate target k_t, subject to the common noise ε_t stemming from the rate-setting process. The unemployment rate U_t is informative about k_t because the central bank, facing the downward-sloping Phillips curve, is tempted to inject inflation in order to reduce the unemployment rate toward its preferred target k_t (i.e., $\mu < 0$). Accordingly, both the equilibrium inflation rate and the equilibrium unemployment rate become dependent on k_t so that observing the unemployment rate reveals information about k_t. This model extension establishes the endogenous source of the common noise in the private signals exogenously specified in the base model.

5.5 A Microfounded Phillips Curve

In our main analysis, to maintain tractability, we have assumed a reduced-form Phillips curve (2) in line with the standard New Keynesian Phillips curve when information is complete. Nonetheless, recent work by Angeletos and Lian (2018) and Angeletos and Huo (2021) suggest that when information is incomplete, the form of the Phillips curve varies from the standard one. Importantly, both studies show that, considering incomplete information, the current inflation depends on the entire future path of average expectations about the output gap and the inflation rate, instead of only the average expectation about the next-period inflation. Angeletos and Huo (2021, p. 1174) demonstrate that analyzing the equilibrium under the modified Phillips curve is considerably more complicated, as "the relevant set of higher-order beliefs is significantly richer." In light of this complexity, we have imposed a simple reduced-form Phillips curve in our main analysis to focus on deriving implications for how central bank transparency/disagreement affects aggregate volatilities. To assess the robustness of our analysis, we now analyze a variant of our model with a microfounded Phillips curve.

Specifically, consider a setting in which the Calvo (1983) friction lasts only for one period, i.e., a firm that has reset its price p_t^j in period t is restricted from resetting its price in period $t+1$ with probability $\theta \in (0,1)$, whereas a firm restricted from resetting its price in period t gains full price-resetting flexibility in period $t+1$. For simplicity, in this extension, we assume that the cost-push shocks $u_t = 0$. Consider a firm j that has the opportunity to reset its price in period t. Following similar steps as in Angeletos and Lian (2018, Equation (32)), the optimal reset price, denoted by p_t^{j*}, can be derived as

$$p_t^{j*} = \frac{1}{1+\beta\theta}\left(mc_t^j + p_t\right) + \frac{\beta\theta}{1+\beta\theta} E_t^j\left[mc_{t+1}^j + p_{t+1}\right], \tag{51}$$

where mc_t^j denotes firm j's real marginal cost in period t. Conditional on the firm being restricted from resetting the price in period $t+1$, it gains full price-resetting flexibility in period $t+2$. Accordingly, the optimal reset price p_t^{j*} only depends on the firm's expectation about the marginal cost and the aggregate price in period $t+1$. Aggregating p_t^{j*} over the population of the firms gives

$$p_t^{*} = \frac{1}{1+\beta\theta}\left(mc_t + p_t\right) + \frac{\beta\theta}{1+\beta\theta} \bar{E}_t\left[mc_{t+1} + p_{t+1}\right], \tag{52}$$

where $\bar{E}_t[\cdot]$ denotes the average expectation of the firms. As shown in the proof of Proposition 9, at the steady state, in each period t, a fraction $\frac{1}{1+\theta}$ of the firms can reset the price, whereas the remaining firms are restricted from resetting the price. This, in turn, gives the inflation rate:

$$\pi_t = \frac{p_t^* - p_{t-1}}{1+\theta}. \tag{53}$$

Using (52) and (53) and applying the usual condition that the real marginal cost is proportional to the output gap (Clarida, Gali, and Gertler 1999), i.e., $mc_t = \kappa x_t$, we obtain the following condition for the level of inflation in period t:

$$\pi_t = \lambda x_t + \beta' \bar{E}_t\left[\kappa x_{t+1} + \pi_{t+1}\right], \tag{54}$$

where $\beta' \equiv \frac{\beta}{1+\beta+\beta\theta} \in (0,1)$ and $\lambda \equiv \frac{\kappa}{\beta\theta+\theta(1+\beta\theta)} > 0$. Note that (54) is identical to the reduced-form Phillips curve (2) assumed in the main analysis except that in the microfounded Phillips curve (54), the inflation rate π_t also depends on the expectation about the future output gap in period $t+1$. To assess how this change affects the robustness of our analysis, we now derive the equilibrium inflation rate π_t^{**}. Note that under the modified Phillips curve, the optimal discretionary monetary policy remains intact, i.e.,

$$x_t = -\frac{\lambda}{\alpha}\pi_t + k_t, \tag{55}$$

because the central bank still takes the firms' future expectations about x_{t+1} and π_{t+1} as given. Substituting the optimal discretionary monetary policy into the Phillips curve (54) gives

$$\pi_t = \frac{\alpha\lambda}{\alpha+\lambda^2}k_t + \frac{\alpha\beta'\kappa}{\alpha+\lambda^2}\bar{E}_t\left[k_{t+1}\right] + \frac{\left(\alpha-\lambda\kappa\right)\beta'}{\alpha+\lambda^2}\bar{E}_t\left[\pi_{t+1}\right]. \tag{56}$$

Comparing the modified law of motion (56) for the inflation π_t with its counterpart (17) in the main analysis yields two insights. First, the inflation π_t depends not only on the central bank's inflation incentive k_t in period t but also on average expectations about the future inflation incentive k_{t+1}. The latter result is driven by the modified Phillips curve (54), where the inflation π_t depends on average expectations about the future output gap x_{t+1} and, accordingly, the central bank's future inflation incentive k_{t+1} due to the discretionary rule (55). Second, the inflation π_t continues to depend on average expectations $\bar{E}_t\left[\pi_{t+1}\right]$, although the coefficient before $\bar{E}_t\left[\pi_{t+1}\right]$, $\frac{\beta'(\alpha-\lambda\kappa)}{\alpha+\lambda^2}$, can be negative. The reason is that when the expectation about the future inflation π_{t+1} increases, firms rationally anticipate that the central bank will choose a policy in period $t+1$ to reduce the future output gap x_{t+1}. Such policy, in turn, dampens the current inflation through the Phillips curve (54). This economic force goes against the usual force that the central bank responds positively to changes in $\bar{E}_t\left[\pi_{t+1}\right]$ in setting π_t, and can even make π_t respond negatively to $\bar{E}_t\left[\pi_{t+1}\right]$ when π_t is more dependent on the future output gap x_{t+1} (i.e., $\lambda\kappa = \frac{\kappa^2}{\beta\theta+\theta(1+\beta\theta)} > \alpha$). We show that, as long as κ is sufficiently small so that π_t responds positively to $\bar{E}_t\left[\pi_{t+1}\right]$, all the implications from our main analysis remain valid under the modified Phillips curve. To see this, iterating (56) gives

$$\pi_t^{**} = \frac{\alpha\lambda}{\alpha+\lambda^2}k_t + \left\{\sum_{l=1}^{\infty}\left(\frac{\left(\alpha-\lambda\kappa\right)\beta'}{\alpha+\lambda^2}\right)^{l-1}\frac{\alpha\beta'}{\alpha+\lambda^2}\frac{\alpha\left(\lambda+\kappa\right)}{\alpha+\lambda^2}\bar{E}_t^l\left[k_{t+l}\right]\right\}. \tag{57}$$

Under the modified Phillips curve (54), the equilibrium inflation rate π_t^{**} continues to be a function of the sum of the higher-order beliefs about $\{k_{t+l}\}_{l=0}^{l=\infty}$, similar to that characterized in Proposition 1. The only difference is that the "discounting factor" before

$\bar{E}_t^l\left[k_{t+l}\right]$ is $\frac{(\alpha-\lambda\kappa)\beta'}{\alpha+\lambda^2}$ instead of $\frac{\alpha\beta}{\alpha+\lambda^2}$. Accordingly, we show that as long as $\frac{(\alpha-\lambda\kappa)\beta'}{\alpha+\lambda^2} > 0$, the implications regarding the effect of the informational properties on the volatilities of inflation and output are similar to the ones in our main analysis. We formally state these results in the following proposition.

PROPOSITION 9. *Under the modified Phillips curve (54), the equilibrium inflation rate π_t^{**} depends on the sum of the higher-order beliefs about $\{k_{t+l}\}_{l=0}^{l=\infty}$, i.e.,*

$$\pi_t^{**} = \frac{\alpha\lambda}{\alpha+\lambda^2}k_t + \left\{\sum_{l=1}^{\infty}\left(\frac{\beta'\left(\alpha-\lambda\kappa\right)}{\alpha+\lambda^2}\right)^{l-1}\frac{\alpha\beta'}{\alpha+\lambda^2}\frac{\alpha\left(\lambda+\kappa\right)}{\alpha+\lambda^2}\bar{E}_t^l\left[k_{t+l}\right]\right\}. \tag{58}$$

If $\frac{\kappa^2}{\beta\theta+\theta(1+\beta\theta)} < \alpha$, information properties (m, n) influence the volatilities of inflation and output as follows:

(i) Volatilities increase strictly in n, i.e., more agreement always increases volatilities.

(ii) There exists a unique $\hat{n}''$, such that volatilities decrease strictly in m if and only if $n < \hat{n}''$, i.e., more accurate average forecast decreases volatilities when disagreement is sufficiently high.

6. Conclusion

With its simplicity, our paper makes a core argument for the inclusion of information diversity among agents in monetary policy discussions. A direct implication of our model is on the explanation and characterization of the observed inflation dynamics. Our model would suggest that the precision of the aggregate estimation of future inflation is a determinant of current inflation. In this regard, our paper is related to the voluminous macroliterature on inflation trend (Goodfriend and King 2012 and Ascari and Sbordone 2014). In these studies, firms are more sophisticated in their understanding of the inflation trend and adjust their pricing behavior (such as indexing). In an extension, we verify that our main qualitative results survive in a more general model (e.g., Woodford 2008) in which an inflation trend term is inserted into the New Keynesian Phillips curve.

More broadly, we view our paper as an attempt at constructing a positive understanding of the macroeconomy under the information imperfections about the incentives of an authority player. Our paper is not directly concerned about how these imperfections emerge endogenously from the *information production* of each player in the model, but any such studies should take into account the results of our paper. Recent interests in studying the communication strategies of the central bank are evidence of its perceived importance (see, e.g., Rudebusch and Williams 2008).

Aside from the information flow from the central bank to the marketplace, a more organic environment would also feature active private information activities. As shown in the Fed-watch literature, individual agents are motivated to acquire relevant information in anticipated information management by the central bank.

Finally, our paper raises issues that future studies could blend with other important considerations related to information and coordination. They include other coordination problems in macroeconomics (Cooper and John 1988; Kiyotaki and Wright 1989; Baxter and King 1991), robust policies by Hansen and Sargent (2007), and the information role of the financial market (King 1982; Baxter and King 1991).

Appendix. Proofs

Proof of Proposition 1

This can be verified by iterating (17).

Proof of Proposition 2

For our convenience, we define the vectors of firm i's demeaned signals at period t and the vector of the demeaned average signals at period t as

$$S_t^i = \begin{bmatrix} k_t - \bar{k} \\ s_{t+1}^i - \bar{k} \\ s_{t+2}^i - \bar{k} \end{bmatrix}, \bar{S}_t = \begin{bmatrix} k_t - \bar{k} \\ \bar{s}_{t+1} - \bar{k} \\ \bar{s}_{t+2} - \bar{k} \end{bmatrix}, \tag{A.1}$$

the variance of S_t^i as

$$Var\left(S_t^i\right) = \Sigma = \begin{bmatrix} \frac{1}{q_k} & \frac{\rho_k}{q_k} & \frac{\rho_k^2}{q_k} \\ \frac{\rho_k}{q_k} & \frac{1}{q_k} + \frac{1}{m} + \frac{1}{n} & \frac{\rho_k}{q_k} \\ \frac{\rho_k^2}{q_k} & \frac{\rho_k}{q_k} & \frac{1}{q_k} + \frac{1}{m} + \frac{1}{n} \end{bmatrix}, \tag{A.2}$$

and the covariance between $\bar{S}_{t+1}$ and S_t^i as

$$Cov\left(\bar{S}_{t+1}, S_t^i\right) = \Omega = \begin{bmatrix} \frac{\rho_k}{q_k} & \frac{1}{q_k} & \frac{\rho_k}{q_k} \\ \frac{\rho_k^2}{q_k} & \frac{\rho_k}{q_k} & \frac{1}{q_k} + \frac{1}{m} \\ \frac{\rho_k^3}{q_k} & \frac{\rho_k^2}{q_k} & \frac{\rho_k}{q_k} \end{bmatrix}. \tag{A.3}$$

In particular, we define the first row of Ω, the covariance between k_{t+1} and S_t^i, as

$$\Omega^{\text{row1}} = L = \begin{bmatrix} \frac{\rho_k}{q_k} & \frac{1}{q_k} & \frac{\rho_k}{q_k} \end{bmatrix}. \tag{A.4}$$

We now derive the hierarchy of higher-order beliefs. In the first-order belief in period t, each firm's forecast of k_{t+1} is

$$E_t^i\left[k_{t+1}\right] = \bar{k} + L\Sigma^{-1}S_t^i, \tag{A.5}$$

and the average forecast is

$$\bar{E}_t\left[k_{t+1}\right] = \bar{k} + L\Sigma^{-1}\bar{S}_t. \tag{A.6}$$

Building on the first-order expectation, now move to the second-order belief. For firm i, its period-t belief about the aggregate period $t+1$ belief about the central bank's period $t+2$ incentive becomes

$$E_t^i\left[\bar{E}_{t+1}\left[k_{t+2}\right]\right] = \bar{k} + E_t^i\left[L\Sigma^{-1}\bar{S}_{t+1}\right] = \bar{k} + L\Sigma^{-1}E_t^i\left[\bar{S}_{t+1}\right], \tag{A.7}$$

where $E_t^i\left[\bar{S}_{t+1}\right] = \Omega\Sigma^{-1}S_t^i$. Therefore, the average second-order belief becomes

$$\bar{E}_t\left[\bar{E}_{t+1}\left[k_{t+2}\right]\right] = \bar{k} + L\Sigma^{-1}\Omega\Sigma^{-1}\bar{S}_t. \tag{A.8}$$

Notice that the law of iterated expectation fails, i.e., $\bar{E}_t^2\left[k_{t+2}\right] \neq \bar{E}_t\left[k_{t+2}\right] = \bar{k} + \begin{bmatrix} \frac{\rho_k^2}{q_k} & \frac{\rho_k}{q_k} & \frac{1}{q_k} \end{bmatrix} \Sigma^{-1}\bar{S}_t$, since $L\Sigma^{-1}\Omega = \begin{bmatrix} \frac{\rho_k^2}{q_k} & \frac{\rho_k}{q_k} & \frac{1}{q_k} \frac{\left(\frac{1}{m}+\frac{1}{q}\right)\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)+\left(\frac{1}{mq}+\frac{1}{mn}+\frac{2}{nq}+\frac{1}{n^2}\right)\rho_k^2}{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2} \end{bmatrix} \neq \begin{bmatrix} \frac{\rho_k^2}{q_k} & \frac{\rho_k}{q_k} & \frac{1}{q_k} \end{bmatrix}$ for $n \neq \infty$ (i.e., there is some disagreement among firms). In particular, we verify that in $\bar{E}_t^2\left[k_{t+2}\right]$, the signal $\bar{s}_{t+2}$ is weighted less than in $\bar{E}_t\left[k_{t+2}\right]$. Moreover, for the third-order belief, firm i's period-t belief about the aggregate period $t+1$ belief about the aggregate period $t+2$ belief about the central bank's period $t+3$ incentive becomes

$$\begin{aligned} E_t^i\left[\bar{E}_{t+1}\left[\bar{E}_{t+2}\left[k_{t+3}\right]\right]\right] &= \bar{k} + L\Sigma^{-1}\Omega\Sigma^{-1}E_t^i\left[\bar{S}_{t+1}\right] \\ &= \bar{k} + L\Sigma^{-1}\Omega\Sigma^{-1}\Omega\Sigma^{-1}S_t^i, \end{aligned} \tag{A.9}$$

and thus the average third-order belief becomes

$$\begin{aligned} \bar{E}_t\left[\bar{E}_{t+1}\left[\bar{E}_{t+2}\left[k_{t+3}\right]\right]\right] &= \bar{k} + L\Sigma^{-1}\Omega\Sigma^{-1}\Omega\Sigma^{-1}\bar{S}_t \\ &= \bar{k} + L\Sigma^{-1}\left(\Omega\Sigma^{-1}\right)^2\bar{S}_t. \end{aligned} \tag{A.10}$$

Keeping iterating $\bar{E}_t^3\left[k_{t+3}\right]$ characterizes the entire hierarchy of higher-order beliefs with

$$\bar{E}_t^l\left[k_{t+l}\right] = \bar{k} + L\left(\Sigma^{-1}\Omega\right)^{l-1}\Sigma^{-1}\bar{S}_t. \tag{A.11}$$

To derive $\bar{E}_t^l\left[k_{t+l}\right]$, we make an eigenvalue decomposition on $\Sigma^{-1}\Omega$, such that

$$\Sigma^{-1}\Omega = Q\Lambda Q^{-1}, \tag{A.12}$$

where Λ is a diagonal matrix with the eigenvalue of $\Sigma^{-1}\Omega$ on its diagonal, i.e.,

$$\Lambda = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \frac{\frac{1}{n}\frac{\rho_k}{q}}{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2} & 0 \\ 0 & 0 & \rho_k \end{bmatrix}, \tag{A.13}$$

and Q is the associated matrix of eigenvectors. As a result,

$$\begin{aligned}\bar{E}_t^l\left[k_{t+l}\right] &= \bar{k} + LQ\Lambda^{l-1}Q^{-1}\Sigma^{-1}\bar{S}_t \\ &= \bar{k} + LQ\begin{bmatrix} 0 & 0 & 0 \\ 0 & \left[\frac{\frac{1}{n}\frac{\rho_k}{q}}{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2}\right]^{l-1} & 0 \\ 0 & 0 & \rho_k^{l-1}\end{bmatrix} Q^{-1}\Sigma^{-1}\bar{S}_t,\end{aligned} \tag{A.14}$$

and can be simplified into

$$\bar{E}_t^l\left[k_{t+l}\right] = \bar{k}+\rho_k^{l-1}\left\{\left[1-w\left(l\right)\right]\bar{E}_t\left[k_{t+1}-\bar{k}|s_{t+1}^i,k_t\right]+w\left(l\right)\frac{\bar{s}_{t+2}-\bar{k}}{\rho_k}\right\}, \tag{A.15}$$

where $\bar{E}_t\left[k_{t+1}-\bar{k}|s_{t+1}^i,k_t\right] = \frac{q}{q+\frac{mn}{m+n}}\rho_k\left(k_t-\bar{k}\right)+\frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\left(\bar{s}_{t+1}-\bar{k}\right)$ and

$$\begin{aligned} w\left(l\right) = &\frac{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}}{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2} \\ &+\frac{\begin{aligned}&\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)\frac{1}{q}\left[\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)\left(\frac{1}{m}+\frac{1}{q}\right)\right] \\ &\times\left\{1-\left[\frac{\frac{1}{n}\frac{1}{q}}{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2}\right]^{l-1}\right\}\end{aligned}}{\begin{aligned}&\left[\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2\right]^2 \\ &\times\left\{1-\left[\frac{\frac{1}{n}\frac{1}{q}}{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2}\right]\right\}\end{aligned}}.\end{aligned} \tag{A.16}$$

Notice that since $\frac{\frac{1}{n}\frac{1}{q}}{\left(\frac{1}{m}+\frac{1}{n}\right)\frac{\rho_k^2}{q}+\left(\frac{1}{m}+\frac{1}{n}+\frac{1}{q}\right)^2} < 1$, then $w\left(l\right)$ is strictly increasing in l.

Proof of Proposition 3

Substituting the expressions for the higher-order-beliefs terms into the expression for the inflation specified in Proposition 1, we have

$$\pi_t^{**} = \frac{\alpha u_t}{\alpha(1-\beta\rho_u)+\lambda^2} + \frac{\alpha\lambda}{\alpha(1-\beta)+\lambda^2}\bar{k} + \frac{\alpha\lambda}{\alpha(1-\beta\rho_k)+\lambda^2}\left(k_t-\bar{k}\right) + \frac{\alpha\lambda}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1}\left(\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}} - \rho_k\left(k_t-\bar{k}\right)\right), \tag{A.17}$$

where the "demeaned" higher-order beliefs are

$$\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}} - \rho_k\left(k_t-\bar{k}\right) = \left[1-w(l)\right]\left[\bar{E}_t\left[k_{t+1}-\bar{k}|s_{t+1}^i,k_t\right]-\rho_k\left(k_t-\bar{k}\right)\right] + w(l)\left(\frac{\bar{s}_{t+2}-\bar{k}}{\rho_k}-\rho_k\left(k_t-\bar{k}\right)\right) \tag{A.18}$$

with

$$\bar{E}_t\left[k_{t+1}-\bar{k}|s_{t+1}^i,k_t\right]-\rho_k\left(k_t-\bar{k}\right) = \frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\left(\bar{s}_{t+1}-\bar{k}-\rho_k\left(k_t-\bar{k}\right)\right) = \frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\left(\eta_{t+1}+\nu_{t+1}\right), \tag{A.19}$$

and

$$\frac{\bar{s}_{t+2}-\bar{k}}{\rho_k}-\rho_k\left(k_t-\bar{k}\right) = \frac{\nu_{t+2}+\rho_k\nu_{t+1}+\eta_{t+2}}{\rho_k}. \tag{A.20}$$

By the first-order condition, the equilibrium output is

$$
\begin{aligned}
x_t^{**} &= -\frac{\lambda}{\alpha}\pi_t^{**} + k_t \\
&= -\frac{\lambda u_t}{\alpha(1-\beta\rho_u)+\lambda^2} + \frac{\alpha(1-\beta)}{\alpha(1-\beta)+\lambda^2}\bar{k} + \frac{\alpha(1-\beta\rho_k)}{\alpha(1-\beta\rho_k)+\lambda^2}\left(k_t - \bar{k}\right) \\
&\quad - \frac{\lambda^2}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1}\left(\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}} - \rho_k\left(k_t-\bar{k}\right)\right).
\end{aligned} \tag{A.21}
$$

The equilibrium nominal interest rate can be derived by substituting the pair (π_t^{**}, x_t^{**}) into the IS curve (7):

$$
i_t^{**} = \frac{E_t^H x_{t+1}^{**}}{\phi} + \frac{g_t}{\phi} - \frac{x_t^{**}}{\phi} + E_t^H \pi_{t+1}^{**}, \tag{A.22}
$$

where, given the information set of the household, $I_t^H = \left\{\{u_\tau\}_{\tau=0}^t, \{g_\tau\}_{\tau=0}^t, \{k_\tau\}_{\tau=0}^t\right\}$,

$$
\begin{aligned}
E_t^H x_{t+1}^{**} &= -\frac{\lambda\rho_u u_t}{\alpha(1-\beta\rho_u)+\lambda^2} + \frac{\alpha(1-\beta)}{\alpha(1-\beta)+\lambda^2}\bar{k} \\
&\quad + \frac{\alpha(1-\beta\rho_k)}{\alpha(1-\beta\rho_k)+\lambda^2}\rho_k\left(k_t-\bar{k}\right), \\
E_t^H \pi_{t+1}^{**} &= \frac{\alpha\rho_u u_t}{\alpha(1-\beta\rho_u)+\lambda^2} + \frac{\alpha\lambda}{\alpha(1-\beta)+\lambda^2}\bar{k} \\
&\quad + \frac{\alpha\lambda}{\alpha(1-\beta\rho_k)+\lambda^2}\rho_k\left(k_t-\bar{k}\right),
\end{aligned} \tag{A.23}
$$

and as a result,

$$
\begin{aligned}
i_t^{**} &= \frac{g_t}{\phi} + \frac{u_t}{\alpha(1-\beta\rho_u)+\lambda^2}\left[\alpha\rho_u + \frac{\lambda(1-\rho_u)}{\phi}\right] + \frac{\alpha\lambda}{\alpha(1-\beta)+\lambda^2}\bar{k} \\
&\quad + \frac{\alpha\left(k_t-\bar{k}\right)}{\alpha(1-\beta\rho_k)+\lambda^2}\left[\lambda\rho_k - \frac{(1-\beta\rho_k)(1-\rho_k)}{\phi}\right]
\end{aligned}
$$

$$+ \ \frac{1}{\phi}\left[\frac{\lambda^2}{\alpha+\lambda^2}\frac{\alpha\beta}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}\right)^{l-1}\right.$$
$$\left.\times\left(\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}}-\rho_k\left(k_t-\bar{k}\right)\right)\right]. \qquad \text{(A.24)}$$

Proof of Proposition 4

Notice that $x_t^{**} = -\frac{\lambda}{\alpha}\pi_t^{**} + k_t$. Thus

$$Var_t\left(x_t^{**}\right) = \frac{\lambda^2}{\alpha^2}Var_t\left(\pi_t^{**}\right) + Var_t\left(k_t\right) - \frac{2\lambda}{\alpha}Cov_t\left(\pi_t^{**}, k_t\right), \quad \text{(A.25)}$$

where $Var_t\left(k_t\right) = \frac{1}{q(1-\rho_k)^2}$ and $Cov_t\left(\pi_t^{**}, k_t\right) = \frac{\alpha\lambda}{\alpha(1-\beta\rho_k)+\lambda^2}\frac{1}{q(1-\rho_k)^2}$ are both independent of m and n. Therefore, the effects of (m, n) on $Var_t\left(x_t^{**}\right)$ are the same as their effects on $Var_t\left(\pi_t^{**}\right)$.

One can verify $\frac{\partial Var_t(\pi_t^{**})}{\partial n} > 0$ by directly computing the derivative. For the sign of $\frac{\partial Var_t(\pi_t^{**})}{\partial m}$, first, one can verify that at $n = 0$, $\frac{\partial Var_t(\pi_t^{**})}{\partial m} = 0$, $\frac{\partial}{\partial n}\left(\frac{\partial Var_t(\pi_t^{**})}{\partial m}\right) = 0$, and $\frac{\partial^2}{\partial n^2}\left(\frac{\partial Var_t(\pi_t^{**})}{\partial m}\right) = -\frac{2\left(1+\rho_k^2\right)}{m^2q^2\left(1-\frac{\alpha\beta}{\alpha+\lambda^2}\rho_k\right)^2} < 0$. As a result, for n close to 0, $\lim_{n\to 0^+}\frac{\partial Var_t(\pi_t^{**})}{\partial m} < 0$. Second, at $n = \infty$,

$$\frac{\partial Var_t\left(\pi_t^{**}\right)}{\partial m}$$
$$= \frac{\left(1+\left(\frac{\alpha\beta}{\alpha+\lambda^2}\right)^2\right)(m+q)^2 + 2\frac{\alpha\beta}{\alpha+\lambda^2}\rho_k\left(q^2-m^2\right) + \left(\left(\frac{\alpha\beta}{\alpha+\lambda^2}m\right)^2+q^2\right)\rho_k^2}{\left(1-\frac{\alpha\beta}{\alpha+\lambda^2}\rho_k\right)^2\left[(m+q)^2+mq\rho_k^2\right]^2} > 0. \quad \text{(A.26)}$$

Therefore, by the intermediate value theorem, there exists an $\hat{n} > 0$, such that $\frac{\partial Var_t(\pi_t^{**})}{\partial m} = 0$. Lastly, we verify that such an $\hat{n}$ is also unique. More specifically, we verify that $\frac{\partial Var_t(\pi_t^{**})}{\partial m} = 0$ can be reduced into $P\left(n\right) = 0$ and $P\left(n\right)$ is a fourth-order polynomial of n,

$$P\left(n\right) = \kappa_1 n^4 + \kappa_2 n^3 + \kappa_3 n^2 + \kappa_4 n^3 + \kappa_5, \qquad \text{(A.27)}$$

where the expressions of the coefficients $\{\kappa_i\}_{i=1}^{5}$ are available upon request. We verify that $\kappa_1 > 0$, $\kappa_2 > 0$, $\kappa_5 < 0$, and the signs of κ_3 and κ_4 are ambiguous. However, it is impossible to have $\kappa_3 < 0$ and $\kappa_4 > 0$ at the same time. As a result, there can be the following three possible scenarios of the signs of $\{\kappa_i\}_{i=1}^{5}$:

κ_1	κ_2	κ_3	κ_4	κ_5
+	+	+	+	−
+	+	+	−	−
+	+	−	−	−

,

where "+" means positive and "−" means negative. Notice that for the polynomial $P(n)$, there is one sign change in its coefficients. Therefore, by Descartes's rule of signs, the polynomial $p(n)$ has a unique positive root. That is, there exists a unique $\hat{n}$ that makes $\frac{\partial Var_t(\pi_t^{**})}{\partial m} = 0$. As a result, $\frac{\partial Var_t(\pi_t^{**})}{\partial m} < 0$ if and only if $n < \hat{n}$.

Proof of Proposition 5

We only analyze the inflation volatility $Var_t(\pi_t^{**})$, as from Proposition 4 the results regarding the output $Var_t(x_t^{**})$ are the same. We denote the (inverse) total precision $T \equiv \frac{1}{m} + \frac{1}{n}$. Hence $m = \frac{1}{T - \frac{1}{n}}$. Since $m \geq 0$, $n > \frac{1}{T}$. Substituting this into the expression of $Var_t(\pi_t^{**})$ and taking the derivative of $Var_t(\pi_t^{**})$ with respect to n gives

$$\frac{\partial Var_t(\pi_t^{**})}{\partial n} = \frac{Q(n)}{n^2\left(1 - \frac{\alpha\beta}{\alpha+\lambda^2}\rho_k\right)^2\left(n + 2nqT + nq^2T^2 - \frac{\alpha\beta}{\alpha+\lambda^2}q\rho_k + nqT\rho_k^2\right)^3}. \tag{A.28}$$

The denominator is positive because $2nqT - \frac{\alpha\beta}{\alpha+\lambda^2}q\rho_k > 2nqT - q\rho_k > 2q - q\rho_k > 0$. The first step uses $\frac{\alpha\beta}{\alpha+\lambda^2} < 1$, the second step uses $\frac{1}{n} < T$, and the last step uses $\rho_k < 1$. The numerator $Q(n)$ is a third-order polynomial of n. We now prove that for $n > \frac{1}{T}$,

$Q(n) > 0$. First, replacing n in $Q(n)$ with $n = b + \frac{1}{T}$ (where $b > 0$), we obtain

$$Q(b) = \delta_1 b^3 + \delta_2 b^2 + \delta_3 b + \delta_4, \tag{A.29}$$

where the expressions of the coefficients $\{\delta_i\}_{i=1}^4$ are available upon request. Hence we need to prove that $Q(b) > 0$ for any $b > 0$. After some tedious algebra, we can verify that all δ_is are positive. Thus $Q(b) > 0$ for any $b > 0$. This, in turn, proves that $Q(n) > 0$ and $\frac{\partial Var_t(\pi_t^{**})}{\partial n} > 0$. In addition, note that fixing $T \equiv \frac{1}{m} + \frac{1}{n}$, an increase in n is the same as a decrease in m. Hence $\frac{\partial Var_t(\pi_t^{**})}{\partial m} < 0$.

Proof of Proposition 6

We only analyze the inflation volatility $Var_t(\pi_t^{**})$, as from Proposition 4 the results regarding the output $Var_t(x_t^{**})$ are the same. The third term in Equation (34) represents the fundamental volatility stemming from the innovations in the central bank's future inflation incentive $\{\nu_{t+1}, \nu_{t+2}\}$, defined below:

$$Var_t^F(\pi_t^{**}) \equiv \frac{\left[W_{\bar{s}_{t+1}}(m,n) + \rho_k W_{\bar{s}_{t+2}}(m,n)\right]^2 + \left[W_{s_{t+2}}(m,n)\right]^2}{q}. \tag{A.30}$$

The fourth term in Equation (34) is the non-fundamental volatility stemming from the noises in firms' signals, i.e., $\{\eta_{t+1}, \eta_{t+2}\}$, defined below:

$$Var_t^{NF}(\pi_t^{**}) \equiv \frac{\left[W_{\bar{s}_{t+1}}(m,n)\right]^2 + \left[W_{s_{t+2}}(m,n)\right]^2}{m}. \tag{A.31}$$

One can verify parts (i) and (ii) of the proposition, i.e., $\frac{\partial Var_t^F(\pi_t^{**})}{\partial n}, \frac{\partial Var_t^{NF}(\pi_t^{**})}{\partial n}, \frac{\partial Var_t^F(\pi_t^{**})}{\partial m} > 0$ by directly computing the derivatives.

For the sign of $\frac{\partial Var_t^{NF}(\pi_t^{**})}{\partial m}$ in part (iii) of the proposition, we can show that, after some tedious algebra,

$$\frac{\partial Var_t^{NF}\left(\pi_t^{**}\right)}{\partial m} = \frac{n^2 H\left(n\right)}{\left(1-\frac{\alpha\beta}{\alpha+\lambda^2}\rho_k\right)^2 \left[\left(mn+nq+mq\right)^2 + mnq\rho_k\left(\left(m+n\right)\rho_k - \frac{\alpha\beta}{\alpha+\lambda^2}m\right)\right]^3}. \tag{A.32}$$

It can be verified that the denominator of $\frac{\partial Var_t^{NF}(\pi_t^{**})}{\partial m}$ is positive, so the sign of $\frac{\partial Var_t^{NF}(\pi_t^{**})}{\partial m}$ is determined by $H\left(n\right)$, which is a fourth-order polynomial of n,

$$H\left(n\right) = \mu_1 n^4 + \mu_2 n^3 + \mu_3 n^2 + \mu_4 n^3 + \mu_5, \tag{A.33}$$

where the expressions of the coefficients $\{\mu_i\}_{i=1}^5$ are available upon request. We verify that $\mu_3 < 0$, $\mu_4 < 0$, $\mu_5 < 0$, and the signs of μ_1 and μ_2 are ambiguous. As a result, there can be the following four possible scenarios of the signs of $\{\mu_i\}_{i=1}^5$:

μ_1	μ_2	μ_3	μ_4	μ_5
+	+	−	−	−
+	−	−	−	−
−	+	−	−	−
−	−	−	−	−

,

where "+" means positive and "−" means negative. In the first three cases, notice that for the polynomial $H\left(n\right)$, there is one sign change in its coefficients. Therefore, by Descartes's rule of signs, the polynomial $H\left(n\right)$ has a unique positive root. Denote this unique root as $\hat{n}'$, where $H\left(\hat{n}'\right) = 0$. For $n < \hat{n}'$, $H\left(n\right) < 0$ since at $n = 0$, $H\left(0\right) = \mu_5 < 0$ whereas for $n > \hat{n}'$, $H\left(n\right) > 0$. This, in turn, implies that $\frac{\partial Var_t^{NF}(\pi_t^{**})}{\partial m} < 0$ if and only if $n < \hat{n}'$. In the last case, all μ_is are negative, so $H\left(n\right) < 0$ for all $n > 0$. In this case, $\frac{\partial Var_t^{NF}(\pi_t^{**})}{\partial m} < 0$. Without loss of generality, define $\hat{n}' = \infty$ in this case. In sum, we have shown that there exists a unique $\hat{n}'$ such that $\frac{\partial Var_t^{NF}(\pi_t^{**})}{\partial m} < 0$ if and only if $n < \hat{n}'$.

Proof of Proposition 7

Note that, in (34), the first two terms in the inflation volatility $var\left(\pi_t^{**}\right)$ are both strictly increasing in α. In addition, the weight α affects the last two terms of $var\left(\pi_t^{**}\right)$ only through affecting $W_{\bar{s}_{t+1}}(m,n)$ and $W_{\bar{s}_{t+2}}(m,n)$ (i.e., the sensitivity effect) in (31) and (32), respectively. It is straightforward to verify that both $W_{\bar{s}_{t+1}}(m,n)$ and $W_{\bar{s}_{t+2}}(m,n)$ are strictly in α. Hence $var\left(\pi_t^{**}\right)$ is strictly increasing in α.

Proof of Proposition 8

Collecting the first-order condition (45) and matching the coefficients in the equilibrium unemployment rate in Equation (43) with firms' conjecture in Equation (38) yields the following set of equations:

$$\mu = -\frac{1}{b}\left(1 + \frac{a\mu}{1 - a\mu\chi}\right)\frac{a}{1 - a\mu\chi}, \tag{A.34}$$

$$\delta = -\frac{a\mu}{1 - a\mu\chi}, \tag{A.35}$$

$$\lambda = \frac{a}{1 - a\mu\chi}, \tag{A.36}$$

$$\chi \equiv \frac{\delta\tau_\varepsilon}{\delta^2\tau_\varepsilon + \lambda^2\tau_k}. \tag{A.37}$$

Note first that $\delta = 0$ or $\mu = 0$ cannot be an equilibrium. To see this, assume by contradiction that there exists an equilibrium of $\delta = 0$. Hence $\mu = 0$ from Equation (A.35). In addition, $\chi = 0$ from Equation (A.37). Plugging $\mu = 0$ and $\chi = 0$ into Equation (A.34) gives its left-hand side as 0 but the right-hand side as $-\frac{a}{b}$, which implies that $\mu = 0$ does not solve the equation. This is a contradiction.

To determine the equilibrium, substituting Equation (A.35) into Equation (A.36) gives

$$\lambda = -\frac{\delta}{\mu}. \tag{A.38}$$

Substituting Equation (A.35) and Equation (A.36) into Equation (A.34) and using Equation (A.38) yields

$$\mu^2 = \frac{1}{b}\left(1-\delta\right)\delta. \tag{A.39}$$

This implies that $\delta \in (0,1)$ since $\mu^2 > 0$. Substituting Equation (A.38) and Equation (A.39) into Equation (A.37) yields

$$\chi = \frac{\left(1-\delta\right)\tau_\varepsilon}{\left(1-\delta\right)\delta\tau_\varepsilon + b\tau_k}. \tag{A.40}$$

Substituting Equation (A.40) into Equation (A.35) gives

$$\delta = -\frac{a\mu b\tau_k}{\left(1-\delta\right)\delta\tau_\varepsilon + b\tau_k}. \tag{A.41}$$

Since $\delta \in (0,1)$, the denominator of the right-hand side of Equation (A.41) must be positive, i.e., $\left(1-\delta\right)\delta\tau_\varepsilon + b\tau_k > 0$. Hence $\mu < 0$. Using Equation (A.39), we obtain

$$\mu = -\sqrt{\frac{1}{b}\left(1-\delta\right)\delta}. \tag{A.42}$$

Substituting Equation (A.42) into Equation (A.41) yields

$$\delta = \frac{a\sqrt{b\left(1-\delta\right)\delta}\,\tau_k}{\left(1-\delta\right)\delta\tau_\varepsilon + b\tau_k}. \tag{A.43}$$

The equilibrium value of $\delta \in (0,1)$ is determined by Equation (A.43). Using the implicit function theorem, one can show that δ is strictly increasing in a. Substituting (A.42) into Equation (A.38) gives

$$\lambda = \sqrt{\frac{b\delta}{1-\delta}}. \tag{A.44}$$

Since $\delta \in (0,1)$, $\lambda > 0$.

Proof of Proposition 9

This proof complements the steps given in the main text. First, we prove that at the steady state, in each period t, a fraction $\frac{1}{1+\theta}$ of the firms can reset the price whereas the remaining firms are restricted from resetting the price. Denote as Θ_t the fraction of the firms in period t that can reset the price. Hence

$$\Theta_{t+1} = 1 - \Theta_t + \Theta_t (1 - \theta), \tag{A.45}$$

since the fraction $1 - \Theta_t$ of the firms that are restricted from resetting the price in period t gains full price-resetting flexibility in period $t+1$ and the fraction Θ_t of the firms that reset the price in period t can reset the price in period $t+1$ with probability $1-\theta$. In addition, the initial condition for iterating Θ_t is that $\Theta_0 = 1$ (i.e., all firms can reset the price at the beginning). Solving the difference, Equation (A.45) yields

$$\Theta_t = \frac{1}{1+\theta} - \frac{(-\theta)^{t+1}}{1+\theta}. \tag{A.46}$$

At the steady state, $\lim_{t\to\infty} \Theta_t = \frac{1}{1+\theta}$.

Next, we briefly explain that when $\frac{(\alpha-\lambda\kappa)\beta'}{\alpha+\lambda^2} > 0$, the effect of the informational properties (m, n) on the volatilities of inflation and output is qualitatively the same as that in the main analysis. First, note that the hierarchy of the higher-order beliefs $\bar{E}_t^l [k_{t+l}]$ remains the same under the modified Phillips curve, i.e.,

$$\bar{E}_t^l [k_{t+l}] = \bar{k} + \rho_k^{l-1} \left\{ [1 - w(l)] \bar{E}_t \left[k_{t+1} - \bar{k} | s_{t+1}^i, k_t \right] + w(l) \frac{\bar{s}_{t+2} - \bar{k}}{\rho_k} \right\}, \tag{A.47}$$

where $\bar{E}_t \left[k_{t+1} - \bar{k} | s_{t+1}^i, k_t \right] = \frac{q}{q+\frac{mn}{m+n}} \rho_k \left(k_t - \bar{k} \right) + \frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}} \left(\bar{s}_{t+1} - \bar{k} \right)$. Second, substituting the expressions for $\bar{E}_t^l [k_{t+l}]$ into (57) yields

$$\pi_t^{**} = \frac{\alpha\lambda}{\alpha+\lambda^2} k_t + \sum_{l=1}^{\infty} \left(\frac{(\alpha - \lambda\kappa)\beta'}{\alpha+\lambda^2} \right)^{l-1} \frac{\alpha\beta'}{\alpha+\lambda^2} \frac{\alpha(\lambda+\kappa)}{\alpha+\lambda^2} \bar{k}$$

$$
+\left\{\sum_{l=1}^{\infty}\left(\frac{(\alpha-\lambda\kappa)\,\beta'\rho_k}{\alpha+\lambda^2}\right)^{l-1}\frac{\alpha\beta'\rho_k}{\alpha+\lambda^2}\frac{\alpha\,(\lambda+\kappa)}{\alpha+\lambda^2}\right\}\left(k_t-\bar{k}\right)
$$

$$
+\left\{\sum_{l=1}^{\infty}\left(\frac{(\alpha-\lambda\kappa)\,\beta'\rho_k}{\alpha+\lambda^2}\right)^{l-1}\frac{\alpha\beta'}{\alpha+\lambda^2}\frac{\alpha\,(\lambda+\kappa)}{\alpha+\lambda^2}\right.
$$

$$
\left.\times\left[\frac{\bar{E}_t^l\left[k_{t+l}-\bar{k}\right]}{\rho_k^{l-1}}-\rho_k\left(k_t-\bar{k}\right)\right]\right\}, \tag{A.48}
$$

where

$$
\frac{\bar{E}_t^l\left[k_{t+l}\right]-\bar{k}}{\rho_k^{l-1}}-\rho_k\left(k_t-\bar{k}\right)=\left[1-w\left(l\right)\right]\frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\left(\nu_{t+1}+\eta_{t+1}\right)+w\left(l\right)\left(\frac{\nu_{t+2}+\rho_k\nu_{t+1}+\eta_{t+2}}{\rho_k}\right). \tag{A.49}
$$

Define

$$
W'_{\bar{s}_{t+1}}=\frac{\frac{mn}{m+n}}{q+\frac{mn}{m+n}}\frac{\alpha\beta'}{\alpha+\lambda^2}\frac{\alpha\,(\lambda+\kappa)}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{(\alpha-\lambda\kappa)\,\beta'\rho_k}{\alpha+\lambda^2}\right)^{l-1}\left[1-w\left(l\right)\right], \tag{A.50}
$$

$$
W'_{\bar{s}_{t+2}}=\frac{\alpha\beta'}{\alpha+\lambda^2}\frac{\alpha\,(\lambda+\kappa)}{\alpha+\lambda^2}\sum_{l=1}^{\infty}\left(\frac{(\alpha-\lambda\kappa)\,\beta'\rho_k}{\alpha+\lambda^2}\right)^{l-1}\frac{w\left(l\right)}{\rho_k}. \tag{A.51}
$$

Dropping the terms that are independent of (m,n), the inflation volatility $Var\left(\pi_t^{**}\right)$ can be expressed as

$$
\frac{\left[W_{\bar{s}_{t+1}}+\rho_k W_{\bar{s}_{t+2}}\right]^2+W_{\bar{s}_{t+2}}^2}{q}+\frac{W_{\bar{s}_{t+1}}^2+W_{\bar{s}_{t+2}}^2}{m}. \tag{A.52}
$$

Note that (A.52) is the same as the expression (34) for the inflation volatility in the main analysis, except that under the modified Phillips curve, the discounting factor in $\left\{W'_{\bar{s}_{t+1}},W'_{\bar{s}_{t+2}}\right\}$ is $\frac{(\alpha-\lambda\kappa)\beta'\rho_k}{\alpha+\lambda^2}$ instead of $\frac{\alpha\beta\rho_k}{\alpha+\lambda^2}$. However, since $\frac{(\alpha-\lambda\kappa)\beta'\rho_k}{\alpha+\lambda^2}\in(0,1)$, the derivations in the proof of Proposition 4 for the effect of the informational properties (m,n) on the volatilities of inflation and output apply equally. Accordingly, the comparative statics results remain valid.

References

Angeletos, G., and Z. Huo. 2021. "Myopia and Anchoring." *American Economic Review* 111 (4): 1166–1200.

Angeletos, G., and J. La'O. 2013. "Sentiments." *Econometrica* 81 (2): 739–79.

———. 2020. "Optimal Monetary Policy with Informational Frictions." *Journal of Political Economy* 128 (3): 1027–64.

Angeletos, G., and C. Lian. 2016. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." In *Handbook of Macroeconomics*, Vol. 2, ed. J. B. Taylor and H. Uhlig, 1065–1240. Elsevier.

———. 2018. "Forward Guidance without Common Knowledge." *American Economic Review* 108 (9): 2477–2512.

Angeletos, G., and A. Pavan. 2004. "Transparency of Information and Coordination in Economies with Investment Complementarities." *American Economic Review* 94 (2): 91–98.

———. 2007. "Efficient Use of Information and Social Value of Information." *Econometrica* 75 (4): 1103–42.

Ascari, G., and A. Sbordone. 2014. "The Macroeconomics of Trend Inflation." *Journal of Economic Literature* 52 (3): 679–739.

Barro, R., and D. Gordon. 1983a. "A Positive Theory of Monetary Policy in a Natural Rate Model." *Journal of Political Economy* 91 (4): 589–610.

———. 1983b. "Rules, Discretion and Reputation in a Model of Monetary Policy." *Journal of Monetary Economics* 12 (1): 101–21.

Baxter, M., and R. King. 1991. "Productive Externalities and Business Cycles." Discussion Paper No. 53, Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis.

Benhabib, J., P. Wang, and Y. Wen. 2015. "Sentiments and Aggregate Demand Fluctuations." *Econometrica* 83 (2): 549–85.

Brunner, K. 1981. "The Art of Central Banking." Working Paper No. GPB 81-6, Center for Research in Government Policy and Business, University of Rochester.

Calvo, G. A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Clarida, R., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707.

Cooper, R., and A. John. 1988. "Coordinating Coordination Failures in Keynesian Models." *Quarterly Journal of Economics* 103 (3): 441–63.

Cukierman, A., and A. Meltzer. 1986. "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information." *Econometrica* 54 (5): 1099–1128.

Dotsey, M., and R. King. 1986. "Informational Implications of Interest Rate Rules." *American Economic Review* 76 (1): 33–42.

Feltham, G. 1973. *Information Evaluation.* Studies in Accounting Research No. 5. Sarasota, FL: American Accounting Association.

Galí, J. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework.* Princeton, NJ: Princeton University Press.

Geraats, P. 2002. "Central Bank Transparency." *Economic Journal* 112 (483): 532–65.

Goodfriend, M. 1986. "Monetary Mystique: Secrecy and Central Banking." *Journal of Monetary Economics* 17 (1): 63–92.

Goodfriend, M., and R. King. 2012. "The Great Inflation Drift." In *The Great Inflation: The Rebirth of Modern Central Banking*, ed. M. D. Bordo and A. Orphanides, 181–209. University of Chicago Press.

Hansen, L., and T. Sargent. 2007. "Recursive Robust Estimation and Control without Commitment." *Journal of Economic Theory* 136 (1): 1–27.

Hellwig, C., and L. Veldkamp. 2009. "Knowing What Others Know: Coordination Motives in Information Acquisition." *Review of Economic Studies* 76 (1): 223–51.

Holthausen, R., and R. Verrecchia. 1990. "The Effect of Informedness and Consensus on Price and Volume Behavior." *Accounting Review* 65 (1): 191–08.

Ijiri, Y., and R. Jaedicke. 1966. "Reliability and Objectivity of Accounting Measurements." *Accounting Review* 41 (3): 474–83.

King, R. 1982. "Monetary Policy and the Information Content of Prices." *Journal of Political Economy* 90 (2): 247–79.

———. 1983. "Interest Rates, Aggregate Information, and Monetary Policy." *Journal of Monetary Economics* 12 (2): 199–234.

King, R., and A. Wolman, eds. 2022. *Essays in Honor of Marvin Goodfriend: Economist and Central Banker*. Federal Reserve Bank of Richmond.

Kiyotaki, N., and R. Wright. 1989. "On Money as a Medium of Exchange." *Journal of Political Economy* 97 (4): 927–54.

Kydland, F., and E. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (3): 473–91.

Liang, P., and G. Zhang. 2019. "On the Social Value of Accounting Objectivity in Financial Stability." *Accounting Review* 94 (1): 229–48.

Mailath, J., and L. Samuelson. 2006. *Repeated Games and Reputations.* Oxford: Oxford University Press.

McCallum, B. 1990. "Could a Monetary Base Rule Have Prevented the Great Depression?" *Journal of Monetary Economics* 26 (1): 3–26.

Mock, M. 1976. *Measurement and Accounting Information Criteria.* Studies in Accounting Research No. 13. Sarasota, FL: American Accounting Association.

Morris, S., and H. Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.

———. 2005. "Central Bank Transparency and the Signal Value of Prices." *Brookings Papers on Economic Activity* 2: 1–66.

Myatt, D., and C. Wallace. 2012. "Endogenous Information Acquisition in Coordination Games." *Review of Economic Studies* 79 (1): 340–74.

Paciello, L., and M. Wiederholt. 2014. "Exogenous Information, Endogenous Information, and Optimal Monetary Policy." *Review of Economic Studies* 81 (1): 356–88.

Phelps, E. 1983. "The Trouble with 'Rational Expectations' and the Problem of Inflation Stabilization." In *Individual Forecasting and Aggregate Outcomes: 'Rational Expectations' Examined*, ed. R. Frydman and E. S. Phelps, 31–45. Cambridge: Cambridge University Press.

Rogoff, K. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (4): 1169–89.

Rudebusch, G., and J. Williams. 2008. "Revealing the Secrets of the Temple: The Value of Publishing Interest Rate Projections." In

Asset Prices and Monetary Policy, ed. J. Y. Campbell. University of Chicago Press.

Sims, C. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.

Stein, J., and A. Sunderam. 2018. "The Fed, the Bond Market, and Gradualism in Monetary Policy." *Journal of Finance* 73 (3): 1015–60.

Townsend, R. 1983. "Forecasting the Forecasts of Others." *Journal of Political Economy* 91 (4): 546–88.

Winkler, B. 2002. "Which Kind of Transparency? On the Need for Effective Communication in Monetary Policy-making." *Ifo Studies* 48 (3): 401–27.

Woodford, M. 2001. "Imperfect Common Knowledge and the Effects of Monetary Policy." NBER Working Paper No. 8673.

———. 2008. "How Important is Money in the Conduct of Monetary Policy?" *Journal of Money, Credit and Banking* 40 (8): 1561–98.

Yun, T. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles." *Journal of Monetary Economics* 37 (2): 345–70.