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 $v_t \sim N(0, \sigma_v^2)$, 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

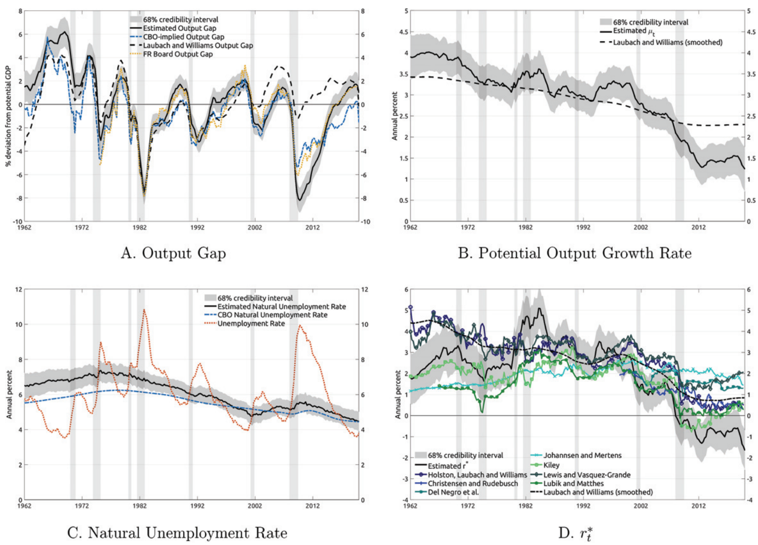

Note: Shaded vertical areas indicate NBER recession periods. Smoothed estimates are reported, except for the r_t^* estimates of other studies in the bottom right panel, which are the filtered estimates.

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

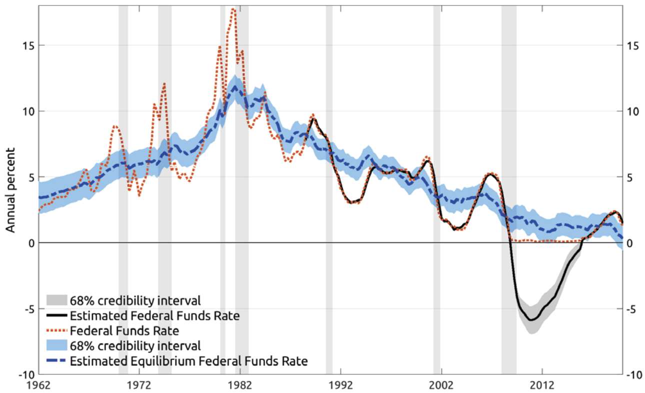

Note: In our estimation, data on the federal funds rate are treated as missing prior to 1987:Q3.

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^*

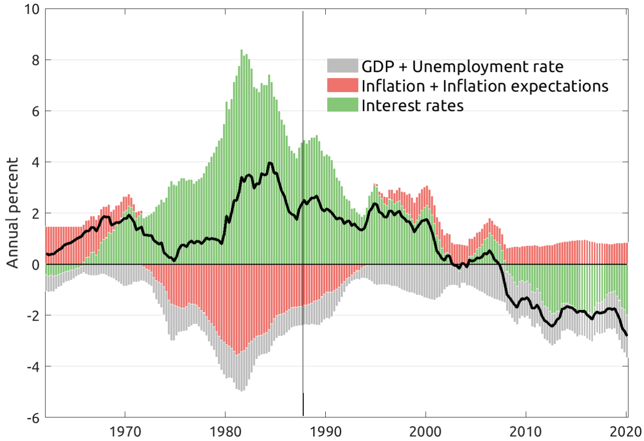

Note: The contributions of GDP and the unemployment rate (GDP + Unemployment rate) have been added together. The same is true for the inflation rate and PTR (Inflation + Inflation expectations), and for the federal funds rate and the 10-year Treasury yield (Interest rates). The gray vertical line indicates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

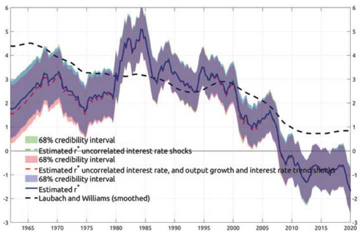

A. Correlation Coefficients

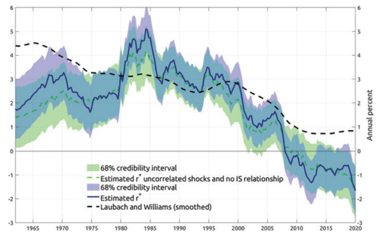

B. IS Relationship

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

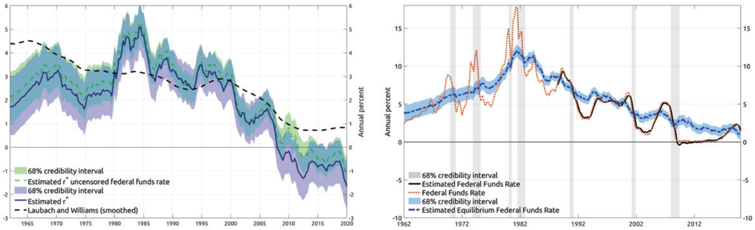

A. r_t^*

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

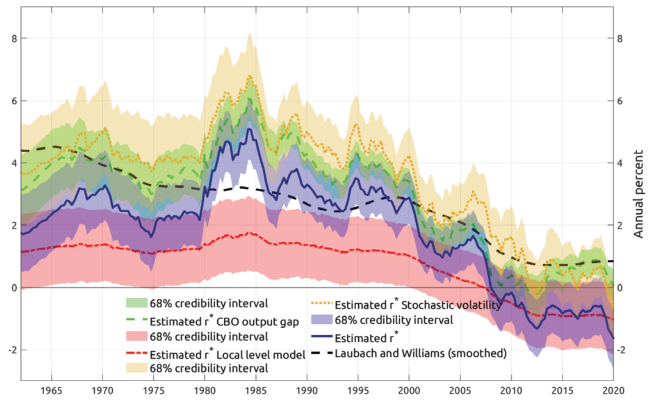

The estimated r_t^* associated with this model specification appears in Figure 6 along with that from our baseline model. Conditioning the estimation of the model on the CBO measure of the output gap yields an estimate of r_t^* about 1 percentage point higher than in our baseline specification, on average. 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

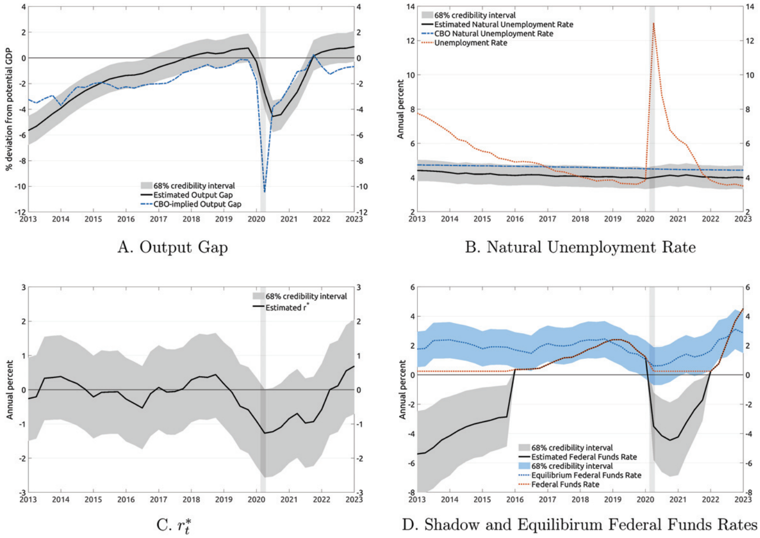

A. Output Gap

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 $\sigma^2_{\eta^R}$ 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 $\sigma^2_{x_0}(\bar{x_0}/s^2_{x_0} + x_1/\sigma^2_{\eta^x})$ and posterior variance $\sigma^2_{x_0} = 1/(1/s^2_{x_0} + 1/\sigma^2_{\eta^x})$, for $x = y^*, \mu, u^*, r^*$, π^*, and p^{10}, where $\bar{x_0}$ and $s^2_{x_0}$ are the prior mean and variance, respectively.

4. Sample $\sigma^2_{\eta^x}$ for $x = y^*, u^*, \pi^*$, and p^{10} from an inverse-gamma distribution with shape coefficient $\underline{a}_{\sigma^2_{\eta^x}} + 0.5 * T$ and rate coefficient $\underline{b}_{\sigma^2_{\eta^x}} + 0.5 * \hat{\eta}^{x\intercal}\hat{\eta}^x$, where $\hat{\eta}^x$ is the vector of residuals obtained from $x_t - x_{t-1}, t = 1, 2, \ldots, T$, and $\underline{a}_{\sigma^2_{\eta^x}}$ and $\underline{b}_{\sigma^2_{\eta^x}}$ are the prior shape and rate coefficients.

5. Sample $\sigma^2_{\eta^\mu}$, $\sigma^2_{\eta^{r*}}$, 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

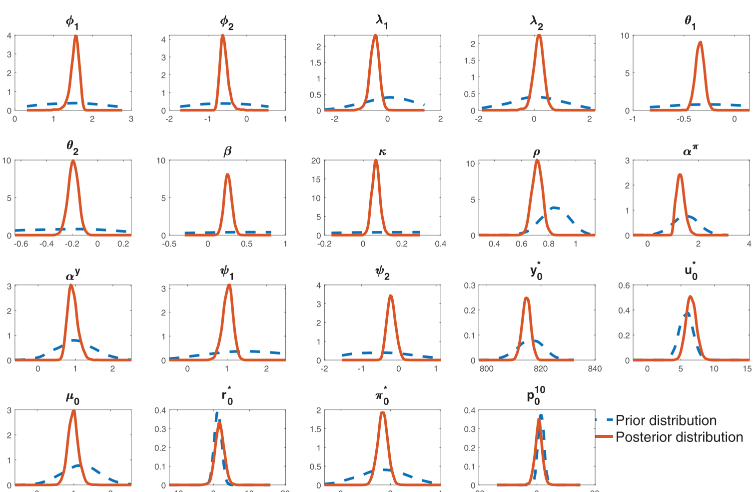

Figure F.2. Prior and Posterior Distributions: Variance Parameters

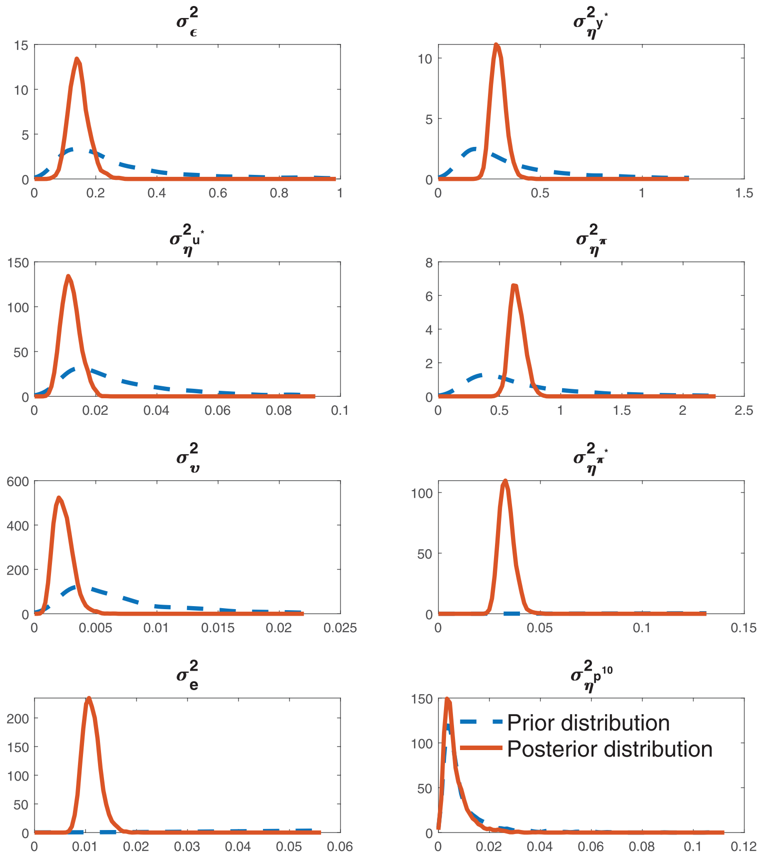

Figure F.3. Prior and Posterior Distributions: Variance and Correlation Parameters

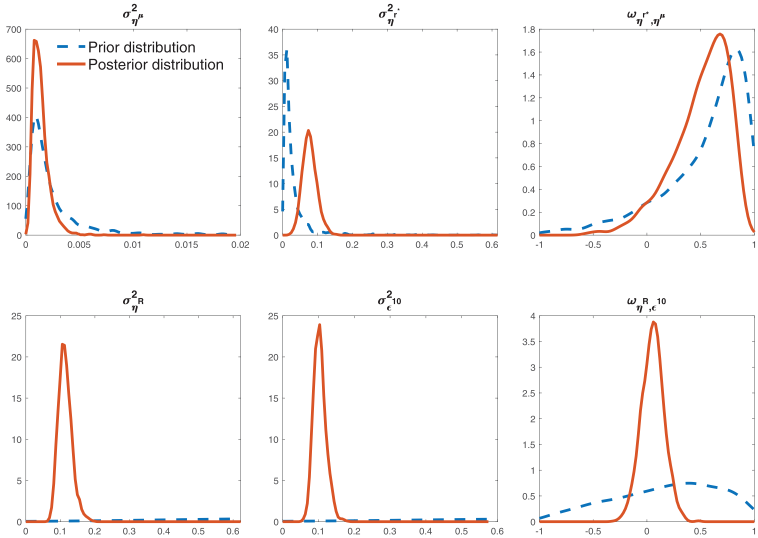

Appendix G. Data Decomposition

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

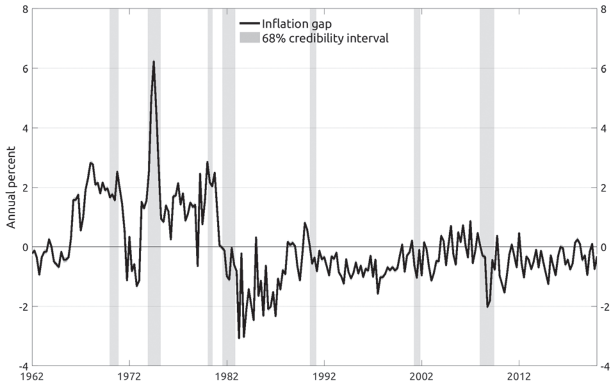

Note: Shaded vertical areas indicate NBER recession periods. Smoothed estimates are reported. The inflation gap is defined as the posterior mean of $\pi_t - \pi_t^*$.

Figure G.2. Historical Data Decomposition of the Output Gap

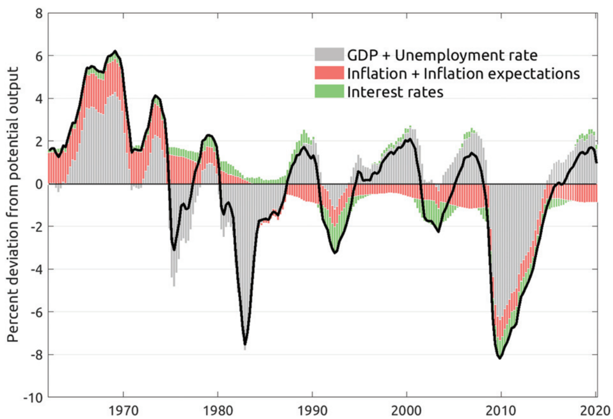

Note: The contributions of GDP and the unemployment rate (GDP + Unemployment rate) have been added together. The same is true for the inflation rate and PTR (Inflation + Inflation expectations), and for the federal funds rate and the 10-year Treasury yield (Interest rates).

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

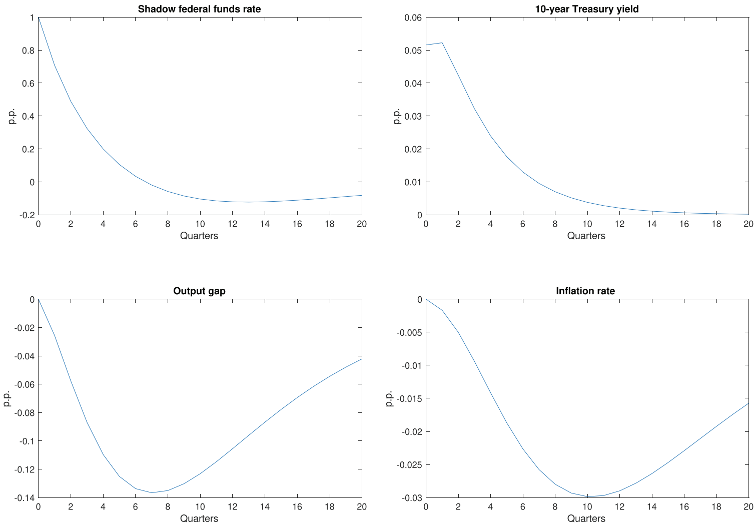

A. Shock to the Federal Funds Rate

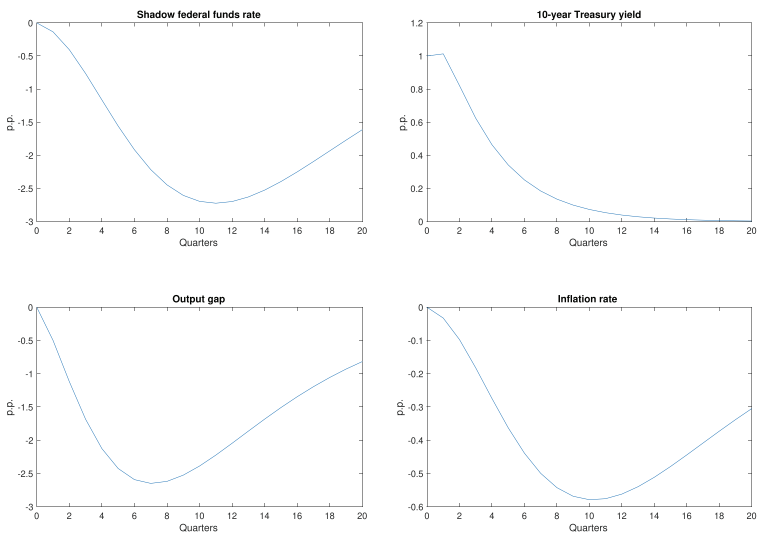

B. Shock to the Long-Term Interest Rate Cycle

(*continued*)

Figure H.1. (Continued)

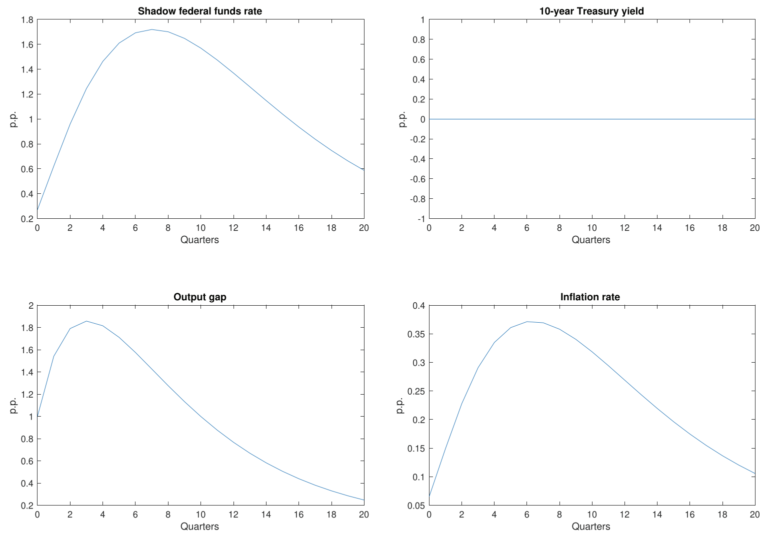

C. Shock to the Output Gap

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

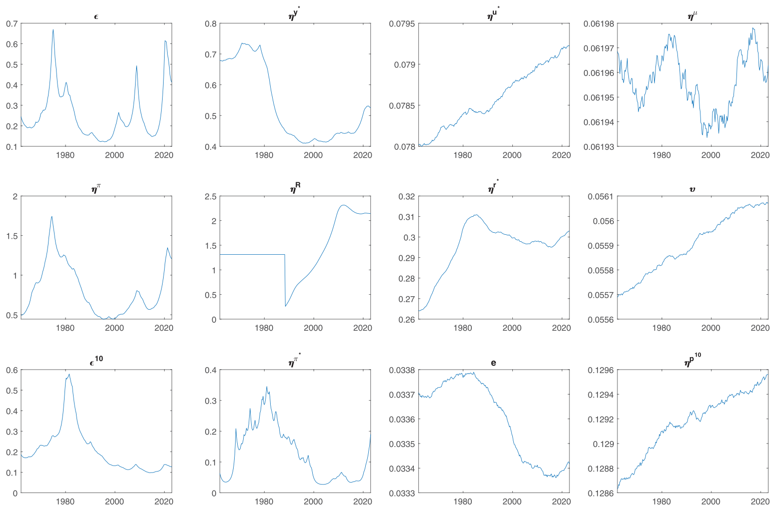

Figure I.2. Parsing of the COVID-19 Pandemic Period

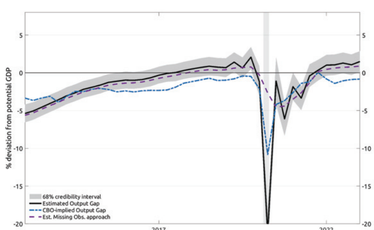

A. Output Gap

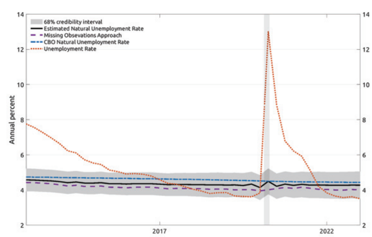

B. Natural Unemployment Rate

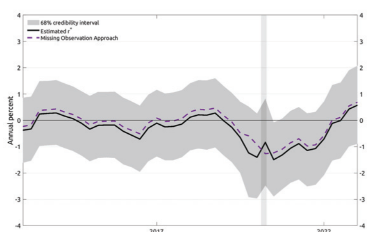

C. r_t^*

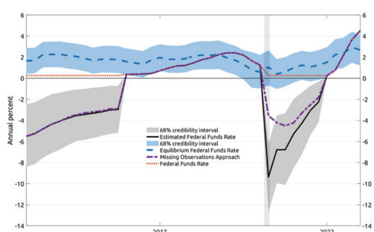

D. Shadow Federal Funds Rate

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

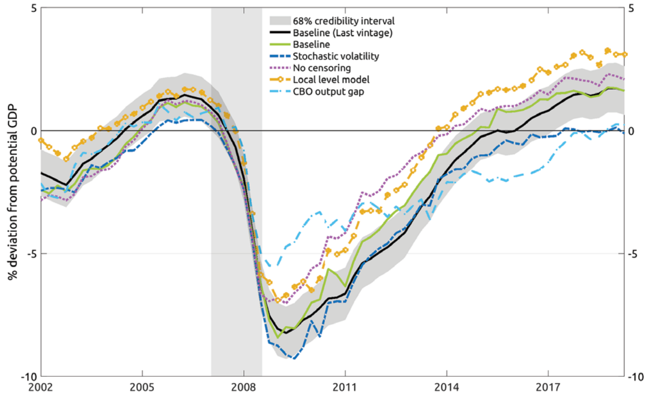

Note: Shaded vertical area indicates NBER recession period.

Figure J.2. Pseudo-Real-Time Estimates of r_t^* across Models

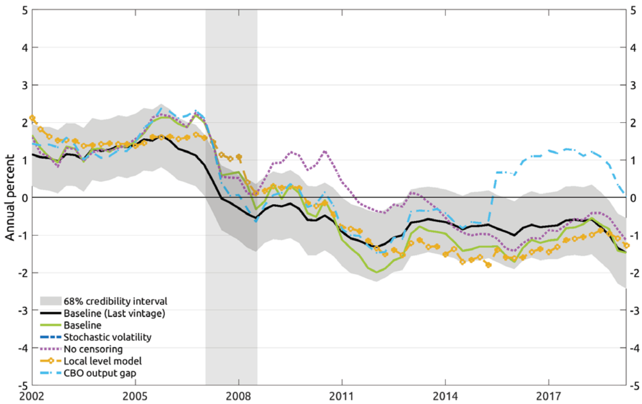

Note: Shaded vertical area indicates NBER recession period.

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.