BIS Working Papers
No 1349

Learning monetary policy strategies at the effective lower bound with sudden surprises

by Spencer Krane, Leonardo Melosi and Matthias Rottner

Monetary and Economic Department

May 2026

JEL classification: E52, C63, E31

Keywords: new framework, central bank's communications, deflationary bias, inflation surprises, asymmetric average inflation targeting, imperfect credibility, liftoff, Bayesian learning

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Learning Monetary Policy Strategies at the Effective Lower Bound with Sudden Surprises*

Spencer Krane
FRB Chicago

Leonardo Melosi
European University Institute & CEPR

Matthias Rottner
BIS & Deutsche Bundesbank

May, 2026

Abstract

We examine how private sector agents might learn a new monetary strategy introduced while policy rates are at their effective lower bound (ELB) in an environment with large inflationary and deflationary shocks. We consider the adoption of a new asymmetric average inflation targeting rule aimed at countering the disinflationary bias imparted by the ELB. The most crucial time for learning runs from when rates would be near liftoff under the old strategy through early liftoff under the new rule. Recessionary shocks during this time could delay learning while large inflationary shocks could outright stop it, inhibiting the ability of the new strategy to address the costs associated with the ELB. Using the US post-Covid experience as an example, we also find that the monetary policy shocks can have important feedback on the learning process.

Keywords: New framework, central bank's communications, deflationary bias, inflation surprises, asymmetric average inflation targeting, imperfect credibility, liftoff, Bayesian learning.
JEL classification: E52, C63, E31

*Emails: spencer.krane@chi.frb.org, lmelosi@gmail.com, matthias.rottner@bis.org. We thank seminar participants at the Workshop on Central Bank Communication organized by the University of Oxford and the CEPR, Federal Reserve Bank of Chicago, and Deutsche Bundesbank. The views in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Bank for International Settlements (BIS), Deutsche Bundesbank, the Federal Reserve Bank of Chicago, the Eurosystem or the Federal Reserve System.

1 Introduction

The past 20 years have featured major changes to the monetary policy environment. With the onset of the Great Recession and ensuing slow economic recovery, the effective lower bound (ELB) on policy rates went from being an isolated curiosity to a prominent feature of the monetary policy environment. Equilibrium policy rates fell below where they were in the 1980s and 1990s in many advanced economies, making the ELB a threat with even a modest recession. Meanwhile, inflation remained stubbornly subdued for an entire decade.

In response, central banks adopted new strategies aimed at offsetting the risks posed by the ELB. Notably, ELB risks were stated as an important motivating factor in the Federal Reserve's revised statement on Longer-Run Goals and Monetary Policy Strategy it adopted in August 2020 and the ECB's Monetary Policy Strategy Statement issued in July 2021. For both central banks, these changes were made when policy rates were still at the ELB during the Covid-19 pandemic. These economies recovered quite swiftly following the pandemic, but then a mix of large shocks suddenly caused inflation to reach levels unseen in the past 40 years.

In this paper, we investigate how private sector agents might learn such new monetary strategies introduced at the ELB in a stochastic environment in which large inflationary and deflationary shocks may occur unexpectedly. The rule we consider is a form of asymmetric average inflation targeting. The exercise highlights several important considerations in implementing such strategies emanating from interactions between the learning process and economic outcomes. First, the period around when the liftoff of rates from the ELB would have occurred under the old rule is key for learning the new rule. Second, shocks that either return the economy to the ELB or that cause a surge in inflation that quickly satisfies the inflation averaging goal can curtail agents learning the new rule and prevent the economy from capitalizing on its long-run benefits. Third, monetary policy shocks– perhaps due to discretionary actions by the central bank–come with a set of costs and benefits due to their interaction with the learning process.

The theoretical foundations for strategies aimed at countering the ELB center on research showing that under traditional monetary strategies the limits on reducing policy rates presented by the proximity of the ELB imparts a downward bias to inflation and inflation expectations relative to the central bank's target. The expectational bias is a feature of the distribution of inflation over time; it's not just an occurrence associated with a large negative shock to the economy. In turn, this bias impinges on the central bank's ability to stabilize both output and inflation.[1]

This literature has proposed a range of alternative monetary policy frameworks to address this inflation bias and improve overall macroeconomic performance. These include average inflation targeting, price level targeting, shadow rate makeup polices, running "dovish policies" away from the ELB, and asymmetric reaction functions and target ranges. These policies can reduce the odds of returning to the ELB and the length of ELB episodes that do occur. Some can be calibrated to eliminate the downward inflation bias altogether.[2] However, these theoretical solutions generally involve strong assumptions: completely rational and forward-looking agents; complete credibility of the monetary authority; and no adjustment lags or other inertia in the economy. Under such conditions, these models offer clear policy prescriptions and their implementation works perfectly. Rational agents immediately align their views with the bias-corrected distribution of inflation generated by the new policy rule, substantially reducing the costs of ELB episodes.

Of course, the real world is not so clean. New monetary policy strategies need to be learned and their credibility established. This process can be complicated even away from the ELB as agents need to estimate new policy parameters, and it can take time to hone in on precise values given the noisy relationships between policy rates and economic outcomes.[3] But there are added challenges if a new policy rule is introduced while rates are mired at the ELB because there is no co-movement between economic variables and interest rates to inform agents. The central bank's proclamations may lead agents to form priors, but there is little they can do in terms of statistical inference until they actually observe different policy outcomes from what would have occurred under the old strategy. The events surrounding the Covid-19 pandemic highlight other challenges to learning the new rule. A new recession may force the economy back to the ELB, or a sharp rise in inflation may dictate an aggressive policy tightening. If either event occurs before the new rule is learned completely, the benefits of the alternative framework will be curtailed, at least for a time.

[1] See Adam and Billi (2007), Nakov (2008) and Hills et al. (2019).

[2] Discussions of average inflation and price level targeting are found in Evans (2012), Bernanke (2017), Hebden and López-Salido (2018), and Bernanke et al. (2019). Shadow rate make up policies were proposed in Reifschneider and Willams (2000). Dovish policies are discussed in Mertens and Williams (2019). Bianchi et al. (2021) examine asymmetric reaction functions and target ranges. Similar asymmetric strategies are advocated by Caballero and Simsek (2023) during temporary supply contractions when aggregate demand has inertia. Arias et al. (2020) study the performance of many of these rules using stochastic simulations of the FRB/US model.

[3] Erceg and Levin (2003) for instance study a model in which agents have to learn the inflation target of the central bank. Barnett and Ellison (2013) study how the decision of a central bank to disinflate the economy influences the learning of economic agent and that of the central bank itself.

We examine these issues in the context of a standard New Keynesian framework in which the economy is driven to the ELB by large negative demand shocks. The central bank then announces it is adopting a new policy rule to counter the ELB. But we also assume the central bank does not announce all of the elements of the new policy function, or cannot do so with full credibility. Instead, agents must learn them.

The new rule we consider is a form of asymmetric average inflation targeting, which adds a term in the weighted average of past shortfalls of inflation from target to the central bank's old policy function. The weighted average inflation shortfall is known to agents, but the intensity of the central bank's response to the shortfall is not. Importantly, our framework features two regions in which learning effectively comes to a halt: a disinflationay recession that drives rates to the ELB, and at high inflation, where the old and the new policy rules coincide. In the first region, learning is impaired because agents cannot observe the central bank's monetary stance. In the second, the model captures a key feature of makeup strategies: their objective is to address the deflationary bias associated with ELB risk, rather than to induce additional tightening following past inflation overshoots.[4]

The new rule contains two policy shocks, the typical shock that entered the old rule plus a new one that is operative only when past inflation shortfalls are influencing policy. Although modeled as random shocks, these terms could capture discretionary deviations from the policy rule by the central bank. Agents know such shocks can occur, but have to learn about their volatility.

Agents form priors over the unknown parameters when the rule is announced. We then roll the problem sequentially forward through time, with agents using Bayesian updating to revise their estimates of the unknown parameters each period. The process is constructed so that the updated estimates are consistent with equilibrium output and inflation on a period-by-period basis. We let the negative demand shocks dissipate over time, so that after some time the economy is able to exit the ELB and then eventually converge to its stochastic steady-state.

Little can be learned about the parameters of the new policy function as long as economic conditions are such that the shadow rate under the old monetary policy rule is well below the ELB - for all agents know, the central bank is engaging in cheap talk, and there is no reason to move away from their priors. The information structure evolves as

[4]This appears consistent with the 2020 Federal Reserve and 2021 ECB monetary policy strategy statements, which discuss overshooting following a period of below-target inflation but did not mention seeking undershooting following a period of above-target inflation.

the economy improves enough that rates would be near liftoff from the ELB under the old policy rule, and then changes radically when the old rule would prescribe raising rates but the central bank keeps rates at the ELB. Agents can now use the difference between the prescriptions of the old rule and the ELB to make more accurate inferences about the new rule's unknown parameters. And once liftoff from the ELB occurs, agents can hone in more precisely on the new policy function. This means that the key time for policymakers to establish the credibility of the new strategy is around the period when the old rule would have called for liftoff and in the early stages of rate increases under the new regime.

Technically, this evolution of the learning process reflects the fact that the likelihood function component of the posterior distribution of the unknown parameters resembles a Tobit regression in which the modification to the interest rate rule is the dependent variable, the explanatory variable is the average inflation shortfall, the error term is the policy shocks, and the censoring point —below which observed rates are at the ELB— is where the new monetary policy rule prescribes liftoff. Deep in recession, when there is little chance of being near the censoring point, the likelihood is flat and the posterior puts almost all its weight on the prior. As a positive rate becomes increasingly probable, the slope of the likelihood steepens and moves the posterior away from the prior; the information content of the data increases further once liftoff occurs and non-censored observations enter the likelihood function.

In terms of macroeconomic outcomes, the new policy rule is calibrated to offset the disinflationary bias in the ELB and produce long-run inflation expectations that are anchored at target in a full information environment. In the learning problem, and under a baseline scenario with no shocks to the policy function, agents' long-run inflation expectations are stuck below target until just before the old policy function would have called for a liftoff in rates. Inflation expectations subsequently rise towards target, with the pace quickening once liftoff would have occurred under the old rule. This path of expectations convergence contrasts the immediate jump to the inflation target seen in the full-information analysis. But it is still a marked improvement relative to the old policy regime, in which long-run inflation expectations would remain stuck below target. Correspondingly, the paths for output and inflation generated by the new rule with learning are an improvement over outcomes under the old regime, though under our baseline parameter assumptions they still fall short of the full information economy.

We next consider what happens when the economy experiences either another round of recessionary shocks that drive it back to the ELB or a set of shocks that generates a bout of inflation, with the shocks hitting before the learning process is complete.[5] In the case of recessionary shocks, the new rule performs better than the old rule to the degree that the new parameters have been learned prior to the shocks. Learning stalls during the return to the ELB, but then re-emerges as the economy makes its second exit. In the case of inflationary shocks, the new model will generate higher inflation than the old rule until the average inflation shortfall is eliminated, after which the inflation paths are similar. But at this time as the outcomes of the two rules are almost identical, learning about the new rule's new parameters also ceases. Two important lessons emerge. First, because high inflation will eventually drive the average past inflation shortfalls to zero, there is a natural limit to the "extra" inflation the new rule can generate. Second, if the inflation shortfall is driven to zero before the new rule is learned completely, then the economy will still suffer extra losses in output and inflation when the next ELB-inducing recessionary shocks hit.

The detrimental outcomes associated with the inflation shock may lead the central bank to consider purposefully raise rates less aggressively in order to emphasize its commitment to the new policy rule and spur the learning process. Alternatively, it might consider a more aggressive response to the inflation shock to insure that the adoption of the new rule is not misperceived as a lack of commitment to both sides of the inflation target. We find that both alternatives can have their intended effects, but they also come with the cost of increased uncertainty over the policy function that the central bank would need to weigh when considering such interventions.

As an empirical example, we look at how our highly stylized model interprets the post-Covid experience in the US, when inflation surged soon after the adoption of flexible average inflation targeting by the Federal Reserve. We find that matching the data requires large monetary policy shocks that hold rates much lower than the new policy rule would prescribe based on output and inflation alone. In turn, these shocks feed back on to the learning process, accelerating agents' learning the new policy rule's response to the inflation shortfall, but also causing them to overestimate the volatility of monetary policy shocks. This improves the ability of the new rule to offset the disinflationary bias from the ELB, but comes at the costs associated with agents' misperception of volatility. We also look at how policy shocks that fight inflation more aggressively interact with the learning

[5]If learning is completed by the time the shocks hit, by construction the learned-rule performs just like the full information model, with much small declines in output and inflation than under the old rule during recessions. And because the past average inflation shortfall has been eliminated by the time learning is completed, the new and old rules essentially look alike during an inflation shock.

process.

To examine tradeoffs, we calculate the gains from learning through stochastic simulations of a standard quadratic loss function of deviations over time in output and inflation from their steady states. This indicates that even a modest amount of learning about the policy response to the shortfall in inflation has large returns in terms of reduced incidence of ELB episodes and stabilization. Losses due to perceived policy noise appear smaller, but this might be an artifact of our model, in which the costs of hitting the ELB dominate higher frequency volatility in output and inflation.

The paper is related to a growing literature on learning new policy strategies and the macroeconomic implications of the learning process. Andolfatto and Gomme (2003) and Erceg and Levin (2003) develop structural models with monetary policy regime changes in which private sector agents have imperfect information about the regimes. Schorfheide (2005) estimates this type of models using likelihood methods. Bianchi and Melosi (2016, 2018) develop a class of models where agents are uncertain about the persistence of policy regimes. These papers, however, do not study an environment that includes the ELB, and therefore do not consider how agents' learning problems might proceed when their information sets are censored.

Gust et al. (2018) study a learning problem with the ELB constraint in which agents must infer the intercept of a monetary rule that is governed by an exogenous discrete Markov-switching process in which all states have a chance of occurring in any period. Similar to the learning problem in our model, episodes at the ELB generate a Tobit-style likelihood function. But in our model agents face a brand new policy rule introduced at the ELB and they need to learn the central bank's response to past misses of inflation, which are endogenous. This endogeneity combined with the asymmetry of the makeup strategy give rise to a second type of censoring in which sudden inflation surprises outweigh past inflation shortfalls and abruptly stop the learning process.

Bodenstein et al. (2022) study learning in a model in which a binding ELB constraint impairs private agents' ability to learn newly introduced make-up strategies. By contrast, our framework features an additional region in which learning also comes to a halt, namely at high levels of inflation where the new rule coincides with the old rule. Moreover, our paper breaks new ground by providing an empirical evaluation of a model in which agents learn about a newly introduced policy strategy at the ELB, a setting motivated by recent historical episodes in the United States and other economies. Finally, their learning is sequential, so agents' beliefs about the parameters of the new rule are predetermined each period.[6] Hebden et al. (2020) use the FRB/US model to consider how learning is delayed if a new makeup policy is introduced while at the ELB compared with full information or if the new policy is introduced before the ELB episode, with resulting negative effects on macroeconomic outcomes. They limit their analysis to shocks that generate the initial ELB episode.

Other recent papers study mechanisms to explain the surge and subsequent decline in inflation following the pandemic through the lens of agents' beliefs about policymakers' behavior. A common theme in this work is that inflation dynamics depend on how private agents interpret the interaction between monetary and fiscal policy, in the tradition of the fiscal theory of the price level (FTPL). Bassetto and Miller (2025) introduce an information channel in which agents endogenously choose whether to acquire costly information about future fiscal surpluses backing government debt. This generates state-dependent inflation dynamics, with low sensitivity to fiscal news in normal times but potentially sharp increases in inflation responsiveness and even discontinuous jumps when beliefs shift toward fiscal stress and agents begin acquiring information. Bianchi et al. (2023) show that 2021 US fiscal stimulus significantly shifted expectations about the degree of fiscal backing of U.S. federal debt, accounting for a substantial portion of the subsequent surge in inflation. Barro and Bianchi (2026) provide empirical evidence that post-pandemic inflation across OECD countries is strongly related to a theoretically derived measure of fiscal expansion that scales spending increases by initial debt levels and debt maturity.

Other studies evaluate the macroeconomic and welfare implications of introducing makeup strategies under the assumption that agents are bounded rational and, therefore, have an imperfect knowledge of the economy they live in (Amano et al., 2020; Budianto et al., 2023; Honkapohja and McClung, 2023). In our study (as well as in Gust et al., 2018 and Bodenstein et al., 2022) agents know the structure of the economy except for certain policy parameters. Limiting the scope of the problem in this manner allows us to concentrate on mechanisms associated with learning the new rule and avoid complications arising from other sources of misspecification or uncertainty.

Another related strand of literature studies optimal monetary policy in environments where private-sector expectations are not fully model-consistent, emphasizing the role of imperfect knowledge in shaping policy design. For instance, Adam and Woodford (2012)

[6] This timing facilitates solving the learning problem analytically with Kalman filtering techniques. However it also means agents' parameter updating does not have a contemporaneous general equilibrium impact on economic outcomes, where in our model and in Gust et al. (2018), parameter learning, output, and inflation are simultaneously determined outcomes. This, for example, would influence the trade offs facing the central bank when making discretionary policy actions.

study optimal monetary policy when private-sector expectations may deviate from full rationality, deriving policies that are robust to near-rational belief distortions without modeling learning explicitly. Mele et al. (2020) instead consider an economy in which agents learn the mapping from shocks to outcomes, showing that gradual belief updating can make price level stabilization suboptimal by creating incentives to exploit transition dynamics. By contrast, our paper studies the learning of a newly introduced monetary policy strategy. Learning is state-dependent—impaired by censoring at the ELB and potentially halted when inflation is high—and interacts directly with shocks and policy deviations.

2 The model

We start with an off-the-shelf New Keynesian model, but work with its nonlinear formulation and not the log-linearization. The model features a standard Euler equation and Phillips curve (using Rotemberg pricing), where π_t is inflation, y_t is output, π and y are their deterministic steady state values, mc_t are marginal costs, and $\Lambda_{t,t+1}$ is the stochastic discount factor:

$$1 = \beta R_t \mathbb{E}_t \left[\left(\frac{\epsilon_{d,t+1}}{\epsilon_{d,t}} \right) \left(\frac{c_t}{c_{t+1}} \right)^{\sigma} \frac{1}{\pi_{t+1}} \right] \tag{1}$$

$$\varphi \left(\frac{\pi_t}{\pi} - 1 \right) \frac{\pi_t}{\pi} = (1 - \epsilon) + \epsilon mc_t + \epsilon_{\mu,t} + \varphi \mathbb{E}_t \left[\Lambda_{t,t+1} \left(\frac{\pi_{t+1}}{\pi} - 1 \right) \frac{\pi_{t+1}}{\pi} \frac{y_{t+1}}{y_t} \right] \tag{2}$$

There are two sources of uncertainty in this part of the model: the preference shock in the Euler equation, $\epsilon_{d,t}$, and the markup shock, $\epsilon_{\mu,t}$, in the Phillips curve. Both follow an AR(1) process. All of the parameters in these equations and the shock process are known to the agents and the central bank.

The new features of the model involve the announcement by the central bank that it will now follow an asymmetric average inflation targeting monetary policy rule given by:

$$R_t = \max \left[1, R \left(\frac{\pi_t}{\pi} \right)^{\phi_\Pi} \left(\frac{y_t}{y} \right)^{\phi_Y} + \sigma_M \epsilon_{M,t} + \eta_t^z \cdot \mathbf{1}_{z_t < 0} \right] \tag{3}$$

$$\eta_t^z = \xi_z z_t + \sigma_\eta \epsilon_{\eta,t} \tag{4}$$

$$z_t = \rho_z z_{t-1} + \phi_z \left(\pi_{t-1} - \pi \right) \tag{5}$$

R_t is the (gross) nominal interest rate, and it equals the maximum of the ELB, which we assume is the zero lower bound of 1, and the interest rate determined by the new rule.[7]

The first term in the new rule is the usual Taylor prescription, where R is the equilibrium real rate of interest and the second term is the usual monetary policy shock, $\epsilon_{M,t} \sim \mathcal{N}(0,1)$.[8] This is the policy rule followed prior to the introduction of the new rule. The parameters ϕ_Π, ϕ_Y, and σ_M are all know to the agents and the central bank. If we just stopped here, the model would display the familiar features of the simple three-equation New Keynesian model. Notably, future events would have a powerful influence on the present, so that the chance of reaching the ELB sometime in the future induces a downward bias to inflation expectations that can impede economic performance today.

Our model differs from the usual specification through the third term in (3), $\eta_t^z \cdot \mathbf{1}_{z_t<0}$, the adjustment to the policy rule announced by the central bank, where $\mathbf{1}_{z_t<0}$ is an indicator function for when $z_t < 0$. This adjustment is a function of the weighted average of past deviations (positive or negative) of inflation from target, z_t, when that average represents a net shortfall from target, the intensity of the central bank's response to the shortfall, $\xi_z \geq 0$, and a random shock, $\epsilon_{\eta,t} \sim \mathcal{N}(0,1)$, and is in place as long as z_t is negative. Confining the adjustment to periods when average inflation is below target reflects the new rule's explicit mission to counter the downward inflation bias generated by the ELB. We calibrate the new rule so that it exactly offsets the ELB inflation bias in a full-information setting. We include a second policy shock, $\epsilon_{\eta,t}$, to account for the possibility that the new rule could come with implementation issues over and above those experienced in the past, potentially making policy more unpredictable. We assume $\epsilon_{M,t}$ and $\epsilon_{\eta,t}$ are independent and define the total monetary policy shock, $\epsilon_{R,t} \sim \mathcal{N}(0, \sigma_{R,t})$, as:

$$\epsilon_{R,t} = \epsilon_{M,t} + \epsilon_{\eta,t} \cdot \mathbf{1}_{z_t<0} \tag{6}$$

$$\sigma_{R,t} = \sigma_M + \sigma_\eta \cdot \mathbf{1}_{z_t<0} \tag{7}$$

The functional form of the new rule is known to the public, as are the parameters determining the weighted average of inflation shortfalls, ρ_z and ϕ_z. However, we assume that even though the central bank can announce a new policy, it cannot do so with full credibility. This is captured by assuming agents do not know the intensity of the central bank's response to inflation shortfall, ξ_z, nor the standard deviation, σ_η, of the additional error term in the policy function adjustment, $\epsilon_{\eta,t}$. Agents use Bayesian updating

[7] Our framework is general enough to incorporate a negative lower bound, including one that is endogenously determined within a structural model as in Darracq Pariès et al. (2023).

[8] The shock is introduced additively to ease the solution of the model.

to estimate all unknown parameters, and then act as if policy is dictated by their resulting perception of the monetary policy rule, $\hat{R}_t$:

$$\hat{R}_t = \max\left[1, R\left(\frac{\pi_t}{\pi}\right)^{\phi_\Pi}\left(\frac{y_t}{y}\right)^{\phi_Y} + \sigma_M \hat{\epsilon}_{M,t|t} + \hat{\eta}_t^z \cdot \mathbf{1}_{z_t<0}\right] \tag{8}$$

$$\hat{\eta}_t^z = \hat{\xi}_{z|t} z_t + \hat{\sigma}_{\eta|t}\hat{\epsilon}_{\eta,t|t} \tag{9}$$

$$z_t = \rho_z z_{t-1} + \phi_z\left(\pi_{t-1} - \pi\right) \tag{10}$$

where $\hat{\xi}_{z|t}$, $\hat{\sigma}_{\eta|t}$, $\hat{\epsilon}_{M|t}$ and $\hat{\epsilon}_{\eta|t}$ are the estimates agents make at time t for ξ_z, σ_η, $\epsilon_{M,t}$ and $\epsilon_{\eta,t}$.

3 Estimating the new monetary policy rule

Suppose $z_t < 0$ but that the ELB is not binding under the new rule, so $R_t > 1$ is observed. Then agents also can observe η_t^z. Given $\epsilon_{R,t} \sim \mathcal{N}(0,1)$, the likelihood of this observation is:

$$\frac{1}{\sigma_{R,t}}\phi\left(\frac{\eta_t^z - \xi_z z_t}{\sigma_{R,t}}\right)$$

where ϕ is the standardized normal density.

When the ELB is binding under the new rule, R_t is censored at 1 and η_t^z is unobservable. All agents know is that η_t^z is small enough to keep rates from lifting off, and hence is less than an upper bound, $\bar{\eta}_t^z$. The likelihood of this event occurring is given by:

$$\Phi\left(\frac{\bar{\eta}_t^z - \xi_z z_t}{\sigma_{R,t}}\right) \tag{11}$$

where

$$\bar{\eta}_t^z = 1 - R\left(\frac{\pi_t}{\pi}\right)^{\phi_\Pi}\left(\frac{y_t}{y}\right)^{\phi_Y} > \eta_t^z \tag{12}$$

and Φ is the standardized cumulative normal density. Note that $\bar{\eta}_t^z$ is the difference between the ELB and the shadow interest rate under the old monetary policy rule. The bound thus reveals how much work η_t^z has to do to prevent liftoff under the old rule given the observed values for output and inflation.

With both censored and uncensored observations, the likelihood function is a Tobit regression in which η_t^z is the dependent variable, z_t is the explanatory variable, the unknown parameters are ξ_z and $\sigma_{R,t}$ and the censoring point is $\bar{\eta}_z$. However, our problem is different than the standard Tobit, because we assume the central bank announces the new policy framework at time $t = 1$ when $R_t = 1$, so at the start of the learning process agents only have censored observations. Accordingly, maximum likelihood estimation is infeasible during this time.[9]

Instead, we assume at the time of the announcement, agents form the perceived monetary policy rule, equation 10, and a set of priors over the unknown parameters, $\Omega(\xi_{z|P}, \sigma_{R|P})$. The agents then engage in sequential Bayesian updating as they move forward through time. Specifically, in every period when $z_t < 0$, they choose $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R|t}$ to maximize the posterior distribution:

$$\prod_{\tau \leq t} \left\{ \left[\frac{1}{\hat{\sigma}_{R,t|t}} \phi \left(\frac{\eta_\tau^z - \hat{\xi}_{z|t} z_\tau}{\hat{\sigma}_{R,t|t}} \right) \right] \mathbb{I}\left(R_\tau > 1\right) \cdot \left[\Phi \left(\frac{\bar{\eta}_\tau^z - \hat{\xi}_{z|t} z_\tau}{\hat{\sigma}_{R,t|t}} \right) \right] \mathbb{I}\left(R_\tau = 1, z_\tau < 0\right) \right\} \cdot \Omega\left(\xi_{z|P}, \sigma_{R|P}\right) \tag{13}$$

where $\mathbb{I}(\cdot)$ is the indicator function.

Once $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R,t|t}$ are in hand, if the ELB is not binding $\hat{\epsilon}_{R,t|t}$ can be estimated as the residual that sets $\hat{R}_t = R_t$:

$$\hat{\epsilon}_{R,t|t} = \hat{\epsilon}_{\eta,t|t} + \hat{\epsilon}_{M,t|t} = \frac{\eta_{z,t} - \hat{\xi}_{z|t} z_t}{\hat{\sigma}_{R,t|t}} \tag{14}$$

or if the ELB is binding as the expected value of the standard truncated normal:

$$\hat{\epsilon}_{R,t|t} = \hat{\epsilon}_{\eta,t|t} + \hat{\epsilon}_{M,t|t} = \mathbb{E}\left(\epsilon_{z,t} | \hat{\xi}_{z|t}, \hat{\sigma}_{R,t|t}\right) = -\hat{\sigma}_{R,t|t} \frac{\phi\left(\frac{\bar{\eta}_{z,t} - \hat{\xi}_{z|t} z_t}{\hat{\sigma}_{R,t|t}}\right)}{\Phi\left(\frac{\bar{\eta}_{z,t} - \hat{\xi}_{z|t} z_t}{\hat{\sigma}_{R,t|t}}\right)} \tag{15}$$

Since $\epsilon_{M.t}$ and $\epsilon_{\eta,t}$ are independent normals, they can be estimated as:

$$\hat{\epsilon}_{\eta,t|t} = \hat{\epsilon}_{R,t|t} \frac{\hat{\sigma}_{\eta|t}}{\hat{\sigma}_{R,t|t}}, \quad \text{and} \quad \hat{\epsilon}_{M,t|t} = \hat{\epsilon}_{R,t|t} \frac{\sigma_M}{\hat{\sigma}_{R,t|t}} \tag{16}$$

Agents stop updating their beliefs when inflation has been running above the central bank's target enough that z_t reaches zero. When $z_t \geq 0$ and $\eta_t^z = 0$ there is nothing agents can learn about the unknown parameters, ξ_z and $\sigma_{R,t}$, because the policy rate R_t is no

[9] With only censored values, agents are maximizing the standard normal cdf, which is done with $\hat{\xi}_{z|t} \to \infty$ so $\Phi\left(\frac{\bar{\eta}_\tau^z - \hat{\xi}_{z|t} z_\tau}{\hat{\sigma}_{R|t}}\right) \to 1$

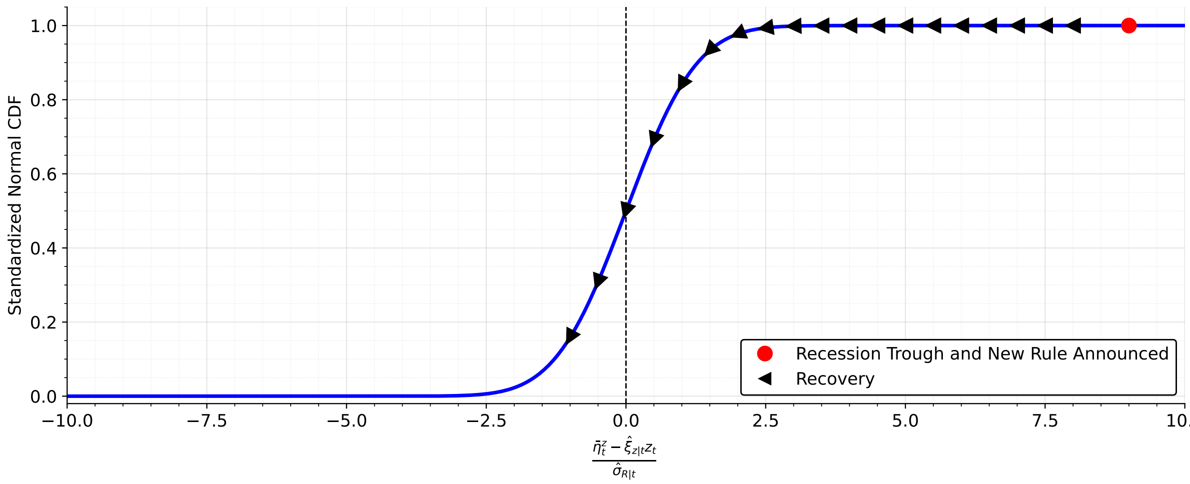

Figure 1: The likelihood function of the censored data. The likelihood (vertical axis) becomes informative as the economy improves, as indicated by movement of $\left(\bar{\eta}^z_t - \hat{\xi}_{z|t} z_t\right)/\hat{\sigma}_{R|t}$ (horizontal axis) to the left (black arrows). The starting point is the red dot.

longer dependent on those parameters. Note, too, that when $z_t > 0$ agents are able to infer $\epsilon_{M,t}$ without error.

Learning proceeds as follows. The central bank announces the new policy rule when the economy is deep in recession and $\bar{\eta}^z_1 >> 0$ and $z_1 << 0$. At this time the likelihood function is uninformative. This can be seen in Figure 1, which plots the likelihood when the data are censored, with the red dot to the right depicting the $t = 1$ starting point. The likelihood here is essentially 1 one for many values of $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R|t}$, and so maximization of the posterior distribution doesn't move the parameters from their priors. As the economy improves, $\bar{\eta}^z$ falls and z_t rises, moving $\left(\bar{\eta}^z_t - \hat{\xi}_{z|t} z_t\right)/\hat{\sigma}_{R|t}$ to left (the black arrows). The likelihood becomes informative, and maximization of the posterior will begin to change the values of $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R|t}$ from the priors.

Specifically, consider the first order conditions for the log posterior with respect to $\hat{\xi}_{z|t}$ when there are only censored data (so $R_t = 1$ for all $\tau < t$):

$$\sum_{\tau \leq t} \frac{-\phi\left(\frac{\bar{\eta}^z_\tau - \hat{\xi}_{z|t} z_\tau}{\hat{\sigma}_{R|t}}\right)}{\Phi\left(\frac{\bar{\eta}^z_\tau - \hat{\xi}_{z|t} z_\tau}{\hat{\sigma}_{R|t}}\right)} \frac{z_\tau}{2\hat{\sigma}_{R|t}} - \omega\left(\frac{\hat{\xi}_{z|t} - \xi_{z|P}}{\sigma_{R|P}}\right) = 0$$

where ω is the derivative of $\ln \Omega$. As long as $\left(\bar{\eta}^z_t - \hat{\xi}_{z|t} z_t\right)/\left(\hat{\sigma}_{R|t}\right)$ is very large, $\phi\left[(\bar{\eta}^z_t - \hat{\xi}_{z|t} z_t)/\left(\hat{\sigma}_{R|t}\right)\right]$ is near zero and nearly all of the weight in estimation is on the prior; but as $\bar{\eta}^z_t$ falls and z_t increases, so does $\phi\left[(\bar{\eta}^z_t - \hat{\xi}_{z|t} z_t)/\left(\hat{\sigma}_{R|t}\right)\right]$, the posterior begins putting meaningful weight on the likelihood, and the parameter estimates will change.

Note that the likelihood provides information even before the old rule would have dictated liftoff conditional on $\{y_t, \pi_t\}$. This is because as $\bar{\eta}_t^z$ falls but rates stay at the ELB, agents put lower weight on: 1) values of $\sigma_{R,t}$ that would admit a meaningful probability on large positive realizations of $\epsilon_{R,t}$ that would accelerate liftoff; and 2) small values of ξ_z that would imply little response to the observed z_t. More information is gleaned when liftoff would have occurred under the old rule ($\bar{\eta}_t^z$ turns negative) but rates remain at the ELB. Once liftoff occurs, non-censored observations on η_t^z enter the calculations, and the likelihood portion of the estimation is now a standard Tobit regression. Learning slows as $z_t \to 0$ and $\partial\phi\left[\left(\eta_t^z - \hat{\xi}_{z|t} z_t\right) / \left(\hat{\sigma}_{R,t|t}\right)\right] \ / \ \partial\xi_z \ \to \ 0$ for all ξ_z and stops completely when $z_t \ = \ 0$ and policy reverts back to being determined by the old rule.

4 Solving and simulating the model

We solve the model with global methods using a time iteration algorithm with linear interpolation based on Richter et al. (2014). We do so to allow the model's nonlinearities to influence economic outcomes. If we solved a linearized version of the model, the only feedback of learning about the monetary policy rule onto current behavior would occur through the rule's effect on expectations of hitting the ELB in the future.[10] But learning about the new rule also reshapes agents' views about the future volatility of inflation and output and hence the welfare risks posed by the economy returning to the ELB. These higher order effects could be important; the multiplicative interactions between variables in the model have the potential to amplify the effects of shocks relative to solutions from linear approximations, and so could provide a richer description of the effects of the ELB than linearized models. Such factors are particularly relevant because the new rule is specifically aimed at eliminating the non-linear downside risks imposed by the ELB constraint. Solving the nonlinear model in which the entire distribution of agents' beliefs influences their current behavior is necessary to capture them.

Using global methods means for each exercise we run, we need to solve the model every period for many different values of the structural parameters agents are learning, ξ_Z and σ_η, and with three stochastic shocks as state variables. The computational burden renders typical solution approaches intractable. We instead solve the model using extended decision rules following Kase et al. (2022). This method essentially expands the state vector

[10] In the standard New Keynesian analysis, higher odds that the ELB will bind and reduce consumption tomorrow induce expected utility smoothing agents to lower consumption today.

by including ξ_z and σ_η as "pseudo-state" along with the stochastic shocks and constructs decision rules for y_t and π_t that are functions of this expanded state vector. These decision rules are then used to quickly solve the model in subsequent simulations. So while the initial calculation of the decision rules takes several hours, subsequent simulations are solved in less than a second.[11] Appendix B describes this solution method in detail.

We calibrate the model by first taking the specification absent the asymmetric monetary policy piece and matching a selected set of moments from 2000:Q1 to 2019:Q4, a period in which the risk of the ELB was particularly acute. We target the standard deviations of GDP per capita growth, core PCE inflation, and the federal funds rate.[12] The central bank inflation target is set to 2.0%. In this symmetric specification average long-run inflation expectations are 1.8% due to the presence of the ELB. This deflationary bias of 20 bps is in line with Bianchi et al. (2021) and implies that the ELB binds in around 14% of periods.[13] We calibrate the asymmetric portion of the policy rule in order to eliminate the deflationary bias in the full information economy. We set the standard deviation of the monetary policy shock component associated with the new rule, σ_η, at zero and the persistence of the shortfalls, ρ_z, at 0.875.[14] With the response to shortfalls ξ_z normalized at 1, the deflationary bias is eliminated with ϕ_z equal to 0.1. The parameters are shown in Table 1:

For the priors, we picked means $\mathbb{E}_{\text{Prior}}(\xi_Z) = 0.1$ and $\mathbb{E}_{\text{Prior}}(\sigma_\eta) = 0.1\%$ that to be far away from the true values of 1 and 0. This means agents are taking the central bank's announcement with a substantial grain of salt; accordingly, they will have a good deal of learning to do to discover the new policy framework. Note, too, that $\mathbb{E}_{\text{Prior}}(\sigma_\eta) = 0.1\%$ is in terms of the quarterly gross interest factor, and so translates into about a 0.4 percentage point standard deviation in the annualized policy rate. Shocks of such size can have a meaningful impact on economic outcomes in the model when interest rates are low, as they then have a good chance of generating an episode at the ELB. The effects, however, are small and transitory when interest rates are far from the ELB given the shock is iid. We explore some alternative priors in the appendix.

To simulate the model, we first set a sequence of shocks to the demand curve, the

[11] The initial calculation repeatedly simulates the model within each period of the sample for each stochastic draw of the shocks, and so requires solving the model several thousand times.

[12] Our model predicts a standard deviation of 0.5 (vs. 0.6 in the data) for output, 0.8 (vs. 0.6 in the data) for inflation and 1.8 (vs. 1.8 in the data) for the interest rate.

[13] Note that we are using a low persistence for the shocks, which helps to match a higher frequency of periods with a binding ELB.

[14] Setting $\sigma_\eta = 0$ means the central bank faces no new implementation issues with the new rule. The choice of ρ_z means that 80 percent of the impact of $\phi_z(\pi_{t-1} - \pi)$ on z_t will be completed at 12 quarters.

Parameters	Sign	Value	Parameters	Sign	Value
		Model Parameters (Baseline)			
Discount factor	β	0.998	Production function	α	1
Relative risk aversion	σ	1	Inverse Frisch elasticity	η	1
Disutility of labor	χ	0.87	Price elasticity of demand	ϵ	7.67
Rotemberg pricing	φ	79.41	Response Inflation	κ_Π	2.5
Response Output	κ_Y	0.125	Inflation Target	$4\log(\Pi)$	2%
Persistence Pref. Shock	ρ_ζ	0.6	Std. Dev. Pref. Shock	σ_ζ	0.91%
Persistence Markup Shock	ρ_μ	0.6	Std. Dev. Markup Shock	σ_μ	5.0%
Std. Dev. MP Shock	σ_M	0.25%			
		Average Inflation Targeting			
Response to AIT	ξ_Z	1	Deviations AIT	ϕ_Z	0.1
Persistence AIT	ρ_Z	0.875	Std. Dev. of AIT Shock	σ_η	0
		Priors			
Prior Mean ξ_Z	$\mathbb{E}_{\text{Prior}}(\xi_Z)$	0.1	Prior Std. Dev. ξ_Z	$\text{Std}_{\text{Prior}}(\xi_Z)$	0.5
Prior Mean σ_η	$\mathbb{E}_{\text{Prior}}(\sigma_\eta)$	0.1%	Prior Stad. Dev. σ_η	$\text{Std}_{\text{Prior}}(\sigma_\eta)$	0.05%

Table 1: Model calibration.

markup, and the monetary policy function for the entire simulation sample period, $\{\epsilon_{d,t}, \epsilon_{\mu,t}, \epsilon_{\eta,t}, \epsilon_{M,t}\}_t^T$. The current values of $\epsilon_{d,t}$ and $\epsilon_{\mu,t}$ are observed by the agents at time t, but their future values are unknown. As noted earlier, agents can exactly recover $\epsilon_{M,t}$ if $z_t > 0$, but $\epsilon_{\eta,t}$ and $\epsilon_{M,t}$ (and so $\epsilon_{R,t}$) are unknown if $z_t < 0$. Each period we initialize $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R,t|t}$ at the values estimated in period t-1 and make a guess for $\hat{\epsilon}_{R,t}$. We then calculate outcomes based on the agents' perceived policy rule; that is, we plug these parameter estimates into equations (1), (2), and (10) to generate y_t and π_t. We then plug these values for y_t and π_t into the actual monetary policy rule (3) to obtain the actual interest rate, R_t. If $z_t < 0$, we then re-estimate $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R,t|t}$ by maximizing the posterior distribution (13) and recalculate $\{\hat{\epsilon}_{R,\tau|t}\}$ using equations (14) or (15). If the changes in $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R,t|t}$ are small enough to meet our convergence criteria, we go on to period t+1. Otherwise, we repeat the process until convergence is achieved.

This general equilibrium approach means that every period economic outcomes are jointly determined with agents' parameter estimates and the central bank's execution of the policy rule. There are no surprises for either agents or the central bank to evaluate at the end of the day.

Figure 2 shows our baseline simulation. Starting in t=1, the economy is hit with a series of negative preference shocks, $\epsilon_{d,t}$, which drive interest rates to the. In period 2, the central bank announces the new policy rule. The depth of the policy shortfall can be

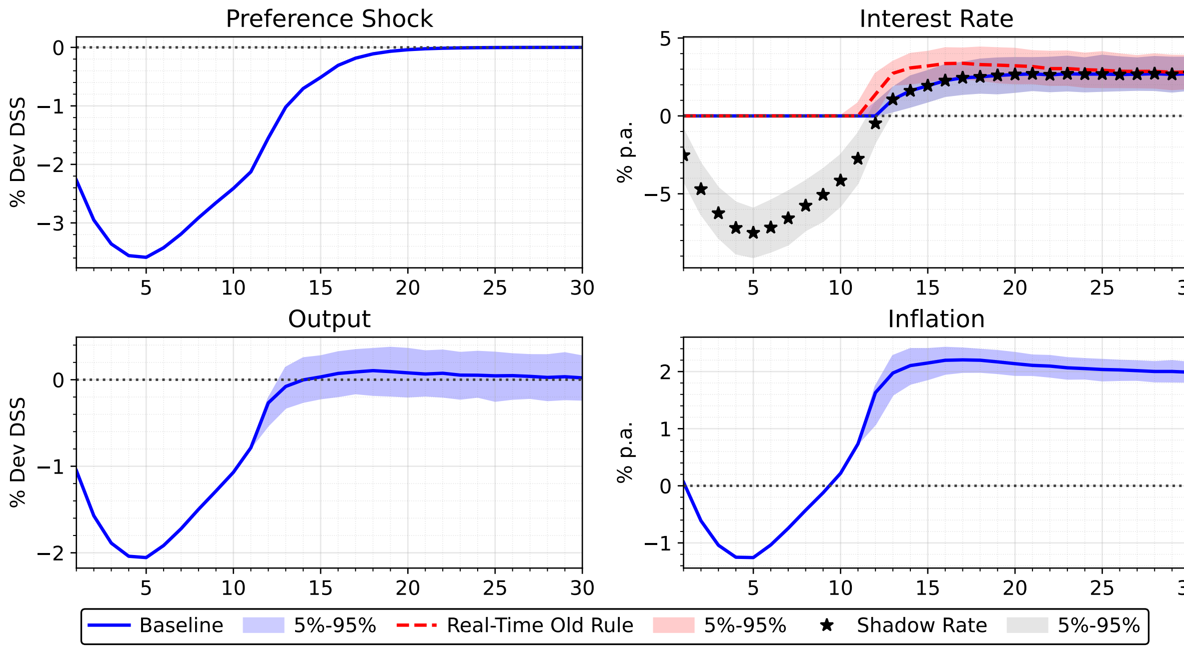

Figure 2: Baseline simulation for the new monetary policy strategy introduced during a recession (blue lines). The economy is hit by a series of negative preference shocks starting in period 1 (upper left); the new rule is introduced in period 2. Keeping the sequence of preference shocks fixed, we simulate the economy 1,000 times using alternative random draws for the sequence of monetary policy shocks. The solid lines show the mean across simulations, and the shaded areas indicate the 90 percent range of outcomes. The dynamics of interest rates (upper right) are compared with the shadow rate in the absence of the ELB (black dots) and the real-time old rule without η_t^z (black dashed line). Output and inflation are shown in the bottom panels.

seen by the negative shadow rate under the new rule as well as the declines in output and inflation well below their steady state values. For comparison, we also show the rate path under the old policy rule given the output and inflation rates generated by the learning model (dashed red line). To highlight the role of uncertainty over the policy function, we repeat the exercise for 1000 draws of $\{\epsilon_{M,t}\}_t^T$ (with the same $\epsilon_{d,t}$); the solid lines are the means from these draws and the shaded areas the 90 percent range of outcomes.

As the demand shocks dissipate, the economic situation improves. By period 12, the old monetary policy rule would dictate liftoff, with rates rising quickly to somewhat above their long-run level of 2-1/2 percent before gradually returning to steady state. In contrast, policy under the new rule delays liftoff until period 13, and rates then slowly asymptote to the steady state. The new policy also is consistent with modest overshooting of both output and inflation above their long-run targets.

Figure 3 describes the learning process. The blue solid lines show agents' parameter estimates, while the stars show the true values set by the central bank. The upper panels

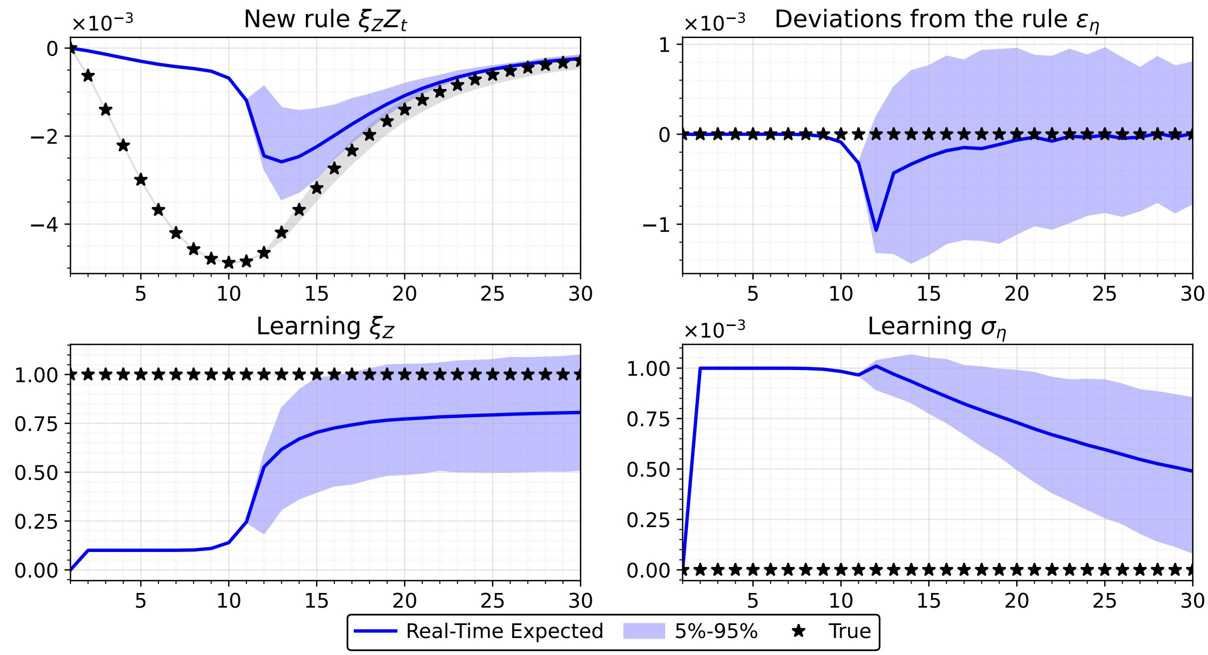

Figure 3: Learning the new rule. The blue lines show the mean of agents' real-time expectations across the simulations, the shaded areas indicate the 90 percent range of outcomes, and the black stars denote the true values. Upper panels: Components of η_t^z: structural effects of the average inflation shortfall (left) and the policy shock (right graph). The lower panels show the parameters agents estimate: intensity of central bank's response to shortfall (left) and the standard deviation of the policy shock (right).

show the components of η_t^z, while the lower panels show the parameters that enter its calculation. A small amount of learning occurs before liftoff would have taken place under the old rule as agents are able to eliminate values of ξ_z and $\sigma_{R,t}$ that would generate early liftoff under the new rule. Learning about ξ_z picks up notably in period 12 as agents observe rates remaining at the ELB instead of lifting off as would have been occurred under the old rule. Once liftoff occurs under the new rule agents can observe η_t^z, but they still do not know its individual components, $\xi_z z_t$ and $\sigma_R \epsilon_{R,t}$, and must continue to estimate them. Indeed, $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{eta|t}$ remain far from their true values for some time. Learning continues at a moderate pace until z_t gets very close to zero and the interest rate prescriptions under the new and old rule are essentially the same.

The new monetary policy rule is designed to eliminate the downward biases presented by the ELB under conventional policy rules. Figure 4 examines how well the learning model performs this job. To do so, we compare it to two alternatives. The first is the old monetary policy regime. In this alternative the central bank never announces a new monetary policy; there is no learning problem, and the path for $\{y_t, \pi_t\}$ is the rational expectations equilibrium generated by equations (1), (2), and (3) with $\xi_z = \sigma_{eta,t} = 0$.[15,16] The second alternative is the new policy rule in the full-information rational expectations environment. This is the case in which agents immediately know all the elements of the new policy rule; again there is no learning, and the path for $\{y_t, \pi_t\}$ is the rational expectations equilibrium generated by equations (1) - (3) with ξ_z and $\sigma_{R,t}$ fixed at their true values.

Early during the ELB period, the learning model and the old policy regime produce similar paths for output, but the learning model generates modestly higher inflation, reflecting the non-zero prior on ξ_z. Bigger differences emerge as agents begin to gain knowledge of ξ_z and $\sigma_{R,t}$. The differences are largest for inflation. Reflecting the powerful biases generated by the ELB, inflation under the old rule asymptotes to a level below target, while soon after liftoff inflation under the new rule is close to the full-information path.[17,18]

The lower-right panel shows the evolution of long-term inflation expectations, which we define as the model's stochastic steady state conditioned on $\{\hat{\xi}_{z|t}, \hat{\sigma}_{R,t|t}\}$.[19] Expectations under the old rule are stuck below the central bank's 2 percent target - this is the well-known downward bias to inflation under the ELB. In contrast, with the new rule, long-run inflation expectations begin to move up a bit before liftoff would have occurred under the old rule, reflecting the small degree of learning that takes place then. Expectations rise quickly after agents see liftoff is delayed, quickly moving to near the 2 percent target. Complete convergence to target is slow, however, reflecting the sluggish final stages of the

[15]The old regime differs from our earlier calculations of the shadow rate under the old monetary policy rule. Those shadow rates, which are key for the agents' inference problem under learning, were calculated by plugging the model's output for $\{y_t, \pi_t\}$ into equation (3) with $\eta_t^z = 0$; those $\{y_t, \pi_t\}$, however, depend on $\{\hat{\xi}_{z|t}, \hat{\sigma}_{R,t|t}\}$ and all the previous periods' data. In contrast, under the old policy regime, ξ_z, σ_η, $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{R,t|t}$ do not enter the solutions for $\{y_t, \pi_t, R_t\}$ at all; that is, the outcomes are what would happen if the central bank never announced a new policy rule.

[16]Note that agents cannot learn anything extra from comparing the output and inflation paths between these two scenarios. This is because the difference between the paths solely reflects $\{\hat{\xi}_{z|t}, \hat{\sigma}_{R,t|t}\}$, and agents already know those values when observing $\{y_t, \pi_t, R_t\}$ under the new regime.

[17]Recall in the New Keynesian model, events in the future have a powerful influence on outcomes today. Importantly, the output and inflation losses under future expected ELB episodes cause lower output and inflation today through the forward expectations terms in the Euler equation and Philips curve equations.

[18]These results are qualitatively similar to those in Hebden et al. (2020) using a version of FRB/US model in which agents form expectations using a VAR, but learning and the convergence of output and inflation to steady state occur much faster in our set up.

[19]Specifically, for the learning rule this inflation rate is calculated on a period-by-period basis as what the model converges to when it is run forward from $\{y_t, \pi_t, R_t\}$ conditional on $\{\hat{\xi}_{z|t}, \hat{\sigma}_{R,t|t}\}$ and with no future shocks to demand or monetary policy. For the old regime, we run only equations (1)-(3) forward with $\eta_t^z = 0$; the full information case runs these equations forward using the true values for $\xi_z = 1$ and $\sigma_\eta = 0$.

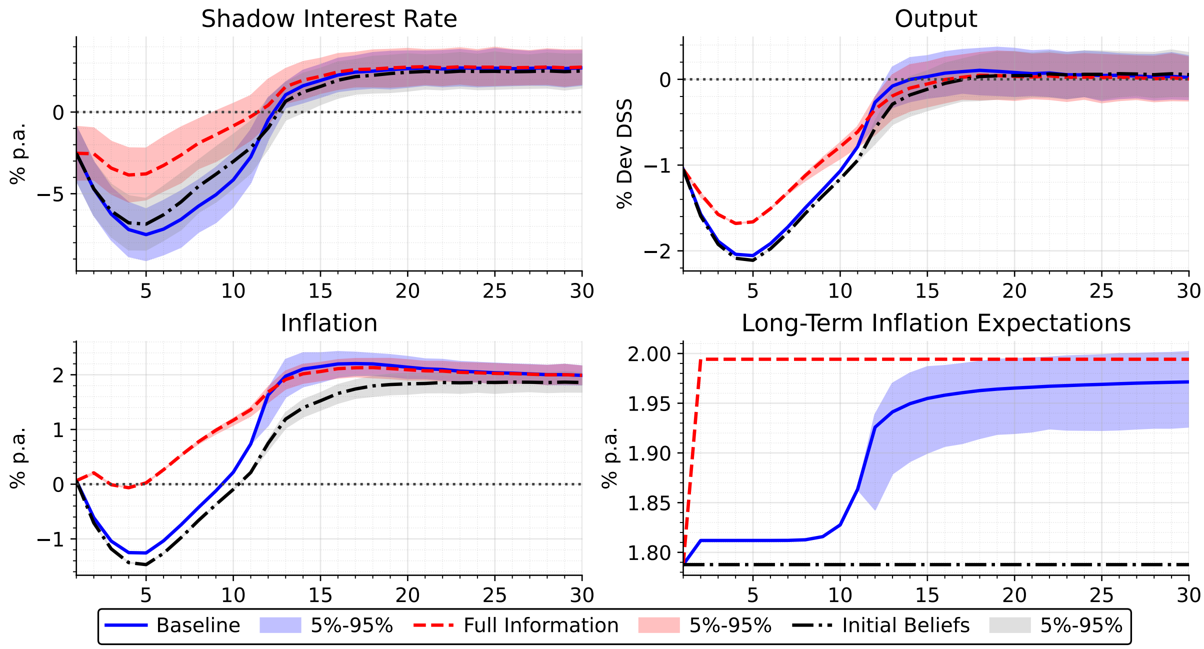

Figure 4: Comparison of outcomes under the baseline simulation of the learning model (blue) with those for the old monetary regime (black) and the new policy rule under full-information, in which agents learn the rule immediately (red). The lines show the means, and the shaded areas indicate the 90 percent range of outcomes across the 1,000 simulations, which differ only in the sequence of monetary policy shocks.

learning process. By design, the full-information model exhibits an immediate jump to target upon adoption of the new rule; furthermore, output and inflation do not fall nearly as much during the ELB episode as in the old regime or the model with learning.

5 Performance over time: Future recessions and inflation

Our analysis so far has focused on the performance of the new rule and the learning process only around the ELB-generating disturbance that motivated its adoption. How well will the rule perform and what happens to learning when the economy is hit by future shocks? To address this question, we consider two different kinds of episodes; a recession that generates another episode at the ELB and a surge in inflation well beyond the central bank's target.

Starting with the recession scenario, we assume a large negative demand shock of roughly the same order of magnitude as the one that caused the first trip to the ELB

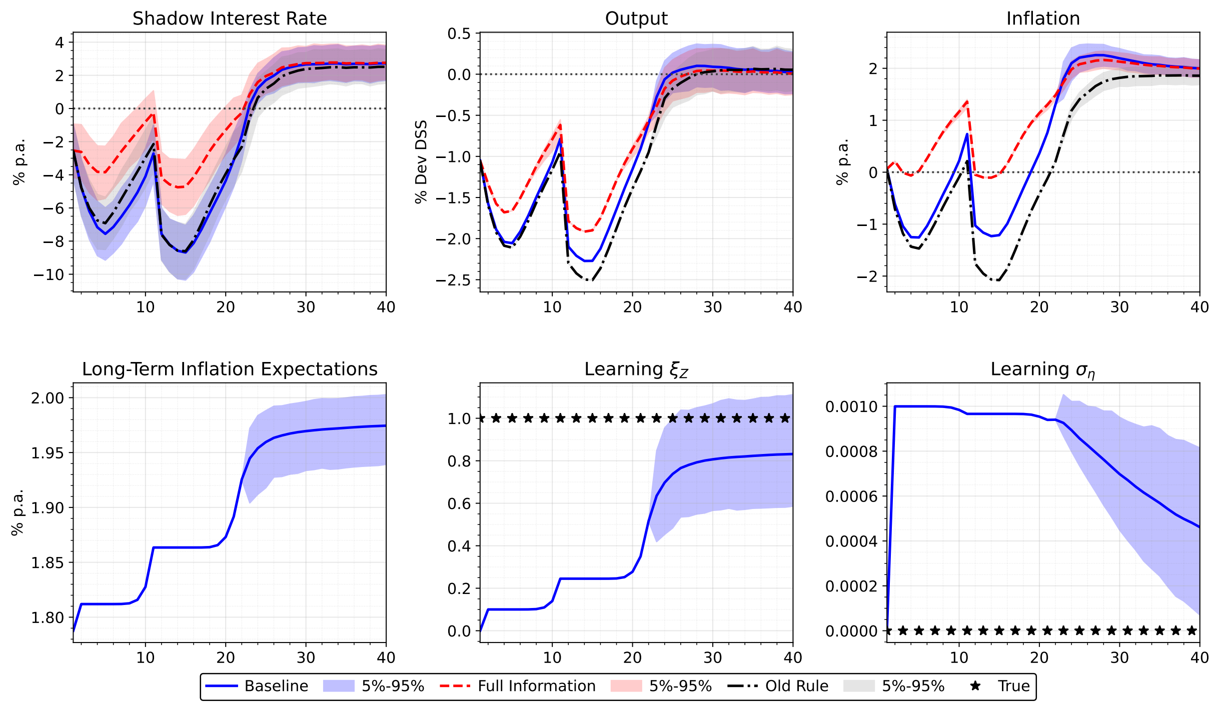

Figure 5: Second recession scenario. The economy is hit additionally by the same series of negative preference shocks starting also in period 12. Comparison of outcomes under baseline simulation (blue) with those for the old monetary regime (black) and the new policy rule under full-information (red). The lines show the means, and the shaded areas indicate the 90 percent range of outcomes across the 1,000 simulations, which differ only in the sequence of monetary policy shocks. The black stars denote the true values.

occurs in period 12, just as the central bank would otherwise be lifting off from the ELB. As seen in Figure 5, even the modest amount of learning that occurred prior to the second shock has a noticeable effect on economic outcomes. Though it still falls well short of the full information economy, the baseline model delivers a somewhat less severe recession and much smaller shortfall of inflation from target than the old policy rule. This occurs even though there has only been a modest amount of learning by the time the shock hits; $\hat{\xi}_{z|12}$ has only reached 0.25 and $\hat{\sigma}_{\eta|12}$ has barely budged from its prior.

The second recession does stall the learning process. The shadow rate falls well below the ELB when the shock hits; learning stops and does not resume when the shadow rate again rises to near the ELB. Conditions improve enough to generate liftoff in period 24. At that time learning quickly accelerates, with $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{\eta|t}$ and output, inflation, and long term inflation expectations soon closely track the full-information model.[20] In sum, the second recession delayed the learning process in the short run, but had no meaningful

[20]Note that $z_t < 0$ and learning is continuing throughout the entire period shown in Figure 5.

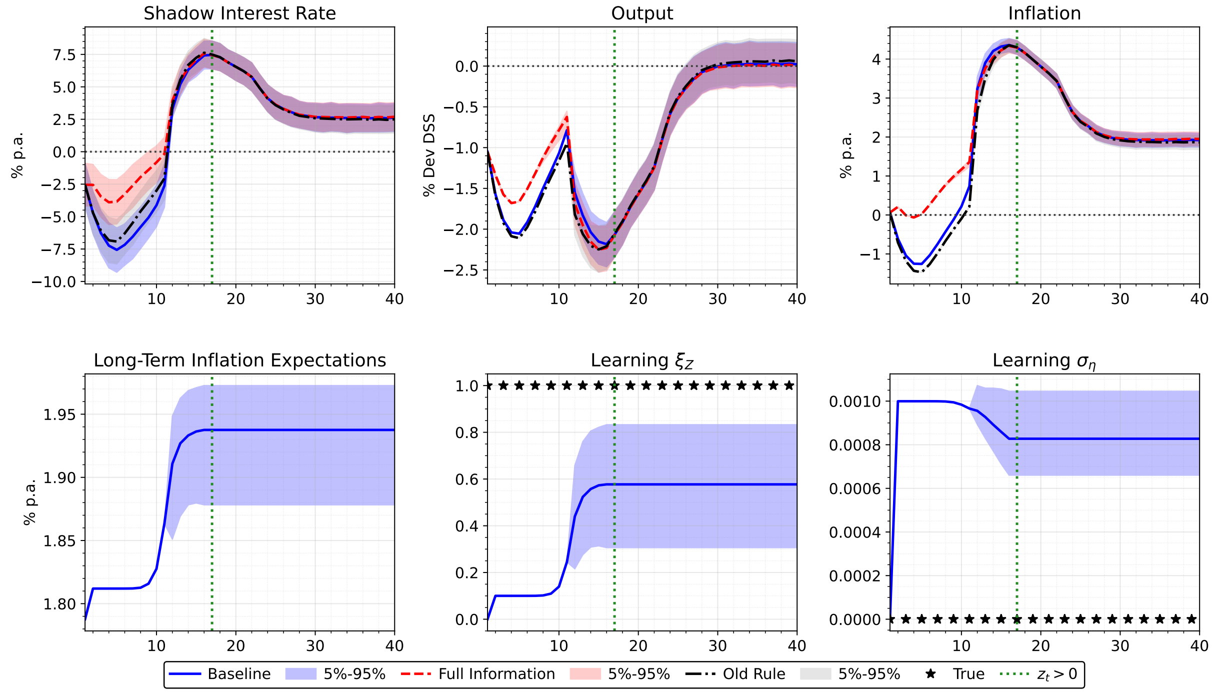

Figure 6: Inflationary scenario. The economy is hit additionally by a series of inflationary markup shocks starting also in period 12. Comparison of outcomes under baseline simulation (blue) with those for the old monetary regime (black) and the new policy rule under full-information (red). The lines show the means, and the shaded areas indicate the 90 percent range of outcomes across the 1,000 simulations, which differ only in the sequence of monetary policy shocks. The black stars denote the true values. The green dotted line indicates when z_t turns positive in the baseline scenario (for the mean).

long-run impact on learning the new rule.

We next consider an inflationary shock. We generate this with a markup shock in period 12 that causes inflation to jump 2 percentage points above target and output to fall another 1-1/2 percentage points below steady state. As seen in Figure 6, rates immediately exit the ELB under the baseline and the old policy rule. The new rule helps cushion the blow to output some, but also generates a bit higher inflation than the old rule. As the inflation shock quickly drives z_t to zero in period 17 (the green vertical line), the policy and economic outcomes are soon virtually identical between all of the rules. Importantly, the new rule has not resulted in a persistently worse outcome for inflation.[21]

The inflation shock, however, has a significant and lasting impact on learning. Because z_t reaches zero quickly, agents have only a short time span to observe differences between

[21] Long run inflation expectations are higher in the full information and baseline rules than under the old rule. However, the effect of this on economic outcomes in this scenario is dwarfed by the dynamics surrounding the markup shock.

the old and new rule. And even before z_t hits zero learning slows to a near halt. This occurs because the interest rate paths under the baseline and old rule are very close, as the impact on rates of the slightly lower output and inflation under the old rule is close to effect of η_t^z in the new rule. There is little agents can learn when the prescriptions of the rules are so close. As a result, $\hat{\eta}_{z|t}$ and long-run inflation expectations are lower and $\hat{\sigma}_{\eta|t}$ is higher than the values in Figure 4. This smaller the reaction to disinflationary shocks and the larger estimated variance of policy shocks mean more and more severe future trips to the ELB.

The inflationary scenario highlights two important features of the new rule. First, because high inflation will drive z_t to zero, there is a limit to the amount of "extra" inflation that can be generated by adoption of the new asymmetric inflation averaging rule. Second, if the inflationary shock hits before the learning process has converged the new rule will be less effective in countering future trips the ELB; ironically, this means that high inflation might be successful in achieving an average inflation target in the short run, but could also generate longer-term costs.

6 Monetary policy shocks and learning

Monetary policy shocks play a central role in the inference problem faced by the agents in our model. These shocks encompass measurement error or other random technical factors that have little economic content. However, in the real world, they also reflect purposeful deviations from simple rules with particular policy goals in mind. A notable example is when policy is tilted in one direction or the other for risk management considerations during periods of heightened uncertainty or asymmetric risks to policy goals.[22]

How might such discrete policy shocks impact our model? The previous section demonstrated how an outsized inflationary shock could quickly shut down the learning process. Recognizing this, the central bank may choose to purposefully raise rates less aggressively in order to emphasize its commitment to the new policy rule and spur the learning process. Alternatively, a large inflationary shock could raise the central bank's concern over its ability to reign in inflation expeditiously enough to keep inflation expectations from rising above target. Here it might consider a more aggressive response to the inflation shock to insure that the adoption of the new rule is not misperceived as a lack of commitment

[22]Evans et al. (2015) cite historical narrative and find some statistical support that the Federal Reserve has deviated from its baseline policy due to uncertainty or asymmetric risks.

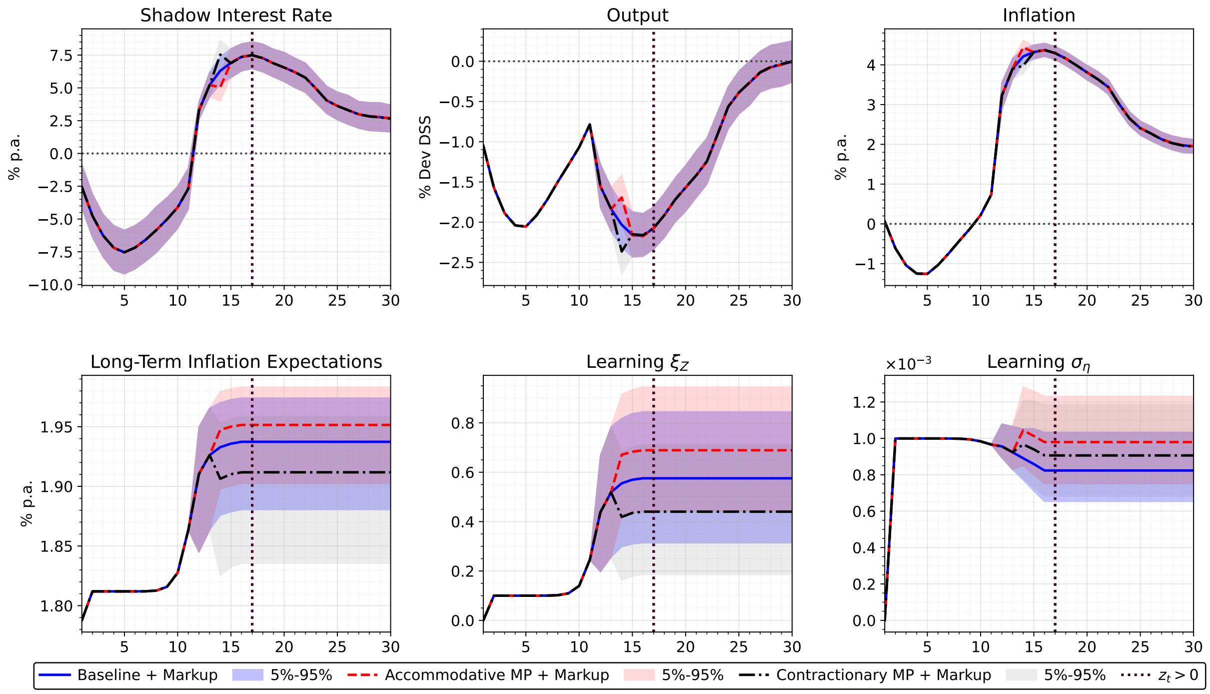

Figure 7: Monetary policy shocks and learning. Comparison of outcomes under baseline simulation (blue) with monetary policy shocks that result in a less aggressive (red) or more aggressive (black) response to inflation. The lines show the means, and the shaded areas indicate the 90 percent range of outcomes across the 1,000 simulations, which differ only in the sequence of monetary policy shocks. The dotted line indicates when z_t turns positive in all scenarios (for the mean).

to both sides of the inflation target.[23]

In Figure 7 we start with the inflation surge from the previous section, when a large markup shock in period 12 soon shuts down the learning process. We first compare this simulation to one in which the central bank purposefully raises rates less aggressively. We do so by introducing a negative $\epsilon_{M,t}$ monetary shock in period 13, so after there has been some time to observe the impact of the markup shock but while z_t is still negative and the new rule is operative. As seen in the red dashed lines, with this shock inflation and output rise faster than in the baseline, but once the monetary policy shock returns to zero the rate paths and economic outcomes are qualitatively the same. However, agents now estimate larger values for $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{\eta|t}$. This occurs because when agents see rates rise less than they expected, they surmise they have witnessed some combination of a reaction to $z_t < 0$ that is bigger than previously estimated or a monetary policy shock from the

[23]If agents believe the central bank might inject discretionary shocks aimed at influencing the learning process, then the equilibrium in our model could be different. Expanding the model to consider these interactions is beyond the scope of this paper.

tail of the estimated distribution. They revise their posteriors according to the inference problem described in Section 3, which increases both $\hat{\xi}_{z|t}$ and $\hat{\sigma}_{\eta|t}$. So this discretionary action will result in offsetting more of the disinflationary bias, but comes at the cost of temporarily higher inflation and increased volatility relative to the baseline.[24]

We next consider the central bank making a more aggressive response to the inflation shock by imposing a positive shock $\epsilon_{M,t}$ shock to the policy function that boosts the policy rate in period 13. As seen in the dashed black lines, this policy shock reduces both output and inflation relative to the baseline and the looser policy alternative. The effects on learning are large, as $\hat{\xi}_{z|t}$ falls back toward the prior and $\hat{\sigma}_{\eta|t}$ increases almost as much as in the looser policy case. Accordingly, this policy intervention may have helped keep current inflation down, but it comes with two costs – a reduced ability of the new rule to counter the economy's disinflationary bias and an increase in policy uncertainty.

These exercises highlight the complex interactions between learning, economic outcomes, and discretionary actions by the central bank. The lesson here is that monetary policy shocks not only have a short-lived impact on output and inflation, but also a longer-lasting effect on learning and hence the future performance of a new the monetary policy rule.

7 The new strategy and the 2021-22 inflation shock

In this section we consider a real world example of learning a new monetary policy rule in the presence of large shocks to the economy. Not long after the Federal Reserve's new framework was introduced in August 2020, inflation in the US rose sharply. The four-quarter change in the core PCE price index increased from 1 percent in 2020:Q2 to 3.4 percent in 2021:Q2 and 5.2 percent in 2022:Q1.[25] But with forecasts that the inflation surge would be transitory and continued concerns about the health of the labor market, the FOMC held the federal funds rate at the ELB until March 2022. The Eurozone experienced a similar similar experience around the time of its adoption of a new monetary policy strategy.

[24] One could also construct examples where even stronger negative deviations from the rule drive $\hat{\xi}_{z|t}$ above one, imparting a long-run inflationary bias that the central bank would need to counter. However, this risk is limited; while the less aggressive response boosts $\hat{\xi}_{z|t}$, it also results in higher inflation and thus pushes z_t to zero faster and shortens the learning period. These trade-offs are not large in the scenarios we consider; still, conceptually the central bank would need to take account if considering such an intervention.

[25] These inflation rates are those published at the time.

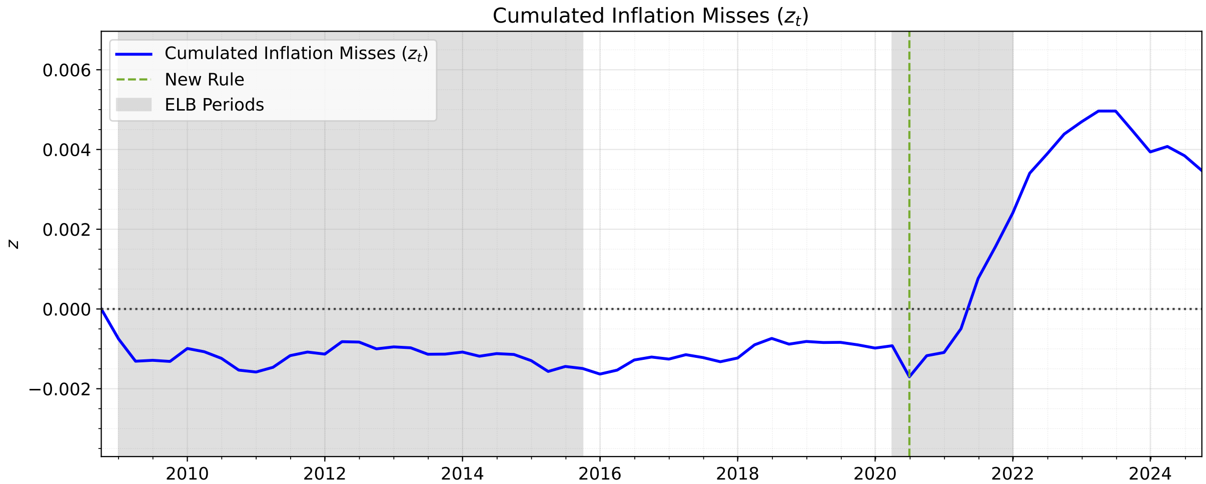

Figure 8: Cumulated inflation misses over time. The blue line shows the evolution of z_t, the accumulated discounted past inflation misses from equation (5), starting from when interest rates in the US first went to the ELB at the end of 2008. The vertical green line marks the announcement of the new rule in 2020:Q3, and the gray shaded areas indicate ELB periods.

We look at this US episode through the lens of our model to see how such an inflationary shock and monetary policy path would influence private sector agents' learning of a new policy rule and to consider how the change in the rule may have influenced output and inflation. To be clear, we are not suggesting that our highly stylized model accurately captures the economic circumstances of the time or the decisions made by the Federal Reserve. Instead, we see it as an interesting calibration exercise for the kinds of shocks that might pose substantial challenges to the learning process and in turn how these might feed back on economic performance.

We start by making our model match the US data. We search for a combination of preference, markup, and monetary policy shocks that explain the observed data on real per capita GDP growth, core PCE inflation, and the federal funds rate from 2020:Q2 to 2023:Q3.[26] Appendix C provides the details how we do so. Given we have 3 shocks, we are able to closely match the data (see Figure 9).[27]

[26] There is no consensus on the source(s) of the post Covid inflation. Bernanke and Blanchard (2023) estimate a large role for shocks to commodity prices and major sectoral shifts in demand running into sector-level capacity constraints; this latter effect is also the focus of theoretical work by Guerrieri et al. (2021). Giannone and Primiceri (2024) estimate a predominant role for unexpectedly strong demand, while Ball et al. (2022) argue nonlinearities in the Phillips curve and the pass through of food and energy price shocks played an important role. With regard to fiscal sources, theoretical work by Bassetto and Miller (2025) show how a discrete public reassessment of fiscal sustainably in light of unfavorable fiscal news could result in a sudden jump in inflation and Bianchi et al. (2023) show how unfunded government debt can help fight downturns but occasionally fuel a persistent rise in inflation.

[27] Our model clearly misses many important features needed to explain the Covid recession and recovery.

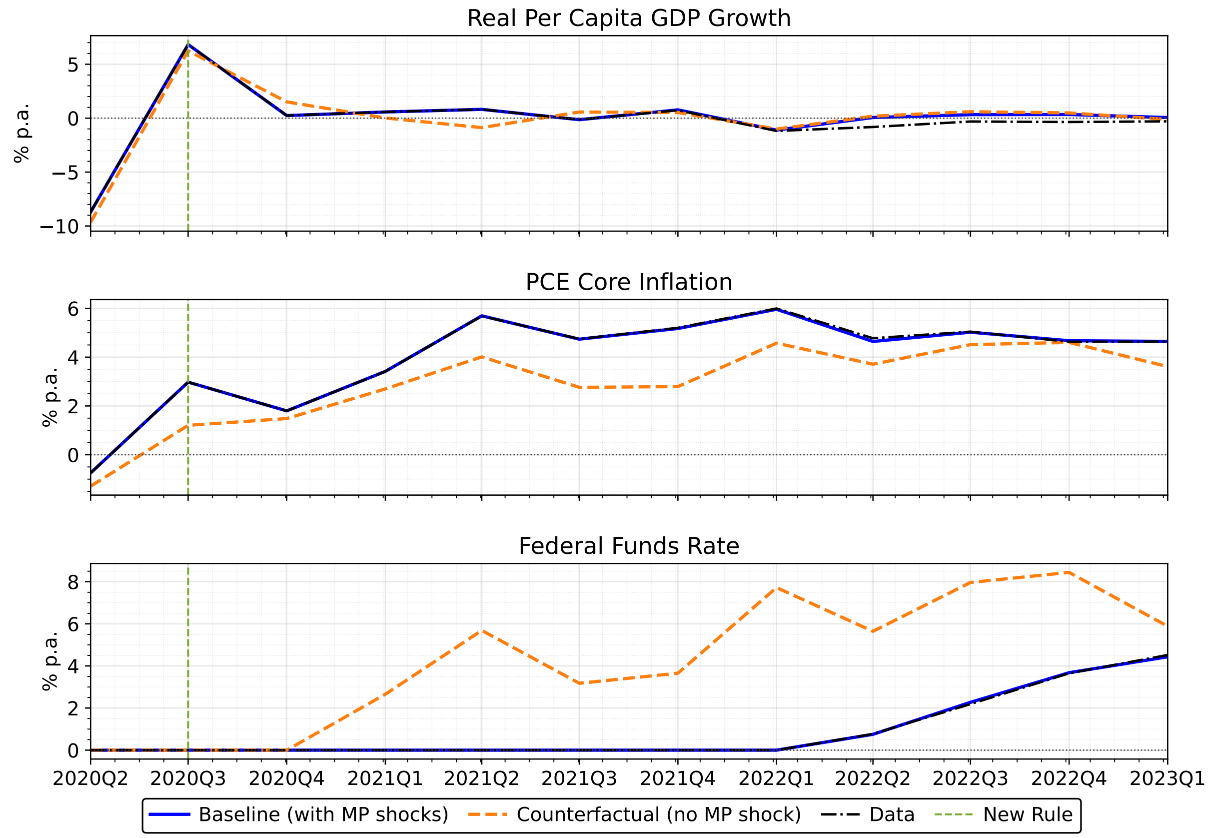

Figure 9: Fit of the model to the data and the role of monetary policy shocks. The blue lines show real GDP growth, core inflation, and the federal funds rate from the baseline fit of the model to the data. The orange lines show the counterfactual path when monetary policy shocks are set to zero, holding the other shocks fixed. The vertical green line marks the announcement of the new rule in 2020:Q3.

Figure 8 shows the evolution of z_t, the accumulated discounted past inflation misses from equation (5), starting from when interest rates in the US first went to the ELB at the end of 2008. The vertical green line in 2020:Q3 is when the new monetary policy rule is announced and agents begin to learn it; z_t is negative then and hence the new rule is operative when introduced. However, it remains so only until 2021:Q2, when the surge in inflation pushes $z_t > 0$. Thus, agents only have 4 quarters in which to learn the new rule. Further complicating the learning problem, as seen by the gray shaded bars, the federal funds rate was held at the ELB for about a year after the inflation shortfall was eliminated. The model fits this feature of the data by positing large negative monetary

Furthermore, our policy rule has no inertia in it, and so it would have difficulty fitting actual policy, even in normal times. Consequently, the shocks we needed to use to match the data were quite large, sometimes exceeding the boundaries of the grids we used when constructing our solution algorithm. We also tightened the prior on σ_η for this exercise.

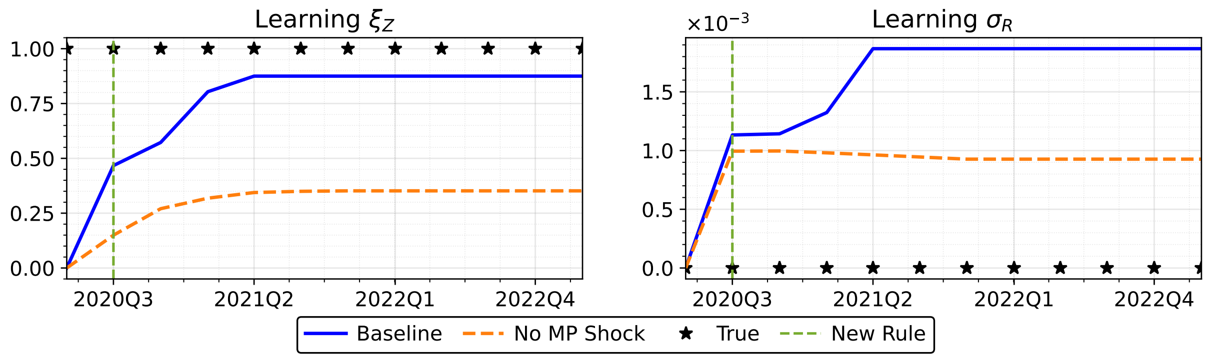

Figure 10: Learning under the fitted shock paths and the role of monetary policy shocks. Comparison of learning outcomes under baseline fitting (blue) and the counterfactual with the absence of monetary policy shocks (orange). The black stars denote the true values. The vertical green line marks the announcement of the new rule in 2020:Q3.

policy shocks that prevent rates from rising. Such shocks also are needed to fit the slope of the funds rate path after liftoff.

The blue lines in Figure 9 show output, inflation, and the federal funds rate from our fit of the baseline model to the data. Since the monetary policy shocks were so important in the fitting exercise, we also show in the orange lines a simulation of our model in which these shocks are set to zero but with the same demand and markup shocks as in the baseline. As can be seen, the policy shocks have a small effect on GDP, but produce a much more accommodative path for policy and boost inflation by a noticeable amount.

The results for the learning process under the two monetary policy shock paths are shown in Figure 10. In the baseline agents boost $\hat{\xi}_{z|t}$ close to its true value of 1 within the short learning window, a substantially faster pace of learning that in our earlier stylized examples. This is accompanied, however, by $\hat{\sigma}_{\eta|t}$ moving far away from its true value. In contrast, in the absence of monetary policy shocks, agents learn very little about the new rule: $\hat{\xi}_{z|t}$ only gets to about 0.35 and $\hat{\sigma}_{\eta|t}$ barely moves from it's prior.

The rationale for these learning results in baseline scenario is qualitatively the same as that for the negative monetary policy shock in the previous section, only the magnitudes are much bigger as the shocks are larger and they are sustained over several periods. In contrast, if there were no monetary policy shocks, policy would soon liftoff from the ELB and then quickly become indistinguishable from the old rule. Hence not much learning takes place in that scenario.

The jump in $\hat{\xi}_{z|t}$ in the baseline will improve prospects for the economy in a subsequent ELB episode. As seen in Figure 11, in terms of long-run inflation expectations, the terminal

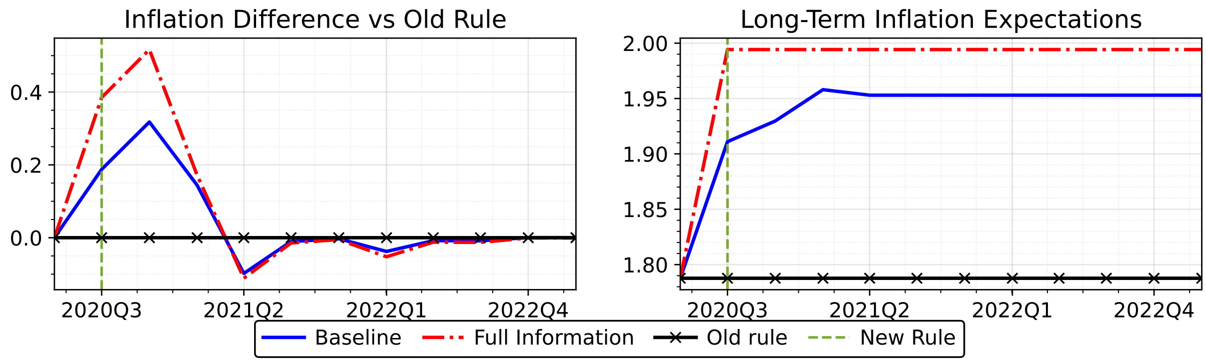

Figure 11: Impact of the rule on inflation and long-term inflation expectations. Comparison of outcomes under the new policy rule with learning (blue), the new policy rule with full information (red), and the old monetary regime (black). The left panel shows the difference in inflation relative to the old monetary regime. The right panel shows long-term inflation expectation fir the different cases. The vertical green line marks the announcement of the new rule in 2020:Q3.

value of $\hat{\xi}_{z|t}$ is large enough to essentially eliminate the disinflationary bias associated with the ELB. But this achievement only occurred because of the large monetary policy shocks, and as noted above, these shocks are not a free lunch.[28]

The new policy rule improves performance of the economy by producing higher inflation than the old rule around ELB episodes. This raises the question of how much of the post-Covid inflation surge does our exercise attribute to the new rule. This is shown in Figure 11. The blue line plots the difference between inflation in our baseline and inflation if the central bank had followed the old policy rule, and the red line plots the difference in inflation from the old rule if the new rule was known with full information. These are subject to the same monetary policy shocks used to construct the baseline. Here we see that little of the run up in inflation is due to the new rule per se; at most it's 1/2 percentage point under full information a quarter after introduction, and the difference is negligible after 2021:Q1 when z_t turns positive and the new and old rules look the same. This is evidence of the natural limit on inflationary presented by the asymmetric inflation averaging term. The larger differences in inflation between the baseline and the no-monetary-shock case seen in Figure 9 are attributable to the effects of the monetary policy shocks, not the effects of operating under the new rule.

[28]This exercise likely underestimates the degree of policy uncertainty induced by the monetary policy shocks. Recall we assume agents know σ_M and only estimate σ_η. Once z_t is zero, only $\epsilon_{M,t}$ is operative, so agents would not have any reason to update $\hat{\sigma}_{\eta|t}$. In contrast, if agents did not know σ_M and were estimating it, rates being kept at the ELB beyond 2021:Q2 would be interpreted as unusually large $\epsilon_{M,t}$ shocks, increasing $\hat{\sigma}_{M|t}$.

8 The costs of incomplete learning

This section provides some illustrative calculations of costs associated with incomplete learning of the new rule and trade-offs that might arise from the central bank taking discretionary actions that influence both $\hat{\xi}_{z,t}$ and $\epsilon_{R,t}$. One way to summarize the relative performance of the different policy functions is to consider a loss function that adds up squared deviations of output and inflation from target. Such a function or close variants of it are often used to evaluate the relative efficacy of different monetary policy rules, see, for example, Reifschneider and Willams (2000) and Bernanke et al. (2019).[29]

Specifically, we consider the stochastic loss, SL. This is a forward looking calculation that takes into account potential future shocks. Starting at steady state values for output and inflation, we simulate the model T periods ahead conditioned on fixed values for $\hat{\xi}_z$ and $\hat{\sigma}_\eta$ and N future draws of $\{\epsilon_{d,t}, \epsilon_{R,t}\}_1^T$, and take the average discounted sum of future squared deviations of output and inflation from target:

$$SL \;=\; \frac{1}{N}\sum_{n=1}^{N}\sum_{\tau=0}^{T}\beta^{-\tau}\;\left[(\pi_{t+\tau,n}-\pi)^2 \;+\; (y_{t+\tau,n}\;-y)^2\right]$$

where $\pi_{t+\tau,n}$ and $y_{t+\tau,n}$ are the values $\pi_{t+\tau}$ and $y_{t+\tau}$ from draw n. Accordingly, the stochastic loss tells us about the average performance the policy regime can be expected to deliver given the entire distribution of future shocks to the economy conditioned on agents' perception of the new rule's parameters.[30]

We start with the economy at steady state, and then calculate SL for a range of values for $\{\hat{\xi}_z, \hat{\sigma}_R\}$. Table 2 shows the results of these experiments. In each entry, the central bank follows the true rule with $\{\xi_\eta = 1, \sigma_\eta = 0\}$ while agents believe and act as if ξ_z and σ_η correspond to the entries in the table and do no learning. The left-hand panel of the table shows the loss relative to the full information case in which the new rule is known, the $\{\hat{\xi}_z = 1, \hat{\sigma}_\eta = 0.0\}$ case in the lower-left corner. As most of the losses are experienced when the economy hits the ELB, the right-hand panel shows the frequency of hitting the ELB in each of the cases.

[29] Many models that calculate optimal monetary policy do so by minimizing such loss functions; see Woodford (2003), Svensson (2010), and the Federal Reserve Board's optimal control Tealbook analysis.

[30] Note that SL does not include a term in the squared change in the policy rate, which enters some commonly used loss functions to capture financial stability concerns associated with large movements in interest rates as well as the empirical fact that central banks usually move policy rates in a smooth manner. Such a loss function is consistent with a monetary policy rule that includes inertia in the policy rate, which is not a feature of the rule in our model.

Policy shock $\hat{\sigma}_\eta$		Stochastic loss 0.00	0.20	0.40	0.60	1.00	ELB probability 0.00	0.20	0.40	0.60	1.00
AIT rule $\hat{\xi}_z$	0.00	1.19	1.20	1.22	1.26	1.38	17.75	18.37	19.09	20.59	24.31
	0.25	1.09	1.10	1.11	1.13	1.17	13.69	14.01	14.40	15.27	17.23
	0.50	1.04	1.05	1.05	1.07	1.10	11.37	11.61	11.91	12.55	13.95
	0.75	1.02	1.02	1.03	1.04	1.06	9.81	9.99	10.21	10.73	11.88
	1.00	1.00	1.00	1.01	1.01	1.03	8.69	8.83	9.00	9.41	10.32

Table 2: Policy shocks, stochastic loss, and ELB probability. $\hat{\sigma}_\eta$ is measured in percentage points of the net interest rate. Stochastic loss is reported relative to $\{\xi_z = 1, \sigma_R = 0\}$. The ELB probability is in %.

The upper left hand corner with $\{\hat{\xi}_z = 0, \hat{\sigma}_\eta = 0\}$ corresponds to agents believing the central bank has made no change to its policy rule; the associated losses are 19 percent larger than if the rule had been enacted with full credibility.[31] As was suggested by our analysis in Section 5, just a small degree of learning can result in much improved macro outcomes, with SL for $\{\hat{\xi}_z = 0.25, \hat{\sigma}_\eta = 0\}$ closing about half of the gap to the full information case and the differences being negligible by $\hat{\xi}_\eta = 0.75$. Correspondingly, the probability of hitting the ELB is shrunk in half if agents learn the new rule with certainty. Still, at 8.7 percent, the odds of hitting the ELB are not trivial even under the new rule.

Increased uncertainty over the reaction function is costly, but not substantially so for modest changes. At the prior for $\hat{\sigma}_\eta$, 40 basis points, holding the incorrect view of variation in the policy rate with $\hat{\xi}_z = 0$ increases the normalized loss by only 3 basis points, with the loss trivial for higher values of $\hat{\xi}_z$. Recall that in most of the scenarios we ran, $\hat{\sigma}_\eta$ at most exceeded the prior by a minute amount. The exception was our exercise matching the 2021-22 inflation surge, when $\hat{\sigma}_\eta$ rose to about 0.8 percent. This would result in more meaningful losses for low values of $\hat{\xi}_z$, when agents do not believe the central bank had made a meaningful change in its policy rule.

Taken as a whole, these results suggest that in our model inducing a modest degree of uncertainty into agents' perception of the interest rate rule may be a cost worth taking if it improves learning ξ_z. This may reflect the well known fact that policy rules using common parameterizations such as in our old rule will not replicate optimal policy that minimizes a loss function such as SL. Rules with stronger reactions to inflation and output gaps

[31] Since the central bank follows the new rule in all of the entries on the table, the difference between the upper left and lower left entries does not give the value of the new rule known with certainty relative to the old rule. That calculation would indicate a 17 percentage point improvement in the new rule over the old one with the probability of hitting the ELB falling from 15 to 8.7 percent

generally do a better job at approaching optimal policy, and these effects are first order. Hence a stronger reaction to inflation (larger values of $\hat{\xi}_z$) can be expected to result in lower losses even if they are accompanied by modest changes in the perceived volatility of policy.

We don't want to push this result to far. Alternative parameterizations would influence the trade-offs. Greater risk aversion would boost precautionary saving motives, making increased uncertainty over the policy function more costly. Greater inertia, for example through interest rates smoothing in the policy function or backward-looking price setting behavior, would prolong the effects of policy shocks. Either of these could increase the relative costs of missing estimates of σ_η.

9 Conclusions

Under traditional monetary policy strategies, the limits on reducing policy rates presented by the proximity of the ELB impart a downward bias to inflation and inflation expectations relative to the central bank's target. This bias impinges on the central bank's ability to stabilize output and inflation. Economic theory has provided a number of alternative policy strategies to combat these limitations. In this paper we consider a form of asymmetric average inflation averaging, which adds a term in the weighted average of past shortfalls of inflation from target to the central bank's old policy function. This policy successfully offsets the disinflationary bias.

For this or any other alternative strategy to be effective, agents need to understand a regime change has taken place and adjust their behavior accordingly. This happens automatically in an abstract economy in which agents have complete information about the parameters of the new policy and the central bank has full credibility. However, in the real world, agents need to be convinced that the new policy process is indeed in place and then must learn its parameters. We show how this occurs through observation of the central bank's policy settings, with the period around the time that liftoff from the ELB would have occurred under the old policy regime being particularly crucial for learning process. The interactions between policy decisions, economic outcomes, and agents' learning process are important determinants of the performance of the new policy rule.

Shocks inevitably will hit the economy during this time, and they can present important impediments to the learning process. Recessionary shocks may delay learning by inducing a second trip to the ELB; however, learning will resume and potentially be more powerful upon the second ELB exit. In contrast, inflationary shocks that cause the inflation averaging goal to be achieved before agents fully absorb the new policy rule can derail learning entirely. But as long as credibility is maintained, the costs of any additional above-target inflation generated by the new rule are limited as past average shortfalls are driven to zero and policy reverts back to the same inflation-fighting prescriptions the central bank had always followed.

The central bank could consider deviating from its new rule at times in order to deliver better near term economic outcomes or to influence the learning process and better establish the new policy. But doing so comes with costs, including inducing additional uncertainty about future monetary policy actions. Such trade offs are evident when looking through the lens of our model at the surge in inflation following introduction of a new policy rule in the US.

References

Adam, K., Billi, R.M., 2007. Discretionary monetary policy and the zero lower bound on nominal interest rates. Journal of Monetary Economics 54, 728 – 752.

Adam, K., Woodford, M., 2012. Robustly optimal monetary policy in a microfounded new keynesian model. Journal of Monetary Economics 59, 468–487. Carnegie-NYU-Rochester Conference Series on Public Policy - Robust Macroeconomic Policy at Carnegie Mellon University on November 11-12, 2011.

Amano, R., Gnocchi, S., Leduc, S., Wagner, J., 2020. Average Is Good Enough: Average-Inflation Targeting and the ELB. Technical Report. Federal Reserve Bank of San Francisco Working Paper 2020-21.

Andolfatto, D., Gomme, P., 2003. Monetary policy regimes and beliefs. International Economic Review 44, 1–30.

Arias, J., Bodenstein, M., Chung, H., Drautzburg, T., Raffo, A., 2020. Alternative Strategies: How Do They Work? How Might They Help? Technical Report. Federal Reserve Board Finance and Economics Discussion Series Working Paper 2020-068.

Ball, L., Leighand, D., Mishra, P., 2022. Understanding us inflation during the coivd 19 era. Brookings Papers on Economic Activity 2022, 1–54.

Barnett, A., Ellison, M., 2013. Learning by disinflating. Journal of Money, Credit and Banking 45, 731–746.

Barro, R.J., Bianchi, F., 2026. Fiscal influences on inflation in oecd countries, 2020â2023. The Economic Journal 136, 626–654.

Bassetto, M., Miller, D.S., 2025. A monetary-fiscal theory of sudden inflations. The Quarterly Journal of Economics 140, 1959–2000.

Bernanke, B., 2017. Monetary Policy in a New Era. Technical Report. Conference on Rethinking Macroeconomic Policy, Peterson Institute for International Economics, October 12, 2017.

Bernanke, B., Blanchard, O., 2023. What Caused the US Pandemic-era Inflation? Technical Report. NBER Workikng Paper 31417.

Bernanke, B.S., Kiley, M.T., Roberts, J.M., 2019. Monetary policy strategies for a low-rate environment. AEA Papers and Proceedings 109, 421–26.

Bianchi, F., Faccini, R., Melosi, L., 2023. A fiscal theory of persistent inflation. Quarterly Journal of Economics 138, 2127–2179.

Bianchi, F., Melosi, L., 2016. Modeling the evolution of expectations and uncertainty in general equilibrium. International Economic Review 57, 717–756.

Bianchi, F., Melosi, L., 2018. Constrained discretion and central bank transparency. Review of Economics and Statistics 100, 187–202.

Bianchi, F., Melosi, L., Rottner, M., 2021. Hitting the elusive inflation target. Journal of Monetary Economics 124, 107–122.

Bodenstein, M., Hebden, J., Winkler, F., 2022. Learning and misperception of makeup strategies. Journal of Economic Dynamics and Control 139, 104417.

Budianto, F., Nakata, T., Schmidt, S., 2023. Average inflation targeting and the interest rate lower bound. European Economic Review 152, 104384.

Caballero, R.J., Simsek, A., 2023. A note on temporary supply shocks with aggregate demand inertia. American Economic Review: Insights 5, 241–258.

Darracq Pariès, M., Kok, C., Rottner, M., 2023. Reversal interest rate and macroprudential policy. European Economic Review 159, 104572.

Engen, E.M., Laubach, T., Reifschneider, D.L., 2015. The Macroeconomic Effects of the Federal Reserve's Unconventional Monetary Policies. Finance and Economics Discussion Series 2015-5. Board of Governors of the Federal Reserve System (U.S.).

Erceg, C.J., Levin, A.T., 2003. Imperfect credibility and inflation persistence. Journal of Monetary Economics 50, 915–944.

Evans, C., Fisher, J., Gourio, F., Krane, S., 2015. Risk management for monetary policy near the zero lower bound. Brookings Papers on Economic Activity 46, 141–219.

Evans, C.L., 2012. Monetary policy in a low-inflation environment: Developing a state-contingent price-level target. Journal of Money, Credit and Banking 44, 147–155.

Giannone, D., Primiceri, G., 2024. The Drivers of Post-Pandemic Inflation. Technical Report. ECB Central Banking Forum.

Guerrieri, V., Lorenzoni, G., Straub, L., Werning, I., 2021. Monetary Policy in Times of Structural Reallocation. Technical Report. Federal Reserve Bank of Kansas City Jackson Hole Syposium.

Gust, C., Herbst, E., López-Salido, D., 2018. Forward guidance with bayesian learning and estimation. Finance and Economics Discussion Series 2018-072, Board of Governors of the Federal Reserve.

Hebden, J., Herbst, E., Tang, J., Topa, G., Winkler, F., 2020. How Robust are Makeup

Strategies to Key Assumptions? Technical Report. Federal Reserve Board Finance and Economics Discussion Series Working Paper 2020-069.

Hebden, J., López-Salido, D., 2018. From Taylor's Rule to Bernanke's Temporary Price Level Targeting. Technical Report. Finance and Economics Discussion Series 2018-051, Board of Governors of the Federal Reserve System.

Hills, T.S., Nakata, T., Schmidt, S., 2019. Effective lower bound risk. European Economic Review 120, 103321.

Honkapohja, S., McClung, N., 2023. On robustness of average inflation targeting. Technical Report.

Kase, H., Melosi, L., Rottner, M., 2022. Estimating nonlinear heterogeneous agents models with neural networks. Centre for Economic Policy Research.

Mele, A., Molnar, K., Santoro, S., 2020. On the perils of stabilizing prices when agents are learning. Journal of Monetary Economics 115, 339–353. URL: https://www.sciencedirect.com/science/article/pii/S0304393219301382, doi:https://doi.org/10.1016/j.jmoneco.2019.08.006.

Mertens, T.M., Williams, J.C., 2019. Tying down the anchor: Monetary policy rules and the lower bound on interest rates. Federal Reserve Bank of San Francisco, Working Paper 2019-14, 2019.

Nakov, A., 2008. Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate. International Journal of Central Banking 4, 73–127.

Reifschneider, D., Willams, J.C., 2000. Three Lessons for Monetary Policy in a Low-Inflation Era. Journal of Money, Credit and Banking 32, 936–966.

Richter, A.W., Throckmorton, N.A., Walker, T.B., 2014. Accuracy, speed and robustness of policy function iteration. Computational Economics 44, 445–476.

Schorfheide, F., 2005. Learning and monetary policy shifts. Review of Economic Dynamics 8, 392–419.

Svensson, L.E., 2010. Inflation targeting, in: Handbook of Monetary Economics. Elsevier. volume 3, pp. 1237–1302.

Woodford, M., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press.

A The role of prior information

The results of our learning model are sensitive to the prior distribution of the unknown parameters. As noted earlier, our baseline case uses prior means that are far from the actual values for ξ_z and σ_R. Figure 12 shows some alternatives for the prior mean and standard deviation of ξ_z, $\mathbb{E}_{\text{Prior}}(\xi_Z)$ and $\text{Std}_{\text{Prior}}(\sigma_R)$, respectively. The simulations follow the analysis in Section 4, and so do not include subsequent recessionary or inflationary shocks.

As would be expected, increasing the prior mean by itself results in a higher shadow rate and faster learning of ξ_z and σ_η – and hence higher long term inflation expectations – than in the baseline. This can be seen by comparing the blue baseline scenario with the red dashed line case, which has the same $\text{Std}_{\text{Prior}}(\sigma_R)$ but larger $\mathbb{E}_{\text{Prior}}(\xi_Z)$. The differences are less pronounced between the black and orange lines, which have the same pair of $\mathbb{E}_{\text{Prior}}(\xi_Z)$ as the blue and red lines, but a larger common $\text{Std}_{\text{Prior}}(\sigma_R)$.

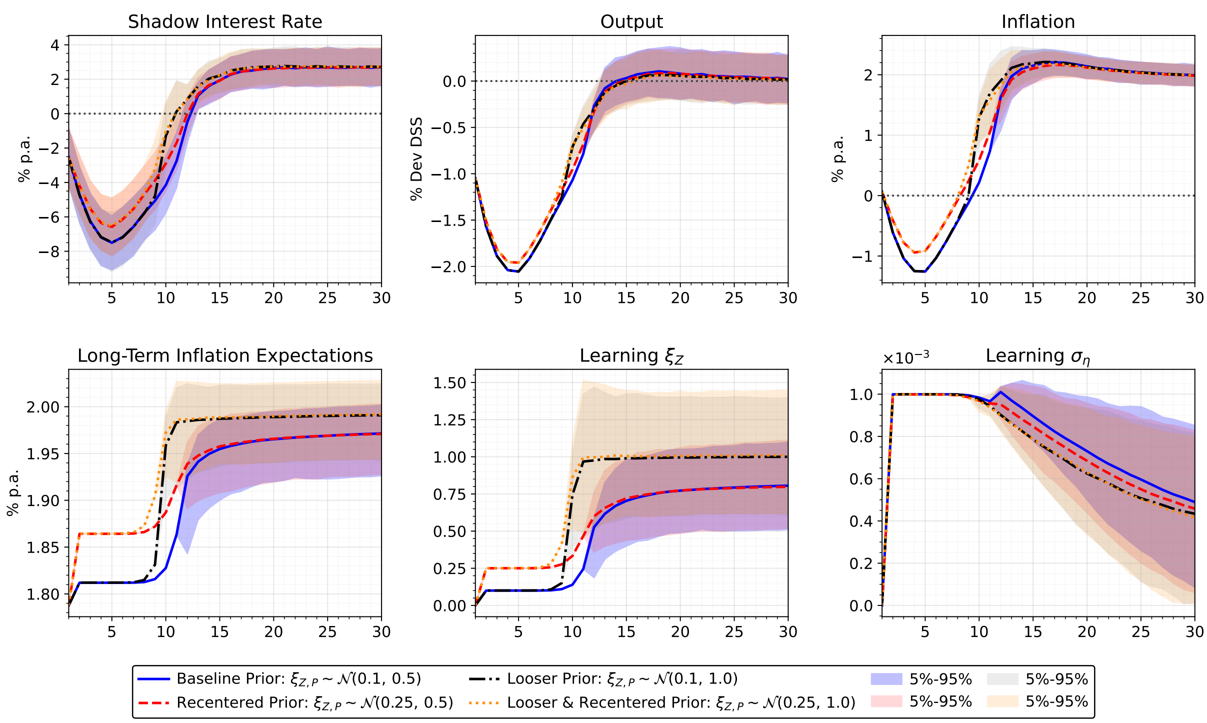

Figure 12: Impact of alternative priors for the structural components ξ_t^z of the new rule in the baseline simulation. The prior of the baseline case (blue) is compared to a scenario with a higher prior mean (red), a larger prior standard deviation (black) and a combination of higher mean and larger standard deviation (orange). The lines show the means, and the shaded areas indicate the 90 percent range of outcomes across the 1,000 simulations, which differ only in the sequence of monetary policy shocks.

A larger prior standard deviation also speeds learning relative to the baseline, because agents are more receptive to incoming information moving their views away from their prior. This is seem by comparing the black with the blue line and the orange with the red lines. However, this can work in two directions – information that is not supportive of the new policy would more readily push perceptions away from the correct ξ_z as well. In particular, if an $\varepsilon_{R,t}$ shock occurs that hastens liftoff relative to the baseline, agents would more readily lower their views of ξ_z if they had a larger prior standard deviation.[32]

These results point to the benefits of the central bank trying to move agents' priors closer to the true policy parameters or make them more receptive to the idea that a change in the policy rule has taken place. Clear communications about the new policy rule at the announcement and while shadow rates are still well below the ELB would be helpful. So would be concrete actions. For example, as discussed in Engen et al. (2015), large scale asset purchases may have provided monetary accommodation in the U.S. not just through portfolio balance effects, but also by the signal they sent about the Federal Reserve's commitment to operating a lower-for-longer interest rate policy.

[32] In the examples we chose, the effects of starting the learning process closer to the truth – a larger $\mathbb{E}_{\text{Prior}}(\xi_Z)$ – is less important than having a looser prior – a larger $\text{Std}_{\text{Prior}}(\sigma_R)$. We do not analyze how robust this result would be to alternative shock scenarios.

B Non-linear solution method with parameters as pseudo state variables

We solve the model in its non-linear form using global methods to allow the model's nonlinearities to influence economic outcomes. Therefore, the possibility of hitting the zero lower bound in the future affects potentially the equilibrium outcome in times of unconstrained monetary policy.

One challenge we face is that for every experiment we run, we need to solve the model in every period at many different values for the structural parameters, ξ_Z and σ_R, that agents are learning. As the model features three stochastic shocks, the computational burden is significant, rendering usual global solution approaches intractable.[33] To overcome this issue, we speed up the process by treating these parameters as pseudo-state variables and solving the model conceptually in line with Kase et al. (2022). We compute decision rules that are conditioned not only on the time-varying state variables, but also on the constant parameters. We use this step to front-load the costs of solving the model. When we need to simulate the model, we can rely on our pre-solved solution and quickly simulate the model forward. While our initial model solution with these decision rules takes several hours to compute, a model simulation during the likelihood evaluation step is solved in less than a second.

We use time iteration with piecewise linear interpolation of the decision rules as in Richter et al. (2014), which has been used for instance in Bianchi et al. (2021) to handle the non-linearities associated with zero lower bound.We adapt this solution approach by treating the parameters as pseudo-state variables. Expectations are calculated using numerical integration based on Gauss-Hermite quadrature.

In the parlance of this literature, the pseudo state variables $\mathbb{P}_t$ are ξ_Z and σ_η, the state variables $\mathbb{X}_t$ are z_t, $\epsilon_{d,t}$, $\epsilon_{\mu,t}$, and $\epsilon_{R,t}$, and the control variables are inflation π_t and output y_t.[34] The control variables are determined by an extended decision rule, g:

$$\begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = g\left(\xi_Z, \sigma_\eta, z_t, \epsilon_{d,t}, \epsilon_{\mu,t}, \epsilon_{R,t}\right) \tag{17}$$

where $g : R^6 \to R^2$. We approximate this unknown extended decision rule with a piecewise

[33]The model actually features four shocks. However, we can combine the two components of the monetary policy shock to a single monetary policy shock when solving the model.

[34]Note that $\epsilon_{R,t} = \epsilon_{M,t} + \epsilon_{\eta,t} \cdot \mathbf{1}_{z_t<0}$.

linear function $\tilde{g}$ that can be written as:

$$\begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = \tilde{g}\left(\xi_Z, \sigma_\eta, z_t, \epsilon_{d,t}, \epsilon_{\mu,t}, \epsilon_{R,t}\right) \tag{18}$$

The solution procedure for the extended decision rule is summarized as follows:

1. Define a discretized grid for the pseudo-state variables $\left\{\left[\underline{\xi}_Z, \overline{\xi}_Z\right], \left[\underline{\sigma}_R, \overline{\sigma}_R\right]\right\}$, states $\left\{\left[\underline{z}, \overline{z}\right], \left[\underline{\epsilon}_d, \overline{\epsilon}_d\right], \left[\underline{\epsilon}_\mu, \overline{\epsilon}_\mu\right], \left[\underline{\epsilon}_R, \overline{\epsilon}_R\right]\right\}$ and the integration nodes $\epsilon = \left\{\left[\underline{e}_d^I, \overline{e}_d^I\right], \left[\underline{e}_\mu^I, \overline{e}_\mu^I\right], \left[\underline{\epsilon}_R^I, \overline{\epsilon}_R^I\right]\right\}$.

2. Guess the piece-wise linear extended decision rule $\tilde{g}\left(\xi_Z, \sigma_R, z_t, \epsilon_{d,t}, \epsilon_{\mu,t}, \epsilon_{R,t}\right)$.

3. Solve for all endogenous time t variables for a given pseudo state variables vector ξ_Z, σ_R and state vector $z_t, \epsilon_{d,t}, \epsilon_{\mu,t}, \epsilon_{R,t}$. The control variables are:

$$\begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = \tilde{g}\left(\xi_Z, \sigma_R, z_t, \epsilon_{d,t}, \epsilon_{\mu,t}, \epsilon_{R,t}\right) \tag{19}$$

   The interest rate is then given as

$$R_t = \max\left[1, R\left(\frac{\pi_t}{\pi}\right)^{\phi_\Pi}\left(\frac{y_t}{y}\right)^{\phi_Y} + \mathbf{1}_{z_t<0} \cdot \xi_z z_t + \epsilon_{R,t}\right] \tag{20}$$

   where we do not need to distinguish whether the monetary shock comes from the usual component or from the new average inflation targeting component. We can then calculate consumption, wage and marginal costs:

$$c_t = y_t\left(1 - 0.5\varphi(\pi_t/\pi - 1)^2\right) \tag{21}$$

$$w_t = \chi y_t^\eta c_t^\sigma \tag{22}$$

$$mc_t = w_t \tag{23}$$

4. Move to time t+1. Update the term that captures the inflation misses,

$$z_{t+1} = \rho_z z_t + \phi_z(\pi_t - \pi) \tag{24}$$

   We then calculate the the monetary shock at each integration node i, which is

$$\epsilon_{R,t+1}^i = \begin{cases} \epsilon_{m,t+1}^i + \epsilon_{\nu,t+1}^i, & \text{if } z_{t+1} < 0 \\ \epsilon_{m,t+1}^i, & \text{else} \end{cases} \tag{25}$$

where we exploit that $\sigma_R = \sigma_M + \sigma_\eta \cdot \mathbf{1}_{z_t<0}$.

Next, calculate the remaining exogenous state variables for period $t+1$ at each integration node i:

$$\epsilon^i_{d,t+1} = 1 - \rho_d + \rho_d \epsilon_{d,t} + e^i_{d,t+1} \tag{26}$$

$$\epsilon^i_{\mu,t+1} = 1 - \rho_\mu + \rho_\mu \epsilon_{\mu,t} + e^i_{\mu,t+1} \tag{27}$$

For each integration node, calculate the control variables:

$$\begin{bmatrix} \pi^i_{t+1} \\ y^i_{t+1} \end{bmatrix} = \tilde{g}\left(\xi_Z, \sigma_R, z_{t+1}, \epsilon^i_{d,t+1}, \epsilon^i_{\mu,t+1}, \epsilon^i_{R,t+1}\right) \tag{28}$$

Afterwards, calculate consumption at each integration node:

$$c^i_{t+1} = y^i_{t+1}\left(1 - 0.5\varphi(\pi^i_{t+1}/\pi - 1)^2\right) \tag{29}$$

Calculate the errors for the Euler Equation and the New Keynesian Phillips curve

$$err_1 = 1 - \beta R_t E_t\left[\frac{\epsilon_{d,t+1}}{\epsilon_{d,t}}\left(\frac{c_t}{c_{t+1}}\right)^\sigma \frac{1}{\pi_{t+1} g_{t+1}}\right] \tag{30}$$

$$\begin{aligned} err_2 = {} & \varphi\left(\frac{\pi_t}{\pi} - 1\right)\frac{\pi_t}{\pi} - (1-\epsilon) - \epsilon mc_t - \epsilon_{\mu,t} \\ & - \beta\varphi E_t\left[\frac{\epsilon_{d,t+1}}{\epsilon_{d,t}}\left(\frac{c_t}{c_{t+1}}\right)^\sigma \left(\frac{\pi_{t+1}}{\pi} - 1\right)\left(\frac{\pi_{t+1}}{\pi}\right)\frac{\tilde{y}_{t+1}}{\tilde{y}_t}\right] \end{aligned} \tag{31}$$

where the expectations are numerically integrated across the integration nodes. The nodes and weights are based on Gaussian-Hermite quadrature.

5. Minimize the errors for the equations by adjusting the decision rules with a numerical root finder.

6. Update the decision rules until the errors at each point of the discretized state space (including the pseudo state variables) are sufficiently small.

7. Once $\tilde{g}\left(\xi_Z, \sigma_R, z_{t+1}, \epsilon^i_{d,t+1}, \epsilon^i_{\mu,t+1}, \epsilon^i_{R,t+1}\right)$ is determined, linearly interpolate between between discretized points to get the continuous decision rule $g\left(\xi_Z, \sigma_\eta, z_t, \epsilon_{d,t}, \epsilon_{\mu,t}, \epsilon_{R,t}\right)$.

We discretize the two endogenous pseudo state variables ξ_Z and σ_R in 12 and 11, respectively, evenly spaced points with bounds of $[0, 1.1]$ and $[0.00001, 0.002]$. The endogenous state variable z_t is discretized using 15 evenly spaced points with bounds of ± 0.012. The preference, markup and (combined) monetary policy shocks are discretized in 11 evenly-spaced points, where the bounds are chosen to be $1 \pm 6\sigma_d$, $\pm 5\sigma_\mu$, and $\pm 16\sigma_R$, respectively. The total number of nodes is $2,635,380$. Gauss-Hermite quadrature nodes provide the integration nodes and the corresponding weights used for numerical integration. We use 7 integration nodes for each shock, so that expected values are computed using 343 weighted points.

C Empirical analysis of the 2020 framework review

This appendix section describes our approach to map our model to the data over 2020-2023. This allows us to evaluate the learning process of the Federal Reserve's new framework through the lens of the model.

We start with a measurement equation that connects the observed time series that we want to match to our model. Specifically, we want to map the dynamics of real GDP per capita growth, PCE core inflation, and the federal funds rate. The measurement equation is

$$\begin{bmatrix} \text{GDP Growth Per Capita} \\ \text{PCE Core Inflation Rate} \\ \text{Federal Funds Rate} \end{bmatrix} = \begin{bmatrix} 100\dfrac{Y_t - Y_{t-1}}{Y_{t-1}} \\ 400\left(\Pi_t - 1\right) \\ 400\left(R_t - 1\right) \end{bmatrix} \tag{32}$$

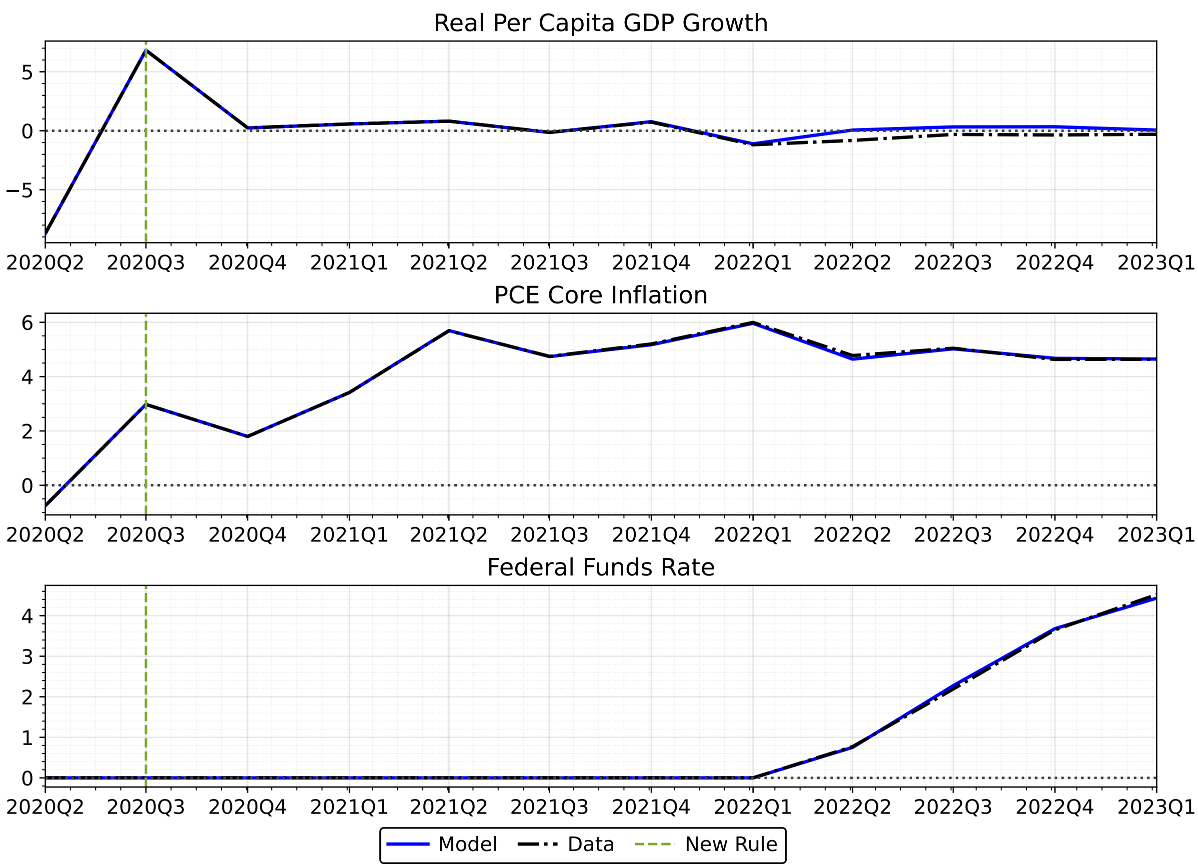

Figure 13: Fit of the model to the data. The blue lines show real GDP growth, core inflation, and the federal funds rate from the baseline fit of the model to the data. The black dash-dotted lines show the data. The vertical green line marks the announcement of the new rule in 2020:Q3.

We search for a sequence of shocks for our demand, markup, and monetary policy shocks that minimizes the distance in the measurement equation. We start with a random guess of 100 draws for each shock in each period, and then, in each period, use the draw that minimizes the squared distance between the data and the model. We then use the best draw in each period as our new starting point and draw again 100 shocks around this new point. The procedure is then continued until a sufficiently close fit of the model and the data.

The fit of the model and data can be seen in Figure 13 for the three series. The results indicate that our method achieves a strong fit. While the provided solution is not unique, as we have four shocks and the model is non-linear, such a poor man's filter serves as a very reasonable reference point.