university of groningen

faculty of economics and business

2026004-EEF

Disagreement at the Zero Lower Bound: The Case for Forward Guidance

Viktor Marinkov

FEBRI is the research institute of the Faculty of Economics & Business at the University of Groningen. FEBRI has seven programmes:

- Accounting
- Economics, Econometrics and Finance
- Global Economics & Management
- Innovation & Organization
- Marketing
- Operations Management & Operations Research
- Organizational Behaviour

FEB Research Institute (FEBRI)
Faculty of Economics & Business
University of Groningen

Visiting address:
Nettelbosje 2
9747 AE Groningen
The Netherlands

Postal address:
P.O. Box 800
9700 AV Groningen
The Netherlands

T +31 50 363 9090/7068

www.rug.nl/research/research-feb/

university of groningen / faculty of economics and business

Disagreement at the Zero Lower Bound: The Case for Forward Guidance

Viktor Marinkov
University of Groningen, Faculty of Economics and Business, Department of Economics, Econometrics and Finance
viktor.marinkov@rug.nl

Disagreement at the Zero Lower Bound: The Case for Forward Guidance

Viktor Marinkov*
University of Groningen
September 2026

Abstract

At the ZLB, the policy rate stops revealing the central bank's shadow stance, censoring the forecast errors needed for learning. I show that this distorts agents' perceived policy law of motion, producing earlier expected lift-off than the central bank under an unchanged policy rule - a pattern documented in Swedish and U.S. forecasts. In a New Keynesian model, agents expect lift-off seven quarters early, and repeated premature forecasts can push perceived policy persistence beyond the stability boundary. Communicating the rule-consistent lift-off date reduces disagreement, contains belief drift, and improves macroeconomic outcomes, providing a case for parsimonious, informational, commitment-free forward guidance.

Keywords: Delphic forward guidance; unanchored expectations; central bank communication; zero lower bound; adaptive learning;

JEL Classifications: E43; E52; E58; D84;

*I thank Guido Ascari, Roberto Billi, William Branch, Laura Coroneo, Jakob de Haan, Juan Dolado, John Duffy, Martin Ellison, Stefano Eusepi, George W. Evans, Andrea Ferrero, Gaetano Gaballo, Emanuel Gasteiger, Christopher G. Gibbs, Alex Grimaud, Richard Harrison, Klodiana Istrefi, Joep Lustenhower, Ramon Marimon, Bruce McGough, Michael McMahon, Francesca Monti, Luba Peterson, Bruce Preston, Christian R. Proaño, Kate Reinold, Luis Rojas, Thomas Sampson, Sergey Slobodyan, Christiaan van der Kwaak, Matt Waldron, and Francesco Zanetti for helpful comments and discussions. I also thank participants in Ramon Marimon's reading group and at the annual Behavioral Macroeconomics Workshop, Expectations in Dynamic Macroeconomic Models, and the Barcelona Summer Forum groups on expectations and behavioral macroeconomics.
An earlier version of this paper, under the title "Communication at the Zero Lower Bound: The Case for Forward Guidance," circulated as University of Oxford Department of Economics Working Paper No. 923.

1 Introduction

During recent low interest rate episodes, private policy rate forecasts often exceeded central bank projections, implying earlier expected lift-off. Can the lower bound itself generate this disagreement when the policy rule is unchanged? When the policy rate is constrained by the zero lower bound (ZLB)[1], its realization ceases to reveal the stance that monetary policy would take absent the constraint. This paper shows that, after agents have learned normal-times dynamics, the resulting data censoring leads adaptive learners to expect lift-off too early. It also creates a role for forward guidance that does not rely on a promise to deviate from future policy: communicating the central bank's rule-consistent lift-off projection supplies a missing signal about the policy path.

The mechanism is straightforward. Above the ZLB, agents can compare their policy rate forecasts with realizations and update their perceived law of motion. At the ZLB, they expect and observe the lower bound. The resulting zero policy rate forecast error prevents them from learning that their latent shadow rate path is too high. That is the case because the ZLB changes how shocks propagate to output and inflation: adverse shocks contract activity more when policy cannot provide its normal accommodation, which implies a deeper shadow rate path under an unchanged policy rule. Agents gradually learn the output and inflation dynamics from observable errors, but the censored shadow rate error leaves their expected policy path uncorrected and too high. When they prematurely forecast lift-off and the policy rate remains at the floor, the subsequent negative forecast errors tend to raise their perceived persistence of the shadow rate. With long-horizon expectations, repeated errors can therefore unanchor the expected policy path and feed back into output and inflation. The disagreement therefore reflects an informational advantage of the central bank.[2] Although both sides observe the shocks, the bank knows the structural model and can compute how the binding ZLB changes the

[1] I use "ZLB" for expositional convenience. The mechanism applies to an effective lower bound at any level.

[2] This differs from Andrade et al. (2019), where private forecasters largely agree about the expected low-rate path but disagree about its macroeconomic interpretation. Here, disagreement is instead between the representative private sector and the central bank over the policy rate path itself.

aggregate law of motion and the rule-consistent shadow-rate path, whereas private agents must learn that mapping from censored observations. In this broad sense, the mechanism is an information effect, but its source is asymmetric model knowledge rather than the private news about fundamentals usually meant by the "Fed information effect" (Bauer and Swanson, 2023).

I first document the same state-dependent disagreement under two markedly different transparency regimes. The Riksbank has published its repo-rate path in real time since 2007, allowing me to match these projections with contemporaneous money-market surveys through 2026. By contrast, Federal Reserve staff funds-rate projections were confidential and became public only after a five-year lag; I match the released vintages with contemporaneous *Wall Street Journal* forecasts for the 2008 - 2015 and 2020 ZLB episodes. Despite this institutional contrast, both comparisons yield the same pattern: at the lower bound, private forecasters expect earlier lift-off than the central bank. The U.S. pattern is consistent with evidence that market participants expected lift-off within roughly four quarters through 2010 (Swanson and Williams, 2014) and that investors over-predicted short rates in downturns and Federal Reserve easings (Cieslak, 2018). These comparisons motivate the direction and state-dependence of the interest rate disagreement pattern studied in the model, but they do not by themselves identify the censoring mechanism presented here.

I then embed constant-gain adaptive learning with long-horizon expectations (Preston, 2005) in a canonical New Keynesian model. Agents are additionally able to understand the lower bound, track a notional shadow rate, and truncate their policy rate forecasts at the floor. They are therefore more sophisticated than learners who simply extrapolate the observed constrained rate, but they do not immediately internalize how the binding constraint changes the aggregate law of motion. I show analytically that, after beliefs have converged to the normal-times law of motion, a sufficiently adverse demand shock that binds the ZLB makes agents' output, inflation, and shadow rate forecasts too high relative to the central bank's model-consistent projections. Agents consequently predict an earlier lift-off even though the policy is unchanged.

The numerical experiments assess the size and consequences of this wedge. Following a common Great-Recession-sized shock, the median agent without communication expects lift-off seven quarters earlier than the central bank. Across 10,000 matched 500-quarter simulations, 45.3% of no-communication paths attempt to push perceived shadow rate persistence beyond the unit circle stability boundary. Depending on how agents incorporate the announcement, lift-off guidance lowers this incidence to 7.4 - 11.2% and, over the 40-quarter crisis window, mean stabilization loss falls by 23 - 40%. The gains are economically meaningful but heterogeneous across histories.

In every binding period, the central bank truthfully reports its expected remaining ZLB duration, which identifies its expected lift-off date; in the learning treatments, it also reports the projected rate at lift-off. Agents may impose the communicated duration on their interest rate forecasts, use the lift-off projection to update their perceived law of motion, or do both. The announcement reveals the bank's model-consistent policy projection rather than a deviation from its policy rule. It is therefore Delphic forward guidance in the terminology of Campbell et al. (2012), rather than the Odyssean commitment studied by Eggertsson and Woodford (2003) and Krugman et al. (1998).[3] The lift-off communication thereby conveys information about the otherwise unobserved shadow policy trajectory. By allowing agents to either align their expected ZLB duration or informing revisions to their perceived law of motion, it reduces forecast disagreement and improves macroeconomic outcomes. The results support the simplicity of the message as they show the biggest gains are received when the ZLB duration disagreement is closed, rather than the precise gradient of the policy rate.

A close econometric counterpart is Aruoba et al. (2022), who develop an occasionally binding SVAR in which the observed policy rate is censored and shock propagation changes when the constraint binds. Whereas they estimate regime-dependent dynamics and recover a shadow rate, this paper asks how private agents learn the changed dynamics when the shadow rate is unobserved. Agents maintain a linear perceived law of motion for a latent shadow rate and truncate the implied policy-rate path at the floor. The resulting

[3] Campbell et al. (2017) find that FOMC announcements contain a substantial Delphic component while leaving room for Odyssean guidance.

censoring generates early-lift-off disagreement and creates a role for communication.

Existing learning models generate liquidity or deflation traps through pessimistic expectations (Benhabib et al., 2014; Evans et al., 2008), while Bullard and Mitra (2002) and Evans and Honkapohja (2003) study stability under alternative Taylor rules. Most closely, Eusepi (2010) shows that policy-rule transparency stabilizes learning in a non-linear liquidity-trap model, and Eusepi and Preston (2010) show that communicating the policy function anchors long-horizon expectations. Here, the rule is unchanged and known; communication instead reveals a shadow policy path hidden by the constraint.

The policy implication is state dependent. In normal times, the model provides no systematic informational case for policy-path communication because the observed rate reveals the policy stance. At the lower bound, however, routine and truthful publication of the bank's rule-consistent policy projection can reduce premature lift-off beliefs and the associated drift in perceived policy persistence. Section 2 presents the evidence, Section 3 develops the model and analytical result, and Section 4 reports the quantitative experiments.

2 Evidence of forecast disagreement at the ZLB

The Great Recession and the long ZLB spell that followed produced historically high disagreement among private forecasters (Andrade et al., 2019) and systematic forecast errors relative to empirical policy rules. Engen et al. (2015) show that over 2008-2014 Blue Chip forecasts of the Treasury bill exceeded the path implied by standard Taylor rules. They attribute this to a misperceived Fed reaction function and excessive optimism about the speed of recovery.

I ask whether the same optimism appears when private forecasts are compared with central banks' internal interest rate projections rather than with stylized rules. That comparison requires both sides of the forecast. Throughout this section I define the policy rate disagreement gap as the private forecast minus the central bank's internal forecast. A positive value means that private agents expect a higher policy rate, which

Table 1: One-year-ahead disagreement on the Swedish repo rate

Based on money-market 1-year-ahead repo rate

	count	mean	se(mean)	min	max
Low	41	0.0756	0.0168	-0.2157	0.2566
High	61	-0.1026	0.0405	-0.8958	0.6688

Based on Riksbank 1-year-ahead repo rate

	count	mean	se(mean)	min	max
Low	44	0.0877	0.0178	-0.2157	0.3677
High	58	-0.1210	0.0410	-0.8958	0.6688

Notes: Gap = money-market expectation minus Riksbank path. Positive: the market expects a higher repo rate. Low (High): one-year-ahead expectation below (above) 0.25. Primary pairs, February 2007 - June 2026.
Source: Riksbank Monetary Policy Reports and Updates; Prospera/Origo Money Market Players.

at the effective lower bound this is the direction associated with earlier expected lift-off.

The Riksbank is unusually transparent among major central banks in publishing its own policy rate path together with a survey of money-market participants. Own-path publications begin with the February 2007 Monetary Policy Report. I reconstruct the comparison from official Riksbank vintages and Prospera/Origo Money Market Players releases through June 2026.[4] Figure 1 plots one- and two-year-ahead repo-rate forecasts. When the one-year-ahead rate is near the lower bound, the market path lies above the Riksbank path most of the time. Table 1 reports the one-year numerical disagreement. Low (High) denotes a one-year-ahead expectation below (above) 0.25 percent. In the low-rate regime the mean disagreement is 0.076 percentage points when the split uses the market forecast and 0.088 percentage points when it uses the Riksbank forecast (standard errors 0.017 and 0.018). The market is higher in 31 of 41 and 34 of 44 observations, respectively. Appendix A.2 presents further horizon and cutoff robustness, and a snapshot case study of 2009-10. The results support the idea that markets tend to have higher medium-run expectations of the policy rate relative to the monetary authority in low-rate regimes such as the lower bound.

On the opposite side in terms of transparency, the Federal Reserve did not use to publish a routine funds-rate path over 2003-2014. The Tealbook (formerly Greenbook)

[4] Appendix A.1 describes sources, pairing, and horizon alignment. I match each Riksbank publication to the latest survey dated on or before it and drop three pairs whose survey is more than 45 days old. Among the 102 retained pairs the median lag is 14 days and the largest is 29 days.

Figure 1: Forecasts of the Swedish repo rate

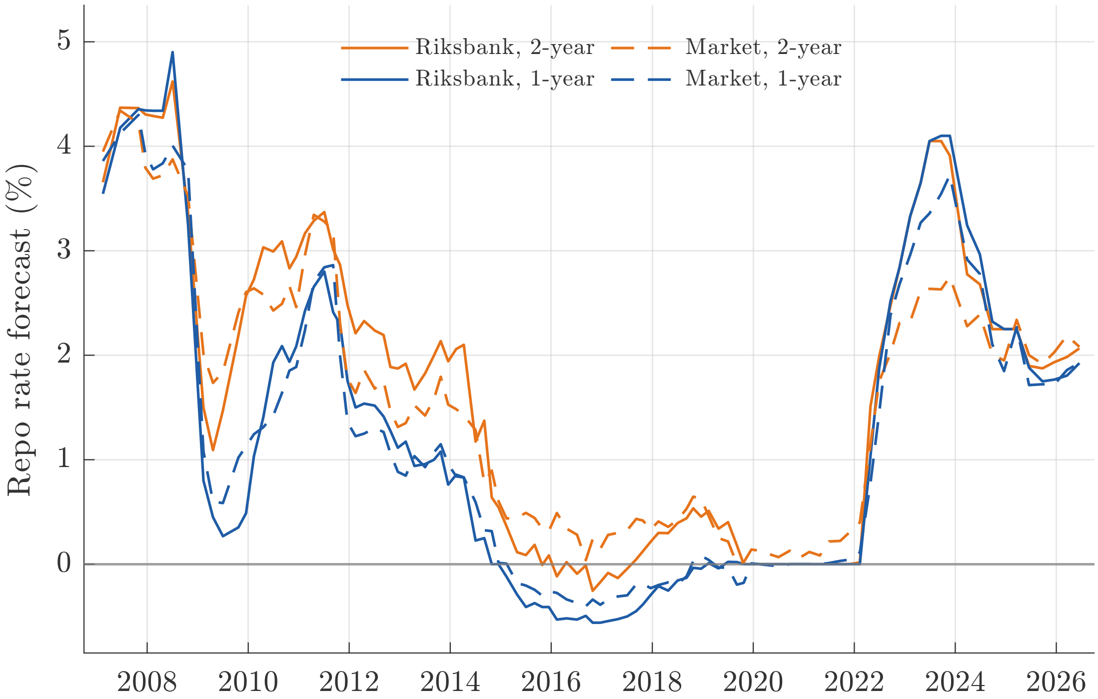

Notes: Solid lines: Riksbank. Dashed lines: money-market expectations (Prospera/Origo). Source: Riksbank Monetary Policy Reports and Updates, 2007–2026.

records Board staff projections prepared before each FOMC meeting and released with a five-year lag. Staff funds-rate assumptions are available from January 1981 through December 2014 at eight meetings per year, with horizons of up to eight quarters.[5] Private forecasts come from the Wall Street Journal Economic Forecasting Survey, available monthly since January 2003. Targets are end-June and end-December; from 2009 - 2011 the horizon sometimes extends to three years. Matching by calendar month of the Tealbook publication yields the observation counts in Table A.2.

Figure 2 plots the WSJ mean minus the FED staff projection at horizons of one to eight quarters. Positive values mean that the survey expects a higher funds rate than staff, the same convention as for Sweden.

During periods at the zero lower bound, notably 2008 - 2011, the survey lies above staff, and the gap is larger at longer horizons. The gap shrinks with the 2005 - 2007

[5] See Appendix A.1. The staff series is the Philadelphia Fed Real-Time Data file "Output Gap and Financial Assumptions from the Board of Governors."

Figure 2: Federal funds rate forecast disagreement, 2003 - 2014

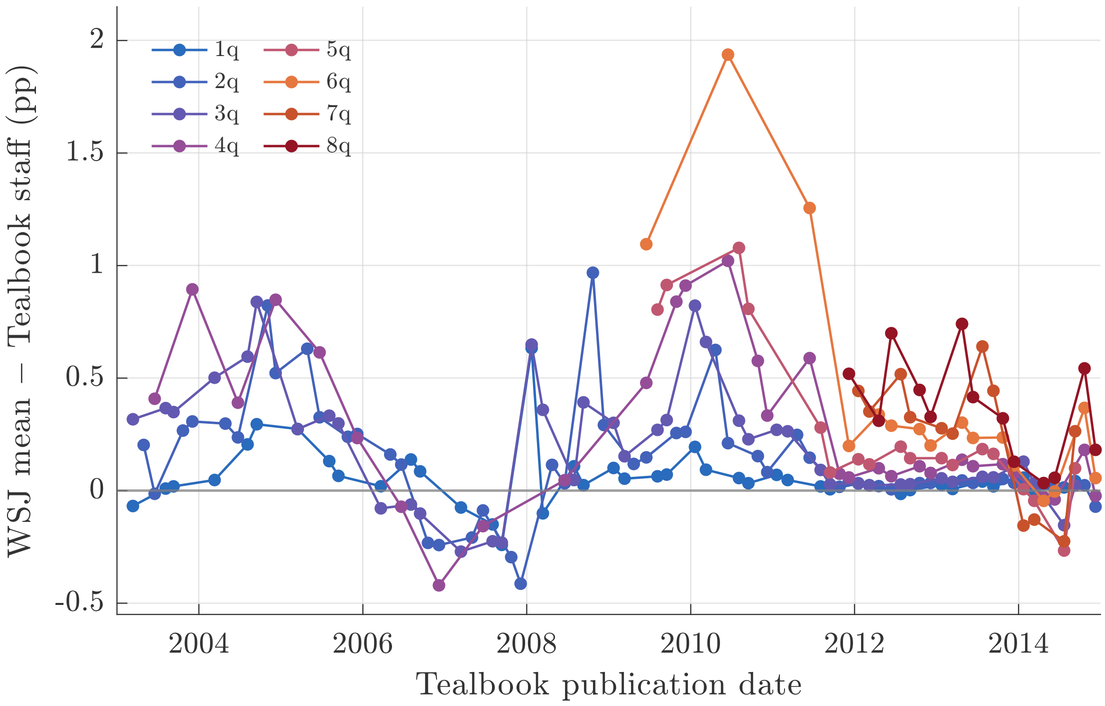

Notes: Positive values mean private forecasters expect a higher funds rate than Board staff. Source: Federal Reserve Tealbook and Wall Street Journal Economic Forecasting Survey.

hiking cycle and after calendar- and threshold-based forward guidance from early 2012.[6] Through mid-2013 the survey still expects a higher funds rate six to eight quarters ahead. The pattern matches the Riksbank evidence: in low-rate environments private agents tend to expect a higher policy rate path than the central bank. Daly (2025) notes that during the Global Financial Crisis markets regularly expected rates to rise faster than FOMC participants did, and that more explicit communication helped clarify the path. Bernanke (2012) and Yellen (2011) likewise argue that guidance can correct expectations of an earlier increase than the FOMC anticipated. This wedge is the empirical counterpart of the information channel modeled below.

Using Blue Chip Financial Forecasts, Cieslak (2018) finds that private funds-rate forecast errors are large and negative in easing episodes, especially 2007 - 2009, and more predictable than the future funds rate itself. Tealbook staff forecasts are more accurate than the market's but still exhibit predictable errors. That paper studies forecast errors

[6]Quantitative easing is outside the scope of this paper; see Bhattarai et al. (2022).

rather than staff - survey disagreement; it is consistent with the difficulty of forecasting the funds rate as the lower bound approaches.

The Philadelphia staff funds-rate file ends in 2014.[7] To study the more recent pandemic ZLB period, I code the Tealbook A Long-Term Outlook Q4 staff path from the eight 2020 issues and match each Q4 point to the same-month WSJ mean for that year's December.[8] Figure 3 shows the two paths at each Tealbook date. In January and March, staff still projected a rising funds rate; the WSJ year-end means lie below that path. From April the staff path is at the ELB through 2022 (later through 2023 - 24), while the survey lies above staff at longer year-end horizons - earlier expected lift-off, as in 2008 - 2011. Figure A.2 reports the same comparison as a gap by Tealbook date. The end-2020 gap is near zero once both sides are at the ELB; the end-2021 through end-2023 gaps are positive and then shrink through the rest of 2020.

Using probabilities implied by federal-funds futures, Janson and Jia (2020) report that the probability of remaining in the ELB target range rose from 88 to 100 percent for the December 2020 meeting and from 86 to 100 percent for the September 2021 meeting around the June 2020 FOMC meeting. The accompanying SEP placed all 17 participants at the ELB through end-2021. The pre-announcement wedge has the same direction as in 2008-2014: markets attached positive probability to earlier lift-off than the subsequently revealed FOMC path (and the then-unreleased staff forecasts). Its disappearance after the SEP dots were published is consistent with guidance affecting the market's policy rate beliefs.

The relevant disagreement need not appear in modal forecasts. In the June 2020 Surveys of Primary Dealers and Market Participants the median modal path remained at 0.13 percent through end-2022, but combined respondents placed about 52 percent probability on an end-2022 midpoint above 0.25 percent. That probability fell to 40 percent in July and 31 percent in September, while the probability on the 0 - 0.25 percent range rose from about 44 to 65 percent (Federal Open Market Committee, 2020). By December 2020, SEP and New York Fed modal paths coincide at six-month, one-year, and

[7]Upon inquiry, no expansion of the funds-rate releases was planned past December 2014.
[8]See Appendix A.1.

Figure 3: Tealbook A staff Q4 path and WSJ year-end (December) mean, 2020

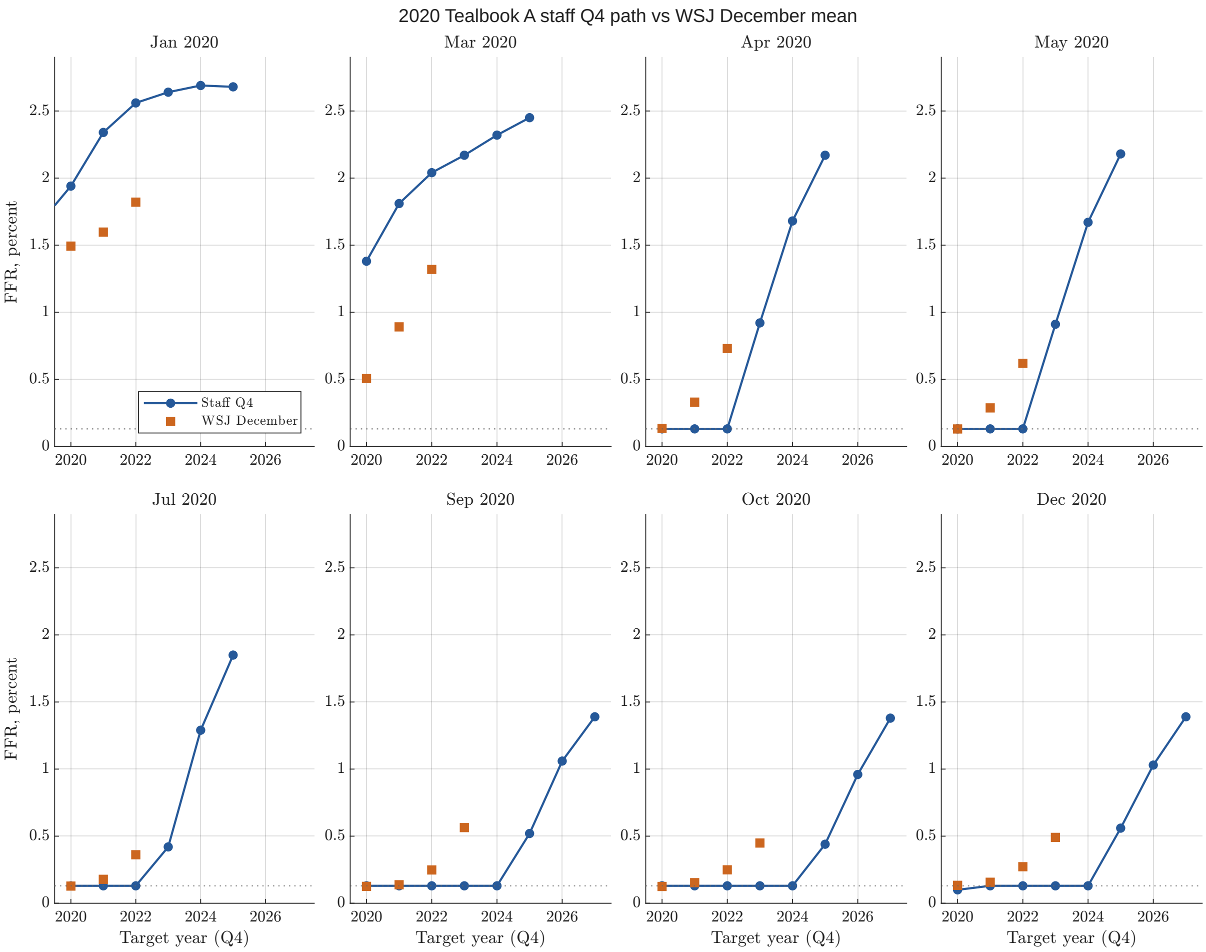

two-year horizons, though private distributions still assigned 29 - 33 percent probability to a midpoint above 0.25 percent at two years (Board of Governors of the Federal Reserve System, 2020; Federal Reserve Bank of New York, 2020a,b).

Abstracting from disagreement, Campbell et al. (2012) find that future monetary policy tightening lowers unemployment expectations and raises inflation expectations in the US, contrary to New Keynesian predictions. They interpret this as Delphic forward guidance. Campbell et al. (2017) classify FOMC private information by the difference between Tealbook forecasts of inflation, GDP growth and unemployment and Blue Chip private forecasts. The four-quarter-ahead futures rate is positively correlated with policymakers' GDP-growth forecast exceeding the market's (and negatively with unemployment). Meeting announcements therefore transmitted the committee's private information about the outlook - supporting a Delphic forward guidance interpretation.

3 Model

To explain why disagreement between the central bank and the private agents arises at the zero lower bound I build a simple New Keynesian model featuring adaptive learning as an expectation formation framework for the agents. The model environment is the canonical New Keynesian model with a representative consumer and monopolistically competitive firms subject to Calvo pricing. The novel contribution is to augment the agents' mental model by keeping track of a fictional shadow rate. This allows the agents to form expectations of future binding ZLB spells and their lift-offs. Moreover, I extend the long-horizon learning framework by allowing the agents to interpret central bank communication about the future path of the policy rate, i.e. forward guidance. This section begins by outlining the model under rational expectations before introducing long-horizon adaptive learning as the expectations formation framework of the agents.

3.1 Rational Expectations

As extensively discussed in Woodford (2003) and Galí (2015), under rational expectations (RE) the linearized aggregate economy of the canonical New Keynesian model can be summarized by the following two equations:

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t\right) \tag{1}$$

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t \pi_{t+1} + u_t \tag{2}$$

with shock processes

$$r_t = \rho_r r_{t-1} + \varepsilon_t^r, \quad \varepsilon_t^r \sim N(0, \sigma_r^2) \tag{3}$$

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u, \quad \varepsilon_t^u \sim N(0, \sigma_u^2) \tag{4}$$

where x_t is the current output gap, defined as the difference between output and its natural rate in an economy with fully flexible prices; π_t denotes the inflation rate; i_t the nominal interest rate; β is the discount factor; σ is the coefficient of relative risk aversion; and κ is a composite of structural parameters. All endogenous variables are expressed as log-deviations from their steady state values. Thus, in steady state $x = \pi = i = 0$. Finally, r_t and u_t stand for exogenous natural rate and cost-push shocks, respectively, and follow $AR(1)$ processes.

The model is closed with a monetary policy (MP) rule subject to the zero lower bound (ZLB):

$$s_t = \rho_i s_{t-1} + (1 - \rho_i)(\chi_\pi \pi_t + \chi_x x_t)$$

$$i_t = \max\{i^*, s_t\} \tag{5}$$

where s_t is the hypothetical unconstrained shadow rate and i_t is the realized policy rate subject to the lower bound i^*. Note that above the ZLB this monetary policy rule implies that the policy and shadow interest rates coincide - that is $s_t = i_t$ if $s_t \geq i^*$, while i_t is constrained by the ZLB where s_t is not - $i_t = i^*$ if $s_t < i^*$.

The reaction parameters satisfy $\chi_\pi > 1$ and $\chi_x > 0$, and the interest rate smoothing - $\rho_i \in (0,1)$. The constant $i^* = 1 - 1/\beta < 0$ represents the effective lower bound on interest rates since, otherwise, as explained in Eggertsson and Woodford (2003) agents would choose to hold all their assets in cash. I will refer to it as the ZLB to be consistent with the arguments in the Introduction and with real world analogies.

The rational expectations minimum state-variable solution to the model in (1) - (5) is of the form:

$$\underset{3\times 1}{Y_t} \equiv \begin{bmatrix} x \\ \pi \\ i \end{bmatrix}_t = \underset{3\times 4}{\Gamma} \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} \equiv \Gamma Z_t \tag{6}$$

where Y_t and Z_t contain the endogenous and state variables, respectively, and Γ is the rational expectations equilibrium policy matrix[9]. Note that due to the smoothing in the Taylor rule, the lagged interest rate becomes a state variable.

3.2 Expectations formation

The specification of expectations employed is adaptive learning (ADL). In particular, agents do not know the true structure of the economy and make forecasts as econometricians using simple regression models.[10] They observe both the realizations of the endogenous (Y_t) and state variables (Z_t) and try to estimate the transition matrix Γ as more data becomes available.[11]

As such, adaptive learning is inherently a linear updating procedure, which when iterated forward produces forecasts for the endogenous variables which are also linear in the states. On the other hand, the ZLB creates a non-linearity in the path of the interest

[9]Note that (6) represents the solutions of the model under RE without a ZLB. If the ZLB is respected, when binding the solution of the model will be piece-wise linear featuring a sequence of different policy transformations Γ^i for every period i when the ZLB is binding.

[10]This follows the "consistency principle" of Evans and Honkapohja (2001).

[11]This updating is usually done using recursive least squares, constant gain learning or gradient descent.

rate by truncating it as in (5). Thus, to account for that the agents must understand the interest rate cannot go below i^* both in realizations and in their forecasts. To get around this issue I model the agents as forming expectations about the shadow interest rate and then truncating at the ZLB to obtain their interest rate forecasts consistent with the monetary policy rule in (5). During a period of binding ZLB the agents use their forecasted shadow rate (s_t) to form lift-off expectations based on the past observed policy prior to the ZLB.

Keeping mental track of both the policy rate and an estimated shadow rate is important to avoid a forecast bias for the learners. When $s_t < i^*$ if they were to use $i_{-1} = i^*$ as a basis for expectations for $t = 0, 1...$ in (6), they would have an upward bias in their projected paths for the interest rate because the ZLB i^* is higher than the shadow rate at $t = -1$. Hence, as outlined below, I assume that above the ZLB agents rely solely on realized prices (i.e. actual interest rate i_t through the states vector Z_t) for updating their expectations and forecasts. When the ZLB binds, on the other hand, due to a lack of exact observable data on the policy rate, they rely on their shadow rate projections (s_t) for keeping track of the full path the policy instrument below the ZLB. The states vector at the ZLB becomes $S_t \equiv \begin{bmatrix} 1 & u_t & r_t & s_{t-1} \end{bmatrix}'$. Thus, even though the use of a shadow rate complicates the notation, this dichotomy is necessary for more realistic and sophisticated expectations. In this sense the imperfect knowledge of the adaptively learning agents here is conservative.

As shown next, as additional data become available each period the agents update their forecasting model by updating the coefficients to their perceived transition matrix Φ_t using a recursive constant gain algorithm. As is common in the literature, I assume that they observe the disturbances r_t and u_t and know their autoregressive coefficients.[12]

Adaptive Learning

Just like in the rational expectations solution in (6) the learning agents use the state variables vector and a transition matrix to forecast the endogenous variables vector.

[12]Eusepi and Preston (2010) show that this assumption can be dispensed with and instead agents would estimate those coefficients. For simplicity, it is maintained.

Unlike RE, however, they do not know the correct transition matrix Γ from (6) and instead use their estimated perceived 3-by-4 transition matrix Φ_{t-1} from the end of period $t-1$:

$$\Phi_{t-1} \equiv \begin{bmatrix} \phi_{x,t-1} \\ \phi_{\pi,t-1} \\ \phi_{i,t-1} \end{bmatrix} \equiv \begin{bmatrix} \phi_{xc,t-1} & \phi_{xu,t-1} & \phi_{xr,t-1} & \phi_{xi,t-1} \\ \phi_{\pi c,t-1} & \phi_{\pi u,t-1} & \phi_{\pi r,t-1} & \phi_{\pi i,t-1} \\ \phi_{ic,t-1} & \phi_{iu,t-1} & \phi_{ir,t-1} & \phi_{ii,t-1} \end{bmatrix} \tag{7}$$

The rows of Φ_{t-1} - $\phi_{x,t-1}, \phi_{\pi,t-1}, \phi_{i,t-1}$, contain the perceived laws of motion for the output gap, inflation, and the shadow interest rate, respectively. Its columns contain the intercept and the responses to the cost-push shock, the natural rate shock, and the lagged observed policy rate. In particular, $\phi_{ii,t-1}$ is the agents' perceived persistence of the interest rate. Note that above the ZLB[13], when $s_t = i_t \Rightarrow S_t = Z_t$, we could simply express the agents' beliefs as:

$$\hat{\mathbb{E}}_t \begin{bmatrix} x_t \\ \pi_t \\ i_t \end{bmatrix} = \Phi_{t-1} Z_t = \begin{bmatrix} \phi_{x,t-1} \\ \phi_{\pi,t-1} \\ \phi_{i,t-1} \end{bmatrix} Z_t \tag{8}$$

However, given the presence of a lower bound constraint on the interest rate and following the discussion above, I assume agents use the vector of realized prices $Z_t \equiv \begin{bmatrix} 1 & u_t & r_t & i_{t-1} \end{bmatrix}'$ to form expectations of the output gap and inflation, but they use $S_t \equiv \begin{bmatrix} 1 & u_t & r_t & s_{t-1} \end{bmatrix}'$ to forecast the shadow rate, and hence - the interest rate by respecting the ZLB. The timing is such that at the beginning of period t the agents observe the states (Z_t and S_t) and use their perceived transition matrix Φ_{t-1} to form nowcasts for x_t, π_t, i_t, s_t as follows:

[13] Or in a model with no binding ZLB.

$$
\begin{aligned}
\hat{\mathbb{E}}_t \begin{bmatrix} x_t \\ \pi_t \end{bmatrix} &= \begin{bmatrix} \phi_{x,t-1} \\ \phi_{\pi,t-1} \end{bmatrix} Z_t \\
\hat{\mathbb{E}}_t s_t &= \phi_{i,t-1} S_t \\
\hat{\mathbb{E}}_t i_t &= \max\left\{ i^*,\ \hat{\mathbb{E}}_t s_t \right\}, \qquad j \geq 0
\end{aligned} \tag{9}
$$

where $\hat{\mathbb{E}}$ is the expectations operator for the learners. Longer-run expectations are also formed at the beginning of period t by iterating forward the state vectors $\mathbb{E}_t Z_{t+j}$ and $\mathbb{E}_t S_{t+j}$ as in (10) and applying again the current perceived transition matrix Φ_{t-1} as in (13).

$$
\begin{aligned}
\hat{\mathbb{E}}_t S_{t+j} &= \begin{bmatrix} 1 & \rho_u^j u_t & \rho_r^j r_t & \hat{\mathbb{E}}_t s_{t+j-1} \end{bmatrix}' \\
\hat{\mathbb{E}}_t s_{t+j} &= \phi_{i,t-1} \hat{\mathbb{E}}_t S_{t+j} \\
\hat{\mathbb{E}}_t i_{t+j} &= \max\left\{ i^*,\ \hat{\mathbb{E}}_t s_{t+j} \right\}, \qquad j \geq 0 \\
\hat{\mathbb{E}}_t Z_{t+j} &= \begin{bmatrix} 1 & \rho_u^j u_t & \rho_r^j r_t & \hat{\mathbb{E}}_t i_{t+j-1} \end{bmatrix}'
\end{aligned} \tag{10}
$$

Note that because the interest rate is a state variable, the agents first need to forecast the interest rate up to period $t+j-1$ before they can forecast output gap or inflation at period $t+j$. To actually obtain a forecast for the interest rate at $t+j-1$, the agents first forecast the shadow rate at the same horizon and apply the ZLB constraint as in (10) above. Therefore, interest rate forecasts depend on recursively forming one-step-ahead forecasts for the shadow rate and updating the forecasted state vectors $\mathbb{E}_t Z_{t+j}$ and $\mathbb{E}_t S_{t+j}$.

To express this recursive forecasting procedure more compactly, consider first the case in which no binding lower bound is expected. Then $i_{t+j} = s_{t+j}$ and hence $Z_{t+j} = S_{t+j}$. Combining the evolution of the exogenous states in equation (10) with the perceived law of motion for the shadow rate in equation (9) gives

$$\hat{\mathbb{E}}_t S_{t+1} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \rho_u & 0 & 0 \\ 0 & 0 & \rho_r & 0 \\ \phi_{ic,t-1} & \phi_{iu,t-1} & \phi_{ir,t-1} & \phi_{ii,t-1} \end{bmatrix}}_{\Psi_{t-1}} S_t. \tag{11}$$

Holding agents' period-t beliefs fixed over the forecast horizon, their expected future state vector is therefore

$$\hat{\mathbb{E}}_t S_{t+j} = \Psi_{t-1}^j S_t. \tag{12}$$

Conditional on computed forecasts for the states, the agents are able to form expectations about the aggregate endogenous variables at any horizon:

$$\begin{aligned} \hat{\mathbb{E}}_t \begin{bmatrix} x_{t+j} \\ \pi_{t+j} \end{bmatrix} &= \begin{bmatrix} \phi_{x,t-1} \\ \phi_{\pi,t-1} \end{bmatrix} \hat{\mathbb{E}}_t Z_{t+j} \\ \hat{\mathbb{E}}_t s_{t+j} &= \phi_{i,t-1} \hat{\mathbb{E}}_t S_{t+j} \\ \hat{\mathbb{E}}_t i_{t+j} &= \max \left\{ i^*,\ \hat{\mathbb{E}}_t s_{t+j} \right\}, \qquad j \geq 0 \end{aligned} \tag{13}$$

Note that the only autoregressive process in the expectations formation is that of the shadow rate - s_{t+j}, and its autoregressive coefficient is $\phi_{ii,t-1}$ - namely, the perceived persistence of the shadow rate.[14] This AR(1) coefficient is then the only unknown eigenvalue of the agents' PLMs for endogenous and exogenous variables - Φ_{t-1} and Ψ_{t-1}, respectively[15]. The anchoring of long-run expectations and the stability of the learning system will depend mainly on whether $\phi_{ii,t}$ stays within the unit circle as the agents update their beliefs. Section 4.3 further highlights its importance in a numerical exercise and discusses its relationship to interest rate forecast disagreement between the agents

[14] In simpler terms, this can also be understood as the perceived persistence of the interest rate but due to the presence of an occasionally binding ZLB it is technically correct to refer to the unconstrained shadow rate.

[15] Remember that, as is standard in the literature, it is assumed that the agents perfectly know the two shocks AR(1) coefficients - ρ_u and ρ_r.

and the central bank (see Section 3.4).

At the end of period t the agents update this perceived law of motion (PLM) by a recursive constant gain algorithm using the difference between their nowcasts of endogenous variables $\hat{\mathbb{E}}Y_t$ and the actual realizations Y_t. Each error correction term is given a constant gain τ against their prior beliefs from $t-1$. They also weigh this forecast error by the historical variance-covariance matrix R_t of the state variables, which they first update in a similar manner, and use the weighted forecast discrepancy for error correction. [16].

$$\Phi_t = \begin{bmatrix} \phi_x \\ \phi_\pi \\ \phi_i \end{bmatrix}_t = \Phi_{t-1} + \tau \left(Y_t - \hat{\mathbb{E}}_t Y_t \right) Z_t' R_t^{-1} \tag{14}$$

$$R_t = R_{t-1} + \tau \left(Z_t Z_t' - R_{t-1} \right) \tag{15}$$

3.3 Bounded Rationality and the Actual Law of Motion

Replacing RE with ADL means that the structural equations of the economy (1)-(2) need to be modified accordingly. For a related class of models Preston (2005, 2006) and Eusepi and Preston (2018) argue that under ADL aggregate expectations $\hat{\mathbb{E}}_t$ are an average of the expectations of heterogeneous households and firms who know only their own objectives, constraints and beliefs and cannot compute aggregate probability laws, i.e. cannot obtain model-consistent expectations like RE. Thus, agents act rationally when it comes to their own objective functions but unlike rational agents fail to anticipate the aggregate laws of motion and resort to econometric learning as in section 3.2. A representative agent can be obtained when a symmetric equilibrium is assumed in which although everyone's problem is identical, no individual is aware of that and as a result the representative agent cannot compute aggregate probability laws. This breaks the law of iterated expectations

[16] Note that here I assume constant gain learning instead of the decreasing gain learning used in Evans and Honkapohja (2001) The reason is that the former is more useful for tracking regime changes, while the latter is useful for studying asymptotic convergence properties of learning models to their RE counterparts. Given the current emphasis on the ZLB, tracking is a necessary feature of the model.

(LIE) for the operator $\hat{\mathbb{E}}$, and hence the recursion from which the aggregate demand (1) and Phillips curve (2) equations are derived. These two equations under ADL and $\hat{\mathbb{E}}$ then depend on a long horizon expectations reading:

$$x_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma} \left(i_T - \pi_{T+1} - r_T \right) \right] \tag{16}$$

$$\pi_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left[\kappa x_T + (1-\alpha)\beta \pi_{T+1} + u_T \right] \tag{17}$$

where $\hat{\mathbb{E}}_t$ again stands for the expectations of the adaptive learners and α is the Calvo probability of not being able to reset prices. I will refer to these two equations as the actual law of motion (ALM) of the economy.

Yet, Honkapohja et al. (2012) point out that assuming a continuum of symmetrical agents as is the case in the used NK model, one could still apply the LIE and resort to one period ahead Euler equation learning. I keep the infinite horizon learning for two reasons.

First, it allows for incorporation of FG as information about future policy rates into the agents' expectations and hence into the law of motion for the aggregate variables, contrary to the Euler equation learning approach. Thus, forward guidance can have real effects as described later in Section 3.5. Second, because agents do not know the structure of the economy, they cannot foresee how the ZLB will change the actual law of motion of the system. Consider equation (1) and apply $\hat{\mathbb{E}}$ instead of the RE operation as the Euler equation learning approach advocated by Honkapohja et al. (2012) would prescribe. Now, if the ZLB is expected to be binding for a few periods ahead, its effect should come through expectations of the output gap ($\hat{\mathbb{E}}_t x_{t+1}$) and inflation ($\hat{\mathbb{E}}_t \pi_{t+1}$). But these are only gradually updated (as described in Section 3.2), implying that although agents respect the ZLB in their forecasts for the interest rate, they are completely oblivious of its future effects on inflation and the output gap when only the one-period-ahead Euler equation (1) is used. In contrast, suppose that agents expect $t = T^{ZLB}$ as the last period of binding ZLB. Then in the long horizon approach (16) they could set $\hat{\mathbb{E}}_t i_{t+s} = i^*$ for all $t + s \leq T^{ZLB}$. Then, the expected duration of the ZLB has an effect on the realizations of the output gap both through the current and future binding periods, which in turn is reflected on future inflation as in (17). Thus, the economy driven by the modified learners with long-run expectations features minimal deviations from the rational expectations economy which are reflected only in the recursively updated $\hat{\mathbb{E}}_{t+s}x_{t+s+j}$ and $\hat{\mathbb{E}}_{t+s}\pi_{t+s+j}$ for $s \geq 1$, $j \geq 0$.

These arguments are related to the discussion of the use of long horizon expectations for anticipated structural changes in Evans et al. (2008).

3.4 Disagreement between the CB and the learning agents

The machinery built so far only describes the expectations formation of the agents and the effects of these expectations on the realizations of output gap and inflation through the ALM in (16) and (17). It notably omits the behavior of monetary policy and the expectations framework of the central bank. Below we will study how the interest rate forecast of a central bank with rational expectations differs from those of the learners and we will argue that this is consistent with the empirical evidence presented in section 2.

Suppose the economy exists for a long enough period with no extreme shocks that bring it to the ZLB. Then, following the forecast and updating procedures from section 3.2 the learning agents have converged to the RE solution of the model in (6). This implies that at some period $t - s,\ s > 1$ the perceived transition matrix has converged to the actual one - $\Phi_{t-s} = \Gamma$.[17] Therefore, the agents have fully learned the model with no binding ZLB and their forecasts (and errors) are indistinguishable of those of rational agents. The period of the Great Moderation is a useful analogy for this scenario.

Now, suppose the economy is hit by a demand shock ε_r at period t which brings the interest rate to the ZLB for at least 2 periods. Since the agents respect the ZLB in their expectations they know that today the interest rate will be at the ZLB - $\hat{\mathbb{E}}i_t = i^*$. Hence, from (15) the forecast error correction term for the interest rate's law of motion is zero and

[17]Note that due to the constant gain learning the PLM Φ_t converges to the rational expectations equilibrium transition matrix Γ in distribution rather than point-wise. See Evans and Honkapohja (2001) for a proof. For the sake of argument here I consider the Great Moderation example as having $\Phi_{t-s} = \Gamma$.

no updating occurs - $\phi_{i,t} = \phi_{i,t-1}$. On the other hand, their perceived LOMs for output gap and inflation ($\phi_{x,t-1}$ and $\phi_{\pi,t-1}$) are the first and second rows of the transition matrix for a world with no binding ZLB ($\Phi_{t-1} = \Gamma$). A model prescribed by Γ is characterized by an active monetary policy rule which accommodates demand shocks. This, however, is no longer true with a binding ZLB which locks the interest rate at an inefficiently high level i^*. Therefore the agents' forecasts for time t of output, inflation, and the interest rate will be based on $t-1$ beliefs of the Great Moderation and will be too optimistic. At the end of period t they will observe the realizations and update their expectations as in (15). Overall, during the expected period of the ZLB the agents will not update their perceived law of motion for the shadow rate because they expect i^* and see i^* at the ZLB, that is - they have no forecast error on the interest rate. Although they might be wrong about the depth of the shadow rate, the ZLB censors the data at i^* and eliminates any interest rate forecast errors. However, they will still make observable forecast errors for output gap and inflation and will update their beliefs for their laws of motion.

Contrary to the adaptively learning agents, let us assume the central bank knows the ALM of the model (16) and (17) and observes agents' expectations for all endogenous variables.[18] Upon observing agents' expectations the CB uses them in the ALM equations (16) and (17) and obtains model-consistent forecasts. Given its nowcasts for output gap and inflation, it uses the MP rule (5) and forms projections for the shadow rate.

Remember that during the ZLB the learners fail to adjust their perceived law of motion for the interest rate ($\phi_{i,t+s} = \phi_{i,t-1}$ if $\hat{\mathbb{E}}_{t+s-1} i_{t+s} = i^*,\ s \geq 0$) because their ZLB expectations in the immediate future are matched by the censored interest rate even if they are wrong about the shadow rate. The CB, on the other hand, uses the actual law of motion of the economy conditional on the learners' expectations, thus obtaining model-consistent forecasts about the shadow rate, the ZLB duration, and the interest rate path. The ZLB acts as an informational curtain for the learning agents and, as explained next, this will prove to be at the heart of the forecast disagreement between them and the rational CB.

[18]Considering the vast amounts of information collected and processed by central banks as well as their sophisticated forecast models this does seem like a realistic assumption.

Due to the inactive monetary policy accommodation, a binding ZLB means a given (negative) shock hits harder output gap and inflation compared to normal times. This naturally necessitates a proportionally lower interest rate trajectory than the same shock would have required above the ZLB. During the ZLB this translates to a deeper shadow rate and a mechanically prolonged ZLB period than what normal-times dynamics would imply. Due to the informational curtain of the ZLB, the agents do not observe the shadow rate and cannot update their shadow rate PLM. They still rely on their "Great Moderation" PLM just before the ZLB which forecasts a shallower interest (shadow) rate response and hence expect an earlier lift-off than the model-consistent forecasts of the CB. The assumption that the central bank knows the true model of the economy implies that the agents are also too optimistic relative to the actual realizations. As discussed in section 2, this is supported by the empirical evidence for policy rate forecasts in the US and Sweden, as well as, in financial markets as shown by Cieslak (2018).

The agents gradually update their output gap and inflation expectations, but the binding ZLB prevents them from understanding how the new ZLB regime changes the dynamics of the MP rule even in the absence of an explicit policy change. The only source of change in the system is the ZLB which affects the propagation of the state variables Z_t to the endogenous variables Y_t, i.e. the transition matrix Γ from (6) is no longer constant but can be approximated by a time-varying piece-wise linear function (see Guerrieri and Iacoviello, 2015).

Proposition 1. *Suppose the economy is brought to the zero lower bound ("Great Recession") after a period of convergence to a rational expectations model with no binding ZLB ("Great Moderation"). Then, the mechanics described above result in a disagreement between the agents and the central bank about the future path of the interest rate even in the absence of any policy change. Namely, the agents expect an earlier lift-off from the ZLB than the central bank.*

Proof: *see Appendix B.2*

Intuitively, as show in (13), agents' period t expectations ($\hat{\mathbb{E}}_t$) depend on estimations up to period $t-1$ (Φ_{t-1}).

Appendices B.1 and B.2 show that the actual law of motion of the economy implies:

$$\begin{aligned}
x_t &= \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma}\left(i_T - \pi_{T+1} - r_T\right) \right] \\
&= (1-\beta) \underbrace{\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} x_{T+1}}_{\mathtt{X_{x,t}}} - \frac{1}{\sigma} \underbrace{\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} i_T}_{\mathtt{X_{i,t}}} + \frac{1}{\sigma} \underbrace{\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \pi_{T+1}}_{\mathtt{X_{\pi,t}}} + \frac{r_t}{\sigma(1-\beta\rho_r)} \\
&= (1-\beta)\mathtt{X_{x,t}} - \frac{1}{\sigma}\mathtt{X_{i,t}} + \frac{1}{\sigma}\mathtt{X_{\pi,t}} + \frac{r_t}{\sigma(1-\beta\rho_r)} = \qquad (18) \\
&= f(\phi_x, \phi_\pi, \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} i_T, Z_t) = \begin{cases} f(\phi_x, \phi_\pi, \phi_i, Z_t), & \text{in normal times} \\ f^{ZLB}(\phi_x, \phi_\pi, \hat{\mathbb{E}}_t \mathtt{X_{i,t}}), & \text{if ZLB expected to bind} \end{cases}
\end{aligned}$$

where the final two cases derive from the agents' understanding of the ZLB. As per (13), they forecast the shadow rate and censor it at the ZLB i^*, thus making $\mathtt{X_{i,t}} \equiv \sum_{T=t}^{\infty} \beta^{T-t} i_T$ highly non-linear in Z_t (Appendix B.2). On the other hand, in normal times the function can be reduced to a linear function of beliefs and states: $f(\phi_x, \phi_\pi, \phi_i, Z_t)$ (see Appendix B.1.1). The agents' period t beliefs (Φ_{t-1}) are consistent with a linear "Great Moderation" economy and correspond to $f(\cdot)$. The proof in Appendix B.2 shows that due to the non-linearity of the ZLB:

$$\hat{\mathbb{E}}_t x_t = \phi_{x,t-1} Z_t = f(\phi_x, \phi_\pi, \phi_i, Z_t) > f^{ZLB}(\phi_x, \phi_\pi, \hat{\mathbb{E}}_t \mathtt{X_{i,t}}) = \gamma_{x,t} Z_t = x_t \qquad (19)$$

where $\gamma_{x,t}$ is the output gap row of one from the sequence of actual transition matrices Γ_t during a binding ZLB spell.[19] Intuitively, this occurs because for a given shock (say r_t) the dynamics of the economy with a binding ZLB imply higher interest rates (at i^* rather than negative as the shadow rate would prescribe), providing less policy accommodation to the shocks. Thus, under $\gamma_{x,t}$ shocks hit harder than what agents were used to in the "Great Moderation" ($\phi_{x,t-1} = \gamma_x$), leading agents to have too optimistic output gap expectations. The same argument applies to the inflation expectations.[20]

[19] As discussed before, Guerrieri and Iacoviello, 2015 show that during a binding ZLB there is no single transition matrix Γ as in the REE during normal times (see (6)), but a continuum of matrices, which can be approximated piece-wise linearly by a sequence of Γ_t's.

[20] Note that Appendix B.2.2 shows that the forecast errors on inflation are lower due to $\kappa, \alpha < 1$.

Given that the CB uses model-consistent expectations (i.e. actual x_t and π_t), and applies its unchanged policy rule (5), then due to the agents too optimistic $\hat{\mathbb{E}}_t x_t$ and $\hat{\mathbb{E}}_t \pi_t$, the CBs desired shadow rate is lower agents' expectation $s_t < \hat{\mathbb{E}}_t i_t = \phi_{i,t-1} Z_t$. This argument can be iterated forward to show that agents expect a higher trajectory for the policy rate and, as a result, a sooner lift-off from the ZLB.

In essence, the ZLB affects actual laws of motion for the endogenous variables through two channels. First, there is the traditional level channel which limits $i_{t+s} \geq i^*$ for the duration of the ZLB spell. Second, as described above, it temporarily changes the mapping of shocks to endogenous variables through a sequence of piece-wise linear matrices $\Gamma_t \neq \Gamma$ ($Y_t = \Gamma_t Z_t$). Unlike the traditional adaptive learning approaches, the learners here are equipped to cope with the first channel by keeping track of a mental shadow rate and censoring it. Yet, unlike rational agents, they do not internalize the second channel. They only adjust their beliefs Φ_t towards Γ_t gradually with observations (and forecast errors), but informational curtain of the ZLB prevents them from doing so for the perceived LOM of the interest rate - ϕ_i. This propagates their forecast disagreement about the lift-off date from the ZLB relative to the rational CB and gives room for policy communication in the form of forward guidance to correct it.

3.5 Forward guidance

I assume that in order to correct disagreement about the future path of interest rates, the CB uses forward guidance by truthfully revealing its expected remaining ZLB duration during every binding period.[21] Because the duration is measured from the current period, it also identifies the expected lift-off date. This is the least informationally dense message the CB can send. Naturally, more sophisticated communication schemes might help the agents update their beliefs even closer to the truth, yet Section 4 shows that even

[21] This assumption abstracts from potential strategic behavior and credibility issues of forward guidance communication. Bassetto (2019) models forward guidance as cheap talk and shows how its informativeness depends on the central bank's private information and pre-existing credibility. Bodenstein et al. (2012) and Haberis et al. (2019) study imperfectly credible policy promises at the ZLB, while Boneva et al. (2018) and Angeletos and Sastry (2021) consider the design of threshold-, calendar-, instrument-, and target-based guidance. See also Campbell et al. (2019) on the quantitative limits imposed by imperfect communication.

such a parsimonious message is effective in guiding the agents and greatly improves macroeconomic outcomes. This communication is inherently conditional on the current state of the economy and does not represent a commitment by the CB, making it a model of Delphic, rather than Odyssean forward guidance.

To interpret this guidance, agents must first estimate their own remaining ZLB duration. In each binding period t, they form long-run expectations $\hat{\mathbb{E}}_t \{x_j, \pi_j, i_j, s_j\}_{j=t}^{\infty}$ as outlined in Section 3.2. Let $T^{ag} \geq 1$ denote the number of quarters for which agents expect the ZLB to remain binding, including the current quarter:

$$T^{ag} \text{ such that } \begin{cases} \hat{\mathbb{E}}_t i_{t+j} = i^*, & j = 0, \dots, T^{ag} - 1, \\ \hat{\mathbb{E}}_t i_{t+T^{ag}} > i^*. \end{cases} \tag{20}$$

Thus, $t + T^{ag}$ is the agents' expected lift-off date.

The central bank is assumed to have model-consistent expectations ($\mathbb{E}_t$), conditioning on the agents' current PLM and their long-run expectations ($\hat{\mathbb{E}}_t Y_{t+j}$). However, it has no choice variable and truthfully reveals its expectations, thus abstracting from strategic behavior. It observes agents' expectations ($\hat{\mathbb{E}}_t Y_{t+j},\ j \geq 0$ and T^{ag}) and uses them to form model-consistent expectations of output gap and inflation according to the model's ALM in (16)-(17):

$$\mathbb{E}_t x_{t+j} = \hat{\mathbb{E}}_t \sum_{T=t+j}^{\infty} \beta^{T-t+j} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma} (i_T - \pi_{T+1} - r_T) \right] \tag{21}$$

$$\mathbb{E}_t \pi_{t+j} = \hat{\mathbb{E}}_t \sum_{T=t+j}^{\infty} (\alpha\beta)^{T-t+j} \left[\kappa x_T + (1-\alpha)\beta \pi_{T+1} + u_T \right] \tag{22}$$

$$\mathbb{E}_t \hat{\mathbb{E}}_t Y_{t+j} = \hat{\mathbb{E}}_t Y_{t+j} \tag{23}$$

where the agents' long-run beliefs evolve as in (13). Then, the CB forecasts its own policy rate i_{t+j} according to the policy rule (5):

$$\mathbb{E}_t s_{t+j} = \mathbb{E}_t \rho_i s_{t+j-1} + (1-\rho_i)(\chi_\pi \mathbb{E}_t \pi_{t+j} + \mathbb{E}_t \chi_x x_{t+j}) \tag{24}$$

$$\mathbb{E}_t i_{t+j} = \max \{ i^*, \mathbb{E}_t s_{t+j} \} \tag{25}$$

In a similar fashion to the agents in (20), the CB obtains its expected remaining ZLB duration, including the current quarter:

$$T^{CB} \text{ such that } \begin{cases} \mathbb{E}_t i_{t+j} = i^*, \quad j = 0, \dots, T^{CB} - 1, \\ \mathbb{E}_t i_{t+T^{CB}} > i^*. \end{cases} \tag{26}$$

Hence, $t + T^{CB}$ is the CB's expected lift-off date.

As per Proposition 1, we would have $T^{CB} > T^{ag}$, because agents' expectations adjust to reflect the new regime brought by the ZLB only gradually through observations by updating the perceived transition matrix Φ_t. Whenever lift-off forecasts disagree, there is scope for forward guidance. I consider four regimes - a non-communication baseline and three ways in which agents might interpret lift-off forward guidance messages:

1. **Baseline *"No FG"* - the central bank does not release any communication.**

   The agents expect lift-off at $t+T^{ag}$ and, even if the central bank expects $T^{CB} > T^{ag}$, it does not communicate that. If the ZLB is still binding at $t + T^{ag}$, agents are surprised by observing $i_{t+T^{ag}} = i^*$ when they had forecast $\hat{\mathbb{E}}_t i_{t+T^{ag}} > i^*$. This forecast error leads them to update their beliefs in the usual manner described in Section 3.2, treating observations of $i_{T^{ag}+j} = i^*, \ j > 0$ as regular interest rate data points.

2. ***"Period FG"* - agents only adopt the communication length of the ZLB spell.**

   At every binding period t, the CB releases its expected remaining ZLB duration T^{CB}. If different from T^{ag}, the agents adopt it outright in their expectations without adjusting their interest rate PLM - $\phi_{i,t-1}$:

$$\hat{\mathbb{E}}_t i_j = i^*, \qquad j \in [t, t + T^{CB} - 1] \tag{27}$$

$$\hat{\mathbb{E}}_t s_j = \phi_{i,t-1} \hat{\mathbb{E}}_t S_j \tag{28}$$

$$\hat{\mathbb{E}}_t i_j = \max\left\{ i^*, \ \hat{\mathbb{E}}_t s_j \right\}, \qquad j \geq t + T^{CB} \tag{29}$$

This expectational mask gets reflected in the aggregate demand equation (16).[22] Note that in this case the law of motion for the interest rate is not updated, so even at the expected lift-off date $t + T^{CB}$ there might still be interest rate nowcast disagreement between the CB and the agents.

3. ***"Learning FG"* - agents adjust their interest rate PLM in response to the CB's lift-off communication**

   The CB announces T^{CB} as in the previous case of *period FG*, along with its first expected interest rate above the ZLB, $\mathbb{E}_t i_{t+T^{CB}}$. Instead of adopting the communicated ZLB duration, agents now treat $\mathbb{E}_t i_{t+T^{CB}}$ as an additional data point for learning. In words, upon receiving a FG message, agents perform a mid-period update on their interest rate PLM ($\phi_{i,t-1}$) by comparing the announced $\mathbb{E}_t i_{t+T^{CB}}$ against their own forecast $\hat{\mathbb{E}}_t i_{t+T^{CB}}$. If there is a forecast error relative to the CB's announcement, they perform a mid-period learning step (as in (15)) solely on their interest rate PLM. That is, they reason how their interest PLM must adjust to fit the CB's announcement. I allow the learning gain λ to differ from the regular gain τ; a lower gain may be interpreted as agents placing less credibility on the message.[23] Namely:

$$\tilde{\phi}_{i,t-1} = \phi_{i,t-1} + \lambda \hat{\mathbb{E}}_t S'_{t+T^{CB}} R^{-1}_{t-1} \left(\mathbb{E}_t i_{t+T^{CB}} - \hat{\mathbb{E}}_t i_{t+T^{CB}} \right) \tag{30}$$

$$\Phi_{t-1} = \left[\phi_{x,t-1}; \phi_{\pi,t-1}; \tilde{\phi}_{i,t-1} \right] \tag{31}$$

4. ***"Period + Learning FG"* - agents both apply the communicated ZLB length and also update their interest rate PLM**

[22] This interpretation is akin to the framework of forward guidance as anticipated shocks by Del Negro et al. (2012).

[23] The baseline calibration sets $\lambda = \tau$, giving an FG announcement the same weight as one quarter of data. Appendix C.4 relaxes this. Learning is performed at the expected lift-off date $t + T^{CB}$ because the duration announcement only reveals that the policy rate is constrained through $t + T^{CB} - 1$; it does not reveal the CB's expected shadow rate in that final binding quarter. Learning instead from i^* at $t + T^{CB} - 1$ would therefore be inaccurate because $\mathbb{E}_t s_{t+T^{CB}-1}$ need not equal i^*. If the CB also communicated that shadow rate, learning could instead take place in the final binding quarter. The update uses the last known variance-covariance matrix R^{-1}_{t-1} (because the FG message arrives before the realizations of the current endogenous variables) and the regressors, including the expected shadow rate, iterated forward to $S_{t+T^{CB}}$.

This interpretation scenario is simply a combination of 2. and 3. It is expected that the largest stability gains will come from eliminating the lift-off disagreement as in *period FG*, yet upon lift-off the updated PLM might be closer to the truth and imply lower forecast errors.

Note that in all FG interpretations (*Period FG*, *Learning FG*, and their combination), forward guidance has real effects on the current output gap (16) and inflation (17) by virtue of affecting the agents' interest rate forecasts - either through the FG mask in *Period FG*, or the updated interest rate PLM (Φ_{t-1}) in *Learning FG*, or both in *Period + Learning FG*. Naturally, the altered output gap and inflation lead the CB to change its policy stance through the policy rule, which next period could lead to a new FG communication. This endogenous interplay is active at every ZLB period and is at the essence of information revealing, Delphic communication, resembling the Riksbank regular internal projections. To the best of my knowledge, this is a novel approach to modeling Delphic FG, that is the use of an endogenous central bank lift-off projection as a communication signal that alters the agents' interest rate expectations, and has real effects.[24]

4 Quantitative experiments: disagreement, learning instability, and communication

4.1 Experimental design and measurement

This section presents the numerical analysis of the learning economy with and without forward guidance. For each communication regime, I simulate 10,000 Monte Carlo (MC) matched paths over 500 periods.[25] All regimes share the matched MC independent draws of the natural rate (r_t) and cost-push (u_t) shocks. To mimic the initial state as coming

[24]Existing adaptive learning analyses generally represent forward guidance as anticipated monetary policy shocks as a commitment to a deviation from the regular policy rule. In Cole (2020a) and Cole (2020b), the announced shocks enter agents' forecasting models and agents recursively learn their macroeconomic effects. These papers therefore already generate contemporaneous effects of forward guidance under adaptive learning, but through an announced policy deviation and commitment (i.e. Odyssean FG) rather than an endogenous, rule-consistent lift-off projection communication (i.e. Delphic FG).

[25]The simulation design is required here because the recursive learning algorithm needs variability in order to reliably estimate the transition matrix Φ.

from a "Great Moderation" past, all simulations are initiated at the rational expectations equilibrium of the linear model with no binding ZLB, i.e. $x_1 = \pi_1 = i_1 = 0,\ \Phi_1 = \Gamma,\ R_1 = \Omega$ where Ω is the variance-covariance matrix of the state vector Z. Then, a common GFC-type negative demand shock ε_2^r pushes the economy to the ZLB for an extended period of time.[26] The calibration follows Table B.3 and is explained in detail in Appendix B.3.

Section 4.2 analyzes the emergent interest rate forecast disagreement between the agents and the central bank as a result of the informational curtain on shadow rates imposed by the binding ZLB. Section 4.3 describes how this disagreement gives rise to learning instability, unanchored expectations, and explosive dynamics. Section 4.4 shows that Delphic forward guidance in the form of lift-off date communication sharply reduces disagreement, instability, and stabilization loss.

Although the model is presented with an intercept for generality, the quantitative experiments impose a zero intercept and estimate the remaining 3×3 coefficient matrix - Φ_t. This restriction removes an additional potential source of forecast drift for long-run expectations and is therefore conservative for the mechanism of the ZLB as an informational curtain.

4.2 No forward guidance - disagreement at the ZLB

In the absence of communication, the ZLB censors shadow rate forecast errors whenever agents expect and observe a binding policy rate. Their interest rate PLM therefore does not adjust to the depth of the latent shadow rate, leading to overly optimistic shadow rate forecasts and premature lift-off forecasts. Across binding ZLB quarters, the agents' median shadow rate nowcast is 7.88 annualized percentage points higher than the central bank's, while their median lift-off forecast is seven quarters early. In contrast, above the ZLB the median shadow rate gap between them is approximately zero.

Figure 4 below shows Monte Carlo fan charts for x_t, π_t, and the shadow rate. These can be interpreted as responses to the large negative natural rate shock in period 2, which produces a prolonged ZLB spell. The thicker solid blue line shows the median

[26] The shock is common across all communication regimes, so their differences are identified from communication rather than from different shock realizations.

end-of-period realizations of output and inflation and the central bank's realized shadow rate, while the thinner solid orange line shows the agents' median beginning-of-period nowcasts. Dashed lines show the corresponding 10 - 90 percentile bounds. Agents are too optimistic about the responses of all three variables. Constant-gain learning quickly narrows the output-gap and inflation nowcast errors. By contrast, while both the expected and observed policy rates remain at the floor, the shadow rate forecast error is censored and cannot correct the interest rate PLM.

Figure 4: No FG: Monte Carlo responses to the common GFC shock

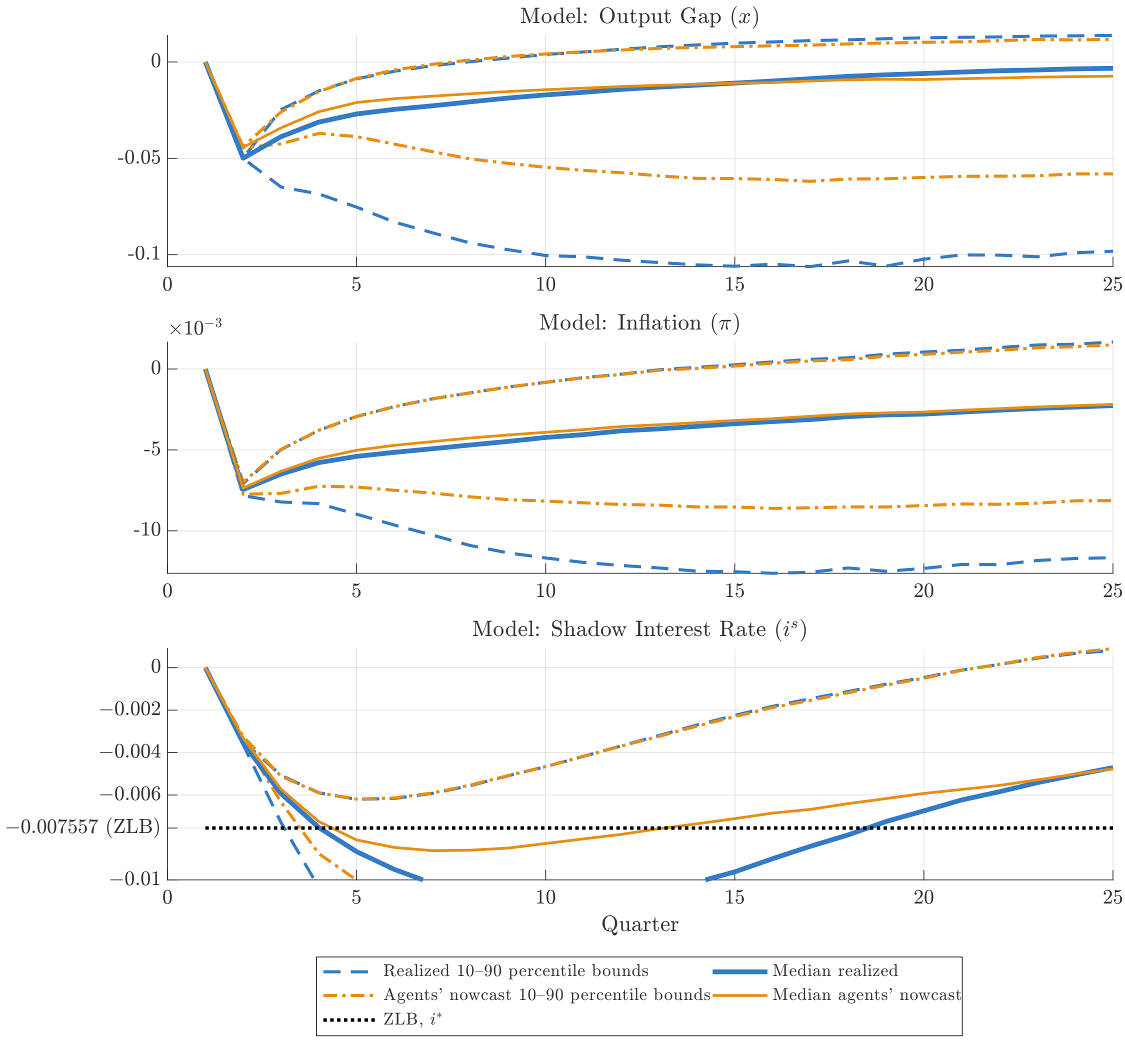

Notes. Interest panel: agents' pre-communication PLM nowcast versus the CB realized shadow rate. Neither series is truncated at the policy-rate floor.

It is particularly striking that the forecast errors are concentrated on the downside: the upper percentile bounds for realizations and nowcasts largely overlap, whereas the lower realization bounds fall well below the corresponding nowcasts. This asymmetry is the result of the informational curtain created by a binding ZLB. Above the ZLB, all variables are observable and agents can correct misspecified beliefs through forecast errors. Moreover, in the upper tail the policy rate is unconstrained, so the initial "Great Moderation" beliefs coincide with the rational-expectations beliefs of the model, $\Phi_0 = \Gamma$. Once the ZLB binds, agents respect the floor in their interest rate forecasts but do not internalize the full change in the aggregate law of motion described in Section 3.4. The lower output-gap and inflation bands gradually converge through adaptive learning, but the informational curtain over the shadow rate leaves interest rate expectations systematically too high.

Figure 5 plots annualized shadow rate disagreement between the agents and the CB, $400(\hat{\mathbb{E}}_t s_t - s_t)$, in normal times (right panel) and at the ZLB (left panel). At the ZLB, agents expect a higher shadow rate in 99.3% of observations, with a median gap of 7.88 annualized percentage points. This is consistent with Proposition 1.[27] Above the ZLB, the median gap is approximately zero and observations are split nearly evenly by sign. Thus, there is no systematic median interest rate forecast disagreement above the ZLB, consistent with the Riksbank and Federal Reserve evidence in Section 2.

The agents' optimism also implies a higher expected policy rate trajectory at the ZLB and, most often, earlier lift-off. Figure 6 (Panel A) plots the agents' and the CB's remaining ZLB durations, T^{ag} and T^{CB}, during the initial episode. Both durations include the current binding quarter, and lift-off is expected at $t + T^{ag}$ and $t + T^{CB}$, respectively. In this episode, agents generally expect three to six fewer binding quarters than the CB.[28]

Because both durations are measured from the same period t, lift-off disagreement is $T^{ag} - T^{CB}$. Figure 6 (Panel B) presents its distribution across all binding path-quarters. Consistent with Panel A, the median agent expects lift-off seven quarters earlier than the CB. Agents expect earlier lift-off in 88.9% of observations, the same date in 2.3%, and

[27]Restricting each path to observations through its own first projection event preserves the pattern: the median ZLB gap is 4.35 annualized percentage points and is positive in 99.3% of observations. The narrower tails indicate that the most extreme full-sample values arise after a path first encounters the projection boundary.

[28]Initially there is no lift-off forecast because the interest rate remains above the ZLB owing to substantial policy inertia. As discussed in Appendix B.3, the one-year delay between the shock and the first ZLB observation is broadly similar to the US experience between the onset of the GFC in 2007Q4 and the first ZLB observation in 2008Q4.

Figure 5: No FG: shadow rate nowcast disagreement

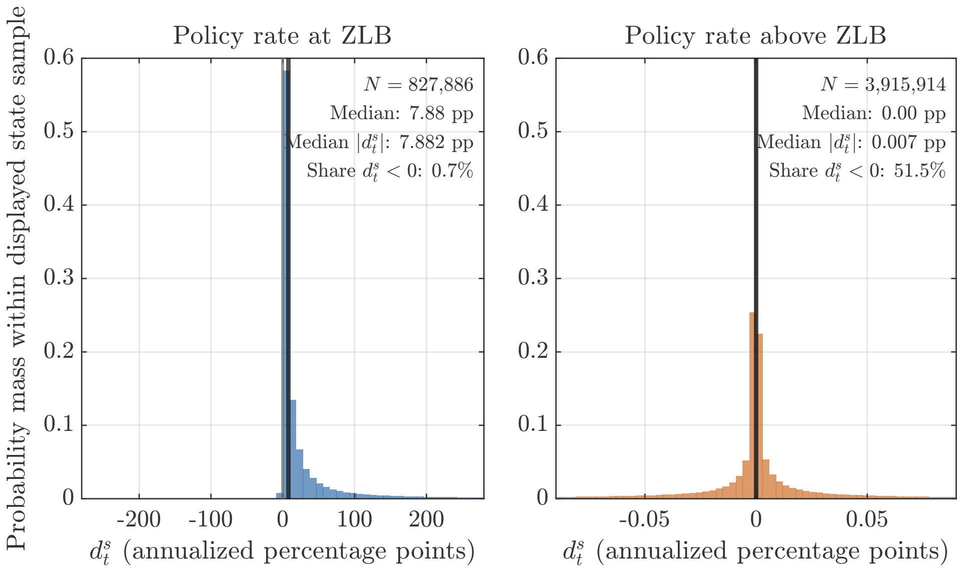

Notes. $d_t^s = 400(s_t^{ag} - s_t^{CB})$ is the agents' internal pre-communication PLM nowcast minus the CB shadow-rate realization, in annualized percentage points.

later lift-off in 8.8%. Figure C.1 further plots T^{CB} conditional on $T^{ag} = 1$: agents then expect the current quarter to be the last at the ZLB and lift-off at $t+1$. The CB expects a longer duration in 93.7% of those observations, implying a forecast error at the agents' anticipated lift-off if the ZLB remains binding.

This state-dependent disagreement offers clear implications for policy communication above and below the ZLB. Above the ZLB, in normal times, there is no systematic median interest rate disagreement between the bank and the agents. There, agents observe all necessary information and the model suggests no systematic role for guidance to correct an informational problem. Once at the ZLB, however, the observed policy rate is censored at the floor and no longer reveals the trajectory of the shadow rate. The resulting disagreement about the future policy rate path creates room for the CB to communicate that its current forecast implies a later lift-off than agents expect. Thus, the ZLB's informational curtain mechanism creates a case for Delphic forward guidance communication at the ZLB.[29]

[29]These implications accord closely with Warsh (2026), who describes forward guidance during the Global Financial Crisis as "essential at the time," while arguing that its role in normal times should be "limited and circumscribed."

Figure 6: No FG: Remaining ZLB duration and lift-off disagreement

Panel A. Remaining ZLB duration

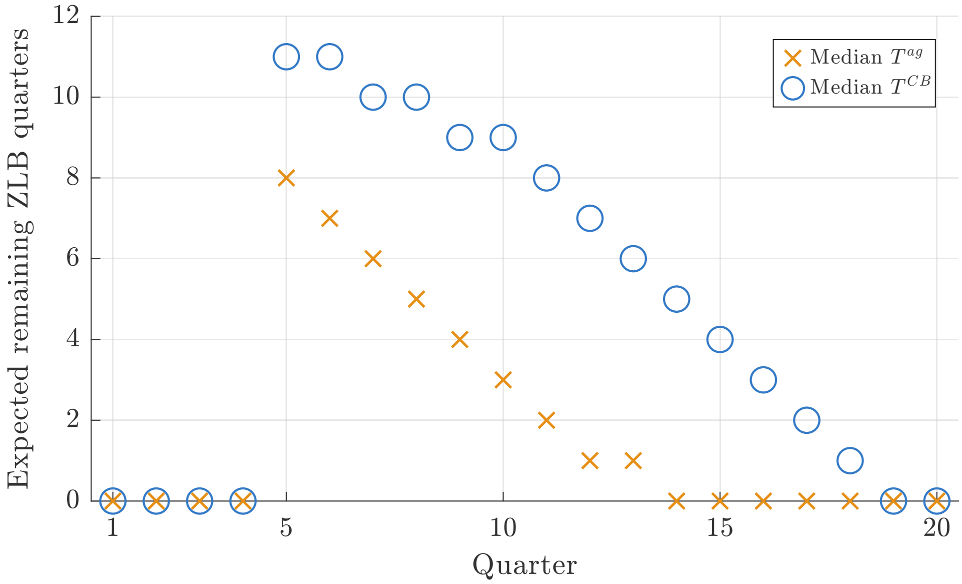

Panel B. Lift-off disagreement

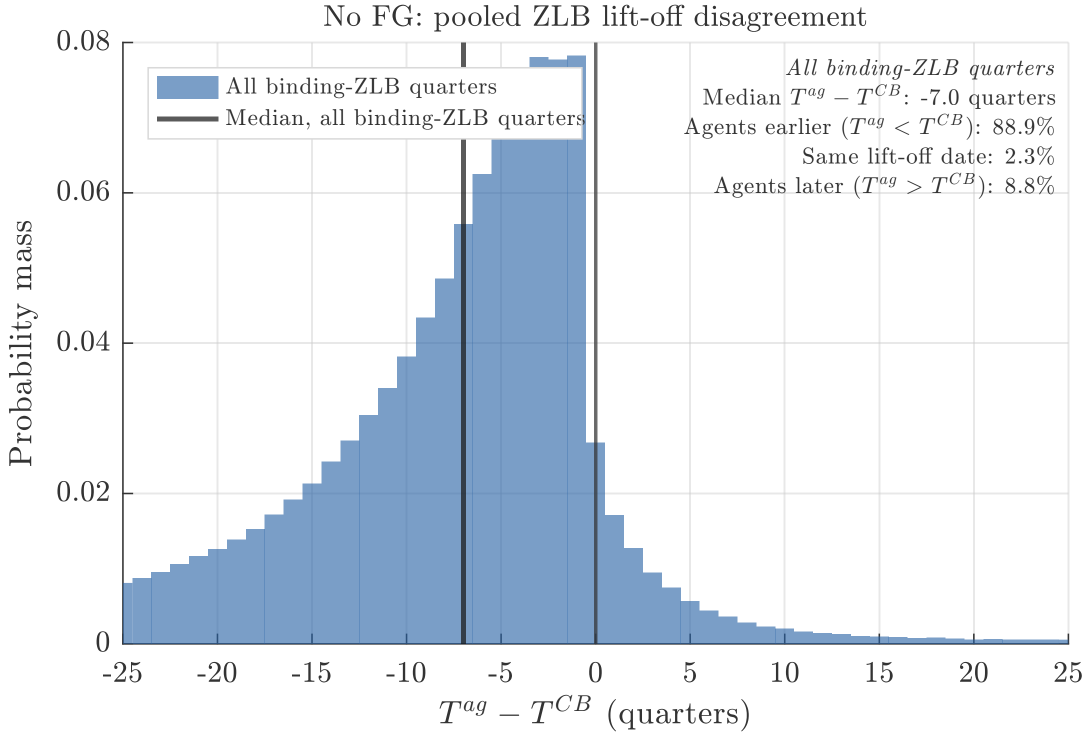

Notes: T^{ag} and T^{CB} denote the expected numbers of remaining binding quarters, including the current quarter. Expected lift-off occurs at $t + T^{ag}$ and $t + T^{CB}$, so lift-off disagreement is $T^{ag} - T^{CB}$.

4.3 From disagreement to instability

Agents' positive shadow rate disagreement at the ZLB becomes an observed policy rate forecast error once they prematurely forecast lift-off. The resulting update to the interest rate PLM changes the perceived persistence of the shadow rate, $\phi_{ii,t}$. Agents' forecasts for the shadow state vector satisfy $\hat{\mathbb{E}}_t S_{t+j} = \Psi_{t-1}^{j} S_t$. As discussed in Section 3.2, because the perceived state-transition matrix is triangular,

$$\text{eig}(\Psi_{t-1}) = \{1, \rho_u, \rho_r, \phi_{ii,t-1}\}\,. \tag{32}$$

The first eigenvalue is associated mechanically with the constant in the state vector, while ρ_u and ρ_r are known and stationary. Consequently, $\phi_{ii,t-1}$ is the only estimated eigenvalue that can destabilize agents' dynamic forecasts.

As $|\phi_{ii,t}|$ approaches unity, the expected shadow rate mean-reverts increasingly slowly. At unity, mean reversion is lost, and beyond unity, iterated future forecasts for the shadow rate, and hence the interest rate, become unanchored and explosive and feed into output gap and inflation through the ALM (16) and (17). Thus, expectations-driven instability is preceded by a region of high perceived interest rate persistence in which relatively small forecast errors can have large effects on the expected path of monetary policy. To deal numerically with such events, the experiments employ a projection facility (PF) bound of $|\phi_{ii,t}| \leq \bar{\phi}_{ii} = 0.98$.[30] The PF bound identifies attempted learning updates that would move the perceived law of motion beyond this stability region. Whenever $|\widetilde{\phi}_{ii,t}| > \bar{\phi}_{ii}$, the entire update step is rejected and beliefs remain at the last admissible configuration, $\Phi_t = \Phi_{t-1}$ and $R_t = R_{t-1}$.[31]

Disagreement about lift-off from the ZLB tends to generate precisely such updates. Let $\widetilde{\phi}_{ii,t}$ denote the candidate coefficient produced by RLS before application of the

[30] Raising the bound to 0.99 reduces PF incidence from 45.28% to 43.17% but slightly increases macroeconomic failures; Section C.3 reports the matched-path comparison.

[31] See Evans and Honkapohja (1998), Marcet and Sargent (1989), and Slobodyan and Wouters (2012) for further details and applications of projection facilities in the adaptive learning literature.

projection facility:

$$\widetilde{\phi}_{ii,t} = \phi_{ii,t-1} + \tau \left[R_t^{-1} Z_t\right]_i \left(i_t - \hat{\mathbb{E}}_t i_t\right). \tag{33}$$

When agents forecast that the economy has already lifted off while the ZLB continues to bind, the interest rate forecast error $(i_t - \hat{\mathbb{E}}_t i_t)$ is negative.[32] During these episodes, the effective gain associated with the lagged interest rate state $(\left[R_t^{-1} Z_t\right]_i)$ is typically also negative. The resulting update to $\phi_{ii,t}$ is therefore positive. Repeated premature lift-off forecasts progressively raise the perceived persistence of the shadow rate and bring agents' forecasting model closer to the stability boundary.

Intuitively, at their expected lift-off date $t + T^{ag}$, agents observe a lower policy rate than expected despite observing $u_{t+T^{ag}}$ and $r_{t+T^{ag}}$. They respond by increasing perceived policy persistence, ϕ_{ii}, and reducing the perceived responses to shocks, ϕ_{iu} and ϕ_{ir}.[33]

This pattern is a quantitative result rather than an analytical restriction, yet it occurs in the vast majority of belief updates when agents expect an earlier lift-off from the ZLB. Table 2 (Panel A) shows that above the ZLB, 53.6% of implemented updates are positive. At the ZLB, 78.6% of updates are zero because agents often also forecast the policy rate at the floor. Conditional on a non-zero at-ZLB update, however, 94.7% are positive. Despite being wrong on the shadow rate (as shown in the previous section), their ability to truncate the interest at the lower bound means at those times they have no forecast error, eliminating any mechanical expectational drifts from learning a fixed variable at the ZLB.

The No FG experiment produces at least one projection facility (PF) event in 4,528 of the 10,000 simulated paths, an unconditional incidence of 45.28%.[34] Of these first events, 1,558 occur above the ZLB and 2,970 occur while the ZLB is binding, with 1,402 during

[32]When agents also expect the ZLB to bind, the realized and forecast policy rates coincide at the floor, so the interest rate PLM is not updated.

[33]Figure C.2 in the Appendix shows the evolution of the components of the agents' PLM, Φ_t. In the No FG regime, the bottom row shows that learning after the GFC-induced ZLB spell lowers $\phi_{ir,t}$ and raises $\phi_{ii,t}$. At longer horizons, updates become more dispersed around the REE initialization (Figure C.3).

[34]Figure C.4 shows the number of PF hits per PF path. Among these paths, 87.2% have more than one event and the median is 12 hits. All 4,528 first projection events are genuine attempted crossings of the interest rate persistence boundary - $\bar{\phi}_{ii}$, and none is classified as a numerical projection trigger. The projection facility is therefore detecting economically generated learning dynamics rather than a failure of the simulation algorithm.

Table 2: No FG: interest rate perceived persistence updates and first projection events

Panel A. All paths until the first PF event		
	At ZLB	Above ZLB
Observations	522,873	2,962,820
Mean, implemented	0.0052	-0.0005
Median, implemented	0	0
P10–P90, implemented	[0,0.0194]	[-0.0014,0.0004]
P1–P99, candidate total	[-0.0002,0.0809]	[-0.0144,0.0069]
Maximum $\|\Delta\|$, candidate total	0.823	36.732
$-$ / 0 / + (%), implemented	1.1/78.6/20.3	46.4/0.1/53.6
Panel B. At the first PF event		
	PF at ZLB	PF outside ZLB
Number of paths	2,970	1,558
Share of all PF paths	65.6%	34.4%
Premature liftoff forecast at the event	100%	0%
Median quarters with premature lift-off forecasts	12	—
Median quarters premature-liftoff streak at PF event	6	—
Median accepted $\Phi_{i,i}$ before update	0.949	0.961
Median attempted change in $\Phi_{i,i}$	0.0639	0.049
Median interest-rate forecast error (annualized pp)	-1.10	0.58

Notes: Panel A: The sample runs from quarter 2 through the path's own first projection-facility event, including the event quarter; never-PF paths remain uncensored. Panel B: The reported interest rate forecast error is realized minus agents' forecast, so a negative value means that agents overpredict the current policy rate.

the initial GFC-induced spell and 1,568 during a later spell. The following event study aligns paths on their first projection event, so repeated events do not receive additional weight.

Table 2 (Panel B) summarizes first PF events at and above the ZLB. Among at-ZLB events, the median path has accumulated 12 quarters with a premature lift-off forecast, including an uninterrupted terminal streak of 6 quarters. The persistence drift is therefore associated with repeated forecast errors generated while the shadow rate remains hidden by the ZLB, rather than merely with the presence of the lower bound itself. The median attempted persistence increase is 0.0639 at an at-ZLB event and 0.049 above the ZLB.[35] At-ZLB events exhibit the familiar overoptimistic pattern, with a median annualized interest rate forecast error of -1.10 percentage points. Above-ZLB events have an opposite-signed median error of 0.58 percentage points.

Figure 7 (Panel A) shows the evolution of the accepted pre-update persistence coefficient around the first PF event. For at-ZLB events, the median increase is concentrated in the final several quarters before the event - characteristic of a sudden uncovered lift-off disagreement; outside-ZLB events exhibit a longer buildup. Panel B reports the share of paths on which agents forecast a policy rate above the ZLB while the realized rate remains constrained. Among paths whose first projection event occurs at the ZLB, the prevalence of premature lift-off forecasts increases markedly as the event approaches and reaches 100% in the event quarter. Figure 8 plots the corresponding interest rate forecast errors. At-ZLB events are driven by premature lift-off expectations, whereas positive errors above the ZLB emerge mainly near the PF event. This supports interpreting the latter as a broader post-ZLB learning instability: the full perceived law of motion has already drifted away from its rational-expectations initialization, even though the monitored interest rate root may remain some distance below the projection boundary until the triggering quarter.

Although the two event types differ, they remain linked to earlier ZLB experience. Appendix C.2 compares the No FG baseline with matched simulations that omit the

[35] Figure C.5 visualizes the cumulative distributions of the accepted pre-event ϕ_{ii} and the post-update candidate.

imposed GFC shock. The GFC raises PF incidence both at and above the ZLB and brings the median first event forward from quarter 270 to quarter 100. This persistence indicates that above-ZLB PF events can reflect belief drift accumulated during earlier ZLB spells.

Projection events are also informative about the subsequent stability of the economy. Table 3 shows that of the 4,528 paths with a PF event, 1,165, or 25.7%, subsequently reach the ad-hoc macroeconomic feasibility bound where the absolute value of the output gap, inflation, or the interest rate reaches 1,000. None of the 5,472 paths without a PF event reaches this bound. Thus, all 1,165 macroeconomic failures occur after a projection event. Among paths experiencing both events, the macroeconomic failure occurs a median of 95 quarters after the first PF event.[36]

Table 3: Macroeconomic volatility through the first projection event

	Target	Stable	Any PF	Macro fail.	PF/no fail
Paths	—	5,472	4,528	1,165	3,363
Output gap	1.00–1.50	1.41	5.57	12.02	4.76
Inflation	0.76–1.24	1.15	2.15	4.38	1.89
Policy rate	—	1.84	1.38	1.04	1.41

Notes: Entries are medians across paths. Output-gap units are percentage points; inflation and policy-rate units are annualized percentage points. For a path with a projection event, its sample standard deviation uses valid observations through and including the first-PF quarter and excludes all later quarters. The event-quarter macro variables realization precedes the rejected learning update.

Table 3 further asks whether PF paths already exhibit greater macroeconomic volatility before the projection intervention. Each PF path is retained through and including its first-event quarter but is censored thereafter. PF paths that do not subsequently fail numerically are already macroeconomically volatile at the time of the first PF event relative to stable paths: their median output-gap volatility is 4.76 rather than 1.41 percentage points, and inflation volatility is 1.89 rather than 1.15 annualized percentage points. A projection event without a later macroeconomic failure should not therefore be treated

[36]These associations are descriptive rather than causal: adverse shock and belief histories can generate both a PF event and a later macroeconomic failure. Moreover, rejecting the candidate update at the first PF changes the subsequent path, so the simulations do not reveal the counterfactual trajectory that would have followed without the intervention. Nevertheless, the concentration of later failures among PF paths shows that the facility identifies economically fragile trajectories rather than innocuous numerical outliers.

Figure 7: Perceived persistence and premature lift-off around the first projection event

Panel A: Perceived interest-rate persistence around the first projection event

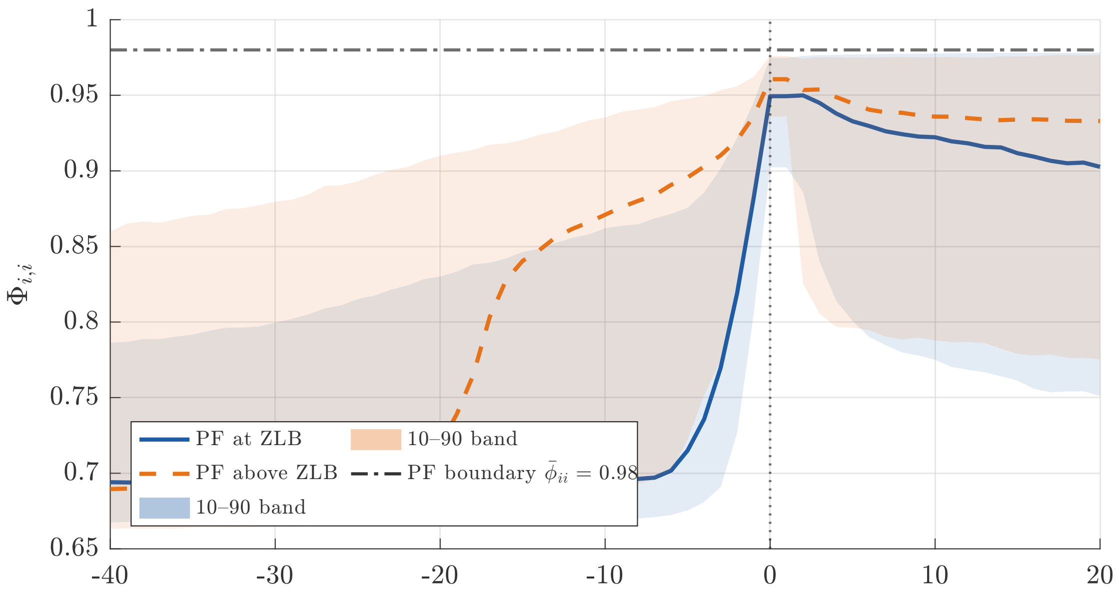

Panel B: Share forecasting premature lift-off around the first projection event

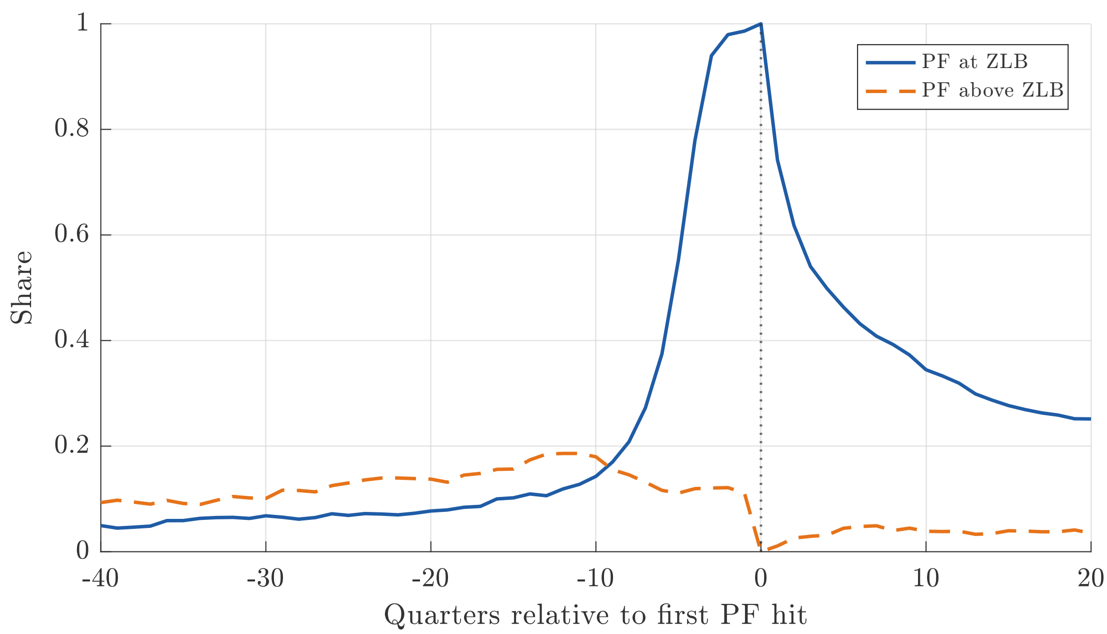

Notes: Event time zero is the quarter of each path's first projection event. Panel A: The rejected candidate update is not plotted. Panel B: The indicator is mechanically zero in the event quarter for paths whose first projection event occurs outside the ZLB.

Figure 8: Agents' interest rate nowcast error before the first projection event

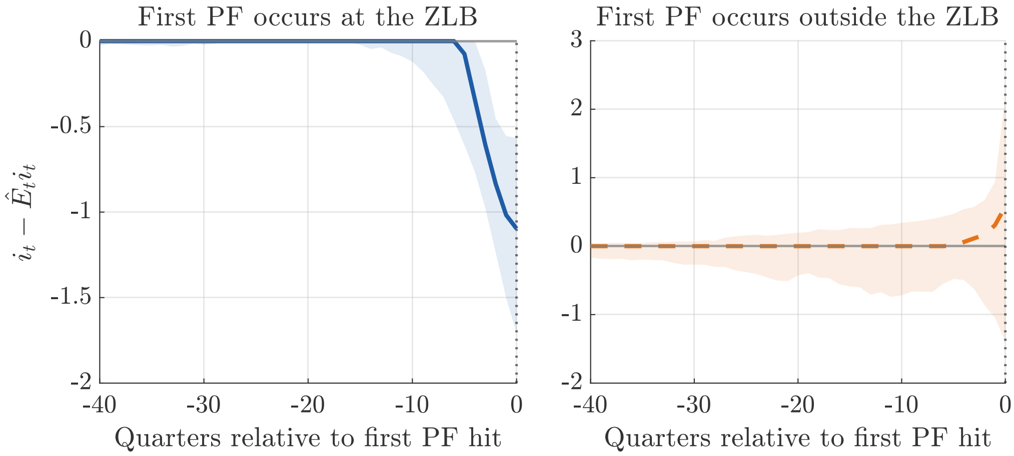

Notes. The annualized interest rate nowcast error, defined as $400(i_t - \hat{E}_t i_t)$. Negative values mean the agents were expecting a higher interest rate than the realization.

automatically as a false positive. The projection intervention removes precisely the candidate update that would have moved the perceived interest rate root beyond the stable region.[37]

In summary, repeated premature lift-off forecasts at the ZLB raise perceived interest rate persistence and can produce PF events. Crossings above the ZLB arise after broader PLM drift and are triggered by opposite-signed forecast errors. The at-ZLB events provide the most direct evidence for the informational-curtain mechanism. PF paths also exhibit substantially higher output-gap and inflation volatility before the projection intervention, consistent with economically consequential fragility rather than innocuous numerical outliers.

4.4 Forward guidance as a stabilizing force for expectations

The preceding sections established a sequence starting from disagreement about lift-off to drift in the perceived persistence of the shadow rate and, eventually, to attempted crossings of the projection boundary. This section asks whether communicating the central bank's expected remaining ZLB duration - and hence its expected lift-off date - is enough to ameliorate the informational curtain and reduce the risk of expectational instability.

[37]Note that interest rate volatility is less discriminating because the lower bound compresses observed variation in the policy rate.

The three FG experiments are again as follows. In *Period FG*, agents impose the communicated duration T^{CB} on their interest rate forecast, but do not update their perceived law of motion in response to the announcement. In *Learning FG*, the announcement instead generates an inter-period PLM update in equation (30), with a learning gain on the message λ (in the baseline $\lambda = \tau = 0.03$), but agents do not impose the communicated ZLB duration mask on their forecasts. *Period + Learning FG* combines the two operations, while *No FG* applies neither.[38]

Figure 9's upper panel shows the evolution of perceived shadow rate persistence, $\phi_{ii,t-1}$, across communication regimes following the common GFC innovation. All four regimes share the initial 4-5 quarters between the crisis shock and hitting the ZLB and have largely similar downward trends in ϕ_{ii}. What is more, Figure C.2 shows that in those periods the agents increase their perceived sensitivity of all three endogenous variables (x, π, i) to the natural rate (r). This is due to the combination of long-run expectations in the ALM and the ability of agents to expect a future binding ZLB as a result of the GFC negative natural rate shock, and accordingly truncate their expectations. An anticipated binding ZLB implies future suboptimally high interest rates. This, of course, leads to a deeper drop in output gap and inflation today, and as a result - in the current policy rate too. Observing these sharp declines, the agents adjust their current PLMs towards higher sensitivities to shocks and lower perceived reduced-form persistence of the policy rate. Overall, this brings their beliefs one step closer to how rational agents would form expectations. In particular, the pattern of initially decreased perceived interest rate persistence renders the following stability results conservative. In the event of more sudden drops to the ZLB (e.g. ad-hoc deviation from the policy rule), the agents wouldn't have the time and observations to correct downwards their PLM persistence, leading to larger lift-off disagreement and more projection facility events. Accordingly, the instability results can be interpreted as a lower bound on the relevance of the informational curtain mechanism.

[38]If the announcement update in the *Period + Learning FG* regime violates the projection facility, both the update and the associated mask are rejected, so the pre-announcement configuration is retained. In a matched-path robustness that retains the mask after rejecting the update, PF incidence is unchanged at 7.89%, while macroeconomic failure falls from 1.53% to 1.42%.

Once at the ZLB, however, the learning dynamics diverge across the FG regimes. Under No FG, median persistence rises to 0.727 in quarter 53. Period FG remains below its initialization throughout the displayed horizon, while Learning FG and Period + Learning FG peak much earlier at 0.708 and 0.706, respectively. [39]

The lower panel further displays the shadow rate disagreement stemming from the different beliefs in each regime. In all regimes disagreement sharply rises after entering the ZLB. Its median peak falls from 1.16 annualized percentage points under No FG to 0.86 under Period FG, 0.47 under Learning FG, and 0.45 under Period + Learning FG. Notably, even Period FG, which has no such beliefs, updating exhibits lower disagreement than the no communication baseline. Interestingly, this effect operates not by bringing the agents' beliefs closer to the truth, but rather by adopting the CB's announced ZLB path in their expectations, agents now expect lower future interest rates. Through the ALM this boost output gap and inflation, bringing them, and by extension - the interest rate, closer to the optimistic views of the agents, thus lowering the nowcast disagreement.

Figures C.6 - C.8 report operative post-communication durations during the initial episode. Accepted announcements align the median duration exactly under Period FG and Period + Learning FG. Learning FG does not impose the duration mask, but reduces the median gap to at most one quarter during the initial spell.

The differences in disagreement and perceived persistence translate into large differences in calibrated tail risk. Table 4 (Panel A) reports the corresponding incidence and event contexts. Without communication, 45.28% of paths experience at least one attempted crossing of the projection boundary. Period FG lowers this incidence by 37.91 percentage points, to 7.37%; Learning FG lowers it to 11.21%, and Period + Learning FG to 7.89%. This shows that eliminating the *duration* or lift-off disagreement in the Period FG variations is much more important for stability than getting the PLM gradient right.

[39] Note that the No FG and Period FG regimes show increasing trends for $\phi_{ii,t}$ at the ZLB. For both the reason is that the plot is showing the median value across all paths, not just those which are strictly at the ZLB at the given period. This means that as the GFC shock and Monte Carlo innovations coexist, some paths enter or exit the ZLB sooner than others. This means the cross-sectional distribution of $\phi_{ii,t}$ is constantly changing, altering which observation occupies its median rank. For any single path in either of these regimes, under a binding and expected ZLB, $\phi_{ii,t}$ is not updated since the agents have zero forecast error for period t.

Figure 9: Perceived shadow rate persistence and disagreement across FG regimes

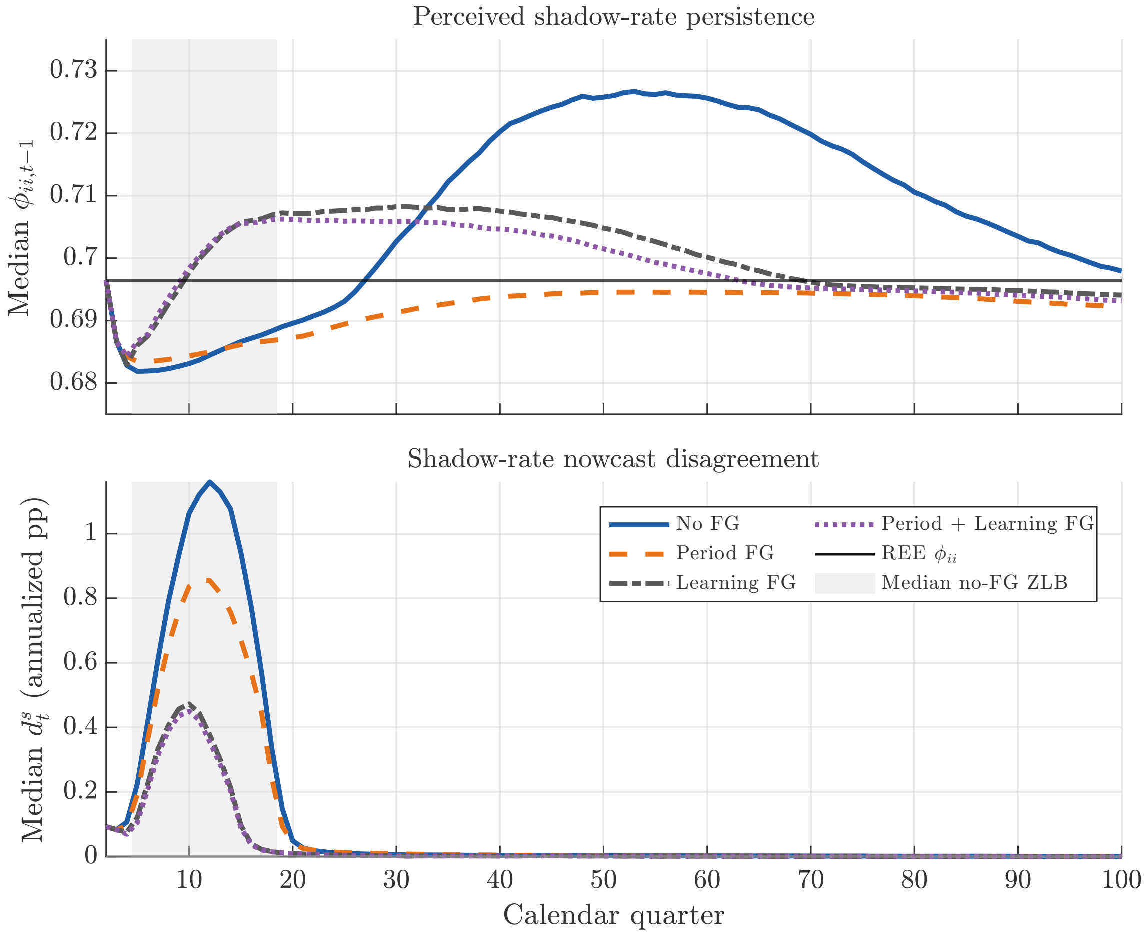

Notes: Each line is the cross-path median over 10,000 matched simulations without PF censoring. Gray areas are quarters in which the median No FG policy rate is at the ZLB. The lower panel plots internal pre-communication shadow rate nowcast disagreement, $d_t^s = 400(s_t^{ag} - s_t^{CB})$, in annualized percentage points.

This happens because if lift-off disagreement is eliminated, any small nowcast disagreement above the ZLB can be corrected properly with the regular learning step. Instead, if agents try to correct their PLM they might not fully close the lift-off disagreement (Figure C.7) or overdo the update step and destabilize their expectations. A communication update generates the first PF event in 4.80% of Learning FG paths and 4.04% of combined-treatment paths. Communication nevertheless delays the median first event from quarter 100 under No FG to quarters 262, 224, and 220, respectively.[40]

Table 4: Stability and stabilization loss across communication regimes

Panel A. Learning stability				
	No FG	Period FG	Learning FG	Per. + Learn. FG
At least one PF	45.28%	7.37%	11.21%	7.89%
First PF at the ZLB	29.70%	1.15%	8.15%	4.77%
First PF outside the ZLB	15.58%	6.22%	3.06%	3.12%
First PF from communication update	—	—	4.80%	4.04%
Median first-PF quarter	100	262	224	220
Median lift-off shadow-rate disagreement	0.69	0.32	0.52	0.24
Macroeconomic failure	11.65%	0.82%	3.05%	1.53%
Panel B. Central-bank stabilization loss over the 40-quarter crisis window				
Mean loss	34.72	20.95	26.63	25.62
Median loss	6.03	5.31	4.96	4.91
Mean loss index	100	60.35	76.71	73.79
P10–P90 path-wise change, %	[0,0]	[-49.7,0]	[-42.1,0]	[-43.4,0]

Notes: Percentages are unconditional shares of 10,000 matched paths. A PF event is a rejected candidate update satisfying $|\widetilde{\phi}_{ii,t}| > 0.98$. Event contexts refer to the path's first PF; the communication row identifies first PF events triggered by the pre-data FG update. Macroeconomic failure is a nonfinite outcome or a crossing of the simulation bound. Lift-off shadow-rate disagreement is the median of $400(s^{ag}_{T^{ag}+1} - s^{CB}_{T^{ag}+1})$, in annualized percentage points, at each path's expected lift-off date. Loss is the discounted average over quarters 2–40; its index normalizes the No FG mean to 100. Path-wise changes compare each FG regime with No FG under the same shocks.

The stability ranking is mirrored in the probability of infeasible macroeconomic outcomes. The No FG economy reaches the macroeconomic feasibility bound in 11.65% of paths. The probability falls to 0.82% under Period FG, 3.05% under Learning FG, and 1.53% under Period + Learning FG. Communication therefore reduces both attempted crossings of the learning-stability boundary and subsequent crossings of the much looser macroeconomic feasibility bound. The result is therefore not merely that communication prevents a candidate coefficient from crossing the expectational stability boundary. It also substantially reduces the frequency with which the simulated macroeconomy subse-

[40]Figure C.9 plots the unconditional cumulative incidence of each regime's first projection event.

quently reaches the much looser feasibility bound.

To approximate a welfare comparison, I use volatility-based loss on a fixed 40-quarter window which spans the initial GFC-induced lower-bound episode and its recovery.For regime s and path n, the central bank's normalized discounted stabilization loss is

$$\bar{L}^s_{40,n} = \frac{\sum_{t=2}^{40} \beta^{t-2} \left[(400\pi^s_{t,n})^2 + 0.5(100x^s_{t,n})^2\right]}{\sum_{t=2}^{40} \beta^{t-2}}. \quad (34)$$

Panel B of Table 4 shows that all three interpretations of forward guidance improve the stabilization criterion. With mean No FG loss normalized to 100, the corresponding mean loss indices are 60.35 under Period FG, 76.71 under Learning FG, and 73.79 under Period + Learning FG. The mean reductions are therefore 39.65%, 23.29%, and 26.21%. The paired median path-wise reductions are 10.54%, 15.09%, and 16.32%, respectively. Period FG delivers the largest reduction in mean loss and the lowest PF incidence, whereas the learning treatments deliver larger improvements for the median path while exposing beliefs to announcement-update risk.

Appendix C.4 varies the announcement gain. Under Learning FG, the loss index declines from 80.65 at $\lambda = 0.01$ to 74.78 at $\lambda = 0.06$, while PF incidence reaches a shallow minimum of 10.86% at $\lambda = 0.02$ and then rises to 23.52%. Under Period + Learning FG, the loss index declines from 75.18 to 67.51; PF incidence falls to 6.75% at $\lambda = 0.04$ before rising to 10.32% at $\lambda = 0.06$. Thus, greater weight on the message improves crisis-window stabilization but increases projection risk at sufficiently high gains; the duration mask moderates this trade-off. The loss distribution is right-skewed, so the substantially larger mean improvements partly reflect protection against adverse tail realizations rather than uniform stabilization across histories.

These results highlight the paper's main argument: the shadow rate censoring by the ZLB creates an informational problem for adaptive learners that generates policy rate forecast disagreement with the central bank. If left unchecked, this disagreement could lead to unanchored expectations and instability. Delphic communication about the expected remaining ZLB duration substantially reduces the associated disagreement and calibrated tail risk. Despite different among themselves, the various FG interpretations all improve on the projection risk over the no-communication case. Yet, the empirical relevance of the different FG interpretations is beyond the scope of this paper and is left for future research.

5 Conclusion

This paper shows that the lower bound can itself generate state-dependent disagreement about the policy path even when the monetary policy rule is unchanged. Above the ZLB, realized policy rates allow private agents to evaluate and update their perceived law of motion for the instrument. At the ZLB, the observed rate instead censors the forecast errors associated with the latent shadow rate. The central bank knows the structural model and can compute the rule-consistent shadow-rate path, whereas private agents must learn that mapping from censored observations. They consequently expect the shadow rate to recover too quickly and predict lift-off earlier than the central bank. The same directional pattern appears in private forecasts relative to both the Riksbank's real-time published paths and the Federal Reserve's historically confidential staff projections, despite the institutions' markedly different past transparency regimes.

The model establishes this disagreement analytically and quantifies its consequences. Following a Great-Recession-sized shock, the median private sector forecast places lift-off seven quarters earlier than the central bank's projection. Learners interpret repeated premature lift-off forecast errors as evidence of a more persistent policy rate that is less responsive to shocks. An excessive increase in the perceived shadow rate persistence can unanchor long-horizon expectations, leading to instability. Across matched simulations, 45.3% of no-communication paths attempt updates beyond the projection boundary. Routine and truthful communication of the central bank's expected ZLB duration lowers this incidence to 7.4 - 11.2% and reduces mean stabilization loss over the crisis window by 23 - 40%. For stability, eliminating disagreement about the duration of the ZLB is more important than correcting the entire perceived policy law of motion.

The policy implication is inherently state dependent. In normal times, the observed policy rate reveals the policy stance, leaving no systematic informational case for publishing a projected path. The paper therefore provides a case for Delphic forward guidance: when a constraint prevents the policy instrument from revealing the underlying policy stance, transparent communication can supply the information that the realized instrument no longer provides, without requiring a promise to deviate from future policy. The analysis assumes that the central bank has better structural knowledge and that its announcements are credible. Allowing the bank itself to learn, introducing imperfect credibility, and identifying empirically how private agents process lift-off announcements are natural directions for future work.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process. During the preparation of this work, the author used ChatGPT (OpenAI) to assist with code development and brainstorming. The model, mathematical derivations, proofs, economic reasoning, intuition, and baseline algorithms were developed independently by the author. ChatGPT was used to assist with automating simulation loops, refactoring large simulation datasets, and generating scripts for the automated export of figures and tables. It was also used to assist with parsing publicly available 2020 Federal Reserve Tealbook PDF files, for which no readily usable structured dataset was available. All AI-assisted outputs were reviewed, validated, and, where necessary, modified by the author. The author takes full responsibility for the content of the work.

References

Adam, K. and R. M. Billi (2006). Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates. *Journal of Money, Credit and Banking* 38.7, pp. 1877–1905.

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon (2019). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics* 11.3, pp. 1–29.

Andreasen, M. M. and A. F. Kronborg (2022). The Extended Perturbation Method: With Applications to the New Keynesian Model and the Zero Lower Bound. *Quantitative Economics* 13.3, pp. 1171–1202. DOI: 10.3982/QE1102.

Angeletos, G.-M. and K. A. Sastry (2021). Managing Expectations: Instruments Versus Targets*. *The Quarterly Journal of Economics* 136.4, pp. 2467–2532. DOI: 10.1093/qje/qjaa045.

Aruoba, S. B., M. Mlikota, F. Schorfheide, and S. Villalvazo (2022). SVARs with occasionally-binding constraints. *Journal of Econometrics* 231.2, pp. 477–499. DOI: 10.1016/j.jeconom.2021.07.013.

Bassetto, M. (2019). Forward Guidance: Communication, Commitment, or Both? *Journal of Monetary Economics* 108, pp. 69–86. DOI: 10.1016/j.jmoneco.2019.08.015.

Bauer, M. D. and E. T. Swanson (2023). An Alternative Explanation for the "Fed Information Effect". *American Economic Review* 113.3, pp. 664–700. DOI: 10.1257/aer.20201220.

Belongia, M. T. and P. N. Ireland (2022). A Reconsideration of Money Growth Rules. *Journal of Economic Dynamics and Control* 135, p. 104312. DOI: 10.1016/j.jedc.2022.104312.

Benhabib, J., G. W. Evans, and S. Honkapohja (2014). Liquidity traps and expectation dynamics: Fiscal stimulus or fiscal austerity? *Journal of Economic Dynamics and Control* 45, pp. 220–238.

Berardi, M. and J. K. Galimberti (2017). Empirical Calibration of Adaptive Learning. *Journal of Economic Behavior & Organization* 144, pp. 219–237. DOI: 10.1016/j.jebo.2017.10.004.

Bernanke, B. S. (2012). *Monetary Policy since the Onset of the Crisis.* Remarks at the Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole, Wyoming. Board of Governors of the Federal Reserve System. URL: https://www.federalreserve.gov/newsevents/speech/bernanke20120831a.htm.

Bernanke, B. S. (2020). The New Tools of Monetary Policy. *The American Economic Review* 110.4, pp. 943–983.

Bhattarai, S., G. B. Eggertsson, and B. Gafarov (2022). Time Consistency and Duration of Government Debt: A Model of Quantitative Easing. *The Review of Economic Studies*, rdac063. DOI: 10.1093/restud/rdac063.

Board of Governors of the Federal Reserve System (2020). *Summary of Economic Projections, December 16, 2020.* URL: https://www.federalreserve.gov/monetarypolicy/fomcprojtabl20201216.htm.

Bodenstein, M., J. Hebden, and R. Nunes (2012). Imperfect Credibility and the Zero Lower Bound. *Journal of Monetary Economics* 59.2, pp. 135–149. DOI: 10.1016/j.jmoneco.2012.01.002.

Boneva, L., R. Harrison, and M. Waldron (2018). Threshold-Based Forward Guidance. *Journal of Economic Dynamics and Control* 90, pp. 138–155. DOI: 10.1016/j.jedc.2018.02.003.

Bullard, J. and K. Mitra (2002). Learning about Monetary Policy Rules. *Journal of Monetary Economics* 49.6, pp. 1105–1129. DOI: 10.1016/S0304-3932(02)00144-7.

Campbell, J. R., C. L. Evans, J. D. Fisher, A. Justiniano, C. W. Calomiris, and M. Woodford (2012). Macroeconomic effects of federal reserve forward guidance [with comments and discussion]. *Brookings Papers on Economic Activity*, pp. 1–80.

Campbell, J. R., J. D. Fisher, A. Justiniano, and L. Melosi (2017). Forward guidance and macroeconomic outcomes since the financial crisis. *NBER Macroeconomics Annual* 31.1, pp. 283–357.

Campbell, J. R., F. Ferroni, J. D. M. Fisher, and L. Melosi (2019). The Limits of Forward Guidance. *Journal of Monetary Economics* 108, pp. 118–134. DOI: 10.1016/j.jmoneco.2019.08.009.

Carvalho, C., F. Nechio, and T. Tristão (2021). Taylor Rule Estimation by OLS. *Journal of Monetary Economics* 124, pp. 140–154. DOI: 10.1016/j.jmoneco.2021.10.010.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2000). Sticky price models of the business cycle: can the contract multiplier solve the persistence problem? *Econometrica* 68.5, pp. 1151–1179.

Cho, S. and A. Moreno (2006). A Small-Sample Study of the New-Keynesian Macro Model. *Journal of Money, Credit and Banking*, pp. 1461–1481.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy* 119.1, pp. 78–121.

Chung, H., J.-P. Laforte, D. Reifschneider, and J. C. Williams (2012). Have We Underestimated the Likelihood and Severity of Zero Lower Bound Events? *Journal of Money, Credit and Banking* 44.s1, pp. 47–82. DOI: 10.1111/j.1538-4616.2011.00478.x.

Chung, H. T., E. Gagnon, T. Nakata, M. O. Paustian, B. Schlusche, J. Trevino, D. Vilan, and W. Zheng (2019). *Monetary Policy Options at the Effective Lower Bound: Assessing the Federal Reserve's Current Policy Toolkit.* Board of Governors of the Federal Reserve System Finance and Economics Discussion Series 2019-003.

Cieslak, A. (2018). Short-rate expectations and unexpected returns in treasury bonds. *The Review of Financial Studies* 31.9, pp. 3265–3306.

Clarida, R., J. Galí, and M. Gertler (2000). Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *The Quarterly Journal of Economics* 115.1, pp. 147–180. DOI: 10.1162/003355300554692.

Coibion, O. and Y. Gorodnichenko (2012). Why Are Target Interest Rate Changes So Persistent? *American Economic Journal: Macroeconomics* 4.4, pp. 126–162. DOI: 10.1257/mac.4.4.126.

Cole, S. J. (2020a). Learning and the Effectiveness of Central Bank Forward Guidance. *Journal of Money, Credit and Banking* n/a.n/a. DOI: 10.1111/jmcb.12696.

Cole, S. J. (2020b). The Limits of Central Bank Forward Guidance under Learning. *63rd issue (September 2020) of the International Journal of Central Banking.*

Daly, M. C. (2025). Dynamic Central Bank Communication. *FRBSF Economic Letter* 2025-15.

Del Negro, M., M. P. Giannoni, and C. Patterson (2012). The forward guidance puzzle. *FRB of New York Staff Report* 574.

Dordal i Carreras, M., O. Coibion, Y. Gorodnichenko, and J. Wieland (2016). Infrequent but Long-Lived Zero Lower Bound Episodes and the Optimal Rate of Inflation. *Annual*

Review of Economics 8, pp. 497–520. DOI: 10.1146/annurev-economics-080315-015306.

Eggertsson, G. B. and M. Woodford (2003). Zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity* 2003.1, pp. 139–233.

Engen, E., T. Laubach, and D. Reifschneider (2015). The macroeconomic effects of the Federal Reserve's unconventional monetary policies.

English, W. B., J. D. Lopez-Salido, and R. J. Tetlow (2015). The Federal Reserve's framework for monetary policy: Recent changes and new questions. *IMF Economic Review* 63.1, pp. 22–70.

Eusepi, S. (2010). Central Bank Communication and the Liquidity Trap. *Journal of Money, Credit and Banking* 42.2-3, pp. 373–397. DOI: 10.1111/j.1538-4616.2009.00291.x.

Eusepi, S. and B. Preston (2010). Central Bank Communication and Expectations Stabilization. *American Economic Journal: Macroeconomics* 2.3, pp. 235–71.

Eusepi, S. and B. Preston (2018). The science of monetary policy: An imperfect knowledge perspective. *Journal of Economic Literature* 56.1, pp. 3–59.

Evans, G. W. and S. Honkapohja (2001). *Learning and expectations in macroeconomics.* Princeton University Press.

Evans, G. W. and S. Honkapohja (2003). Adaptive learning and monetary policy design. *Journal of Money, Credit, and Banking* 35.6, pp. 1045–1072.

Evans, G. W., E. Guse, and S. Honkapohja (2008). Liquidity Traps, Learning and Stagnation. *European Economic Review* 52.8, pp. 1438–1463. DOI: 10.1016/j.euroecorev.2008.05.003.

Evans, G. W. and S. Honkapohja (1998). Convergence of Learning Algorithms without a Projection Facility. *Journal of Mathematical Economics* 30.1, pp. 59–86. DOI: 10.1016/S0304-4068(97)00023-2.

Federal Open Market Committee (2020). *Meeting Presentation Materials, September 15–16, 2020.* See Chart 17, based on the Federal Reserve Bank of New York Surveys of

Primary Dealers and Market Participants. URL: https://www.federalreserve.gov/monetarypolicy/files/FOMC20200916material.pdf.

Federal Reserve Bank of New York (2020a). *Survey of Market Participants: December 2020*. URL: https://www.newyorkfed.org/medialibrary/media/markets/survey/2020/dec-2020-smp-results.pdf.

Federal Reserve Bank of New York (2020b). *Survey of Primary Dealers: December 2020*. URL: https://www.newyorkfed.org/medialibrary/media/markets/survey/2020/dec-2020-spd-results.pdf.

Fuhrer, J. C. (2000). Habit Formation in Consumption and Its Implications for Monetary-Policy Models. *American Economic Review* 90.3, pp. 367–390. DOI: 10.1257/aer.90.3.367.

Furlanetto, F. and A. Lepetit (2025). 08: The Slope of the Phillips Curve. In: *Research Handbook on Inflation*. Edward Elgar Publishing, pp. 167–185.

Gabaix, X. (2020). A Behavioral New Keynesian Model. *American Economic Review* 110.8, pp. 2271–2327. DOI: 10.1257/aer.20162005.

Gagliardone, L., M. Gertler, S. Lenzu, and J. Tielens (2025). Anatomy of the Phillips Curve: Micro Evidence and Macro Implications. *American Economic Review* 115.11, pp. 3941–3974. DOI: 10.1257/aer.20231569.

Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press.

Gáti, L. (2023). Monetary Policy & Anchored Expectations—An Endogenous Gain Learning Model. *Journal of Monetary Economics*. Inflation: Drivers and Dynamics 2022 140, S37–S47. DOI: 10.1016/j.jmoneco.2023.06.009.

Guerrieri, L. and M. Iacoviello (2015). OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics* 70, pp. 22–38.

Haberis, A., R. Harrison, and M. Waldron (2019). Uncertain Policy Promises. *European Economic Review* 111, pp. 459–474. DOI: 10.1016/j.euroecorev.2018.11.003.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2022). The Slope of the Phillips Curve: Evidence from U.S. States*. *The Quarterly Journal of Economics* 137.3, pp. 1299–1344. DOI: 10.1093/qje/qjac010.

Honkapohja, S., K. Mitra, and G. W. Evans (2012). Notes on Agents' Behavioral Rules Under Adaptive Learning and Studies of Monetary Policy.

Janson, W. and C. Jia (2020). Forward Guidance during the Pandemic: Has It Changed the Public's Expectations? *Economic Commentary* 2020-27. DOI: 10.26509/frbc-ec-202027.

Kiley, M. T. and J. M. Roberts (2017). Monetary Policy in a Low Interest Rate World. *Brookings Papers on Economic Activity* 2017.1, pp. 317–396. DOI: 10.1353/eca.2017.0004.

Krugman, P. R., K. M. Dominquez, and K. Rogoff (1998). It's baaack: Japan's slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, pp. 137–205.

Levin, A., D. López-Salido, E. Nelson, and T. Yun (2010). Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound. *International Journal of Central Banking* 6.1, pp. 143–189.

Lubik, T. A. and F. Schorfheide (2004). Testing for indeterminacy: An application to US monetary policy. *American Economic Review* 94.1, pp. 190–217.

Marcet, A. and T. J. Sargent (1989). Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory* 48.2, pp. 337–368. DOI: 10.1016/0022-0531(89)90032-X.

Milani, F. (2007). Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics* 54.7, pp. 2065–2082. DOI: 10.1016/j.jmoneco.2006.11.007.

Milani, F. (2011). Expectation Shocks and Learning as Drivers of the Business Cycle. *The Economic Journal* 121.552, pp. 379–401.

Molnár, K. and S. Santoro (2014). Optimal Monetary Policy When Agents Are Learning. *European Economic Review* 66, pp. 39–62. DOI: 10.1016/j.euroecorev.2013.08.012.

Nakov, A. (2008). Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate. *International Journal of Central Banking*.

Orphanides, A. and J. C. Williams (2007). Robust Monetary Policy with Imperfect Knowledge. *Journal of Monetary Economics*. Carnegie-Rochester Conference Series on Public Policy: Issues in Current Monetary Policy Analysis November 10-11, 2006 54.5, pp. 1406–1435. DOI: 10.1016/j.jmoneco.2007.06.005.

Preston, B. (2005). Learning about Monetary Policy Rules When Long-Horizon Expectations Matter. *International Journal of Central Banking* 1.2.

Preston, B. (2006). Adaptive Learning, Forecast-Based Instrument Rules and Monetary Policy. *Journal of Monetary Economics* 53.3, pp. 507–535. DOI: 10.1016/j.jmoneco.2005.01.008.

Rotemberg, J. J. and M. Woodford (1997). An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy. *NBER macroeconomics annual* 12, pp. 297–346.

Slobodyan, S. and R. Wouters (2012). Learning in an Estimated Medium-Scale DSGE Model. *Journal of Economic Dynamics and Control* 36.1, pp. 26–46. DOI: 10.1016/j.jedc.2011.01.016.

Smets, F. and R. Wouters (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *The American Economic Review* 97.3, pp. 586–606.

Swanson, E. T. and J. C. Williams (2014). Measuring the effect of the zero lower bound on medium-and longer-term interest rates. *American economic review* 104.10, pp. 3154–3185.

Warsh, K. M. (2026). *In Our Time.* Remarks at the Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole, Wyoming. Board of Governors of the Federal Reserve System. URL: https://www.federalreserve.gov/newsevents/speech/warsh20260828a.htm.

Woodford, M. (2003). *Interest and prices: Foundations of a theory of monetary policy.*

Yellen, J. L. (2011). *Unconventional Monetary Policy and Central Bank Communications.* Remarks at the University of Chicago Booth School of Business U.S. Monetary Policy Forum, New York. Board of Governors of the Federal Reserve System. URL: https://www.federalreserve.gov/newsevents/speech/yellen20110225a.htm.

Appendices

A Data

A.1 Construction of the forecast comparisons

Sources are the Riksbank's published policy rate paths (Monetary Policy Reports and Updates, February 2007 - June 2026) and the mean expected repo rate in the Prospera/Origo money-market survey at three, twelve, and twenty-four months. Each vintage is matched to the most recent survey on or before the publication date. Three Updates (September 2008, December 2008, and September 2009) are dropped because those surveys are more than 45 days old, leaving 102 pairs. The Riksbank path is interpolated linearly to the survey's 3-, 12-, and 24-month calendar dates. I evaluate the official Riksbank curve at those calendar dates by linear interpolation between adjacent published path points. The interpolation aligns the two sides on the same date; it is not a substitute for missing vintages. The low-rate (ELB) split in the text uses a one-year-ahead expectation below 0.25 percent.

United States, 2003 - 2014. Internal forecasts are the Board staff short-rate assumptions in the Philadelphia Fed Real-Time Data file "Output Gap and Financial Assumptions from the Board of Governors" (formerly Greenbook, then Tealbook). The funds-rate columns are quarterly averages from the current quarter through nine quarters ahead. The public file ends in December 2014.

Private forecasts are the mean federal funds rate in the Wall Street Journal Economic Forecasting Survey (WSJ), available monthly since January 2003. Targets are always end-June or end-December. I match each Tealbook issue to the WSJ survey in the same calendar month. The horizon in quarters is the difference between the WSJ target quarter and the Tealbook quarter.

United States, 2020. For 2020 I code the Tealbook A Long-Term Outlook Q4 average funds rate from the eight Tealbook A issues dated 17 January, 6 March, 17 April, 29 May, 17 July, 4 September, 23 October, and 4 December. I match each Q4 point to the WSJ mean in the same calendar month for that year's December. As in 2003-2014, the survey's targets are always end-June or end-December of a named year; the eight 2020 vintages therefore contribute year-end forecasts throughout the year, not a single December 2020 survey. For instance, June (Q2) is filled at 0.13 only if that calendar year's printed Q4 remains at the ELB (≤ 0.13). If the printed Q4 already lies above the ELB, I leave the June cell blank rather than interpolate along a rising path.

A.2 Empirical evidence

A.2.1 Robustness of Swedish repo-rate forecast disagreement

Table A.1 reports the market-minus-Riksbank gap at three-, twelve-, and twenty-four-month horizons, using 0.25 and 0.75 cutoffs on the one-year-ahead rate and splitting on either the market or the Riksbank forecast. As in Table 1, in the low-rate regime the mean gap is positive at one and two years (the market expects a higher repo rate). In the high-rate regime the one- and two-year means reverse sign. The two-year high-regime mean is the larger of the two: away from the lower bound, the Riksbank path lies above the market on average. Three-month gaps are small in both regimes.

Figure A.1 illustrates the first spell: in October 2009 the market path lies above the Riksbank path; in 2010, after lift-off, the two coincide. The lower-bound gap persisted even though the Riksbank published an explicit repo-rate path. Publication of the path did not eliminate the wedge.

A.2.2 United States

Table A.1: Disagreement on the Swedish repo rate, 2007 - 2026

Based on money-market expected 1-year-ahead repo rate	$\mathbb{E}_t i_{t+1y} \lessgtr 0.25$		$\mathbb{E}_t i_{t+1y} \lessgtr 0.75$	
	mean	se(mean)	mean	se(mean)
dis_3m_low	0.0130	0.0048	0.0238	0.0095
dis_1y_low	0.0756	0.0168	0.0940	0.0178
dis_2y_low	0.1895	0.0279	0.1859	0.0321
dis_3m_high	0.0122	0.0230	0.0032	0.0240
dis_1y_high	-0.1026	0.0405	-0.1337	0.0414
dis_2y_high	-0.3152	0.0545	-0.3574	0.0535

Based on Riksbank expected 1-year-ahead repo rate	$\mathbb{E}_t i_{t+1y} \lessgtr 0.25$		$\mathbb{E}_t i_{t+1y} \lessgtr 0.75$	
	mean	se(mean)	mean	se(mean)
dis_3m_low	0.0234	0.0084	0.0233	0.0091
dis_1y_low	0.0877	0.0178	0.1174	0.0236
dis_2y_low	0.1714	0.0316	0.1834	0.0310
dis_3m_high	0.0042	0.0235	0.0029	0.0249
dis_1y_high	-0.1210	0.0410	-0.1629	0.0374
dis_2y_high	-0.3276	0.0557	-0.3753	0.0538

Notes: Gap = money-market expectation minus Riksbank path. Primary pairs, February 2007 - June 2026. Low (High): one-year-ahead expectation below (above) the stated cutoff.
Source: Riksbank Monetary Policy Reports and Updates; Prospera/Origo Money Market Players.

Figure A.1: Forecasts of the Swedish repo rate, 2009 - 2010

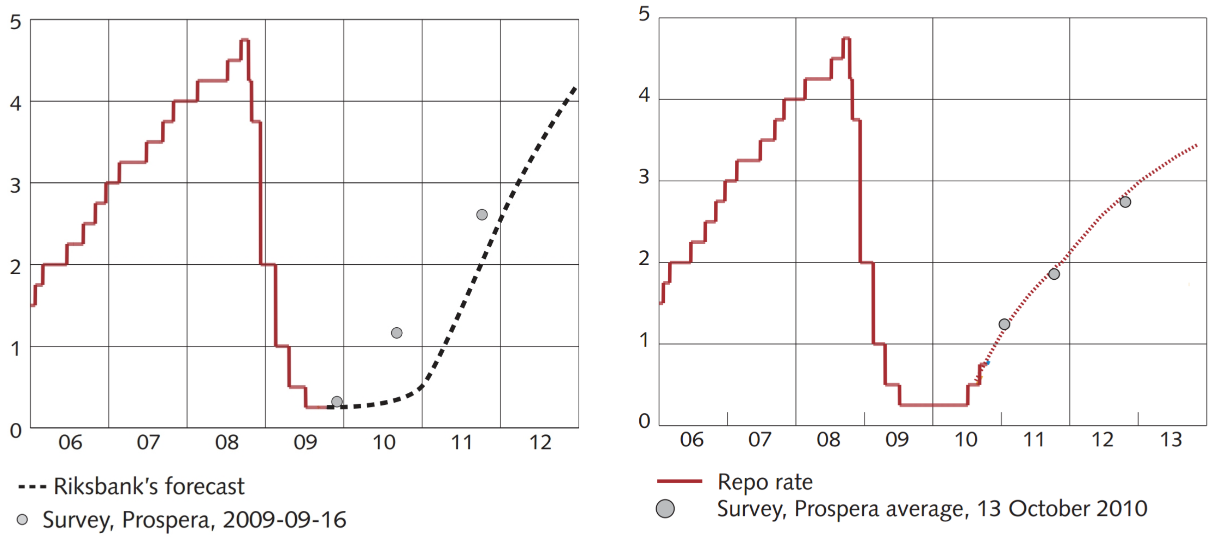

Source: Riksbank Monetary Policy Report, 2009 - 2010.

Table A.2: WSJ federal funds rate forecasts matched to Tealbook staff, by horizon

	1q	2q	3q	4q	5q	6q	7q	8q
count	43	48	43	31	18	16	12	13

Source: Wall Street Journal Economic Forecasting Survey matched to Philadelphia Fed Greenbook/Tealbook short-rate assumptions, March 2003 - December 2014.

Figure A.2: Federal funds rate forecast disagreement, 2020

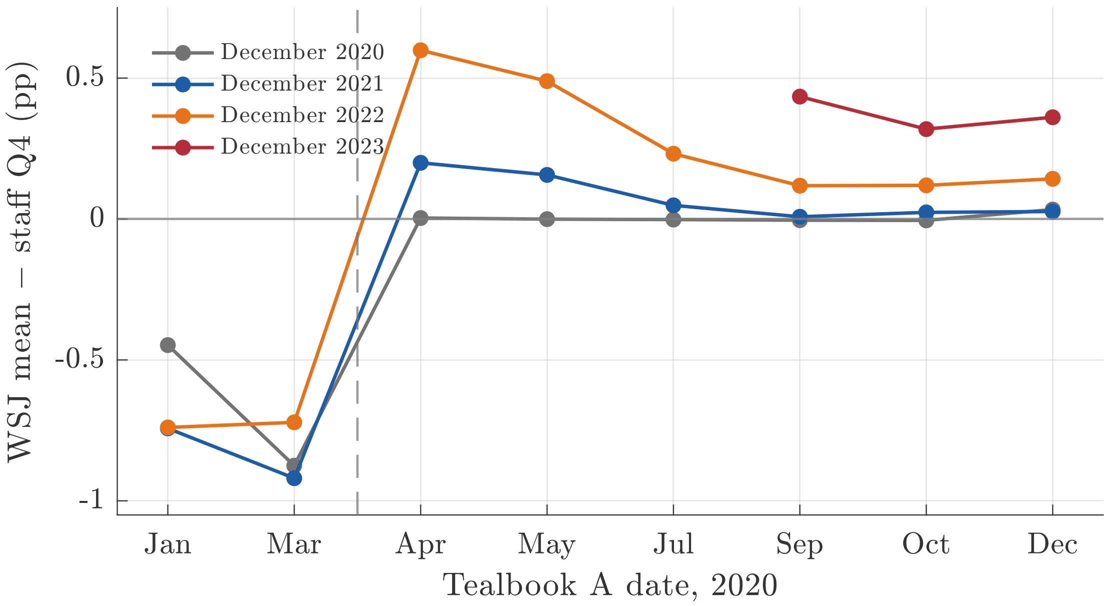

Notes. Gap is the Wall Street Journal December mean minus the Tealbook A Long-Term Outlook Q4 staff path. Positive: private forecasters expect a higher funds rate than staff. Dashed line: 15 March unscheduled FOMC.

B Model

B.1 Actual laws of motion (ALMs)

The ALM is determined by:

$$x_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma} \left(i_T - \pi_{T+1} - r_T \right) \right] \tag{35}$$

$$\pi_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left[\kappa x_T + (1-\alpha)\beta \pi_{T+1} + u_T \right] \tag{36}$$

To compute the values of x and π at any given time we need expressions for the infinite sums of each component in (35) and (36) which can be computed, i.e. are functions of the realized state variables.

B.1.1 ALM - output gap

Taking first the output gap:

$$\begin{aligned}
x_t &= \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma} \left(i_T - \pi_{T+1} - r_T \right) \right] \\
&= (1-\beta) \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} x_{T+1} - \frac{1}{\sigma} \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} i_T + \frac{1}{\sigma} \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \pi_{T+1} + \frac{1}{\sigma} \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} r_T \\
&= (1-\beta) \underbrace{\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} x_{T+1}}_{\mathtt{X_{x,t}}} - \frac{1}{\sigma} \underbrace{\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} i_T}_{\mathtt{X_{i,t}}} + \frac{1}{\sigma} \underbrace{\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \pi_{T+1}}_{\mathtt{X_{\pi,t}}} + \frac{r_t}{\sigma(1-\beta\rho_r)} \\
&= (1-\beta) \mathtt{X_{x,t}} - \frac{1}{\sigma} \mathtt{X_{i,t}} + \frac{1}{\sigma} \mathtt{X_{\pi,t}} + \frac{r_t}{\sigma(1-\beta\rho_r)}
\end{aligned} \tag{37}$$

Since the lagged interest rate is a state variable, the infinite expectational sums for x_T and π_T require infinite sums of forecasts for i_T from the state vectors Z_T. Thus, it will be convenient to first compute the sum for the interest rate - $\mathtt{X_{i,t}}$, before obtaining the ones for $\mathtt{X_{x,t}}$ and $\mathtt{X_{\pi,t}}$.

Let us assume for now that no binding ZLB periods are expected. This assumption will be relaxed in Appendix B.2. Then,

$$\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} i_T = \hat{\mathbb{E}}_t i_t + \beta \hat{\mathbb{E}}_t i_{t+1} + \beta^2 \hat{\mathbb{E}}_t i_{t+2} + \beta^3 \hat{\mathbb{E}}_t i_{t+3} + \beta^4 \hat{\mathbb{E}}_t i_{t+4} + \dots =$$

$$= \phi_{i,t-1} \hat{\mathbb{E}}_t Z_t + \beta \phi_{i,t-1} \hat{\mathbb{E}}_t Z_{t+1} + \beta^2 \phi_{i,t-1} \hat{\mathbb{E}}_t Z_{t+2} + \dots =$$

$$= \phi_{i,t-1} Z_t + \beta \phi_{i,t-1} \Psi_{t-1} Z_t + \beta^2 \phi_{i,t-1} \Psi_{t-1}^2 Z_t + \dots =$$

$$= \phi_{i,t-1} \left[I + \beta \Psi_{t-1} + \beta^2 \Psi_{t-1}^2 + \dots \right] Z_t = \tag{38}$$

$$= \phi_{i,t-1} \left[V_{t-1} I V_{t-1}^{-1} + \beta V_{t-1} D_{t-1} V_{t-1}^{-1} + \beta^2 V_{t-1} D_{t-1}^2 V_{t-1}^{-1} + \dots \right] Z_t = \tag{39}$$

$$= \phi_{i,t-1} V_{t-1} \left[I + \beta D_{t-1} + \beta^2 D_{t-1}^2 + \dots \right] V_{t-1}^{-1} Z_t \tag{40}$$

where we define Ψ_{t-1} as the perceived law of motion of the state variables:

$$\hat{\mathbb{E}}_t Z_{t+1} = \Psi_{t-1} Z_t =$$

$$= \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \rho_u & 0 & 0 \\ 0 & 0 & \rho_r & 0 \\ \phi_{ic,t-1} & \phi_{iu,t-1} & \phi_{ir,t-1} & \phi_{ii,t-1} \end{bmatrix}}_{\Psi_{t-1}} \underbrace{\begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix}}_{Z_t} \tag{41}$$

The transformation between (38) and (39) is an eigendecomposition of the square matrix $\Psi_{t-1} = V_{t-1} D_{t-1} V_{t-1}^{-1}$, where D is a diagonal matrix of Ψ's eigenvalues and V is a matrix featuring its eigenvectors as columns. It can be shown that:

$$D_{t-1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \rho_r & 0 & 0 \\ 0 & 0 & \rho_u & 0 \\ 0 & 0 & 0 & \phi_{ii,t-1} \end{bmatrix} \tag{42}$$

$$V_{t-1} = \begin{bmatrix} \frac{1-\phi_{ii,t-1}}{\phi_{ic,t-1}} & 0 & 0 & 0 \\ 0 & 0 & \frac{\rho_u-\phi_{ii,t-1}}{\phi_{iu,t-1}} & 0 \\ 0 & \frac{\rho_r-\phi_{ii,t-1}}{\phi_{ir,t-1}} & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix} \tag{43}$$

Note that in the cases where $\phi_{ii,t-1} = \rho_u$ or $\phi_{ii,t-1} = \rho_r$ or $\phi_{ii,t-1} = 1$ at least two of the eigenvectors become collinear and the decomposition is not possible. In what follows for the sake of brevity I will omit the subscript $t-1$ on the elements of the law of motion for i: ϕ_i.

Plugging in the equation for D_{t-1} from (42) into (40) and computing the series in the parentheses we obtain $\mathtt{X_{i,t}}$ from (37):

$$\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} i_T = \phi_{i,t-1} V_{t-1} \left[I + \beta D_{t-1} + \beta^2 D_{t-1}^2 + \dots \right] V_{t-1}^{-1} Z_t =$$

$$= \phi_{i,t-1} V_{t-1} \begin{bmatrix} \frac{1}{1-\beta} & 0 & 0 & 0 \\ 0 & \frac{1}{1-\beta\rho_r} & 0 & 0 \\ 0 & 0 & \frac{1}{1-\beta\rho_u} & 0 \\ 0 & 0 & 0 & \frac{1}{1-\beta\phi_{ii}} \end{bmatrix} V_{t-1}^{-1} Z_t =$$

$$= \begin{bmatrix} \frac{\phi_{ic}}{(1-\beta)(1-\beta\phi_{ii})} \\ \frac{\phi_{iu}}{(1-\beta\rho_u)(1-\beta\phi_{ii})} \\ \frac{\phi_{ir}}{(1-\beta\rho_r)(1-\beta\phi_{ii})} \\ \frac{\phi_{ii}}{1-\beta\phi_{ii}} \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} = \boxed{\mathtt{X_{i,t}}} \tag{44}$$

Next we can compute the component for the sum of discounted output gap expectations $\mathtt{X_{x,t}}$ from (37), which similarly to the case of $\mathtt{X_{i,t}}$ above yields:

$$\begin{aligned}
\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} x_{T+1} &= \hat{\mathbb{E}}_t x_{t+1} + \beta \hat{\mathbb{E}}_t x_{t+2} + \beta^2 \hat{\mathbb{E}}_t x_{t+3} + \dots = \\
&= \phi_{x,t-1} \hat{\mathbb{E}}_t Z_{t+1} + \beta \phi_{x,t-1} \hat{\mathbb{E}}_t Z_{t+2} + \beta^2 \phi_{x,t-1} \hat{\mathbb{E}}_t Z_{t+3} + \dots = \\
&= \phi_{x,t-1} \left[\Psi_{t-1} Z_t + \beta \Psi_{t-1}^2 Z_t + \beta^2 \Psi_{t-1}^3 Z_t + \dots \right] = \\
&= \phi_{x,t-1} \Psi_{t-1} \left[I + \beta \Psi_{t-1} + \beta^2 \Psi_{t-1}^2 + \dots \right] Z_t = \\
&= \phi_{x,t-1} \Psi_{t-1} V_{t-1} \left[I + \beta D_{t-1} + \beta^2 D_{t-1}^2 + \dots \right] V_{t-1}^{-1} Z_t = \\
&= \phi_{x,t-1} \Psi_{t-1} V_{t-1} \begin{bmatrix} \frac{1}{1-\beta} & 0 & 0 & 0 \\ 0 & \frac{1}{1-\beta\rho_r} & 0 & 0 \\ 0 & 0 & \frac{1}{1-\beta\rho_u} & 0 \\ 0 & 0 & 0 & \frac{1}{1-\beta\phi_{ii}} \end{bmatrix} V_{t-1}^{-1} Z_t =
\end{aligned} \tag{45}$$

$$\begin{aligned}
&= \begin{bmatrix} \frac{\phi_{xc}}{1-\beta} + \phi_{xi} \frac{\phi_{ic}}{(1-\beta)(1-\beta\phi_{ii})} \\ \frac{\phi_{xu}\rho_u}{1-\beta\rho_u} + \phi_{xi} \frac{\phi_{iu}}{(1-\beta\rho_u)(1-\beta\phi_{ii})} \\ \frac{\phi_{xr}\rho_r}{1-\beta\rho_r} + \phi_{xi} \frac{\phi_{ir}}{(1-\beta\rho_r)(1-\beta\phi_{ii})} \\ \phi_{xi} \frac{\phi_{ii}}{1-\beta\phi_{ii}} \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} = \\
&= \begin{bmatrix} \frac{\phi_{xc}}{1-\beta} \\ \frac{\phi_{xu}\rho_u}{1-\beta\rho_u} \\ \frac{\phi_{xr}\rho_r}{1-\beta\rho_r} \\ 0 \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} + \phi_{xi} \mathtt{X_{i,t}} = \boxed{\mathtt{X_{x,t}}}
\end{aligned} \tag{46}$$

Analogously for the sum of discounted inflation expectations $\mathtt{X_{\pi,t}}$ from (37):

$$
\begin{aligned}
\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \pi_{T+1} &= \hat{\mathbb{E}}_t \pi_{t+1} + \beta \hat{\mathbb{E}}_t \pi_{t+2} + \beta^2 \hat{\mathbb{E}}_t \pi_{t+3} + \cdots = \\
&= \phi_{\pi,t-1} \Psi_{t-1} V_{t-1} \begin{bmatrix} \frac{1}{1-\beta} & 0 & 0 & 0 \\ 0 & \frac{1}{1-\beta\rho_r} & 0 & 0 \\ 0 & 0 & \frac{1}{1-\beta\rho_u} & 0 \\ 0 & 0 & 0 & \frac{1}{1-\beta\phi_{ii}} \end{bmatrix} V_{t-1}^{-1} Z_t = \\
&= \begin{bmatrix} \frac{\phi_{\pi c}}{1-\beta} + \phi_{\pi i} \frac{\phi_{ic}}{(1-\beta)(1-\beta\phi_{ii})} \\ \frac{\phi_{\pi u}\rho_u}{1-\beta\rho_u} + \phi_{\pi i} \frac{\phi_{iu}}{(1-\beta\rho_u)(1-\beta\phi_{ii})} \\ \frac{\phi_{\pi r}\rho_r}{1-\beta\rho_r} + \phi_{\pi i} \frac{\phi_{ir}}{(1-\beta\rho_r)(1-\beta\phi_{ii})} \\ \phi_{\pi i} \frac{\phi_{ii}}{1-\beta\phi_{ii}} \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} = \\
&= \begin{bmatrix} \frac{\phi_{\pi c}}{1-\beta} \\ \frac{\phi_{\pi u}\rho_u}{1-\beta\rho_u} \\ \frac{\phi_{\pi r}\rho_r}{1-\beta\rho_r} \\ 0 \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} + \phi_{\pi i} \mathtt{X}_{\mathtt{i},\mathtt{t}} = \boxed{\mathtt{X}_{\pi,\mathtt{t}}}
\end{aligned} \tag{47}
$$

Using (44), (46), and (47) in (37) we finally obtain the reduced form expression for the law of motion of the output gap:

$$
\begin{aligned}
x_t &= \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma} \left(i_T - \pi_{T+1} - r_T \right) \right] \\
&= (1-\beta) \mathtt{X}_{\mathtt{x},\mathtt{t}} - \frac{1}{\sigma} \mathtt{X}_{\mathtt{i},\mathtt{t}} + \frac{1}{\sigma} \mathtt{X}_{\pi,\mathtt{t}} + \frac{r_t}{\sigma(1-\beta\rho_r)} = x(Z_t)
\end{aligned} \tag{48}
$$

Note that $\mathtt{X}_{\mathtt{x},\mathtt{t}}$ and $\mathtt{X}_{\pi,\mathtt{t}}$ depend negatively on $\mathtt{X}_{\mathtt{i},\mathtt{t}}$ since $\phi_{xi} < 0$ and $\phi_{\pi i} < 0$. Since $\frac{1}{\sigma} > 0$, then x_t depends negatively on $\mathtt{X}_{\mathtt{i},\mathtt{t}}$.

B.1.2 ALM - inflation

Next we repeat the procedure for inflation:

$$\begin{aligned}
\pi_t &= \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left[\kappa x_T + (1-\alpha)\beta\pi_{T+1} + u_T\right] \\
&= \kappa \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} x_T + \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} u_T + (1-\alpha)\beta \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \pi_{T+1} \\
&= \underbrace{\kappa \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} x_T}_{\Pi_{x,t}} + (1-\alpha)\beta \underbrace{\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \pi_{T+1}}_{\Pi_{\pi,t}} + \frac{1}{(1-\alpha\beta\rho_u)} u_t \qquad (49)
\end{aligned}$$

Before computing its components it is useful to follow (44) and first compute:

$$\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} i_T = \begin{bmatrix} \frac{\phi_{ic}}{(1-\alpha\beta)(1-\alpha\beta\phi_{ii})} \\ \frac{\phi_{iu}}{(1-\alpha\beta\rho_u)(1-\alpha\beta\phi_{ii})} \\ \frac{\phi_{ir}}{(1-\alpha\beta\rho_r)(1-\alpha\beta\phi_{ii})} \\ \frac{\phi_{ii}}{1-\alpha\beta\phi_{ii}} \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} = \boxed{\Pi_{i,t}} \qquad (50)$$

Then, continuing with the components of (49) one at a time, the expected future discounted sum of output gaps becomes:

$$
\begin{aligned}
&\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} x_T = \hat{\mathbb{E}}_t x_t + \alpha\beta \hat{\mathbb{E}}_t x_{t+1} + \alpha^2\beta^2 \hat{\mathbb{E}}_t x_{t+2} + \dots \\
&= \hat{\mathbb{E}}_t x_t + \alpha\beta \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} x_{T+1} = \\
&= \phi_x Z_t + \alpha\beta \begin{bmatrix} \frac{\phi_{xc}}{1-\alpha\beta} \\ \frac{\phi_{xu}\rho_u}{1-\alpha\beta\rho_u} \\ \frac{\phi_{xr}\rho_r}{1-\alpha\beta\rho_r} \\ 0 \end{bmatrix}' Z_t + \alpha\beta\phi_{xi}\Pi_{\mathtt{i,t}} \\
&= \begin{bmatrix} \frac{\phi_{xc}}{1-\alpha\beta} \\ \frac{\phi_{xu}}{1-\alpha\beta\rho_u} \\ \frac{\phi_{xr}}{1-\alpha\beta\rho_r} \\ \phi_{xi} \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} + \alpha\beta\phi_{xi}\Pi_{\mathtt{i,t}} = \boxed{\Pi_{\mathtt{x,t}}} \qquad (51)
\end{aligned}
$$

where the result follows analogously following (46) and using (50). Similarly for inflation using (47) and (50):

$$
\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \pi_{T+1} = \begin{bmatrix} \frac{\phi_{\pi c}}{1-\alpha\beta} \\ \frac{\phi_{\pi u}\rho_u}{1-\alpha\beta\rho_u} \\ \frac{\phi_{\pi r}\rho_r}{1-\alpha\beta\rho_r} \\ 0 \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ i_{t-1} \end{bmatrix} + \phi_{\pi i}\Pi_{\mathtt{i,t}} = \boxed{\Pi_{\pi,\mathtt{t}}} \qquad (52)
$$

Plugging in (50), (51), and (52) in (49) we obtain the reduced form expression for the law of motion of inflation:

$$
\begin{aligned}
\pi_t &= \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left[\kappa x_T + (1-\alpha)\beta\pi_{T+1} + u_T\right] \\
&= \kappa\Pi_{\mathtt{x,t}} + (1-\alpha)\beta\Pi_{\pi,\mathtt{t}} + \frac{1}{(1-\alpha\beta\rho_u)} u_t = \pi(Z_t) \qquad (53)
\end{aligned}
$$

Similar to the output gap in the previous section, $\Pi_{\mathtt{x,t}}$ and $\Pi_{\pi,\mathtt{t}}$ depend negatively on

$\Pi_{i,t}$ since $\phi_{xi} < 0$ and $\phi_{\pi i} < 0$. Thus, π_t depends negatively on $\Pi_{i,t}$. Therefore, $x(Z_t)$ in (48) and $\pi(Z_t)$ in (53) can also be expressed as $x(Z_t, \underset{(-)}{X_{i,t}})$ and $\pi(Z_t, \underset{(-)}{\Pi_{i,t}})$, respectively. This will be used in Appendix B.2 to prove the forecast disagreement of interest rates between the central bank and the agents at the ZLB.

B.2 Proof of disagreement at the ZLB

Suppose for a long enough period the economy has always been above the ZLB until period t when a big negative shock is realized. As a result, let us assume at period $t-1$ the learning agents have converged to the REE. That is, after observing the shocks at the beginning of period $t-1$ the agents have model-consistent nowcasts and forecasts: $\hat{\mathbb{E}}_{t-1}x_{t-1} = x_{t-1}$, $\hat{\mathbb{E}}_{t-1}x_{t+j} = E_{t-1}x_{t+j}$ $\forall j$, and $\Phi_{t-2} = \Gamma$. Thus, in normal times no adjustment in beliefs (i.e. the agents' perceived law of motion Φ_{t-1}) should occur between period $t-1$ and t. In more detail:

$$\begin{aligned}\hat{\mathbb{E}}_{t-1}x_{t-1} \equiv &\phi_{x,t-2}Z_{t-1} = x_{t-1} = \\ &= \hat{\mathbb{E}}_{t-1}\sum_{T=t-1}^{\infty}\beta^{T-(t-1)}\left[(1-\beta)x_{T+1} - \frac{1}{\sigma}\left(i_T - \pi_{T+1} - r_T\right)\right]\end{aligned} \tag{54}$$

It will be argued below that the realization of a strong negative shock in period t and the following anticipation of the ZLB by the agents together cause a wedge between the period t expected output gap (55) and realized output gap (56), where the former will be larger (i.e. too optimistic). This means that the agents will underestimate the negative effect of the ZLB on output. After the negative shock, in the presence of an anticipated binding ZLB ($\hat{\mathbb{E}}^{ZLB}$) it will be shown that:

$$\begin{aligned}\hat{\mathbb{E}}_t x_t \equiv &\phi_{x,t-1}Z_t = \\ &\hat{\mathbb{E}}_t^{no\ ZLB}\sum_{T=t}^{\infty}\beta^{T-t}\left[(1-\beta)x_{T+1} - \frac{1}{\sigma}\left(i_T - \pi_{T+1} - r_T\right)\right] > \qquad (55)\\ &\hat{\mathbb{E}}_t^{ZLB}\sum_{T=t}^{\infty}\beta^{T-t}\left[(1-\beta)x_{T+1} - \frac{1}{\sigma}\left(i_T - \pi_{T+1} - r_T\right)\right] = x_t \qquad (56)\end{aligned}$$

where *no ZLB* refers to a world with the same states Z_t but where the ZLB is not a constraint. Note that the empirical expectations of the learners ($\hat{\mathbb{E}}_t x_t \equiv \phi_{x,t-1} Z_t$) are equivalent to deriving the output gap as if no ZLB is expected (equation (55)) because their period t beliefs matrix Φ_{t-1} (i.e. perceived law of motion) has been derived only during normal times without ever hitting the ZLB up until period $t-1$.

The same argument applies to $\pi_t < \hat{\mathbb{E}}_t \pi_t$, yet, as will be seen, to a smaller degree due to the stronger interest rate channel for the output gap. Since the CB uses model-consistent nowcasts in its policy rule (5), these two results together will ensure that $i_t < \hat{\mathbb{E}}_t i_t$, thus showing the interest rate path disagreement between CB and the agents. Where the latter do not anticipate the structural effects of the ZLB on output gap and inflation but only its effects on the interest rate and hence have biased expectations towards a higher shadow rate and an earlier lift-off.

The following sections prove the agents' forecast bias for output gap, inflation, and finally the interest rate.

B.2.1 ALM at the ZLB - output gap

The learning mechanism ensures $PLM \to ALM$ above the ZLB, that is - the nowcast of endogenous variables equals their realizations after observing the shocks:

$$\hat{\mathbb{E}}_t x_t = x_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma} \left(i_T - \pi_{T+1} - r_T \right) \right] \tag{57}$$

To generalize the proof, for what follows it will be assumed that a future anticipated binding ZLB is $m \geq 0$ periods ahead. Suppose further that between some periods $t+m$ and $t+n$, where $n \geq m \geq 0$, the ZLB is expected to hold, thus the anticipated duration of binding ZLB is $n-m+1$. The agents understand that the interest rate cannot be lower than the ZLB i^*, so they form expectations of the shadow rate and censor it at the ZLB to obtain their actual expected future interest rate between $t+m$ and $t+n$. Even though the agents correctly understand the direct effect of the ZLB on the interest rate, they do not anticipate the changes in the LOM for x and π stemming from the ZLB. That is, they understand the interest rate cannot go below the ZLB and will use such censored interest rate for forecasting x and π but they do not understand that the censoring itself causes the LOMs of x and π (i.e. the coefficients in Γ_t) to be different as well. The CB keeps its reaction function unchanged but, because of these changed LOM for x and π, the LOM in terms of state variables of the policy rate itself changes. Without communication at the ZLB the agents have no way of observing this change and amending their forecasts. This creates a forecast disagreement between the CB and the agents with the former expecting a longer ZLB spell. In what follows, we prove the existence and direction of this disagreement and discuss its reliance on the anticipated ZLB spell duration.

In equation (48) in Appendix B.1 we have established that the ALM of x_t depends on Z_t and $\mathtt{X_{i,t}} \equiv \hat{\mathbb{E}}_t \sum\limits_{T=t}^{\infty} \beta^{T-t} i_T$. We have also assumed that the agents know the autoregressive coefficients of u_t and r_t and they can only update the shocks' impact on endogenous variables through observations. Therefore, it is sufficient to establish how $\mathtt{X_{i,t}}$ changes in the presence of the ZLB to understand how the ALM of x_t changes too.

During the anticipated binding ZLB period[41] from $t+m$ to $t+n$ incl., the expectation of the discounted sum of interest rates is:

$$\sum_{T=t+m}^{t+n} \beta^{T-t} i^* = i^* \beta^m \sum_{T=0}^{n-m} \beta^T = i^* \beta^m \frac{1-\beta^{n-m+1}}{1-\beta} \tag{58}$$

By definition, if no binding ZLB is anticipated, this sum is not defined and may be treated as zero. Thus, we may write the expected discounted sum of future interest rates using the shadow rate corrected for a potentially anticipated ZLB period between periods $t+m$ and $t+n$:

$$\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t} i_T = \begin{cases} \hat{\mathbb{E}}_t \sum\limits_{T=t}^{\infty} \beta^{T-t} s_T - \hat{\mathbb{E}}_t \sum\limits_{T=t+m}^{t+n} \beta^{T-t} s_T + \sum\limits_{T=t+m}^{t+n} \beta^{T-t} i^*, & \text{if } n \geq m \geq 0 \\ \hat{\mathbb{E}}_t \sum\limits_{T=t}^{\infty} \beta^{T-t} s_T, & \text{otherwise} \end{cases} \tag{59}$$

[41]Note that due to the linear forecasting rules agents might anticipate at most one uninterrupted duration of binding ZLB. This is in accordance with typical real-world forecasting.

where s_t is the shadow rate such that $\hat{\mathbb{E}}_t i_T = \max(i^*, \hat{\mathbb{E}}_t s_T)$. In words, the first case corrects the infinite sum of discounted expected interest rates by censoring the region where the ZLB is expected to bind. While in the second, when no ZLB is expected, the expected interest rate coincides with the shadow rate. Importantly, notice that increasing the censored region, either by longer duration of the ZLB $(n - m)$ or a deeper recession (shadow rate much lower than i^*), deviates the sum of discounted expected interest rates further from the sum of expected discounted shadow rates. This will prove a crucial element in understanding the source of disagreement between the agents and the CB.

Since we have already computed the infinite sum of discounted interest rates without a binding ZLB - $\mathtt{X_{i,t}}$ in (44) in Appendix B.1, we know the first term on the right, while the third was computed in (58). Thus, we only need to compute the corrective second term:

$$
\begin{aligned}
\hat{\mathbb{E}}_t \sum_{T=t+m}^{t+n} \beta^{T-t} s_T &= \beta^m \hat{\mathbb{E}}_t s_{t+m} + \beta^{m+1} \hat{\mathbb{E}}_t s_{t+m+1} + \dots = \\
&= \beta^m \left[\hat{\mathbb{E}}_t s_{t+m} + \beta \hat{\mathbb{E}}_t s_{t+m+1} + \dots \right] = \\
&= \beta^m \left[\phi_{it-1} \hat{\mathbb{E}}_t S_{t+m} + \beta \phi_{it-1} \hat{\mathbb{E}}_t S_{t+m+1} + \dots \right] = \\
&= \beta^m \phi_{it-1} \left[\Psi_{t-1}^m S_t + \beta \Psi_{t-1}^{m+1} S_t + \dots + \beta^{n-m} \Psi_{t-1}^n S_t \right] = && (60) \\
&= \beta^m \phi_{it-1} \left[V_{t-1} D_{t-1}^m V_{t-1}^{-1} + \dots + \beta^{n-m} V_{t-1} D_{t-1}^n V_{t-1}^{-1} \right] S_t = && (61) \\
&= \beta^m \phi_{it-1} V_{t-1} D_{t-1}^m \left[I + \beta D_{t-1} + \dots + \beta^{n-m} D_{t-1}^{n-m} \right] V_{t-1}^{-1} S_t && (62)
\end{aligned}
$$

Plugging in the equation for D_{t-1} from (42) into (62) and computing the series in the parentheses we obtain:

$$\hat{\mathbb{E}}_t \sum_{T=t+m}^{t+n} \beta^{T-t} s_T = \beta^m \phi_{it-1} V_{t-1} D_{t-1}^m \left[I + \beta D_{t-1} + \cdots + \beta^{n-m} D_{t-1}^{n-m} \right] V_{t-1}^{-1} S_t =$$

$$= \beta^m \phi_{it-1} V_{t-1} D_{t-1}^m \begin{bmatrix} \frac{1-\beta^{n-m+1}}{1-\beta} & 0 & 0 & 0 \\ 0 & \frac{1-(\beta\rho_r)^{n-m+1}}{1-\beta\rho_r} & 0 & 0 \\ 0 & 0 & \frac{1-(\beta\rho_u)^{n-m+1}}{1-\beta\rho_u} & 0 \\ 0 & 0 & 0 & \frac{1-(\beta\phi_{ii})^{n-m+1}}{1-\beta\phi_{ii}} \end{bmatrix} V_{t-1}^{-1} S_t =$$

$$= \phi_{it-1} V_{t-1} \begin{bmatrix} \frac{\beta^m-\beta^{n+1}}{1-\beta} & 0 & 0 & 0 \\ 0 & \frac{(\beta\rho_r)^m-(\beta\rho_r)^{n+1}}{1-\beta\rho_r} & 0 & 0 \\ 0 & 0 & \frac{(\beta\rho_u)^m-(\beta\rho_u)^{n+1}}{1-\beta\rho_u} & 0 \\ 0 & 0 & 0 & \frac{(\beta\phi_{ii})^m-(\beta\phi_{ii})^{n+1}}{1-\beta\phi_{ii}} \end{bmatrix} V_{t-1}^{-1} S_t =$$

$$= \begin{bmatrix} \frac{\phi_{ic}}{1-\phi_{ii}} \left[\frac{\beta^m-\beta^{n+1}}{1-\beta} - \frac{\phi_{ii}\left((\beta\phi_{ii})^m-(\beta\phi_{ii})^{n+1}\right)}{1-\beta\phi_{ii}} \right] \\ \frac{\phi_{iu}}{\rho_u-\phi_{ii}} \left[\frac{\rho_u}{1-\beta\rho_u} \left((\beta\rho_u)^m - (\beta\rho_u)^{n+1}\right) - \frac{\phi_{ii}}{1-\beta\phi_{ii}} \left((\beta\phi_{ii})^m - (\beta\phi_{ii})^{n+1}\right) \right] \\ \frac{\phi_{ir}}{\rho_r-\phi_{ii}} \left[\frac{\rho_r}{1-\beta\rho_r} \left((\beta\rho_r)^m - (\beta\rho_r)^{n+1}\right) - \frac{\phi_{ii}}{1-\beta\phi_{ii}} \left((\beta\phi_{ii})^m - (\beta\phi_{ii})^{n+1}\right) \right] \\ \frac{\phi_{ii}}{1-\beta\phi_{ii}} \left((\beta\phi_{ii})^m - (\beta\phi_{ii})^{n+1}\right) \end{bmatrix}' \begin{bmatrix} 1 \\ u_t \\ r_t \\ s_{t-1} \end{bmatrix} \tag{63}$$

For cross-reference, it can be easily seen that setting $m = 0$ and $n = \infty$, i.e. perpetual expected ZLB starting today, the correction factor in (63) yields exactly the result for the sum of discounted expected interest rates in normal times in (44). This makes sense because the entire shadow rate path needs to be substituted with an effective interest rate peg at i^*. Notice also that the correction vector in (63) is increasing in the ZLB duration $n - m + 1$.[42]

Since the shadow rate reacts positively to all state variables, that is - $\phi_i > 0$, it means that $\hat{\mathbb{E}}_t \sum_{T=t+m}^{t+n} \beta^{T-t} s_T$ will also be a positive function of each state variable in Z_t. This can be verified by considering each term of the correction vector in (63). It can be shown that for either case $\phi_{ii} \lesseqgtr \rho_u$ or $\phi_{ii} \lesseqgtr \rho_r$ all terms are always positive. Therefore, the component containing the sum of discounted expected interest rates in the output gap

[42] This again can be shown to hold for either case $\phi_{ii} \lesseqgtr \rho_u$ or $\phi_{ii} \lesseqgtr \rho_r$.

ALM - $\mathtt{X_{i,t}}$, is corrected negatively for each state variable in the presence of a binding ZLB. Remember that from (48) above the ZLB and the end of Appendix B.1:

$$\begin{aligned} x_t &= \hat{\mathbb{E}}_t^{noZLB} \sum_{T=t}^{\infty} \beta^{T-t} \left[(1-\beta) x_{T+1} - \frac{1}{\sigma} \left(i_T - \pi_{T+1} - r_T \right) \right] \\ &= (1-\beta)\mathtt{X_{x,t}} - \frac{1}{\sigma}\mathtt{X_{i,t}} + \frac{1}{\sigma}\mathtt{X_{\pi,t}} + \frac{r_t}{\sigma(1-\beta\rho_r)} = x(Z_t, \mathtt{X_{i,t}}) \end{aligned} \tag{64}$$

Because the ZLB correction term (63) negatively affects $\mathtt{X_{i,t}}$ (from (59)), it means that the presence of a binding ZLB results in unambiguously larger responses of x_t to all state variables in Z_t, confirming the intuition that once the policy rate cannot fully accommodate shocks, the latter end up hitting harder and lead to a harsher recession. The ZLB correction thus ensures that the inequality in (55) holds, i.e. $\boxed{\hat{\mathbb{E}}_t x_t > x_t}$.

B.2.2 ALM at the ZLB - inflation

From (53) in Appendix B.1.2 we know that:

$$\begin{aligned} \pi_t &= \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left[\kappa x_T + (1-\alpha)\beta\pi_{T+1} + u_T \right] \\ &= \kappa \Pi_{\mathtt{x,t}} + (1-\alpha)\beta \Pi_{\pi,\mathtt{t}} + \frac{1}{(1-\alpha\beta\rho_u)} u_t = \pi(Z_t, \Pi_{i,t}) \end{aligned}$$

The equivalent sum of discounted interest rates here is:

$$\hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} i_T = \begin{cases} \hat{\mathbb{E}}_t \sum\limits_{T=t}^{\infty} (\alpha\beta)^{T-t} s_T - \hat{\mathbb{E}}_t \sum\limits_{T=t+m}^{t+n} (\alpha\beta)^{T-t} s_T + \sum\limits_{T=t+m}^{t+n} (\alpha\beta)^{T-t} i^*, & \text{if } n \geq m \geq 0 \\ \hat{\mathbb{E}}_t \sum\limits_{T=t}^{\infty} (\alpha\beta)^{T-t} s_T, & \text{otherwise} \end{cases} \tag{65}$$

The eigen decomposition of the ZLB correction term is analogous to (63) and similarly the presence of a binding ZLB results in lower reaction of $\Pi_{\mathtt{i,t}}$ to the state variables vector Z_t compared to normal times. Since $\alpha, \beta, \kappa \in (0,1)$ and $\phi_{x,i} < 0$ and $\phi_{\pi,i} < 0$, $\Pi_{\mathtt{i,t}}$ enters negatively in π and hence the ZLB correction means that π_t is more reactive to the state variable vector Z_t. Note, however, that due to $\kappa, \alpha < 1$ the correction in π_t is much smaller than that in x_t meaning that the agents are making larger forecast errors on output gap forecasts than inflation forecasts.

B.2.3 ALM at the ZLB - policy rate

The CB observes the realizations of output gap and inflation and sets its shadow policy rate according to:

$$s_t = \rho_i s_{t-1} + (1 - \rho_i)(\chi_\pi \pi_t + \chi_x x_t) \tag{66}$$

Naturally, $\chi_x, \chi_\pi > 0$, thus the ZLB corrections in x_t and π_t from the previous sections imply that CB shadow rate reacts more strongly to the state variable vector Z_t. In particular, if a negative demand or supply shock drives the interest rate to the ZLB, then the CB would act to adjust its shadow rate and skew its trajectory towards lower values, thus prolonging the a priori expected ZLB duration. In essence, the CB fully understands the indirect impacts of the ZLB on output gap and inflation and chooses lower policy rates when the ZLB is binding despite its policy coefficients χ_x and χ_π remaining unchanged. This creates disagreement between the CB and the agents who expect the policy rate to still follow its pre-ZLB trajectory. That is, the agents expect an earlier lift-off due to the non-linear correction to the ALM as a result of the ZLB. □

B.3 Calibration

B.3.1 Structural parameters

The equilibrium equations of the New Keynesian (NK) model in (1-5) can be fully calibrated by the slope of the Phillips curve - κ, the inverse inter-temporal elasticity of substitution - σ, the discount factor - β, the two shock processes, and the three monetary policy rule parameters - ρ_i, χ_π, χ_x. Finally, the only free learning parameter is the learning gain - τ, in (15).

Common benchmark calibrations in small-scale New Keynesian models place the slope of the quarterly output-gap in the NK Phillips curve between 0.02 and 0.05.Adam and

Billi (2006), Nakov (2008), Levin et al. (2010), and Bodenstein et al. (2012) use $\kappa = 0.024$ as in Rotemberg and Woodford (1997) and Woodford (2003); similarly, Eggertsson and Woodford (2003) use 0.02 and Christiano et al. (2011) use approximately 0.03. Chari et al., 2000 's calibration implies 0.085 (see Gáti (2023)), whilst Hazell et al. (2022) estimate a flat Phillips curve slope to *unemployment* of 0.0062, which assuming an Okun's law coefficient of 2, implies $\kappa = 0.0031$. Recent microeconomic evidence from Gagliardone et al. (2025) implies output-gap slopes of between 0.006 and 0.021. The textbook calibration of a similar small-scale NK model in Galí (2008) implies a relative steep $\kappa = 0.1275$. In a behavioral NK model with cognitive discounting Gabaix (2020) uses $\kappa = 0.11$. Thus, the benchmark calibration here is chosen at $\kappa = 0.024$.For further reading, Furlanetto and Lepetit (2025) provide a recent survey of the identification and estimation of the Phillips curve. They conclude that conventional evidence tends to exaggerate the extent of pre-pandemic flattening of the Phillips curve slope. They also review post-pandemic evidence consistent with nonlinear or state-dependent steepening when inflation is high or capacity constraints bind.[43]

Calibrations of the inverse inter-temporal elasticity of substitution, σ, vary substantially in the literature. Log utility, implying σ=1, is the standard benchmark in canonical New Keynesian models (Galí, 2015; Eusepi and Preston, 2018). Eggertsson and Woodford (2003) use an inter-temporal elasticity of substitution (IES) - $1/\sigma = 0.5$, implying $\sigma = 2$, whilst Lubik and Schorfheide (2004) estimate a mean value of 1.86 with 90% confidence interval of [1.04, 2.64] for the post-Volcker era, and Nakov (2008) uses $\sigma = 4$. Other papers use values which make output much more sensitive to the real interest rate. Namely, Adam and Billi (2006) follow Rotemberg and Woodford (1997) who estimate $\sigma = 0.16$ implying an output sensitive to the real rate of 6.25. The baseline calibration here is set at $\sigma = 1$.

For the monetary policy rule parameters, I focus on the post-Volcker evidence for the

[43]Note that Table 1 in Furlanetto and Lepetit (2025) reports slopes of a Phillips curve relative to *unemployment*. These estimates therefore cannot be compared directly with *kappa*. A quick conversion using Okun's law supports the chosen range for the sensitivity analysis.

United States. Remember that the policy rule is defined in (5) as

$$
\begin{aligned}
s_t &= \rho_i s_{t-1} + (1-\rho_i)(\chi_\pi \pi_t + \chi_x x_t) \\
i_t &= \max\{i^*, s_t\}
\end{aligned}
\tag{67}
$$

Thus, χ_π and χ_x denote the long-run coefficients for the policy reaction to inflation and output gap, respectively, and they are independent from the policy inertia parameter ρ_i.

Empirical estimates generally imply substantial quarterly interest rate smoothing, with policy inertia values around 0.8-0.9. Clarida et al. (2000) and Lubik and Schorfheide (2004) estimate ρ_i at 0.79 and 0.84, respectively. In a similar model to here, Cho and Moreno (2006) use different output gap measurements and across specifications estimate $\rho_i \in [0.84, 0.88]$. In a modern take on OLS of Taylor rules, Carvalho et al. (2021) estimate ρ_i to be 0.74 or 0.81 when using real-time data and internal Federal Reserve Greenbook forecasts, respectively. Andreasen and Kronborg (2022)'s multiple specifications support a wide range of high policy inertia values spanning $\rho_i = 0.80$ for the post-Volcker period and up to $\rho_i = 0.97$ in the full sample since 1959. Similarly Coibion and Gorodnichenko (2012) estimate the inertia between 0.80 and 0.87 depending on the specification, favoring genuine inertial policy making over alternative explanations for the observed persistence of the policy rate. Belongia and Ireland, 2022 estimation of a NK model with various specifications consistently yielded the policy inertia around 0.85. Thus, the benchmark calibration is set at 0.85 as in English et al. (2015), Chung et al. (2019), and Bernanke (2020).

Evidence on the post-Volcker monetary policy reaction to inflation implies a long-run coefficient confidently above 1. Clarida et al. (2000) and Lubik and Schorfheide (2004) estimate χ_π at 2.15 and 2.19, respectively, whilst Cho and Moreno (2006) estimate it between 1.64 and 2.15 depending on the output gap specification. In a larger estimated model Smets and Wouters (2007) find a posterior of 2.03, although their rule also responds to changes in the output gap. More recent structural estimates also support substantial inflation responses. Carvalho et al. (2021) estimate $\chi_\pi = 1.70$ using Greenbook forecasts and 1.24 using real-time data. Using OLS and IV, Andreasen and Kronborg (2022) obtain χ_π around 1.5 but starting from 1985 their estimate rises to 3.66. Belongia and Ireland (2022) obtain an implied long-run response of approximately 1.41 in their baseline lagged-information rule and 1.73 when contemporaneous inflation is used. Accordingly, the benchmark remains $\chi_\pi = 1.5$, as is standard in the literature.

Note that in the model the interest rate and inflation are expressed as quarterly rates, and the output gap is a quarterly log-deviation from potential output. Many empirical studies use interest rate and inflation as annualized rates, whilst retaining the output gap definition. When translating such estimates into the convention used here, the inertia and inflation-response coefficients remain unchanged, whereas the reported output-gap coefficient is divided by four. The estimates of Cho and Moreno (2006), Clarida et al. (2000), and Lubik and Schorfheide (2004) correspond to χ_x values of 0.15, 0.23, and 0.3, respectively, whilst Carvalho et al. (2021) estimate different specifications between 0.20-0.27. Smets and Wouters (2007) obtain a smaller output gap level response of 0.08, but their rule supplements it with a response of 0.22 to changes in the output gap. Using a richer model featuring habit formation, money, etc., Belongia and Ireland (2022) estimate rather high output gap reaction coefficients of 1.47 when using lagged variables in the policy rule and 2.023 when using contemporaneous variables. I therefore retain the standard $\chi_x = 0.25$ benchmark.

In a comprehensive survey, Berardi and Galimberti (2017) show that most estimates for a quarterly single constant gain in the literature lie between 0 and 0.1 with most common estimations (and calibrations) generally landing between 0.01 and 0.05. Specifically, Milani (2007, 2011) estimates values close to 0.02 and Orphanides and Williams (2007) find that gains between 0.01 and 0.04 best reproduce professional forecasts and choose a baseline of 0.02. Molnár and Santoro (2014) settle on a benchmark of 0.03, whilst Gáti (2023) chooses a slightly higher one of 0.05. The benchmark here is set at $\tau = 0.03$.

B.3.2 Empirical moments, target bands, and calibration procedure

I discipline the natural-rate and cost-push processes with moments of ordinary macroeconomic fluctuations and lower-bound episodes. The macroeconomic moments are computed from quarterly U.S. data for 1987Q1 - 2007Q4. This sample begins after the Volcker disinflation and ends before the Great Recession, separating the calibration of routine innovations from the large crisis shock used in the quantitative experiments. Real activity is measured by real GDP relative to CBO potential real GDP (FRED series `GDPC1` and `GDPPOT`), and inflation by the core PCE price index (`JCXFE`). In model units,

$$x_t = \log\left(\frac{\text{GDPC1}_t}{\text{GDPPOT}_t}\right), \qquad \pi_t = \log(\text{JCXFE}_t) - \log(\text{JCXFE}_{t-1}). \tag{68}$$

Thus, x_t is the log output gap and π_t is quarterly, rather than annualized, inflation.

Over the 84-quarter sample, the standard deviations of x_t and π_t are 0.01253 and 0.002421, and their first-order autocorrelations are 0.9247 and 0.8235, respectively. I use bands around these estimates rather than point targets: $\text{sd}(x) \in [0.0100, 0.0150]$, $\text{sd}(\pi) \in [0.0019, 0.0031]$, $\text{AR}(x) \in [0.78, 0.97]$, and $\text{AR}(\pi) \in [0.70, 0.90]$. The tolerance bands accommodate sampling and data-vintage uncertainty and the deliberately parsimonious structure of the model.

The persistence targets receive less weight than the volatility targets. Canonical forward-looking New Keynesian models require additional propagation mechanisms to reproduce the observed inertia of output and inflation (Fuhrer, 2000), while adaptive learning can itself generate persistence (Milani, 2007). Since I calibrate on the rational-expectations model, I place a lower weight on persistence which the learning framework will naturally add.

For lower-bound incidence and duration, I use the post-war international evidence rather than the single U.S. episode. Dordal i Carreras et al. (2016) report an unconditional ZLB frequency of approximately 7.5 percent and an average duration of 14.2 quarters across advanced economies; alternative sample treatments place the frequency near 6-11 percent and mean duration between 11.5 and 17.5 quarters. I accordingly target a

ZLB frequency of 6-10 percent and a mean duration of 12-18 quarters. These are broad benchmarks because lower bound moments depend strongly on the shock process, the steady-state nominal rate, and the policy rule (Chung et al., 2012; Dordal i Carreras et al., 2016; Kiley and Roberts, 2017).

The next subsection describes how these six bands select the ordinary shock processes. The large natural-rate innovation used to initiate the Great Recession experiment is calibrated separately and therefore does not affect the unconditional targets above.

B.3.3 Calibration procedure and results

The natural-rate and cost-push shocks follow

$$r_t = \rho_r r_{t-1} + \varepsilon_t^r, \qquad u_t = \rho_u u_{t-1} + \varepsilon_t^u, \qquad \varepsilon_t^r \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_r^2), \quad \varepsilon_t^u \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_u^2). \tag{69}$$

I first evaluate these processes in the linear rational-expectations version of the model, holding the structural parameters at their benchmark values in Appendix B.3.1. Because the auxiliary model is linear, its rational-expectations solution is exact. The effective lower bound is not imposed in this exercise. Instead, a shadow-ZLB quarter is one in which the unconstrained policy rate satisfies $s_t \leq i^*$, where $i^* = 1 - 1/\beta$ in deviations from the steady-state nominal rate. Consecutive shadow-ZLB quarters form an episode; frequency is the fraction of usable quarters at the bound, and duration is the mean length of completed episodes.

The search grid is

$$\begin{aligned}
&\rho_r \in \{0.80, 0.85, 0.90, 0.91, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97\},\\
&\rho_u \in \{0.20, 0.40, 0.60, 0.80\},\\
&\sigma_r \in \{0.00100, 0.00125, 0.00150, 0.00175, 0.00200, 0.00225, 0.00250\},\\
&\sigma_u \in \{0.00025, 0.00050, 0.00075, 0.00100, 0.00125\},
\end{aligned} \tag{70}$$

for a total of 1,400 candidates. Each candidate is simulated for 100,000 quarters after a 5,000-quarter burn-in. The innovations are kept the same across all candidates which

Table B.1: Empirical targets and auxiliary rational-expectations calibration

Moment	Empirical anchor	Target band	Linear RE model
sd(x)	0.01253	[0.0100, 0.0150]	0.01152
sd(π)	0.002421	[0.0019, 0.0031]	0.003069
AR(x)	0.9247	[0.78, 0.97]	0.7969
AR(π)	0.8235	[0.70, 0.90]	0.8952
ZLB frequency	0.075	[0.06, 0.10]	0.0980
Mean ZLB duration (quarters)	14.2	[12, 18]	12.08

makes comparisons across nearby grid points less sensitive to simulation noise.

Let $m_j(\vartheta)$ denote element j of the simulated moment vector for $\vartheta = (\rho_r, \rho_u, \sigma_r, \sigma_u)'$, and let $[\ell_j, h_j]$ be its target band. Define the normalized distance to the band as

$$D_j(m_j) = \begin{cases} (\ell_j - m_j)/(h_j - \ell_j), & m_j < \ell_j, \\ 0, & \ell_j \leq m_j \leq h_j, \\ (m_j - h_j)/(h_j - \ell_j), & m_j > h_j. \end{cases} \tag{71}$$

The selection criterion is

$$Q(\vartheta) = \sum_{j=1}^{6} w_j D_j(m_j(\vartheta))^2, \qquad (w_1, \ldots, w_6) = (1, 1, 0.2, 0.2, 1, 1), \tag{72}$$

where the moment ordering is $(\text{sd}(x), \text{sd}(\pi), \text{AR}(x), \text{AR}(\pi), f_{\text{ZLB}}, \bar{d}_{\text{ZLB}})$. Thus, moments inside their bands incur no loss, while the two persistence moments receive one fifth of the weight assigned to each remaining target. A normalized distance from the band midpoints breaks ties in $Q(\vartheta)$.

Exactly one grid point lies within all six target bands:

$$\rho_r = 0.96, \qquad \rho_u = 0.20, \qquad \sigma_r = 0.00125, \qquad \sigma_u = 0.00050. \tag{73}$$

Table B.1 reports its fit together with the empirical anchors underlying the bands.

I then refine the innovation standard deviations in the full model with adaptive learning and an occasionally binding lower bound, keeping the two persistence parameters

Table B.2: Full-model validation of the selected shock calibration

Moment	Target band	Model mean	Model median
sd(x)	[0.0100, 0.0150]	0.02066	0.01429
sd(π)	[0.0019, 0.0031]	0.003354	0.002883
AR(x)	[0.78, 0.97]	0.8382	0.8379
AR(π)	[0.70, 0.90]	0.9031	0.9043
ZLB frequency	[0.06, 0.10]	0.0837	0.0751
Pooled ZLB duration (quarters)	[12, 18]	12.91	9

fixed. The refinement uses 100 simulations of 1,000 quarters and common standardized innovations across candidate calibrations. All 100 simulations used for the reported moments remain stable through the full horizon. For volatility, persistence, and ZLB frequency, Table B.2 reports the mean and median of path-specific statistics. The duration entries are instead computed from the 641 completed episodes pooled across simulations.

With $\sigma_r = 0.00100$, mean output and inflation volatility are 0.01375 and 0.00276, but the ZLB frequency falls to 0.0599. Raising σ_r to 0.00110 increases ZLB frequency to 0.0837 while retaining a pooled mean duration of 12.91 quarters. I select the latter value and set $\sigma_u = 0.00025$ to contain inflation volatility. The final calibration is therefore

$$\rho_r = 0.96, \qquad \rho_u = 0.20, \qquad \sigma_r = 0.00110, \qquad \sigma_u = 0.00025. \tag{74}$$

The mean-median gap in macroeconomic volatility reflects a right tail of occasional non-linear amplification under learning and the binding lower bound: the median simulation lies within the volatility bands, whereas the means are pulled above them. The selected process balances this tail risk against an empirically relevant incidence and duration of ZLB episodes. It is therefore a disciplined calibration for the paper's quantitative laboratory, rather than a formal minimum-distance estimate. I hold the parameter vector in equation (74) fixed in all subsequent experiments.

The common Great Recession shock is calibrated in one dimension to reproduce the magnitude of the U.S. output contraction. Setting $\varepsilon_2^r = -0.008$ produces a median output-gap trough of approximately -5% , close to the trough measured from real GDP relative to CBO potential output. The remaining features of the recession are not targeted and therefore provide useful validation. Median inflation falls by approximately 0.75 percentage point at a quarterly rate, closely matching the decline in core-PCE inflation between 2007Q4 and 2008Q4. interest rate smoothing delays arrival at the lower bound until approximately three to four quarters after the shock, consistent with the gradual policy easing between the onset of the recession in 2007Q4 and the federal funds rate reaching its effective lower bound in 2008Q4. Output and inflation subsequently return to their steady states by approximately period 28.

The initial lower-bound spell lasts approximately 10 quarters in the median no-communication simulation. This is shorter than the historical U.S. episode, but a literal duration match would be inappropriate in this parsimonious experiment. The observed episode reflected not only the initial financial crisis shock but also subsequent shocks as well as policy interventions such as quantitative easing and forward guidance. These forces are absent from the no-communication calibration and are likely to have prolonged the duration of the binding lower bound. The calibration instead reproduces the magnitude and timing of the recession and generates a sustained lower-bound episode long enough for the model's disagreement mechanism to operate.

Table B.3: Calibration

parameter	benchmark
β	0.9925
κ	0.024
σ	1
ρ_i	0.85
χ_π	1.5
χ_x	0.25
τ	0.03
ρ_r	0.96
ρ_u	0.20
σ_r	0.0011
σ_u	0.00025
ε_2^r	-0.008

C Numerical simulations

C.1 Figures

Figure C.1: No FG: CB remaining ZLB duration conditional on $T^{ag} = 1$

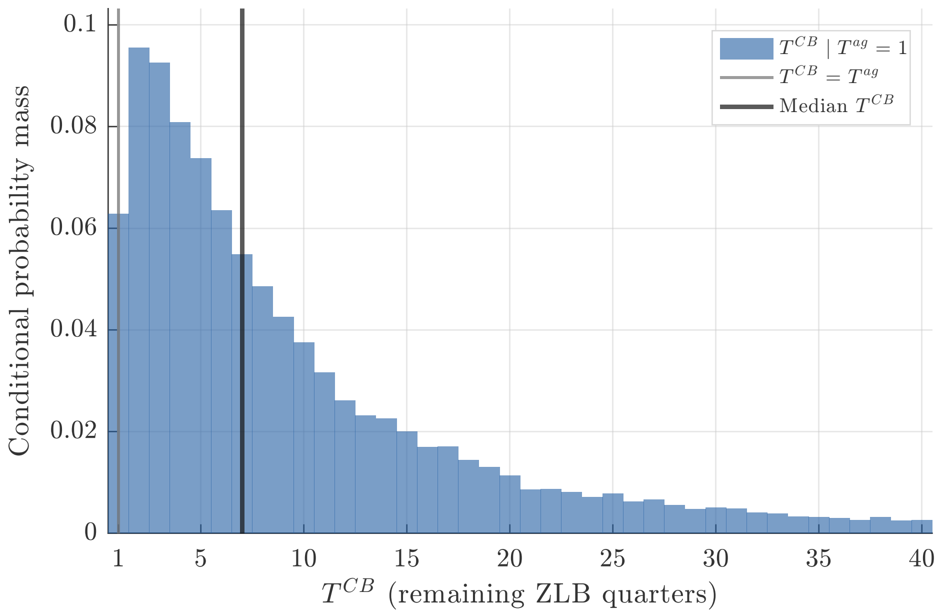

Notes: T^{CB} is the CB's expected number of remaining binding quarters, including the current quarter. The conditioning event $T^{ag} = 1$ means that agents expect lift-off next quarter.

Figure C.2: PLM components evolution across communication regimes

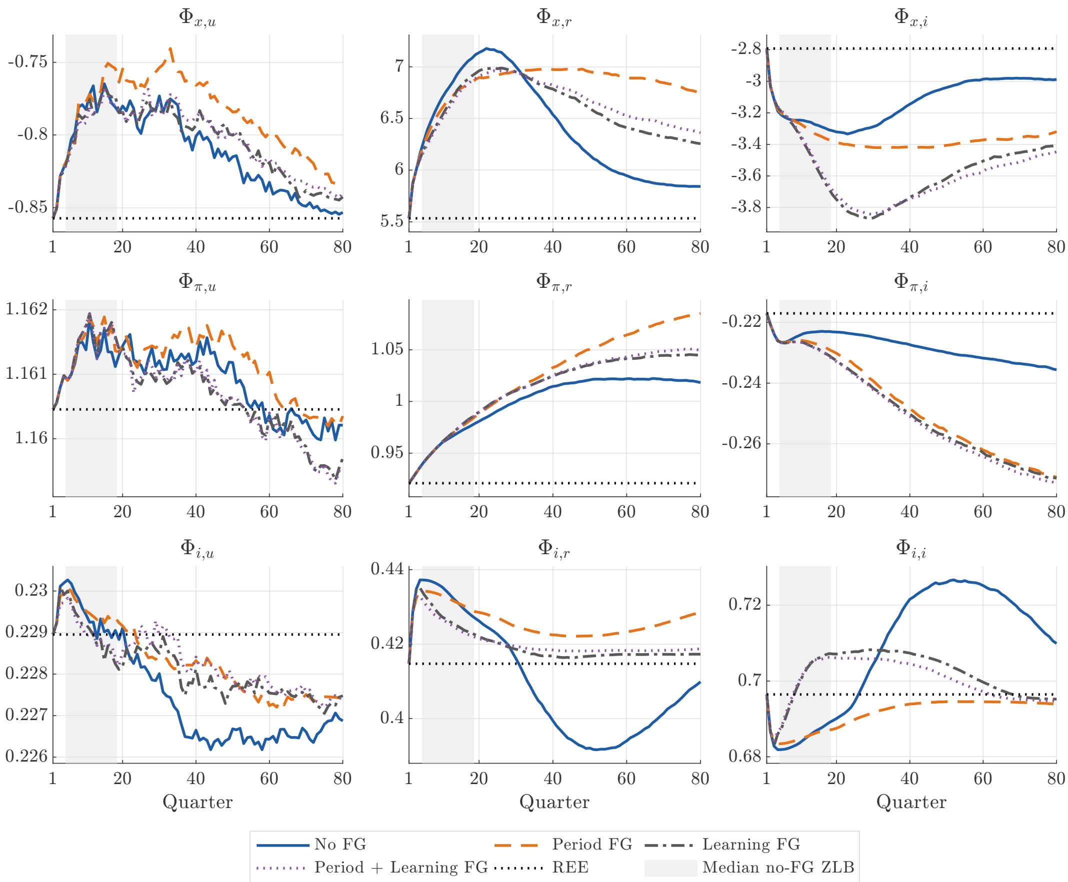

Cross-path medians on the changing risk set. Dotted line: REE initialization. Grey shading: quarters in which the median no-FG policy rate is at the ZLB.

Notes: Lines report cross-path medians of the accepted PLM coefficients over 10,000 matched simulations. All regimes use the same shocks and REE initialization; no PF censoring is imposed.

Figure C.3: Distributional timeline of ϕ_{ii} accepted values

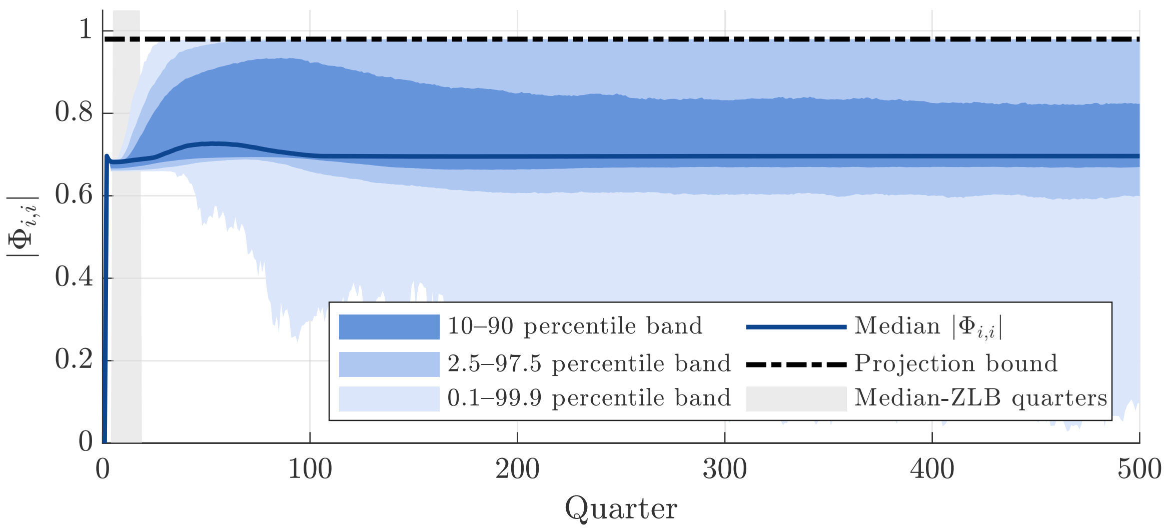

Figure C.4

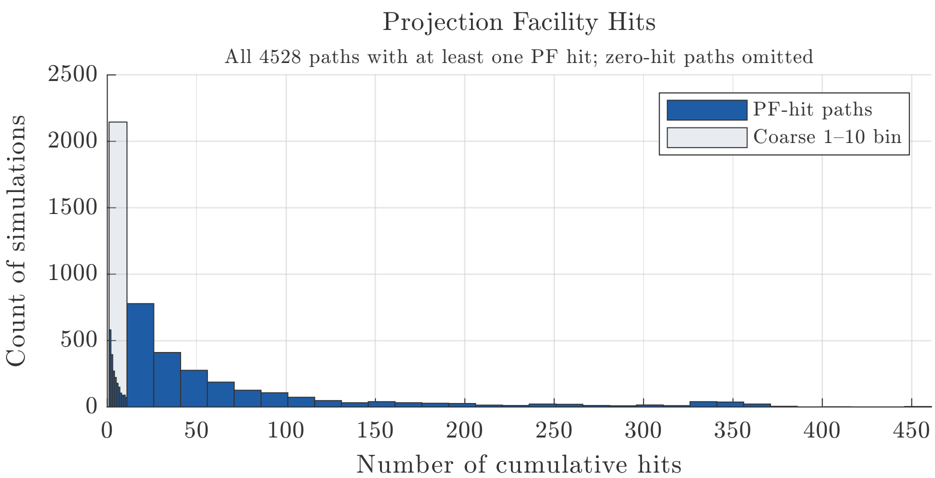

Figure C.5: Accepted persistence and rejected candidates at the first projection event

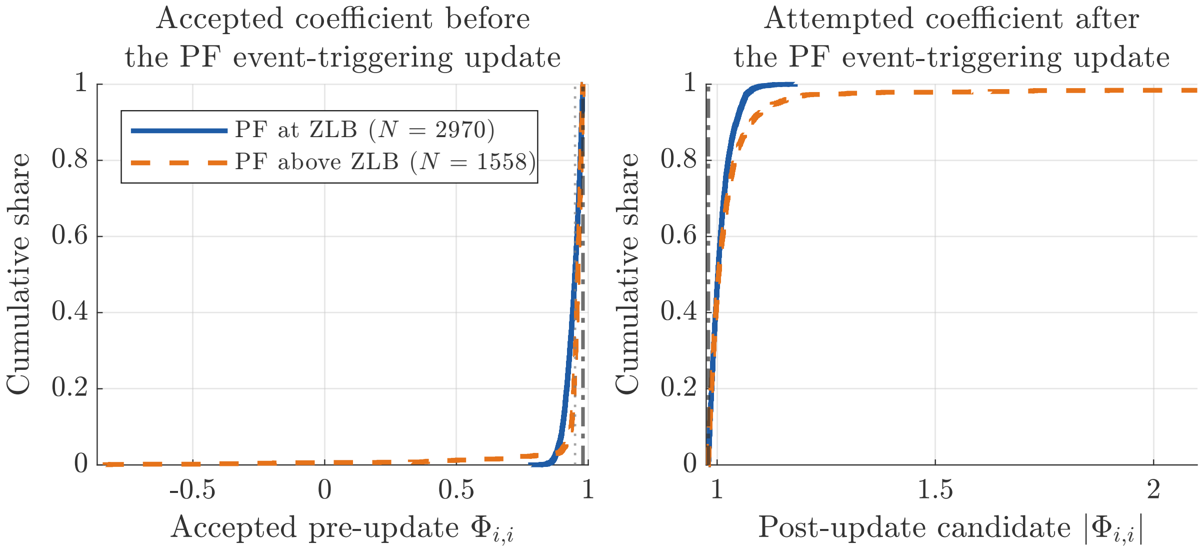

Notes: The left panel reports the empirical distributions of the accepted pre-update coefficient. The right panel reports the absolute value of the event-triggering candidate. The projection test is applied to the latter object. Every finite candidate displayed in the right panel exceeds the bound of 0.98 (dash-dot line).

Figure C.6: Period FG: Remaining ZLB duration

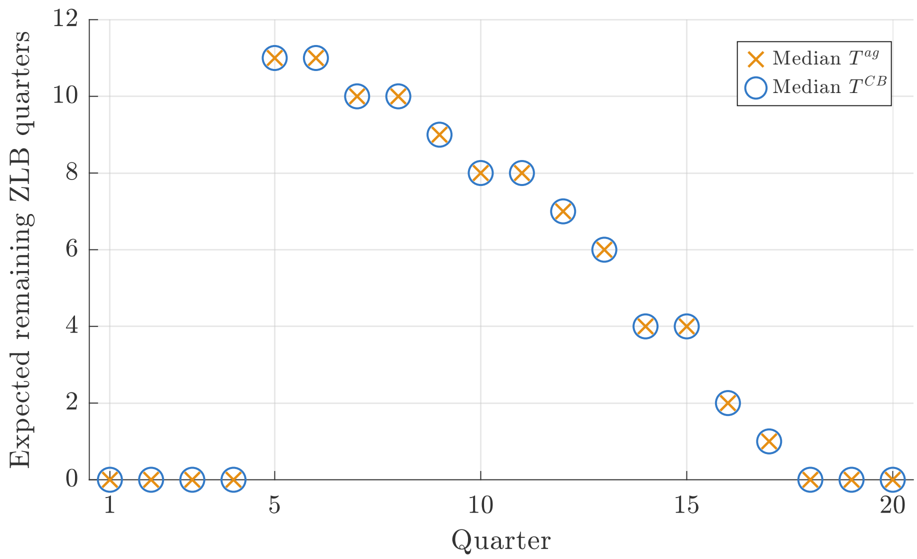

Notes: Medians use all 10,000 matched paths. T^{ag} is the operative post-communication duration, including the current binding quarter; accepted Period FG announcements impose T^{CB} directly.

Figure C.7: Learning FG: Remaining ZLB duration

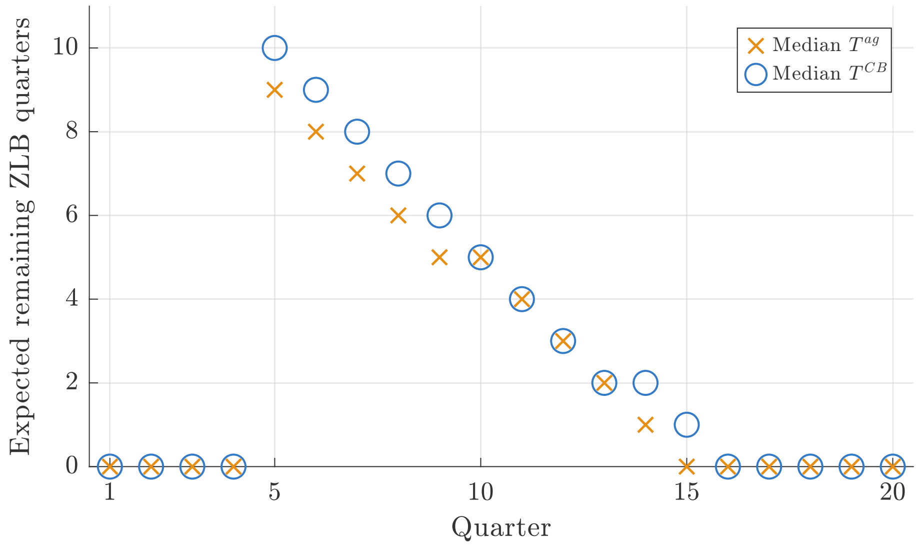

Notes: Medians use all 10,000 matched paths. Learning FG updates the interest rate PLM from the announcement but does not impose the communicated duration mask.

Figure C.8: Period + Learning FG: Remaining ZLB duration

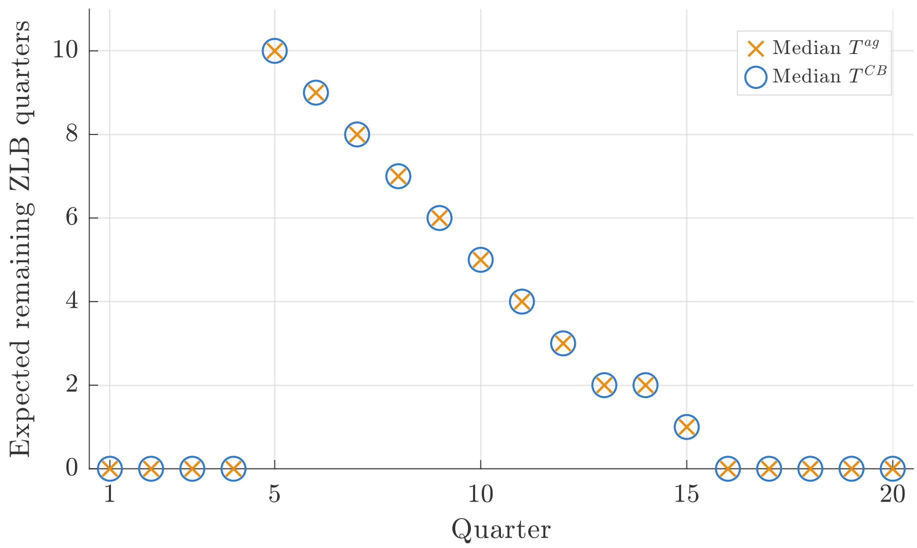

Notes: Medians use all 10,000 matched paths. The regime combines the announcement PLM update with the communicated duration mask whenever the update is accepted.

Figure C.9: Cumulative incidence of the first projection event across communication regimes

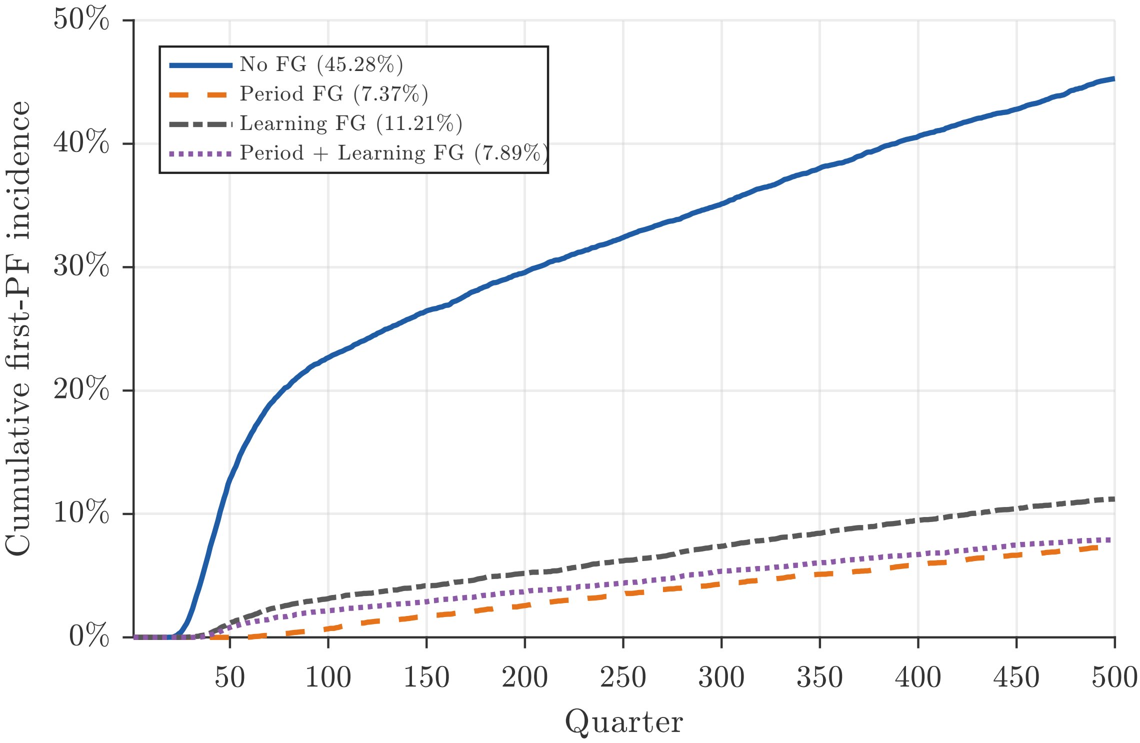

Notes: Each line is the unconditional share of 10,000 matched paths whose first PF event has occurred by the indicated quarter. A PF event is a rejected candidate update satisfying $|\widetilde{\phi}_{ii,t}| > 0.98$.

C.2 Simulation statistics with and without a GFC shock

This section compares the baseline No FG simulations with and without the initial GFC negative demand shock. In both cases agents are assumed to have converged to the REE beliefs $\Phi_1 = \Gamma$. Apart from the initial GFC shock, all simulations are identical and share the baseline calibration. As described in Appendix B.3.3, the GFC shock occurs in the first quarter and induces a 5% drop in the output gap and around 3% drop in the annualized inflation rate.

The imposed GFC raises the probability of at least one projection event from 29.86% to 45.28% (Table C.1). Events at the first ZLB spell rise from 1,080 without the GFC to 1,402 with it, later-ZLB events from 1,133 to 1,568, and outside-ZLB events from 773 to 1,558. The first spells are not the same objects across treatments: of the 1,080 no-GFC first-spell events, 601 occur in a later spell once the GFC is imposed, while 254 paths have first-spell events in both treatments. Thus, the initial GFC shock changes the subsequent evolution of beliefs well beyond the initial ZLB spell.

Table C.1: No FG: projection-facility incidence with and without an imposed GFC shock

	GFC	No GFC
Paths	10,000	10,000
At least one PF	4,528 (45.28%)	2,986 (29.86%)
First PF during the first ZLB	1,402 (14.02%)	1,080 (10.80%)
First PF during a later ZLB	1,568 (15.68%)	1,133 (11.33%)
First PF outside the ZLB	1,558 (15.58%)	773 (7.73%)
Median first-PF quarter	100	270

Notes: Of the PF events, 2,703 paths hit in both treatments, 1,825 hit only with the GFC, and 283 hit only without it. 601 of the 1,080 no GFC first-ZLB projection events are later-ZLB events once the GFC is imposed; 254 paths are first-ZLB projection events in both treatments.

Beliefs and second-moment matrices at entry, reported in Tables C.2 and C.3, explain why a first spell that begins after a long unconstrained history is more fragile than a first spell that begins immediately after the GFC impulse.

Recursive least squares updates

$$R_t = (1-\tau)R_{t-1} + \tau Z_t Z_t', \qquad Z_t = (u_t, r_t, i_{t-1})',$$

do not preserve the rational-expectations initialization R_0. The recursion does not depend on forecast errors. With $\tau = 0.03$, the weight remaining on that initialization after n updates is $(1-\tau)^n$. At the GFC first-ZLB date that weight is still about 0.94; at the no-GFC median of 125 quarters it is 0.024. Over the no-GFC pre-ZLB sample, forecast errors remain negligible, so Φ stays at its rational-expectations value (last column in Table C.2). R_t nonetheless adapts toward the realized second moments of Z_t. Median perceived volatility of the interest rate, R_{ii}, at entry is 64% of the linear-RE variance among all no-GFC paths that reach the ZLB, and 48% among those that subsequently hit the projection facility during that first spell (Table C.3). The first premature-liftoff forecast errors therefore meet a substantially larger RLS gain, $R_t^{-1}Z_t$, than the initialization would have implied.

The GFC treatment produces the opposite entry configuration. The first ZLB arrives after only a few quarters, so R_t remains near its rational-expectations initialization: the median R_{ii} ratio is 0.96 both overall and among paths whose first PF event occurs during that spell. Instead, Φ moves in response to the large natural-rate impulse. At first-ZLB entry, median $\Phi_{i,i}$ is 0.684 rather than its REE value of 0.696, moving initially away from the projection threshold (Table C.2). The ALM nevertheless incorporates agents' expected future ZLB path, so the shock affects output and inflation before the constraint binds and shifts the other interest rate PLM loadings.

Table C.2: Perceived law of motion on entry to the first ZLB

Coefficient	RE	GFC	GFC−RE	No GFC	No GFC−RE
$\Phi_{x,u}$	−0.857	−0.851	0.006	−0.857	0
$\Phi_{\pi,u}$	1.160	1.161	0	1.160	0
$\Phi_{i,u}$	0.229	0.229	0	0.229	0
$\Phi_{x,r}$	5.533	5.978	0.445	5.547	0.013
$\Phi_{\pi,r}$	0.921	0.931	0.010	0.921	0.001
$\Phi_{i,r}$	0.415	0.434	0.019	0.415	0.001
$\Phi_{x,i}$	−2.792	−3.090	−0.298	−2.795	−0.003
$\Phi_{\pi,i}$	−0.217	−0.224	−0.007	−0.217	0
$\Phi_{i,i}$	0.696	0.684	−0.013	0.696	0

Notes: Medians across paths that visit the ZLB at least once (N = 9,868 with the GFC; N = 9,127 without). The reported Φ is the admissible matrix used in the first ZLB quarter. RE denotes the linear rational expectations initialization.

Two further comparisons isolate the role of the initial shock. Conditional on any ZLB

Table C.3: Regressor second-moment matrix on entry to the first ZLB

	RE	GFC/RE	GFC initial-PF/RE	No GFC/RE	No GFC initial-PF/RE
R_{uu}	6.51×10^{-8}	0.98	0.98	0.98	0.98
R_{rr}	1.54×10^{-5}	1.20	1.19	0.80	0.67
R_{ii}	2.60×10^{-5}	0.96	0.96	0.64	0.48
R_{ui}	3.46×10^{-9}	0.97	0.97	0.22	−0.10
R_{ri}	1.85×10^{-5}	0.99	0.99	0.70	0.54

Notes: Entries except the RE column and the last row are ratios of the median R_T at first-ZLB entry to the corresponding linear RE second-moment entry of $Z_t = (u_t, r_t, i_{t-1})'$. "Initial-PF" restricts to paths whose first projection event occurs during that first ZLB spell ($N = 1{,}402$ with the GFC; $N = 1{,}080$ without).

visit, the first-spell PF hazard is 11.83% without the GFC and 14.21% with it (Table C.4). Among spells that begin after quarter 8, the corresponding hazards are 11.83% and 9.81%, indicating that part of the aggregate GFC difference comes from the mass of early first spells. At-ZLB first projection updates themselves are similar across treatments (Table C.5): median accepted $\Phi_{i,i}$ is 0.949 in both, while the median attempted change is 0.0639 with the GFC and 0.0631 without it. Median forecast errors are also close, at -1.10 and -0.99 annualized percentage points. Thus, the GFC changes the incidence and timing of unstable paths but not the underlying premature-lift-off mechanism.

Table C.4: First ZLB visit with and without an imposed GFC

	GFC	No GFC
Share ever at the ZLB (percent)	98.7	91.3
Paths with a first ZLB	9,868	9,127
Median first-ZLB start (quarter)	4	125
10th–90th percentile start	4–73	26–356
Median first-ZLB duration (quarters)	15	11
First-spell PF paths	1,402	1,080
First-spell PF hazard (percent)	14.21	11.83
Hazard if first ZLB starts after quarter 8	9.81	11.83

Notes: The last row restricts to first spells that begin after quarter 8, so it drops the mass of early GFC entries.

C.3 Robustness to a 0.99 projection bound

A matched current-specification comparison raises only the projection threshold from 0.98 to 0.99 while retaining $R_0 = \Omega$, the macroeconomic bound of 1,000, and the same 10,000 shock histories. PF incidence declines from 45.28 percent (4,528 paths) to 43.17

Table C.5: First projection-event learning updates with and without an imposed GFC

	GFC	No GFC
At-ZLB first PF paths	2,970	2,213
Median accepted $\Phi_{i,i}$	0.949	0.949
Median attempted $\Delta\Phi_{i,i}$	0.0639	0.0631
Median |post-update candidate|	1.002	1.001
Median rate forecast error (ann. pp)	−1.10	−0.99
Outside-ZLB first PF paths	1,558	773
Outside-ZLB median accepted $\Phi_{i,i}$	0.961	0.960
Outside-ZLB median |candidate|	1.003	1.004

Notes: At-ZLB rows combine first projection events during the initial and later ZLB spells. The attempted change and post-update candidate are the objects to which the projection test is applied. Forecast errors are realized minus forecast, in annualized percentage points.

percent (4,317 paths), a paired difference of 2.11 percentage points. Every path that hits at 0.99 also hits at 0.98; hence 211 paths are marginal to the lower threshold. Of these, 172 first cross at the ZLB under 0.98 and 39 cross above it. Among paths that hit under both bounds, the median first-event delay is zero quarters and the 90th percentile is one quarter. Across all PF paths, however, the median first event moves from quarter 100 to 105 because the 211 marginal paths exit the risk set. Macroeconomic failure rises from 11.65 to 12.54 percent, while the union of PF and macroeconomic failure falls from 45.28 to 43.20 percent. Thus, the 0.98 bound creates a modest mass of threshold-sensitive events but does not drive the broader instability result. It also provides an interior buffer below the unit-root boundary: persistence of 0.98 already implies a quarterly half-life of about 34 quarters. I therefore retain 0.98 as the baseline projection facility boundary.

C.4 Sensitivity to the weight on forward-guidance messages

The baseline calibration gives the announcement gain the same value as the ordinary learning gain, $\lambda = \tau = 0.03$. I assess the sensitivity of the two learning regimes over $\lambda \in \{0.01, 0.02, 0.03, 0.04, 0.06\}$, so that an announcement receives between one-third and twice the weight of a regular quarter of data in equation (30). All other features of the experiment are held fixed. Table C.6 reports both stabilization outcomes and indicators of learning instability.

Under Learning FG, increasing the announcement gain improves the crisis-window loss

Table C.6: Sensitivity to the forward-guidance message gain

Panel A. Learning FG									
λ	λ/τ	Median first-PF quarter	Any PF	First PF at ZLB	First PF above ZLB	First PF from FG	Macro failure	Loss index	Median ΔL_n (%)
0.01	0.33	244	11.56%	6.36%	5.2%	3.23%	1.8%	80.65	-10.4
0.02	0.67	243	10.86%	7.46%	3.4%	4.83%	2.44%	77.96	-13.6
0.03	1	224	11.21%	8.15%	3.06%	4.8%	3.05%	76.71	-15.1
0.04	1.33	220	13.37%	9.08%	4.29%	4.73%	3.6%	75.98	-15.8
0.06	2	214	23.52%	10.93%	12.59%	5.52%	3.96%	74.78	-16.8
Panel B. Period + Learning FG									
λ	λ/τ	Median first-PF quarter	Any PF	First PF at ZLB	First PF above ZLB	First PF from FG	Macro failure	Loss index	Median ΔL_n (%)
0.01	0.33	253	8.15%	2.91%	5.24%	2.5%	0.95%	75.18	-13.6
0.02	0.67	244	8.47%	4.8%	3.67%	4.07%	1.45%	74.9	-15.2
0.03	1	220	7.89%	4.77%	3.12%	4.04%	1.53%	73.79	-16.3
0.04	1.33	205	6.75%	3.79%	2.96%	3.23%	1.31%	71.99	-17.2
0.06	2	263	10.32%	2.6%	7.72%	2.43%	0.71%	67.51	-18.6

Notes: Each cell uses 10,000 matched paths and the same No FG benchmark. PF and macro-failure entries are unconditional path shares. First-event columns partition the timing of each path's first PF event; the FG column identifies first events triggered by the pre-data message update. The loss index uses the mean discounted loss over quarters 2–40 and normalizes the No FG mean to 100. ΔL_n is the matched path-wise percentage change relative to No FG.

criterion but eventually raises projection risk sharply. The loss index falls monotonically from 80.65 at $\lambda = 0.01$ to 74.78 at $\lambda = 0.06$, and the median path-wise loss reduction increases from 10.4% to 16.8%. PF incidence is initially flat: it falls from 11.56% to 10.86% at $\lambda = 0.02$ and equals 11.21% at the baseline gain. It then rises to 13.37% at $\lambda = 0.04$ and 23.52% at $\lambda = 0.06$. At the largest gain, first events above the ZLB account for 12.59% of paths and first PF events at the ZLB for 10.93%. Macroeconomic failure also rises across the grid, from 1.80% to 3.96%.

The ZLB duration mask moderates this trade-off. Under Period + Learning FG, the loss index falls from 75.18 to 67.51 and the median path-wise loss reduction increases from 13.6% to 18.6%. PF incidence remains between 6.75% and 8.47% through $\lambda = 0.04$, then rises to 10.32% at $\lambda = 0.06$. The increase at the largest gain occurs above the ZLB, where first-event incidence rises to 7.72%; first-event incidence at the ZLB falls to 2.60%. Macroeconomic failure is non-monotone and falls to 0.71% at the largest gain. This difference is not contradictory: a PF event records a rejected candidate update, whereas the macroeconomic-failure measure records the subsequent realized path under the projection facility and, in the combined regime, the duration mask.

Thus, the exercise identifies a stability - stabilization trade-off rather than an optimal gain. Empirically studying λ is beyond the paper's scope and is left for future research.

List of research reports

2021001-OPERA: Baardman, L., K.J. Roodbergen, H.J. Carlo, and A.H. Schrotenboer, A Special Case of the Multiple Traveling Salesmen Problem in End-of-aisle Picking Systems

2021002-EEF: Wiese, R., and S. Eriksen, Willingness to Pay for Improved Public Education and Public Health Systems: The Role of Income Mobility Prospects.

2021003-EEF: Keller, J.T., G.H. Kuper, and M. Mulder, Challenging Natural Monopolies: Assessing Market Power of Gas Transmission System Operators for Cross-Border Capacity

2021004-EEF: Li, X., and M. Mulder, Value of Power-to-Gas as a Flexibililty Option in Integrated Electricity and Hydrogen Markets

2021005-GEM: Rozite, K., J.P.A.M. Jacobs, and D.J. Bezemer, Investor Sentiment and Business Investment

2021006-EEF: Spierdijk, L., and T. Wansbeek, Differencing as a Consistency Test for the Within Estimator

2021007-EEF: Katz, M., and C. van der Kwaak, To Bail-in or to Bailout: that's the (Macro) Question

2021008-EEF: Haan, M.A., N.E. Stoffers, and G.T.J. Zwart, Choosing Your Battles: Endogenous Multihoming and Platform Competition

2021009-I&O: Greve, M., M. Fritsch, and M. Wyrwich, Long-Term Decline of Regions and the Rise of Populism: The Case of Germany

2021010-MARK: Hirche, C.F., T.H.A. Bijmolt, and M.J. Gijsenberg, When Offline Stores Reduce Online Returns

2021011-MARK: Hirche, C.F., M.J. Gijsenberg, and T.H.A. Bijmolt, Promoting Product Returns: Effects of Price Reductions on Customer Return Behavior

2021012-MARK: Hirche, C.F., M.J. Gijsenberg, and T.H.A. Bijmolt, Asking Less, Getting More? The Influence of Fixed-Fee and Threshold-Based Free Shipping on Online Orders and Returns

2021013-I&O: Sorgner, A., and M. Wyrwich, Calling Baumol: What Telephones Can Tell Us about the Allocation of Entrepreneurial Talent in the Face of Radical Institutional Changes

2021014-I&O: Slavtchev, V., and M. Wywich, TV and Entrepreneurship

2021015-EEF: Kate, F. ten, M.J. Klasing, and P. Milionis, Diversity, Identity and Tax Morale

2021016-EEF: Bergemann, A., and R.T. Riphahn, Maternal Employment Effects of Paid Parental Leave

2021017-GEM: Abolhassani, M., Productivity Spillovers of Multinational Enterprises through Worker Mobility: New Evidence for the Netherlands

2021018-GEM: Abolhassani, M., Productivity Spillovers of Superior Firms through Worker Mobility

2022001-GEM: Oosterhaven, J., A Price Reinterpretation of the Leontief Quantity Model

2022002-EEF: Ghaemi, S, X. Li, and M. Mulder, Economic Value of Flexibility Provided by Power to gas Conversion Systems in Low-voltage Distribution Grids

2022003-OB: Meer, P.H. van der, Are All Self-employed Happy?

2022004-EEF: Perey, P., and M. Mulder, International Competitiveness of Low-carbon Hydrogen Supply to the North-west European Market

2022005-OPERA: Kasper, A., M. Land, and R. Teunter, University of Groningen, Faculty of Economics and Business, Department of Operations.

2022006-I&O: Fritsch, M., and M. Wyrwich, Entrepreneurship in the Long-run: Empirical Evidence and Historical Mechanisms.

2022007-EEF: Treurniet, M., and R. Lensink, Belief-based Poverty Traps and the Effects of Material and Psychological Development Interventions.

2022008-EEF: Kwaak, Christiaan van der, Monetary Financing Does Not Produce Miraculous Fiscal Multipliers

2022009-EEF: Beesten, E.R. van, and D. Hulshof, Transmission Capacity Reduction in International Power Systems: Economic Incentives and Welfare Effects

2022010-OPERA: Romeijnders, W., N.D. van Foreest and J. Wijngaard, On Proportionally Fair Solutions for the Divorced-Parents Problem

2022011-I&O: Hipp, A., M. Fritsch, M. Greve, J. Günther, M. Lange, C. Liutik, B. Pfeifer, M. Shkolnykova, and M. Wyrwich, Comprehensive Patent Data of the German Democratic Republic 1949-1990—Technical Report and Dataset Overview

2022012-EEF: Bianchi-Vimercati, R., G. Lecce, and M. Magnaricotte, Persistent Specialization and Growth: The Italian Land Reform

2023001-I&O: Fritsch, M., M. Greve, and M. Wyrwich, Shades of a Socialist Legacy? Innovation Activity in East and West Germany 1877-2014

2023002-EEF: Bahcivan, H., L. Dam and H. Gonenc, New Avenues in Expected Returns: Investor Overreaction and Overnight Price Jumps in US Stock Markets

2023003-I&O: Wyrwich, M. and M. Fritsch, How does Regional Entrepreneurship Transfer over Time? The Role of Household Size and Economic Success

2023004-OPERA: Seepma, A., and M. Spring, The Roles of Technology and Co-location in Overcoming Divergent Logics in Professional Supply Chains: the Case of Criminal Justice

2023005-EEF: Dalo, A., R. Mees, B. Scholtens, and H. Zhao, Corporate Sustainability, Cost of Equity, and Credit Ratings

2023006-EEF: Kwaak, C. van der, and S. van Wijnbergen, Financial Fragility and the Fiscal Multiplier

2024001-EEF: Herrmann, O., A partisan duty to vote?

2024002-EEF: Soetevent, A.R., and G.J. Romensen, The Short and Long Term Effects of In-Person Performance Feedback: Evidence from a Large Bus Driver Coaching Program

2024003-EEF: Stangenberg, L, and A.R. Soetevent, Don't wait on the world to change! How technophilia causes group inaction – an experiment

2024004-GEM: Stumphius, C.J., and D.J. Bezemer, Green Bonds: financial development or financialization? A firm-level analysis of their emission and energy impacts

2024005-OB: Meer, P.H. van der, and R. Wielers, The Value of Meaningful Work

2024006-OB: Meer, P.H. van der: Home-office, Covid-19 and Subjective Well-being

2024007-GEM: Tian, X., The Effect of the Global Financial Cycle on National Financial Cycles: Evidence from BRICS Countries

2024008-EEF: Elgersma, S., Mitigation, Adaptation and Cooperation in Response to Climate Disaster

2024009-MARK: Koster, J., M.C. Leliveld, H. Risselada, and J.W. Bolderdijk, A System Approach to Sustainable Fashion: What Do We Know and Where Do We Go?

2024010-EEF: Moahid, M., N. Hermes, and M. Hudon, The Microfinance Regulation Maze: A Systematic Literature Review

2024011-EEF: Moahid, M., N. Hermes, and M. Hudon, Examining the Implications of Innovative Microfinance Products: A Systematic Literature Review

2024012-EEF: Karan, M.B., and W. Westerman, A narration of banks, economic freedom and liberal democracy

2024013-EEF: Kwaak, C. van der, Bank Risk Taking & Quantitative Easing

2024014-EEF: Kwaak, C. van der, Bank Risk Taking and Central Bank Lending in Financial Crises

2025001-EEF: Romensen, G.J. and A.R. Soetevent, P(l)ay-As-You-Go: Large-Scale Longitudinal Risk-Elicitation in the Field

2025002-I&O: Bergeaud, A., M. Deter, M. Greve, and M. Wyrwich: Migration and Innovation: The Impact of East German Inventors on West Germany's Technological Development

2025003-EEF: Eriksen, S., M. Parkhomenko, A. de Ridder, and B. Scholtens: The Investment Performance of Classical Swedish Painters

2025004-EEF: Heijmans, R.J.R.K., and S. Suetens, Comparing Subsidies to Solve Coordination Failure

2025005-EEF: Oosterhaven, J., On the Math versus Econ of Input-Output Models

2025006-EEF: Romensen, G.J. and A.R. Soetevent, Does Targeting Relative Performance Feedback Improve Worker Productivity? Field Experimental Evidence from Bus Drivers

2025007-EEF: Westerman, W., Betting on the big bucks: greedy bankers beating (?) financial markets (1970's – 2010's)

2025008-GEM: Bezemer, D., Finance for Investment? Explaining Macrofinancial Slippage

2025009-GEM: Kellers, A., An idea-centric perspective on the emerge of belief systems: The case of sustainable investing and high-net-worth individuals

2025010-I&O: Fritsch, M., M. Greve, and M. Wyrwich, The impact of socialist legacy on regional differences in innovation activities and cooperation in Europe: A mediation analysis

2026001-GEM: Schoenauer, A., and T. Trompke: Closing the SME Emissions Data Gap

2026002-I&O: Mitzinneck, B., and F. Noseleit: Wat Werkt Waar: Plaatsgebonden paden voor burgerenergie in Nederland / What Works Where: Place-Based Pathways for Citizen Energy in the Netherlands

2026003-OB: Xu, S.Y., L. Mulder, and F. Rink: Highlighting the Collective Harm: Socially-framed Appeals Decrease Intention to Use Drugs Among People Who Use Drugs

2026004-EEF: Marinkov, V.: Disagreement at the Zero Lower Bound: The Case for Forward Guidance www.rug.nl/feb