Building Central Bank Credibility: The Role of Forecast Performance*

Michael McMahon† Ryan Rholes‡

This Draft: September 2025
Click here for latest draft

Abstract

This paper examines how central banks influence inflation expectations via public signals on inflation, and particularly how their forecast accuracy impacts this effect. We find, using an incentivized experiment, that forecast performance matters. Our main, and novel, finding is the presence of recency bias when subjects evaluate forecast accuracy. This bias, which applies to both short-term and medium-term forecasts, is especially strong after poor forecasting performance. In a New Keynesian model, such biases lead to endogenous forecast credibility, which can increase the persistence of inflation. Importantly, narrative communication can partly mitigate the detrimental effect of recent poor forecasting.

Keywords: Expectation formation, Forecasting, Central Bank Communication

JEL Codes: E52, E58

*We are grateful to seminar participants at the University of Warwick, University of Oxford, University of Mississippi, Indiana University - Bloomington, Simon Fraser University, at the 2022 conference on Monetary Policy: Heterogeneity, Communication and Subjective Inflation Expectations, the Economic Science Association's (ESA) World Meeting 2022 at MIT, the ESA European Meeting 2022 at the University of Bologna, the 2023 meeting of the Southern Europe Experimental Team, the 2023 SUERF-Banca d'Italia-ECB-EIB conference at the Bank of Italy, the 60th SUERF Anniversary 50th OeNB Annual Economic Conference at The Oesterreichische Nationalbank, the 2023 Meeting of the Central Bank Research Association, and at the Workshop on Challenges For Monetary Policy in Times of High Inflation hosted at Leibniz University Hannover. We also thank Luba Petersen and Johannes Abeler for their helpful comments. McMahon gratefully acknowledges financial support from the European Research Council (Consolidator Grant Agreement 819131). The study received IRB approval from the University of Oxford. AEA Trial Registry RCT ID: AEARCTR-0009026. The views presented in the paper are those of the authors alone and do not represent the official views of anyone else. Any errors remain ours alone.

†University of Oxford, CEPR, CfM (LSE). Email: michael.mcmahon@economics.ox.ac.uk
‡University of Mississippi. Email: rarholes@olemiss.edu

1 Introduction

We explore a fundamental issue in monetary economics: (how) can central banks influence inflation expectations by communicating their own inflation projections? Our novel focus is on the endogenous relationship between past forecast performance and the extent to which forecast information is used by agents in forming their expectations, which we will call forecast credibility. We engage thousands of U.S. households in an incentivized forecasting experiment that allows us to isolate a causal relationship between a central bank's own forecasting history and its forecast credibility.

This issue is important because a central tenet of monetary economics, reflected in the widely-adopted inflation targeting framework, is that central banks must anchor inflation expectations (Woodford 2005, King et al. 2008, Candia et al. 2020) because inflation expectations are a key determinant of inflation itself (Clarida et al. 1999, Woodford 2003, Galí 2008, for example). Inflation-targeting central banks use inflation forecasts to guide their policy (inflation-forecast targeting) and typically publish numerical inflation forecasts and contextualizing narratives to communicate these views and influence inflation expectations. While such open-mouth operations have become an integral component of monetary policy, the extent and nature of their impact on expectations remains debated (Haldane and McMahon 2018, Coibion et al. 2020, 2022).

In canonical monetary models populated by fully-informed rational agents, communicating about inflation forecasts is irrelevant since the central bank holds neither an informational nor a processing advantage; the central bank and the agent form coincident, optimal inflation expectations. Thus, central bank communication of these forecasts would be perfectly credible in such theoretical models. In reality, the central bank may hold advantages along either, or both, dimensions and therefore communication of their forecasts should improve the inflation expectations of agents.[1] And central banks worry about establishing and safeguarding credibility, which is considered necessary for the transmission of communication policy (Blinder 2000). If the central bank is highly credible, it can anchor inflation expectations, giving it better control of inflation which reinforces credibility. Instead of such a virtuous cycle, there may exist a vicious cycle; lower credibility could impinge upon the ability of the central bank to manage inflation, which then makes credibility-reducing inflation fluctuations more likely.

Our experiment employs an individual-choice setting where inflation evolves exogenously so that atomistic participants face no strategic uncertainty. As a result, inflation forecasting is not a coordination game allowing us to more cleanly isolate the causal relationship between features of a central bank's forecasting history and participants' perceptions of the bank's forecast credibility. This experimental model aligns with household surveys like the University of Michigan's Survey of Consumers, the New York Federal Reserve's Survey of Consumer Expectations, or the European Central Bank's Consumer Expectations Survey. It contrasts with the often-used learning-to-forecast (LtFEs) experiments, which feature experimental economies that evolve endogenously according to

[1] Eusepi and Preston (2010) consider the role of communicating less-than-full information about the central bank's policy function, which they argue is akin to partially-credible communication. Such communication in their framework is enough to prevent self-fulfilling expectations and restore macroeconomic stability under optimal policy.

the incentivized expectations of participants.[2] We believe that our design combines the best of LtFEs and survey-based, information provision experiments[3]; it combines the use of marginal incentives, experimental control over economic conditions and information, and the large-scale nature of survey-based experiments. This confluence of design features allows us to establish a credible causal link between the central bank's forecasting history and its forecast credibility.

Participants in our experiment provide two sets of one-period-ahead point and range forecasts of inflation in each of three independent decision periods (e.g. 'Initial Forecasts' and 'Updated Forecasts'). We begin each decision period by revealing the three most recent years of a central bank's quarterly inflation forecasts alongside actual inflation. Subjects provide Initial Forecasts (priors) for the next period based on this historical data. We then reveal the central bank's own inflation forecast and allow subjects to update their own density projection (i.e. Updated Forecasts or posterior). Thus, we collect incentivized measures of each participant's initial outlook on inflation, perceived initial inflation forecast precision, and their updated outlook on inflation. Additionally, we control both the central bank's historical forecast precision and its signal. Using these measures, we can precisely quantify the degree to which a participant incorporates the central bank's signal into their inflation outlook using a Bayesian signal processing framework and relate this causally to the realistic historical economic information we reveal at the start of each decision period.

Our design relates most closely to Armantier et al. (2016), who use a Bayesian updating framework to study how inflation expectations respond to historical price information or professional forecasts in an information provision experiment embedded into the Michigan survey. A key difference between our work and theirs is that our design uses experimental control over extraneous economic features and information allowing us to precisely manipulate the central bank's forecasting history to study its relationship to forecast credibility. Also related, Bulutay (2023) uses an information-provision experiment embedded into the Bundesbank's Survey on Consumer Expectations to show that European households typically underestimate the ECB's forecast performance and that correcting this misperception can increase the bank's forecast credibility.

We study how several key features of the central bank's forecasting history relate to its forecast credibility. First, we consider how the central bank's historical forecast precision influences its forecast credibility (*Forecast Performance*). We create a series of economic histories that control for the pattern of historical forecast errors while scaling the magnitude of the historical average absolute forecast error. We find that participants behave qualitatively like Bayesians, in that the central bank's forecast credibility depends significantly on historical forecast precision. However, the link between precision and credibility is not as sharp as theory predicts.[4]

[2]Researchers have used this framework to study the design and efficacy of central bank communication (Kryvtsov and Petersen (2021); Arifovic and Petersen (2017); Cornand and M'baye (2018); Rholes and Petersen (2021); Petersen and Rholes (2022)), expectation formation and equilibria selection (Adam (2007); Bao et al. (2012)), and how various monetary policy rules and targets affect expectation formation (Ahrens et al. (2019); Pfajfar and Žakelj (2014); Pfajfar and Žakelj (2018); Assenza et al. (2013); Hommes et al. (2019); Hommes et al. (2019); Cornand and M'baye (2018)).

[3]See Fuster and Zafar (2023) for an overview of survey-based experiments on economic expectations.

[4]We provide some evidence in A4.1 that this is becuase subjects do not account for their stated forecast precision when evaluating the central bank's signal.

Second, we hold the central bank's historical forecast precision constant on average and ask how variation in the timing of historical inflation forecast errors influences the central bank's forecast credibility (*Timing*). We do this using a set of three core economic histories where historical forecast precision is either consistent (*Consistent*) or mimics the average pattern of forecast errors but has larger forecast errors at the onset (*Early*) or end (*Late*). With such a small sample to assess forecast precision, a rational Bayesian agent should find the central bank equivalently credible across *Early, Late,* and *Consistent*. Our interest is in whether this is true or if, instead, the timing of forecast errors matters. We find that participants rely primarily on recent information to form perceptions of forecast credibility whenever the central bank's forecast precision has changed considerably in the recent past. We refer to the preference for recent information as 'recency bias'.

Importantly, we find that recency bias is significantly more pronounced when recent forecast performance is poor. This is true even though historical forecast precision changes by equivalent speed and magnitude across the *Early* and *Late* histories. This suggests an important asymmetry in how recency bias interacts with forecast performance. Bad forecast performance seems to hold more value to participants and so subjects exhibit less recency bias when the central bank's forecast precision improves than when it deteriorates. This has important implications for the dynamics of credibility; it is much easier for a bank to lose credibility than to rebuild it. Further, the speed with which credibility evaporates suggests that unanticipated shocks leading to poor short-term forecast performance can undermine the efficacy of inflation communication precisely when the monetary authority most needs it as a policy tool.

We explore the robustness of these timing-related results in two ways. First, we show that these results also hold whenever participants forecast long-term average inflation rather than short-term inflation.[5] Second, we introduce a set of treatments wherein we reverse the direction of forecast errors to show that our results hold regardless of whether the bank under- or over-forecasts inflation.

We design our experiment to mitigate the role that inattention, memory and recall issues, and cognitive limitations may have on participant behavior. For instance, we systematically reveal historical information gradually, ensure relevant information is always available to subjects, and communicate information about historical average precision. Beyond these mitigating design features, we take no stand on the underlying source of recency bias in our experiment. Participants may perceive a structural break in historical data, assume a high degree of autocorrelation in central bank forecast errors, or instead be subject to some behavioral tendency for discounting historical information. However, the asymmetry we find is hard to rationalize through the lens of perceived structural changes given the significant differences in recency bias exhibited across different histories. Further, the perception of highly autocorrelated forecast errors fails to rationalize why these core results are robust to long-term average inflation forecasts.

Although understanding the behavioral mechanism for recency bias could be important – for example, by clarifying topics of narrative communication to guard against credibility loss – we leave this for future work. Rather, our goal in this project is to document whether and how forecast credibility relates dynamically to a central bank's forecasting

[5] Weber et al. (2022) find short-term and medium-term expectations are correlated.

history. To explore the implications of our findings, we take an otherwise standard three-equation New Keynesian model but have the central bank communicates with a boundedly rational household to show that accounting for endogenous credibility and recency bias can lead to significant persistence in inflation dynamics. However, a central bank can achieve dynamics similar to those under rational expectations if it successfully maintains forecast credibility following the exogenous shock.

Finally, we test whether narrative communication that reinforces the central bank's outlook and contextualizes its forecasting history can bolster forecast credibility for a bank that finds itself in a position of low forecast credibility (*Contextual Communication*).[6] That is, can a central bank talk its way from a low-credibility position to a better one? To do this, we create a series of text-based communication interventions that contain a forward-looking component, rationalize forecast mistakes as the result of unforeseeable exogenous shocks or endogenous policy errors, and report performing better or worse than peer forecasting institutions. Communication that reinforces the central bank's numerical inflation forecast without providing additional information can significantly improve credibility. Layering on additional information can further increase credibility but the effect is more nuanced. Reporting that the bank under-performed relative to peer institutions reduces credibility sufficiently to eliminate any gains from contextual communication. Reporting that the bank outperformed peer institutions bolsters credibility. We find no difference arising from reporting different sources of forecast errors.

The rest of the paper is organized as follows: Section 2 introduces a simple Bayesian updating framework. Section 3 provides details of our experimental design, Section 4 contains our *Forecast Performance* results, and Section 5 our *Timing* results. Section 6 explores the theoretical implications of our *Forecast Performance* and *Timing* results. Finally, Section 7 explores the role that communication can play in attenuating the effect of poor forecast performance, and Section 8 concludes.

2 Central bank signals and forecast updating

In this section, we introduce a simple Bayesian updating framework (similar to Morris and Shin (2002), for example) to emphasize how central bank signals should influence a Bayesian participant's decision to update her inflation forecasts. This framework, which guides our experimental design, illuminates how an individual should react to inflation forecasts contingent upon his or her perception of the central bank's forecast credibility, providing a Bayesian benchmark for participant behavior. Moreover, the framework provides a precise measure we use to quantify central bank credibility.

We begin a decision period by revealing to participant i the central bank's forecast of inflation alongside realized inflation for the twelve most recent quarters. In response, the participant provides incentivized beliefs about the value of inflation (π_i) for the following quarter and her own forecast precision (α_i^{-1}):

$$\pi_i \sim \mathcal{N}\left(\bar{\pi}_i, \frac{1}{\alpha_i}\right), \tag{1}$$

[6] This relates to the role of narrative (Shiller 2017, Macaulay and Song 2023, Flynn and Sastry 2022).

The central bank then reveals its own point forecast of inflation (π_{cb}):

$$\pi_{cb} = \pi + \tilde{\epsilon}, \quad \tilde{\epsilon} \sim \mathcal{N}\left(0, \frac{1}{\beta}\right). \tag{2}$$

where β^{-1} is related to the precision of the central bank forecast, which i infers from the 12-quarter economic history. We assume the central bank's forecasts are unbiased.[7] We are interested in how the central bank's signal is used to update the prior forecast (π_i). Define the update rate, which we will use to measure forecast credibility, as:

$$u_i \equiv \frac{\mathbb{E}(\pi|\pi_{cb}) - \pi_i}{(\pi_{cb} - \pi_i)} \tag{3}$$

The optimal Bayesian inflation forecast is a precision-weighted, linear combination of π_i and π_{cb}, and the optimal update rate is given by:[8]

$$u_i^* = \frac{\beta}{\alpha + \beta} \tag{4}$$

If $\beta \to \infty$, $\alpha \to 0$, or both, the agent updates fully toward the central bank signal, $u_i^* = 1$.Figure 1 plots this optimal update rate (in percentage terms, $100 \times u_i^*$) for different levels of β and α. There are three main implications:

1. For any given precision of the central bank signal, as the precision of the prior increases, $\alpha \uparrow$, the agent updates less when they receive the central bank signal.

2. For a given prior precision, as the precision of the central bank signal decreases, $\beta \uparrow$, the agent updates less when they receive the central bank signal.

3. The marginal effect of decreasing precision of the central bank signal is larger when the individual's prior is *more* precise.

3 Experimental Design

Our goal in designing this experiment is to isolate the causal relationship between various features of the central bank's forecasting history and its ability to influence inflation expectations. Our interest is in how the historical economic information we reveal to participants influences their perceptions of the central bank's forecast credibility (β^{-1}) by observing how they incorporate (or not) the central bank's forecast into their own inflation outlook. As the experimenters, we directly control β^{-1} and π_{cb}, which we reveal to subjects via economic histories and by announcing the central bank's forecast. Thus, we must collect incentivized measures of π_i and $\mathbb{E}(\pi|\pi_{cb})$ to measure forecast credibility, u_i. We also collect incentivized measures of α^{-1} to allow us to infer $\hat{u}_i$, which we define in Equation (7). The rest of this section provides details on how we do exactly this, how we construct our economic histories, and how we construct *ForecastPerformance* and *Timing* treatments.

[7] If, instead, we assume $\tilde{\epsilon} \sim \mathcal{N}\left(\gamma, \frac{1}{\beta}\right)$, γ represents a systematic bias in the central bank's inflation forecast. We show in Appendix appendix A6 that our results are qualitatively robust to assuming $\gamma \neq 0$.

[8] $\mathbb{E}(\pi|\pi_{cb}) = \frac{\alpha\pi_i + \beta\pi_{cb}}{\alpha+\beta}$.

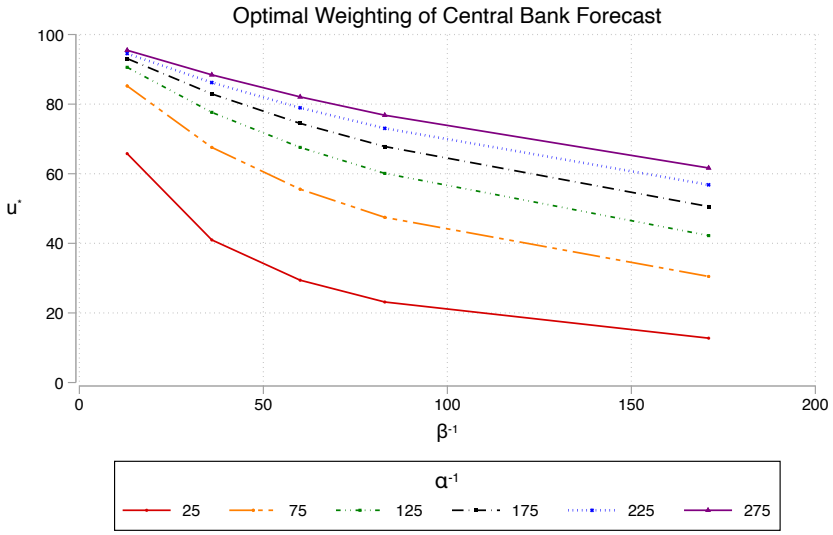


Figure 1: Optimal Updating

Note: This figure shows the optimal update rate in percentage terms (y-axis) prescribed by Equation (4) for different levels of a central bank precision (x-axis). Each line denotes a different level of participant forecast precision ranging from 25 basis points (bottom line) to 275 basis points (top line) in increments of 50 basis points.

3.1 Overall Design and Implementation

We programmed our experiment in oTree (Chen et al. 2016). We conducted our experiment online and recruited participants via Prolific, restricting our sample to experienced users from the United States.

Participants began the experiment by completing a short survey that measured their level of economics knowledge, their level of understanding of and trust in various public institutions, their preferences for obtaining economic information, and their familiarity with prevailing economic conditions. We then provided subjects on-screen instructions that explained the inflation forecasting task, the information available when forming forecasts, how to interact with the available information, how to interact with our software, and how we incentivized their forecasts. These instructions remained available to subjects throughout the experiment via a toggle button on all screens.

Next, subjects completed a comprehension quiz comprising five questions designed to test their understanding of our experiment. Subjects had three attempts to answer all five questions correctly to proceed. Our software ended the experiment early for subjects who failed the quiz more than twice, while subjects who successfully completed the quiz proceeded to the forecasting task.[9]

In the forecasting task, subjects completed three independent decision periods. In each decision period, subjects provided an Initial Forecast and an Updated Forecast. This means that our experiment yields a total of six sets of forecasts, with each set consisting of both point and range forecasts.

[9]In our online appendix appendix A9, we provide questions from the economic literacy quiz, our experimental instructions, and questions from our comprehension quiz.

Following decision periods, we paid subjects based on their performance in one of these sets of forecasts, randomized at the individual level. Participants ended the experiment with a non-compulsory survey of decisions.

3.2 Decision Periods

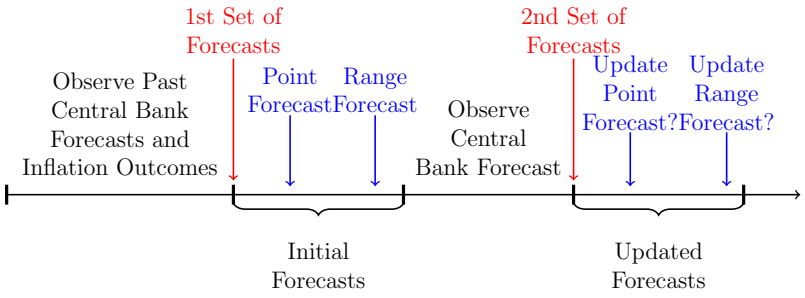


Figure 2: Experimental Timeline: A single decision period

Figure 2 presents the timeline of a decision period. We began each decision period by providing a participant with a 12-quarter economic history of the central bank's inflation forecasts alongside actual inflation. We revealed historical observations sequentially with a one-second lag, ensuring subjects considered each historical observation before forming Initial Forecasts.[10] We displayed this historical data graphically and numerically. All information, once revealed, remained available for the duration of that decision period.

Participants then provided a point forecast of one-period-ahead inflation (i.e. $\mathbb{E}_{i,12}\{\pi_{13}\}$) in percentage terms with two-decimal precision.[11] We incentivized point forecasts according to Equation (5), following the LtFE literature (Rholes and Petersen 2021, Mokhtarzadeh and Petersen 2021, Petersen and Rholes 2022):

$$F_i = 2^{-|\mathbb{E}_{i,12}\{\pi_{13}\}-\pi_{13}|}. \tag{5}$$

Note that a perfect forecast yields $F_i = 1$ and that this forecasting score is reduced by half each percentage point increase in the forecast error.

Participants also submitted a measure of initial forecast uncertainty. Our experimental software randomly generated upper and lower uncertainty bounds bracketing a participant's initial point forecast. The area between these two bounds appeared to participants as a shaded region, denoting a visual representation of the participant's forecast uncertainty. Participants could then change the uncertainty bounds to reflect their true forecast uncertainty.[12] If actual inflation fell outside this range forecast, a participant

[10] This design feature mitigates the role rational inattention might play in our results.

[11] Participants could submit point forecasts two ways. First, they could create a point forecast by clicking on the interactive chart used to display historical economic information. They could subsequently alter this forecast by dragging and dropping this point anywhere inside the forecast region of the graph. Alternatively, participants could type forecasts directly into an available input field. Participants faced no time pressure and could visualize as many forecasts as they desired before submitting the initial point forecast. Once a subject submits the initial point forecast, our software updates to reflect this value graphically and numerically.

[12] They could do this by dragging and dropping the two bounds independently, dragging and dropping both bounds simultaneously, or by typing numbers directly into corresponding input fields. Our software

earned nothing for the forecast. If instead realized inflation fell within a participant's uncertainty bounds, then she earns the amount given by Equation (6), which follows Pfajfar and Žakelj (2016), Rholes and Petersen (2021), Petersen and Rholes (2022).

$$\mathbb{U}_i(r_i) = \begin{cases} 0 & \pi_{13} \notin [\underline{\pi_{i,13}}, \overline{\pi_{i,13}}] \\ \phi\left(\frac{1}{r_i}\right) & \pi_{13} \in [\underline{\pi_{i,13}}, \overline{\pi_{i,13}}]. \end{cases} \tag{6}$$

Here, ϕ is a scalar we can use to calibrate average earnings.[13] $\underline{\pi_{i,13}}$ is the lower-bound of a participant's forecast uncertainty, $\overline{\pi_{i,13}}$ the upper-bound of a participant's forecast uncertainty, and $r_i = \|\overline{\pi_{i,13}} - \underline{\pi_{i,13}}\|$ is the participant's forecast range capturing their uncertainty. The intuition of this measure is that conditional on setting bounds that almost certainly capture inflation, the participant wants to minimize the magnitude of their forecast range.

After collecting Initial Forecasts, we revealed the central bank's point forecast of inflation (i.e. $\mathbb{E}^{CB}_{i,12}(\pi_{13})$) and then allowed the participant to update her forecasts (or not). We provided participants with numerical and graphical information about their Initial Forecasts and emphasized to them in our instructions and with an on-screen reminder that they were not obligated to change their forecasts. If they chose to update, they could update any or all values of $\mathbb{E}_{i,12}(\pi_{13})$, $\underline{\pi_{i,13}}$, $\overline{\pi_{i,13}}$.

If we designate the updated (posterior) estimate as $\mathbb{E}_{i,12}(\pi_{13}|\pi_{cb})$, then these responses allow us to calculate the empirical forecast credibility (Equation (3)) as:

$$\hat{u}_i \equiv \frac{\mathbb{E}_{i,12}(\pi_{13}|\pi_{cb}) - \mathbb{E}_{i,12}(\pi_{13})}{\mathbb{E}^{CB}_{i,12}(\pi_{13}) - \mathbb{E}_{i,12}(\pi_{13})}. \tag{7}$$

We can compare this to a Bayesian benchmark given by

$$u^*_i \equiv \frac{\beta_T}{\alpha_i + \beta_T}. \tag{8}$$

where we estimate α_i as $\frac{1}{r_i}$.[14]

After collecting updated forecast values, our software would reveal to participants the actual value of quarter-13 inflation (π_{13}) alongside their forecasting performance for that decision period. The software then proceeded to the next decision period.

3.3 Creating the economic histories

Differences in the economic histories shown to subjects constitute treatment variation in our experiment. We create these histories by simulating the simple 3-equation New

prevented subjects from inputting values for the upper bound that were below the point forecast and vice versa for lower-bound values. Our software also prevented subjects from visualizing upper and lower bounds that violated these same bounding conditions.

[13] We use $\phi = 1$ for our experiment.

[14] We explore more complex methods of estimating α_i in A8 assuming upper and lower range forecast bounds correspond to upper and lower percentiles in a normal distribution.

Keynesian model in Walsh (2017) linearized around a zero-inflation steady-state, described by Equation (9) through Equation (12). y_t is the output gap (log-deviation of output from the natural rate), π_t is the quarterly rate of inflation between $t-1$ and t, i_t is the nominal interest rate on funds moving between period t and $t+1$, and r_t is the real interest rate. Finally, g_t, p_t, and v_t are demand, cost-push, and monetary policy shocks, respectively.

$$y_t = E_t y_{t+1} - \sigma^{-1}(i_t - \mathbb{E}_t \pi_{t+1}) + g_t \tag{9}$$
$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + p_t \tag{10}$$
$$i_t = \phi_x y_t + \phi_\pi \pi_t + v_t \tag{11}$$
$$r_t = i_t - \mathbb{E}_t \pi_{t+1} \tag{12}$$

We assume each shock follows an AR(1) process

$$\psi_t = \rho_\psi \psi_{t-1} + \varepsilon_t^\psi, \quad \psi \in \{g, p, v\}, \tag{13}$$

where ε_t^ψ are i.i.d. innovations that are uncorrelated both across time and across shocks. The central bank is assumed to form rational expectations, so that for any shock ψ_t the conditional expectation is given by

$$E_t[\psi_{t+1}] = \rho_\psi \psi_t. \tag{14}$$

We calibrate this model using parameters in Table 1 and the inflation gap is then converted to inflation data by assuming a target rate of 2%. We created inflation forecasts and inflation values for the forecast quarter ($\mathbb{E}_{i,12}^{CB}\pi_{13}$ and π_{13}) in each economic history using shocks that roughly preserved the average forecast error of the final year of economic history.

Parameter Values								
β	$\sigma = \eta$	ω	κ	ϕ_π	ϕ_y	ρ_g	ρ_p	ρ_v
.99	1	.8	.104	1.5	0	.5	.5	0

Table 1: Parameter values for simulation exercise

3.3.1 Designing Simulations

Business cycles are a pervasive feature of all economies, arising from unanticipated structural shocks. Such shocks initially generate meaningful forecasts errors, which are autocorrelated over short horizons due to their autoregressive nature of the shocks. Over time, central banks adjust their expectations to account for the autoregressive component of the shock process, and forecast accuracy improves. This dynamic produces a profile in which forecast errors are initially elevated and then gradually diminish, resembling the "bad–early" history described below. Because this mechanism emerges generically from the interaction between shocks, persistence, and expectation formation, it is a challenge that all central banks face over time and thus provides a broadly generalizable basis for our experimental design.

Our goal in constructing the simulated histories is to capture this feature of the business cycle. The *Early* history reflects an economy moving out of a downturn, as forecast performance improves over time. The *Late* history represents the reverse situation—an economy entering a downturn, with forecast performance deteriorating. The *Consistent* histories illustrate the intervening state in which forecast accuracy remains broadly stable. Together, these three core scenarios provide a parsimonious but comprehensive characterization of the forecast-credibility challenges a central bank may face across the business cycle.[15]

To summarize, we simulate the following three economic histories:

- In *Early*, the central bank commits significant forecast errors in the first year (four quarters) of the forecasting history, moderate errors in the second year, and minimal errors in the last year.
- In *Late*, we reverse the pattern of forecast errors observed in *Early*, exactly preserving the absolute average forecast error between *Early* and *Late*. This means that both the magnitude and speed by which historical forecast precision changes are identical across these two histories. All that varies is whether the central bank has recently experienced a spate of poor or great forecast performance.
- For *Consistent*, the central bank exhibits a consistent average annual forecast performance. The key characteristic of *Consistent* is that each of the annual (4-quarter) average absolute forecast errors is the same as the full sample average absolute forecast error.
  - We first generate a version of *Consistent-Bad* so that the annual and sample average absolute forecast errors match the sample average absolute forecast errors of *Late* and *Early*. Next, leaving inflation unchanged, we amplify or moderate the central bank's forecast errors to create the other versions of *Consistent* listed in Table 2. We chose average absolute forecast errors in *Consistent-Great* (*Consistent-Terrible*) to exactly match the average absolute forecast error in the final year of *Consistent-Early* (*Consistent-Late*). Finally, we chose absolute error values for *Consistent-Good* and *Consistent-Moderate* so that they partitioned the performance difference between *Consistent-Great* and *Consistent-Bad*.
  - To isolate the causal relationship between historical forecast precision and forecast credibility (i.e. *ForecastPerformance*), we produce five versions of *Consistent* that preserve to time profile of forecast errors but vary the central bank's historical forecast precision. Precision in these alternative versions of *Consistent* vary from *Consistent-Great* performance, through *Consistent-Good*, *Consistent-Moderate*, *Consistent-Bad* and down to *Consistent-Terrible*.

We summarise forecasting performance in our simulated economic histories in Table 2. We show the economic histories in Figure 4.

[15]For concreteness, we illustrate this dynamic in appendix A2 using the experience of the United Kingdom and the Bank of England between 2010 and 2012, when forecast errors were initially large before declining substantially. Our simulated histories calibrated to that episode, qualitatively preserving the magnitudes and shapes of forecast errors observed over this period.

	Numerical Summary of Economic Histories (bps)				
	Year 1	**Year 2**	**Year 3**	**Full Sample**	$\gamma_{HistAvg}$
			Forecast Performance		
Consistent - Great	13	13	13	13	06
Consistent - Good	36	36	36	36	10
Consistent - Moderate	60	60	60	60	06
Consistent - Bad	83	83	83	83	02
Consistent - Terrible	171	171	171	171	-06
			Timing & Contextual Communication		
Consistent - Bad	83	83	83	83	02
Early	171	65	13	83	-51
Late	13	65	171	83	-52

Table 2: This table provides a numerical summary of our economic histories. Numbers are average absolute forecast errors expressed in basis points. The column labeled $\gamma_{HistAvg}$ lists of historical average forecast errors by economic history, which we explore in appendix A6.

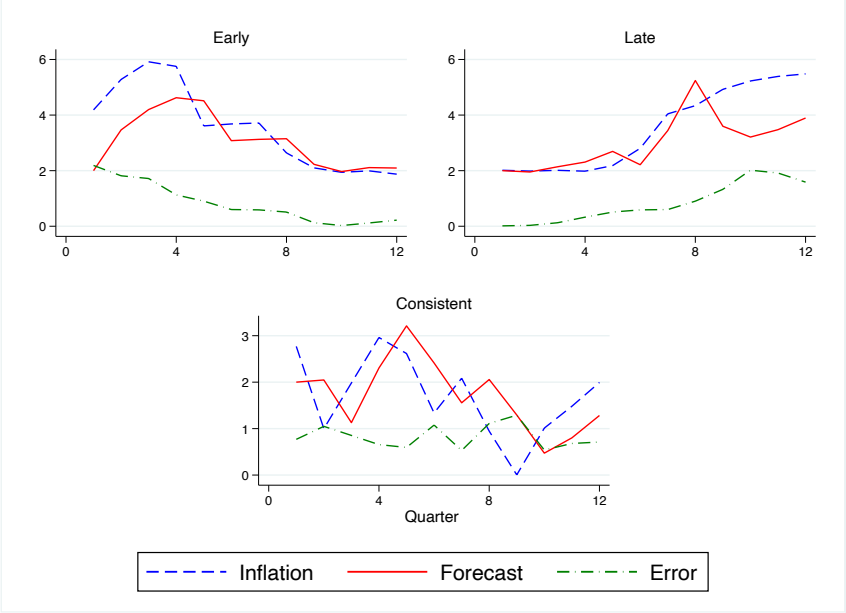


Figure 3: *Early*, *Late*, and *Consistent-Bad* – Timing Histories

Note: This figure presents three subpanels labeled *Early* (top left), *Late* (top right), and *Consistent* (bottom). Each subpanel shows historical inflation trends (blue dashed lines) alongside inflation forecasts (solid red lines) provided to participants in our study for the respective histories. Green dash-dot lines represent the corresponding absolute forecast errors.

All participants completed three independent decision periods consisting of *Early*, *Late*, and some version of *Consistent*. In *Forecast Performance*, subjects first see some ordering of *Early* and *Late*, and then see one of the five possible versions of *Consistent*. In (*Timing*), a participant experienced some ordering of *Early*, *Late*, and *Consistent-Bad*. In *Contextual Communication*, participants see *Early* and *Consistent-Bad*. In their third decision period, these participants see *Late*, where we augment the central bank's forecast with additional written communication. In the following sections we cover the *Forecast Performance*, *Timing*, and *Contextual Communication* treatment waves and,

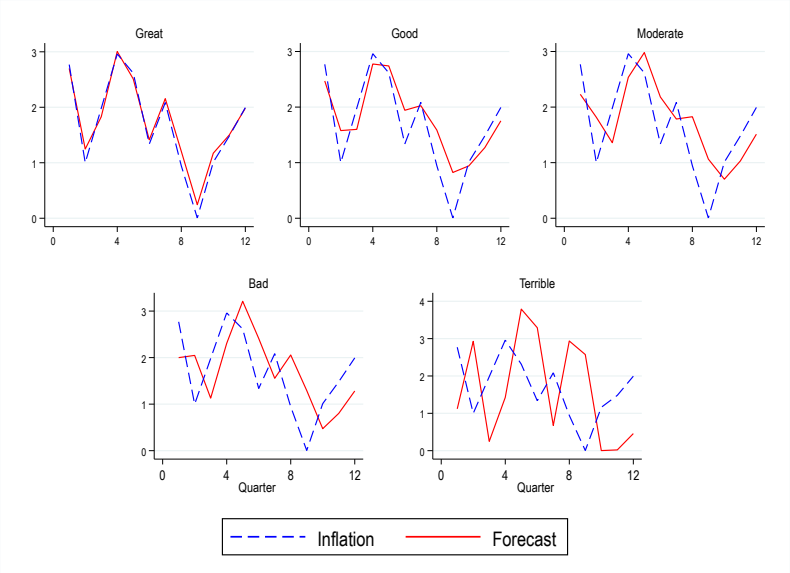


Figure 4: Alternative versions of *Consistent* used in *Forecast Performance*
Note: This figure presents subpanels labeled *Great* through *Terrible* (bottom right) that show historical inflation trends (blue dashed lines) alongside inflation forecasts (solid red lines) provided to participants in our study for the respective histories.

for each, provide additional details regarding treatments and experimental design, state our hypotheses, and detail our results.

4 Does average *Forecast Performance* matter?

In our first experimental wave, which we label *Forecast Performance*, we study how a central bank's historical forecast precision changes the willingness of participants to incorporate the central bank's inflation forecast into their own updated point forecast. As discussed in Section 3, we answer this question using our alternative versions of *Consistent* that vary only in the central bank's historical average absolute forecast errors (i.e. β^{-1}). Recall that we refer to these histories as *Consistent-Great*, *Consistent-Good*, *Consistent-Moderate*, *Consistent-Bad* and *Consistent-Terrible*. For brevity, we drop the *Consistent* prefix from treatment names for the remainder of this section.

More precisely, we randomize *Forecast Performancw* participants into one of ten different possible treatments listed in Table 3. Results are based on a between-subjects comparison of the average perceived forecast credibility across *Great* through *Terrible*. Sample sizes are consistent across treatments, except for *Bad*, where we draw on strictly comparable treatments that arose in the *Timing* waves.

4.1 Hypothesis 1: Forecast performance and credibility

Equation (4) provides a clear hypothesis about the relationship between historical forecast performance and the central bank's forecast credibility, as measured by $\hat{u}_i$. Using the inverse of a history's sample-average absolute forecasting error as a proxy for precision, we have the following:

Treatment Summary: *Forecast Performance*				
	History 1	**History 2**	**History 3**	**Sample Size**
T1a	*Early*	*Late*	*Great*	46
T1b	*Late*	*Early*	*Great*	44
T2a	*Early*	*Late*	*Good*	45
T2b	*Late*	*Early*	*Good*	46
T3a	*Early*	*Late*	*Moderate*	33
T3b	*Late*	*Early*	*Moderate*	39
T4a	*Early*	*Late*	*Bad*	100
T4b	*Late*	*Early*	*Bad*	80
T5a	*Early*	*Late*	*Terrible*	46
T5b	*Late*	*Early*	*Terrible*	50

Table 3: Treatment Summary: Forecast Performance

Note: This table describes our *Forecast Performance* treatments and provides sample sizes for each. Each subject participated in a single treatment, denoted by rows *T1a* through *T5b*. Within this treatment, a subject experienced three histories, the order of which is outlined by the columns denoted *History 1* through *History 3*. *History 3* always presented some version of our *Consistent* forecasting history, which we detail in Table 2.

Hypothesis 1. *A central bank's forecast credibility is decreasing in its historical average absolute forecast error.*

To test this hypothesis, we average over individual-level estimates of perceived forecast credibility, $\hat{u}_i$, to produce estimated average treatment effects (blue dots), which we compare to the treatment-average Bayesian optimal benchmarks (red triangles) in Figure 5.[16] Finally, we note that we Winsorize the top and bottom 5% of our data to mitigate the impact of outliers on our results. Unless otherwise noted, we do this for all results in our text.[17] Additionally, Table 4 presents results from a series of OLS regressions that capture all pairwise treatment-level comparisons of forecast credibility across *Forecast Performance* treatments.

Estimates indicate that the central bank's forecast credibility is negatively correlated with its historical forecast precision, β^{-1}, which constitutes support for hypothesis one – subjects respond to decreases in historical forecast precision like a Bayesian updater. However, the empirical relationship we observe between forecast credibility and forecast precision is flatter than predicted by theory. In fact, estimates in Table 4 suggest there is no difference in credibility if a central bank's forecast deteriorates from *Good* to *Bad*, and there is no statistically significant loss even if performance drops to *Terrible* (though the point estimate does indicate a decline). In some sense, this is good news because the forecast credibility cost of large errors may not be quite as high as the theory predicts. However, we observe much stronger evidence that subjects underuse signals from highly

[16]The treatment average optimal response is given as $u_T^* = \frac{1}{N_T} \sum_{n \in N_T} \frac{\beta_T}{\alpha_n + \beta_T}$, where β_T denotes historical forecast precision in the relevant treatment.

[17]We provide a sensitivity analysis of these cut points in appendix A7.

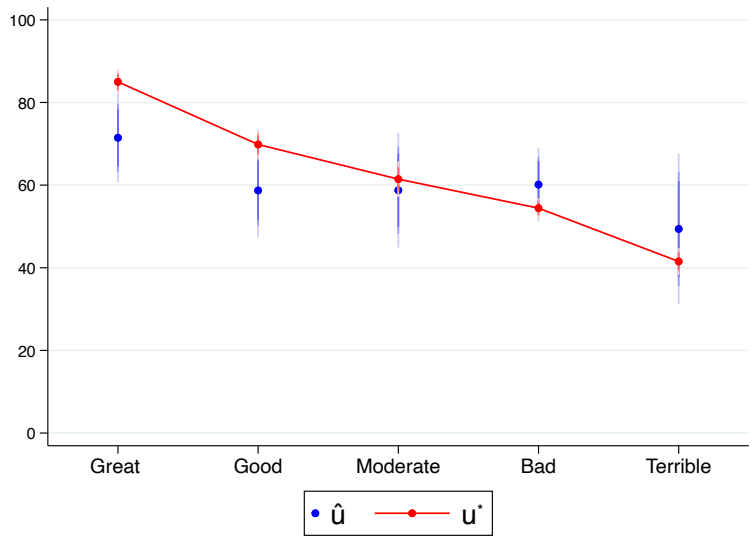


Figure 5: *Forecast Performance* treatment estimates.
Note: This figure presents estimates of central bank forecast credibility in *Forecast Performance* treatments. Blue and red shaded bands surrounding point estimates depict 99% (lightest), 95%, and 90% confidence (darkest) intervals.

precise central banks.[18] A central bank may not fully reap the reward, in terms of forecast credibility, of high forecast precision.

We explore the robustness of these results via regression analysis wherein we project individual-level perceptions of central bank forecast credibility, $\hat{u}_i$, onto a set of indicator variables denoting treatment (i.e. *Great*, *Good*, ..., *Terrible*), demographic characteristics, a subject's own uncertainty regarding future inflation, and controls for economic literacy. We depict the results of this exercise in Table 4. Estimates in column (1) of Table 4 correspond directly to the unconditional mean estimates of $\hat{u}_i$ depicted in Figure 5, column (2) additionally controls for forecast uncertainty, (3) layers in controls for demographics, and (4) includes controls based on our survey of economic literacy and for how much a participant trusts and understands the central bank. The main point of this table is that our baseline result – that subjects qualitatively behave like Bayesians – is robust.

Though outside the scope of our primary question, this result leaves us with an obvious question – why are subjects exhibiting a muted response to forecast precision relative to theoretical predictions? We provide some suggestive evidence in **??** that this results from subjects exhibiting a muted response to their own forecast uncertainty relative to what Equation (4) predicts.

5 Does *Timing* of forecast errors matter?

Our *Forecast Performance* results suggest that the central bank appears to be afforded some leniency when it makes worse forecast errors. This is important from a policy

[18] We show in Table A-7 that deviations from the equal-weighting Bayesian benchmark are statistically significant ($p < .01$) for *Great* and *Good* but not for *Bad* or *Terrible*.

Table 4: Regression Table for *Forecast Performance*

	(1) $\hat{u}_i$	(2) $\hat{u}_i$	(3) $\hat{u}_i$	(4) $\hat{u}_i$
Great	71.16***	71.13***	77.89***	69.74***
	(4.177)	(5.088)	(9.014)	(17.20)
Good	58.37***	58.34***	62.95***	55.18***
	(4.336)	(4.847)	(8.327)	(16.97)
Moderate	58.39***	58.36***	65.13***	58.35***
	(5.434)	(6.114)	(9.404)	(17.57)
Bad	59.70***	59.67***	65.03***	59.16***
	(3.421)	(4.288)	(7.297)	(16.05)
Terrible	47.31***	47.28***	53.49***	46.61***
	(6.967)	(7.665)	(11.06)	(17.62)
Uncertainty		0.000275	0.00126	0.00414
		(0.0241)	(0.0245)	(0.0260)
Demographics			✓	✓
Survey Responses				✓
N	528	528	520	520

Note: This table presents OLS regressions of forecast credibility ($\hat{u}_i$) across *Forecast Performance* treatments. Robust standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

perspective because it informs policymakers about the efficacy of central bank signals following inflation misses. However, these results may depend, quite critically, on the consistency of forecast errors throughout the *Consistent* economic histories. Thus, we next assess the extent to which the timing of forecast errors influences participants' perceptions of central bank forecast credibility using our *Timing* treatments. These treatments comprise all possible permutations of our three core economic histories, *Early*, *Late*, and *Consistent - Bad*. Recall that these economic histories feature identical historical forecast precision but allow for variation in the pattern of historical forecast errors. By comparing our measure of forecast credibility across histories, we learn whether and how the time profile of forecast errors impacts the central banks ability to influence inflation expectations via forecasting.

5.1 Timing-related Hypotheses

Given that the presented sample history is only 12-quarters long, we might expect subjects to use the full history to estimate the central bank's precision. If subjects equally weight all historical information when doing this, they estimate

$$\beta^{-1} = \frac{\sum_{j=1}^{j=12} \left| \mathbb{E}_{j-1}^{CB}\left(\pi_j\right) - \pi_j \right|}{12}. \tag{15}$$

If Equation (15) is correct, then the average level of perceived forecast credibility across *Early, Late*, and *Consistent* ought to be identical because average forecast errors are constant across histories. We summarise this into the following hypothesis:

Hypothesis 2. *Subjects weigh observed histories equally such that the timing of forecast*

errors does not lead to a difference in the average level of forecast credibility across economic histories in the Timing wave.

Though averaging across all three available years of historical performance seems like the natural thing to do (we would fail to reject the null of Hypothesis 2), the results of the previous section suggest that the participants place different weights on very large and very small errors. This does not necessarily translate into timing effects however; if they underweight (overweight) large (small) errors, but the timing does not matter in and of itself, then we might expect that the effects net out over *Early* and *Late* such that $\hat{u}_{Early} = \hat{u}_{Late} = \hat{u}_{Consistent}$.

There is, however, literature that suggests people exhibit time-dependency in economic decision-making in related contexts. Malmendier and Nagel (2016) show that people born at different times weight information differently due to differences in life experiences. Thakral and Tô (2021) show that expectations-based reference points adjust dynamically and exhibit recency bias. In the context of forecast credibility, we might observe significant differences in forecast credibility across these economic histories.

For instance, if the economic agent views the central bank's forecast credibility as ever-changing and accounts for this by more heavily weighting recent performance, then they might calculate β as:

$$\beta^{-1} = \lambda \sum_{j=0}^{j=11} (1-\lambda)^j |\mathbb{E}^{CB}_{t-2-j}(\pi_{t-1-j}) - \pi_{t-1-j}| \tag{16}$$

where the weighting function exhibits exponential decay in time. Figure 6 depicts the implied weighting functions from Equation (16) for different values of λ. This is akin to constant-gain learning models of expectation formation common in the learning literature (Evans et al. 2001).[19] This gives us our second *Timing* hypothesis:

Hypothesis 3. *Subjects exhibit recency bias when forming a perception of central bank credibility.*

5.2 Timing Treatments and Results

We test these hypotheses using a within-subject design that exposes each participant to some ordering of *Early, Late,* and *Consistent - Bad*. A within-subjects design allows us to we implement a full factorial design to nullify concerns about order and learning effects as potential confounds. This yields the *Timing* treatments described in Table 5. Note that, because we only use *Consistent-Bad* in this wave, we will refer to *Consistent - Bad* as *Consistent* throughout the remainder of the *Timing* section.

Figure 7 reports measures of perceived forecast credibility by economic history (blue circles) alongside the Bayesian optimal level of updating (assuming equal weighting of all historical information, red triangles) and the deviation from this Bayesian benchmark

[19] In that context, economists typically motivate these models as a way for an agent to account for structural change in whatever macroeconomic time series an agent is forecasting.

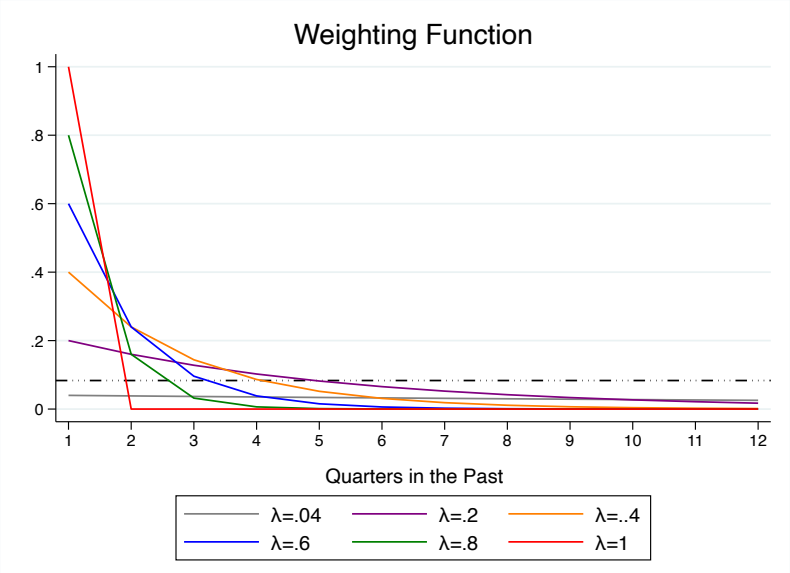


Figure 6: Weighting function (Equation (16)) for different values of λ

Note: This figure depicts hypothetical weighting functions for alternative values of λ in Equation (16). $\lambda = 1$ means a subject cares only about the most recent historical observation. As λ decreases, subjects begin to more equally consider the central bank's full forecasting history. The black horizontal line is the benchmark case, described by Equation (15), where a participant equally weights all twelve quarters of historical forecast information in our experiment.

Table 5: Treatment Summary: *Timing*

Treatment Summary: Timing				
	History 1	**History 2**	**History 3**	**Sample Size**
T6	*Early*	*Late*	*Consistent*	97
T7	*Early*	*Consistent*	*Late*	94
T8	*Late*	*Early*	*Consistent*	76
T9	*Late*	*Consistent*	*Early*	88
T10	*Consistent*	*Late*	*Early*	91
T11	*Consistent*	*Early*	*Late*	79

Note: This table summarizes our *Timing* treatments. Note that T6 and T8 are the same as T4a and T4b in Table 3.

(green diamonds). The mean of $\hat{u}_i$ responses and the standard error in parenthesis for *Early*, *Consistent*, and *Late* are 64.375 (3.523), 58.615 (2.142), and 11.494 (2.468) respectively. Two-sided t-tests comparing means for each pair of histories fails to reject that $\hat{u}$ is different in *Early* and *Consistent* ($p = 0.163$), but clearly rejects equality of the update in *Late* compared to the other two histories ($p < .001$ in both tests).

Participants put more weight on the most recent observations when forming perceptions of central bank forecast credibility in *Early* and *Late*. We shall refer to this as *recency bias*.[20] This is evidenced by the highly significant deviations from the Bayesian optimal benchmark in both economic histories, where deviations are positive in *Early* and negative in *Late*. This matches with subjects placing more emphasis on more recent information in both economic histories. Taken together, this indicates participants re-

[20]Of course, without more information about the process driving inflation, subjects may perceive a structural change in the forecasting performance of the central bank even if none has actually occurred.

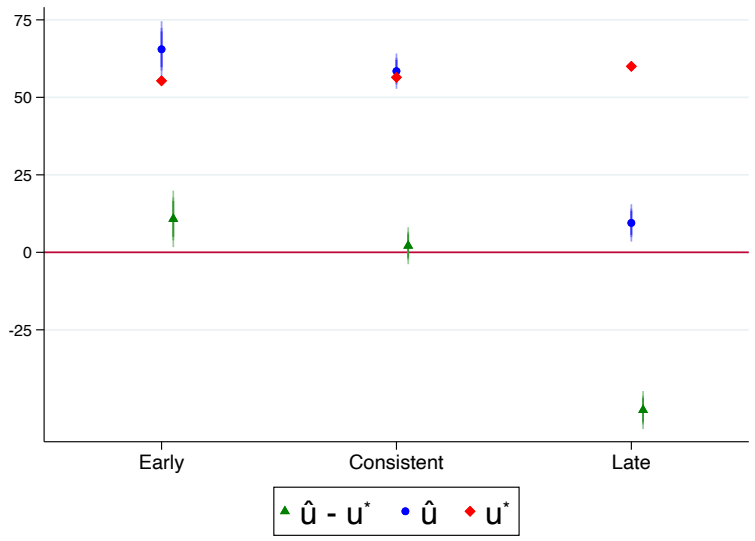


Figure 7: Estimated forecast credibility by economic history in *Timing*

Note: This figure depicts average treatment effects (blue circles) surrounded by confidence intervals ranging from 90 (darkest, thickest lines) to 99 percent (lightest, thinnest lines) from our *Timing* treatments using data Windsorized at the 5% level. Red triangles reflect the Bayesian benchmark weight, while green triangles show participants' average deviations from this benchmark, surrounded by confidence intervals ranging from 90 to 99 percent.

spond strongly to the timing of forecast errors when forming a perception of the central bank's forecast credibility. Thus, we reject Hypothesis 3 but fail to reject Hypothesis 4.

However, we also observe that the magnitude of the deviation from the Bayesian benchmark is about five times larger in *Late* than it is in *Early*. Recall that we design *Early* and *Late* so that forecasting performance changes at the same speed and by the same magnitude across these two economic histories. Thus, the only difference between them is whether the central bank's most recent forecast performance is better or worse than its historical average forecast performance. This sizable difference in deviation away from the Bayesian benchmark, which is based on equally weighting all available information, suggests an asymmetry in recency bias across the two histories. If this is the case, then it isn't just the *change* in forecast performance that drives recency bias. Instead, it could be that poor forecast performance is more salient for subjects whenever forming a perception of forecast credibility.

We also consider a within-subject measure of perceived forecast credibility, which we report in Figure 8. These measures present forecast credibility in *Early* and *Late* relative to *Consistent*. Specifically, for $X \in \{Early,\ Late\}$:

$$\hat{u}_{X,within} = \frac{1}{N_X} \sum_{n \in N} (\hat{u}_{n,X} - \hat{u}_{n,C}) \tag{17}$$

This approach effectively controls for a subject-level fixed effect since idiosyncratic biases are likely to be invariant to forecast history (for example, a participant's risk preferences are equivalent across the three economic histories). We observe the same pattern of relative updating between *Early* and *Late* in our within-subject measure of perceived forecast credibility as we do in Figure 7.

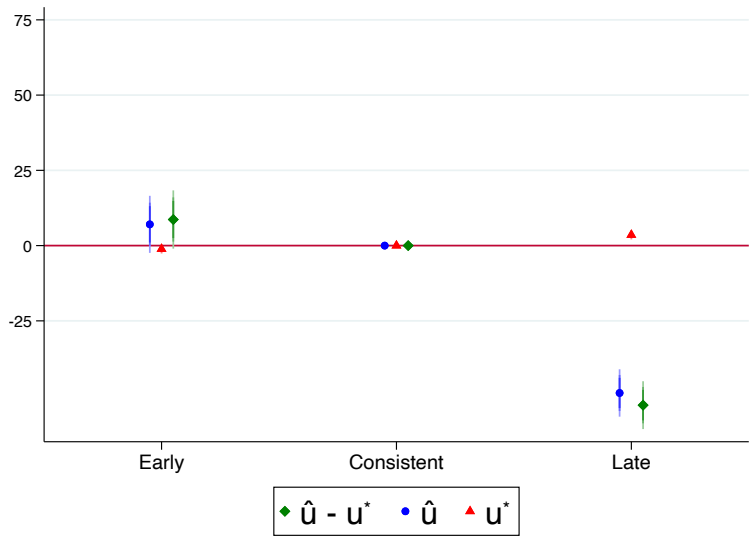


Figure 8: Within-Subject estimated forecast credibility by economic history in *Timing*
Note: This figure presents the average within-subject treatment effects on forecast credibility from the *Timing* treatments, using the 'Consistent' history as a control for idiosyncratic biases. Blue circles illustrate the average treatment effects, bounded by 90 to 99 percent confidence intervals using data Windsorized at the 5% level. Red triangles reflect the Bayesian benchmark weight, while green triangles show participants' average deviations from this benchmark, surrounded by confidence intervals ranging from 90 to 99 percent.

As we did in *ForecastPerformance*, we further explore *Timing* results via regression analysis to understand how perceptions of central bank forecast credibility relate to individual characteristics, to forecast uncertainty, and to our measures of economic literacy. Our approach here is identical to that in *Forecast Performance*. We show these results in Table 6, where column (1) corresponds directly to the treatment effects depicted in Figure 7, column (2) introduces forecast uncertainty as a control, column (3) includes demographic controls, and column (4) includes controls for economic literacy and how much participants trust and understand the central bank. We note that our baseline estimates change very little in terms of magnitude or statistical significance when we introduce controls.

5.3 Exploring the Nature of Recency Bias

Our *Timing* results show that participants exhibit recency bias, more heavily weighting more recent information when forming a perception of the central bank's forecast precision. Further, we observe a large difference in the magnitude of deviations from the Bayesian benchmark between *Early* and *Late*. This leads us to ask: how strong is this recency bias that we observe and how is it different across economic histories? To answer this, we estimate λ in Equation (16) for these two economic histories.[21] We do

[21] Combining $\beta^{-1} = (\alpha \times \hat{u})^{-1}(1-\hat{u})$ with Equation (16), gives:

$$\lambda \sum_{j=0}^{j=11} (1-\lambda)^j |\mathbb{E}^{CB}_{t-2-j}(\pi_{t-1-j}) - \pi_{t-1-j}| - (\alpha \times \hat{u})^{-1}(1-\hat{u}) = 0.$$

which we solve for λ via numerical approximation.

Table 6: Regression Table for *Timing*

	(1) $\hat{u}$	(2) $\hat{u}$	(3) $\hat{u}$	(4) $\hat{u}$
Early	64.38***	56.77***	50.69***	62.34***
	(3.525)	(4.332)	(7.023)	(14.09)
Consistent	58.61***	50.97***	44.40***	56.05***
	(2.143)	(3.515)	(6.637)	(13.73)
Late	11.49***	2.192	-3.639	8.175
	(2.470)	(4.058)	(7.008)	(13.85)
Uncertainty		6.363**	5.908**	5.288**
		(2.523)	(2.547)	(2.515)
Demographics			✓	✓
Survey Responses				✓
N	1548	1548	1518	1518
Clusters	516	516	506	506

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

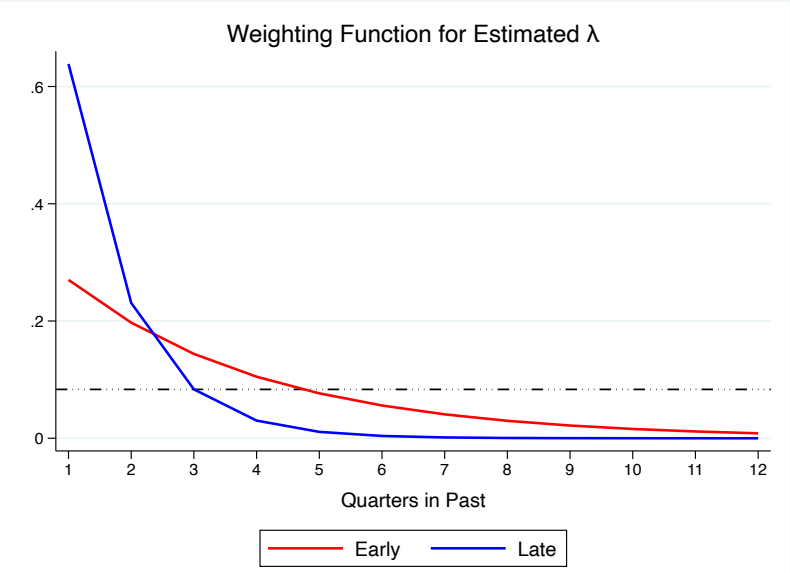


Figure 9: Estimated weighting functions

Note: This figure depicts weighting functions using values of λ that result from estimating λ in Equation (16) using our experimental data. This figure indicates that individuals exhibit significant recency bias and, on average, about threefold more recency bias in *Late* than in *Early*.

not show estimates for *Consistent* since, by design, estimating β^{-1} using the last year of information should be identical to estimating it over the last two years, all years, etc.

The estimated values are $\lambda_{Early} = 0.245$ (standard error is 0.0170) and $\lambda_{Late} = 0.622$ (standard error is 0.0198). This suggests that participants exhibit recency bias in both *Early* and *Late*. However, the degree of recency bias exhibited in *Late* is considerably higher than in *Early*. $\lambda_{Late} = 0.622$ implies that the average participant in *Late* bases approximately 97.95% of her perception of central bank forecast precision on the most

recent four quarters of economic data. By contrast, this same number for the average participant in *Early* ($\lambda_{Early} = 0.245$) is about 67.5%. This difference in λ_{Late} and λ_{Early} is highly significant ($p < .001$, t-test). Using these estimates values of λ, we depict estimated weighting functions for *Early* (solid red line) and *Late* (blue dashed line) alongside an equal-weighting benchmark (black dashed line) in Figure 9.

So far, we have shown that the time profile of forecast errors can significantly change how participants use historical information when forming a perception of central bank forecast credibility. Further, we find a stark asymmetry in the extent of recency bias depending upon whether the central bank's most recent performance is significantly better or worse than its historical average performance. Finally, we find that the magnitude of the average deviation from our equal-weighting Bayesian benchmark is about five-fold larger in *Late* than in *Early*.

Our experimental findings in *Timing* underscore a crucial point: the central bank cannot rest on its laurels in the face of unanticipated shocks that precipitate a sharp decrease in forecast performance, even if confined to short-run forecast performance. Further, such changes in forecast performance undermine the bank's forecast credibility precisely when the benefit of controlling and guiding inflation expectations is highest. Public attention will gravitate toward the recent subpar performance, overshadowing prior successes. Additionally, our results reveal that while a rebound in forecast performance following a negative shock to forecast precision can aid in restoring credibility, the pace of recovery is invariably slower than the rate at which credibility initially eroded. Given this, it is paramount that central banks guard against and react swiftly to negative forecasting episodes to prevent credibility loss because the cost of rebuilding credibility may significantly outweigh the cost required to maintain it.

5.4 Dynamics of Perceived Credibility

Our experimental design also allows us to study the dynamics of perceived forecast credibility by comparing episodes where the central bank's historical forecast precision over the full three-year economic history exactly matches the forecast precision of the most recent year of economic history. We can do this for very bad forecast precision and very good forecast precision. We use the fact that the central bank's historical forecast precision in *Terrible* from *Forecast Performance* is identical to the central bank's forecast precision in the final year of *Late* from *Timing*. Similarly, we can also use the fact that the central bank's historical forecast precision in *Great* is identical to the bank's historical forecast precision in the final year of *Early*.

Turning first to our low-credibility scenario, Figure 10 visualizes the kernel density functions of $\hat{u}$ in *Terrible* (blue line) and *Late* (red dashed line), along with their means and corresponding standard errors. We underscore two crucial observations here.

First, we highlight the counter-intuitive finding that *Terrible* yields a higher average forecast credibility than *Late*, despite historical forecast precision being much higher in *Late*. Second, we observe that *Terrible* induces a higher level of disagreement or variation in $\hat{u}$ than *Late*. These observations combined reveal that the visible decrease in forecast precision in *Late* leads participants to significantly undervalue the central bank's

inflation forecast *systematically* whereas the consistent pattern of forecast errors in *Terrible* yields greater disagreement on how to use the central bank's signal. A considerable number of participants lean heavily on the central bank's forecast when constructing their posterior beliefs about inflation even though the central bank's historical precision is quite low. This aligns with our main result in *Forecast Precision*, where participants tend to overvalue the central bank's inflation forecasts in both *Bad* and *Terrible*.

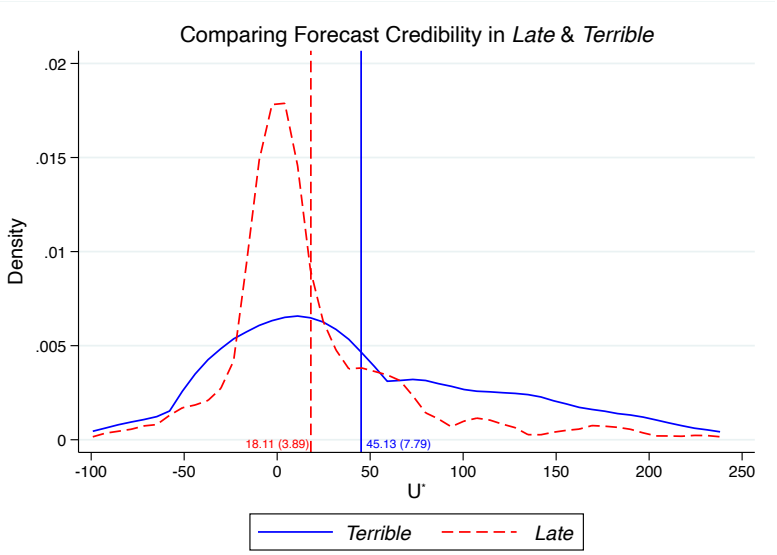


Figure 10: *Consistent-Terrible* vs. *Late*

Note: This figure depicts kernel density estimates of forecast credibility, $\hat{u}$, from *Late* (red, dashed line) and *Terrible* (blue, solid line). The numbers present the mean (corresponding standard errors).

We next consider our high-credibility scenario. Figure 11 visualizes the kernel density functions of $\hat{u}$ in *Great* (blue line) and *Early* (red dashed line). The results here are the opposite of those in our low-credibility scenario. Consistently high forecast precision in *Great* leads to most participants viewing the central bank as highly credible. By contrast, observing that the central bank exhibited poor forecast precision in its recent history leads many subjects to view the bank as less credible. However, the shift from low to high forecast precision in *Early* induces the same sort of disagreement about how to use the central bank's signal observed in *Terrible*.

The overarching message from these findings is that both the consistency and the recentness of the central bank's forecast errors significantly influence participants' perceptions of the bank's forecast credibility. Further, these findings suggest that consistency, even in error, may instill a sense of forecast credibility among participants, which would explain why so many subjects in *Terrible* place considerable emphasis on the central bank's signal when forming updated inflation expectations. However, the shift from low to high forecast precision in *Early*, akin to *Terrible*, introduced disagreements among the participants on how to interpret the bank's signal. In essence, these findings illuminate the nuanced role that consistency, recent performance, and changes in precision play in shaping the credibility of the central bank's forecast. This indicates that both past performance and recent changes matter significantly when it comes to trust in forecasts, and highlights the challenge for central banks in maintaining their credibility.

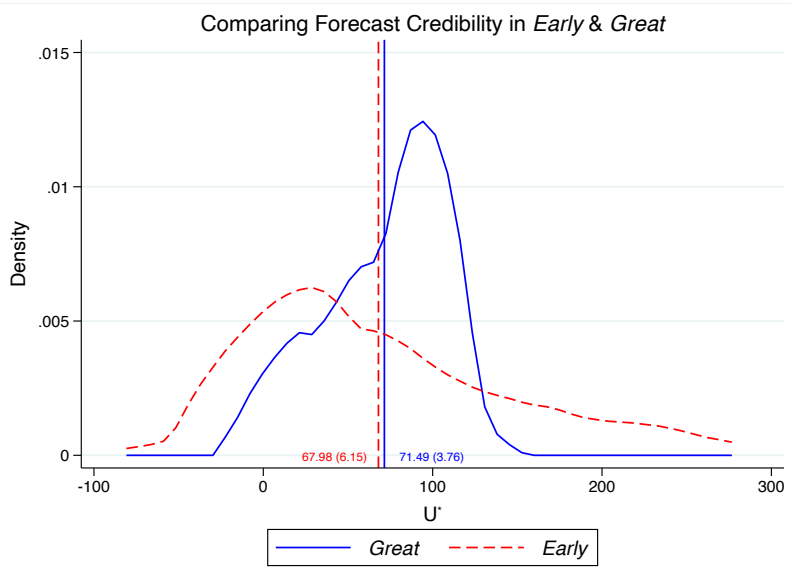


Figure 11: *Consistent-Great* vs. *Early*

Note: This figure depicts kernel density estimates of forecast credibility, $\hat{u}$, from *Early* (red, dashed line) and *Great* (blue, solid line). The numbers present the mean (corresponding standard errors).

5.5 Robustness

5.5.1 Does the direction of forecast error matter?

So far, we have only tested scenarios where the central bank's historical inflation forecasts were too low. A possible concern then is that our results regarding the relationship between the central bank's historical forecast performance and forecast credibility might vary depending on whether the bank over or underestimates inflation. For example, this could be true if people perceive the costs of inflation as asymmetric. An agent who views high inflation as more costly than low inflation might be less forgiving of under-forecasting inflation than over-forecasting inflation. This is because the former case leads to misspecified expectations that are violated by an even more costly scenario than anticipated whereas the latter to a case where expectations are violated but realized inflation is less costly than anticipated. This would align with Guido et al. (2022), which shows that people harshly punish disappointing violations of expectations but not beneficial violations of expectations.

To address this concern, we implement a subset of our original treatments that are identical in every way but reverse the direction of forecast errors. We refer to these treatments as our *Timing: Reversed Shock* treatments. We consider our baseline *Timing* results robust to the direction of forecast errors if we are able to qualitatively reproduce our baseline *Timing* results. We show results from these treatments in Figure 12.

Similar to baseline results, we find that the time profile of forecast errors matters. Our participants place significant emphasis on more recent information in our *Early* and *Late* treatments, as evidenced by the over-updating in *Early* and under-updating in *Late* relative to the Bayesian benchmark.

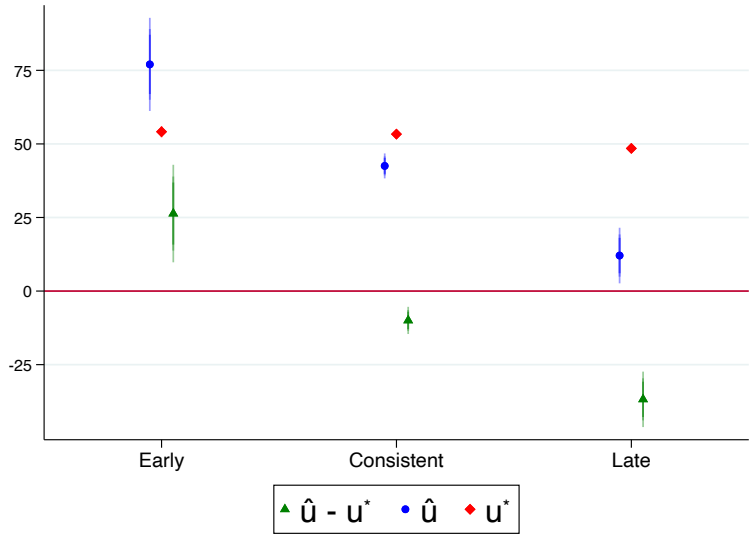


Figure 12: Estimated forecast credibility in *Timing: Reversed Shock* treatments.
Note: This figure depicts average treatment effects (blue circles) surrounded by confidence intervals ranging from 90 (darkest, thickest lines) to 99 percent (lightest, thinnest lines) from our *Timing: Reversed Shock* treatments using Windsorized data (5% level). Red triangles reflect the Bayesian benchmark weight, while green triangles show participants' average deviations from this benchmark, surrounded by confidence intervals ranging from 90 to 99 percent.

5.5.2 Does this hold for the medium term?

Another possible concern deals with the forecasting horizon of our experiment. Up to now, we've considered how perceptions of the central bank's short-term forecast credibility respond to features of its historical short-term forecast performance. While there is evidence that short-term forecasts for inflation typically correlate with longer-term expectations (Weber et al. 2022), a natural robustness check is whether and how our *Timing* results depend on the forecast horizon. Why? Policymakers are often more concerned with managing longer-term expectations than short-term expectations, which allows them to pursue short-term stabilization policies in response to transitory shocks. Therefore, we also consider how historical forecast performance impacts the central bank's ability to control longer-term inflation expectations.

To do this, we implement our *Early*, *Consistent*, and *Late* economic histories from our baseline *Timing* treatments but elicit expectations of average inflation over the next three years. We call these our *Timing: Medium-Term* treatments. Additionally, we provide subjects in these treatments with the central bank's outlook on average inflation over the same three-year horizon between the Initial and Updated forecasts. To do this, we average the short-run forecasts of the central bank while allowing our mean-reverting shocks to dissipate over time. We accompany these changes in the forecast and signal horizon with corresponding changes in our instructions and graphical interface. Finally, we included an additional question in our comprehension quiz to ensure participants understood the forecast horizon.[22]

Our interest is in whether the impact of short-term forecast performance on perceived

[22]We include these instructions, screenshots of the updated graphical interface, and the text of this new question in our appendix.

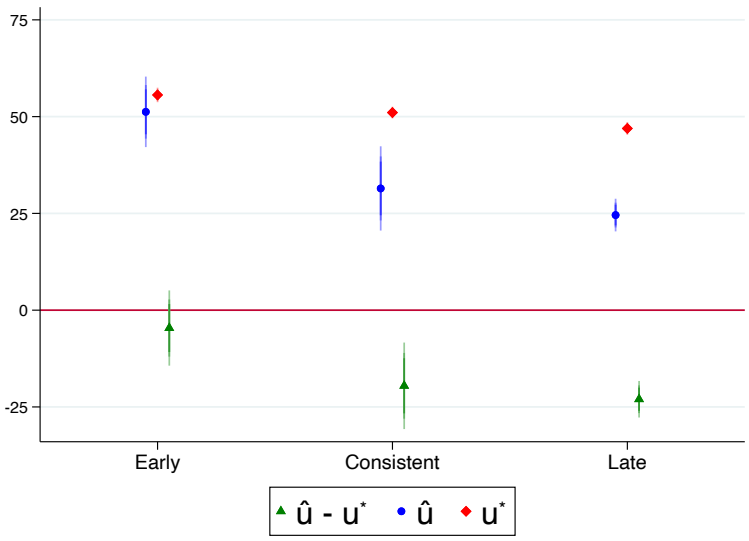


Figure 13: Estimated forecast credibility in *Timing: Medium-Term* treatments.
Note: This figure depicts average treatment effects (blue circles) surrounded by confidence intervals ranging from 90 (darkest, thickest lines) to 99 percent (lightest, thinnest lines) from our *Timing: Medium-Term* treatments using Windsorized data (5% level). Red triangles reflect the Bayesian benchmark weight, while green triangles show participants' average deviations from this benchmark, surrounded by confidence intervals ranging from 90 to 99 percent. These estimates omit individuals for whom $\hat{u} = 0$ and their initial forecast matches the central bank's forecast.

forecast credibility is limited to the short-term or if it affects perceptions of the central bank's longer-term inflation forecast credibility. The idea is similar to Carvalho et al. (2023), which shows that an economic agent's long-term inflation expectations respond endogenously to short-term forecast errors and that the strength of this relationship depends critically on the agent's historical forecast performance. Our litmus test is whether or not we replicate the qualitative findings observed in our baseline *Timing* treatments. That is, we consider our primary results – that people exhibit recency bias when forming a perception of central bank forecast credibility – robust if we qualitatively replicate that result in these treatments.

Note that implementing this litany of changes simultaneously – forecast horizon, signal horizon, instructions, and graphical interface – means that we cannot make causal claims comparing short-term and medium-term *Timing* treatments. Because of this, we are not overly concerned with quantitative differences that may arise across analogous short-term and medium-term treatments.

We show results from these treatments graphically in Figure 13, which depicts a history-contingent pattern of updating that is consistent with our short-term *Timing* treatments.[23] Our results in *Timing* are robust to longer-term forecast horizons.

[23]In our *Early* data from these medium-term treatments, we observe 100 instances where a participant's initial point forecast of inflation perfectly coincides with the central bank's inflation forecast. Of these 100 instances, there are 75 observations where the participant does not update her point forecast of inflation after viewing the central bank's forecast. In Figure 13, we treat these observations as fully credible, which aligns with how we would treat such instances in theory since they coincide perfectly with the central bank's signal. If we instead exclude these observations, the results are qualitatively similar.

6 Theoretical Implications

6.1 Endogenous Credibility In the New Keynesian Model

Our experimental results suggest some ways to modify common features of workhorse theoretical models. First, the central bank's forecast is not necessarily fully credible. Rather, central bank forecast credibility can evolve endogenously. Second, economic agents do not incorporate all available historical information when forming a perception of the central bank's forecast credibility. Instead, agents' perception of central bank credibility places more weight on temporally-recent information. This is especially true following large shocks that lead to large forecasting errors. This asymmetry in the degree of recency bias means that forecast credibility is harder to build than lose. Third, agents tend to exhibit over-precision, which leads them to under-react to their own forecast errors whenever forming a perception of the central bank's forecast credibility. Below, we incorporate our experimental findings into an otherwise standard New Keynesian model to demonstrate how endogenous forecast credibility coupled with recency bias and an overly-precise economic agent can impact aggregate inflation dynamics.

Before proceeding, we first note that we are not the first to consider the impact of endogenous central bank credibility on aggregate dynamics. Hommes and Lustenhouwer (2019) introduces a model of endogenous central bank credibility, whereby a continuum of heuristic-switching agents form either naive inflation expectations or inflation expectations that align with the central bank's inflation target. Credibility in their model is endogenous in the sense that the proportion of agents anchored onto the central bank's inflation target depends on the bank's historical ability to achieve its target. Goy et al. (2022) extend Hommes and Lustenhouwer (2019) to include N-step learning to study the effectiveness of forward guidance in a world populated by heterogeneous, boundedly-rational agents. Relative to these studies, our focus is on the relationship between historical forecast performance and the central bank's ability to influence inflation expectations via the publication of its own subjective outlook on inflation.

6.1.1 A Behavioural Expectations New Keynesian Model

We start with a simple three-equation New Keynesian model linearized around a zero-inflation steady state described by Equation (9) through Equation (12) as was used to calibrate our experiment histories. As in that case, the shocks are assumed to follow AR(1) processes with persistence ρ_g, ρ_p and ρ_v. We calibrate the main model parameters with the same parameter values used in the model to generate our experimental scenarios (described in Table 1). The only deviation from this model that we explore is to allow household inflation expectations to potentially deviate from rationality.

We assume that a Bayesian household[24] forms a composite expectation of π, where some

[24]We refer to a single household for ease of exposition. You could also think of a continuum of households where ξ is the proportion of households forming noisy rational expectations and $1-\xi$ the proportion adhering to a backward-looking, adaptive heuristic.

proportion ξ of the household's expectation is model consistent but noisy

$$\mathbb{E}^{\text{HH}}_{t,\text{RE}}\{\pi_{t+1}\} = \pi^{\text{RE}}$$

The remaining $1-\xi$ of the household's belief follows a backward-looking, adaptive heuristic of the form:

$$\mathbb{E}^{\text{HH}}_{t,\text{BR}}\{\pi_{t+1}\} = \mathbb{E}_{t-2}\{\pi_{t-1}\} + \eta_\pi(\pi_{t-1} - \mathbb{E}_{t-2}\{\pi_{t-1}\})$$

where $\eta_\pi \in [0,1]$ determines how strongly the agent reacts to past forecast errors. Note that $\eta = 1 \implies \mathbb{E}^{\text{HH}}_{t,\text{BR}}\{\pi_{t+1}\} = \pi_{t-1}$, which means the household forms expectations naively in a backward-looking fashion. Alternatively, $\eta_x = 0 \implies \mathbb{E}^{\text{HH}}_{t,\text{BR}}\{\pi_{t+1}\} = \mathbb{E}_{t-1}\{\pi_t\}$, which means the household's expectations are completely anchored. For brevity, we refer to these expectations as AD(1). Intuitively, these forecasting heuristics say that the agent forms her initial expectations in each period t by adjusting her most recent forecasts of inflation and output to account for her most recent forecasting errors.

Thus, the household's prior is a noisy linear combination of a rational component and an adaptive component:

$$\mathbb{E}^{\text{HH}}_{t,\text{prior}}\{\pi_{t+1}\} = \xi\mathbb{E}^{\text{HH}}_{t,\text{RE}}\{\pi_{t+1}\} + (1-\xi)\mathbb{E}^{\text{HH}}_{t,\text{BR}}\{\pi_{t+1}\} + \epsilon_{\text{HH}}, \quad \epsilon_{\text{HH}} \sim \mathcal{N}\left(0, \frac{1}{\alpha}\right).$$

The household holds a belief about her own forecast precision, α. Notice that $\xi = 1$ implies the household forms noisy rational expectations while $\xi = 0$ means the household is boundedly rational.

The central bank in this economy publishes model-consistent forecasts of π in each period, which the household can use to update its belief about π before committing to actions:

$$\mathbb{E}^{\text{CB}}_t\{\pi_{t+1}\} = \pi^{\text{RE}} + \epsilon_{\text{CB}}, \quad \epsilon_{\text{CB}} \sim \mathcal{N}\left(0, \frac{1}{\beta}\right).$$

Given this, the household forms an updated expectation (i.e. a posterior) according to:

$$\mathbb{E}^{\text{HH}}_{t,\text{post}}\{\pi_{t+1}|\mathbb{E}^{\text{CB}}_t\pi_{t+1}\} = (1-\Gamma_{\pi,t})\left(\mathbb{E}^{\text{HH}}_{t,\text{prior}}\{\pi_{t+1}\}\right) + \Gamma_{\pi,t}\left(\mathbb{E}^{\text{CB}}_t\{\pi_{t+1}\}\}\right) \tag{18}$$

where $\Gamma_{\pi,t} = \frac{\tilde{\beta}_{\pi,t}}{\tilde{\alpha}_{i,t}+\tilde{\beta}_{\pi,t}}$.[25] Notice in Equation (18) that $\Gamma_{\pi,t}$ determines how strongly the household's updated inflation expectation depends on the central bank's forecast. Thus, $\Gamma_{\pi,t}$ constitutes a measure of central bank forecast credibility in our model.

$\Gamma_{\pi,t}$ depends on the households belief about its own forecast precision, $\tilde{\alpha}^{-1}$, and the central bank's forecast precision, $\tilde{\beta}^{-1}$. We assume, consistent with our results, that the household forms a subjective outlook on the central bank's actual forecast precision, β, approximating its outlook on the central bank's forecast credibility. This serves as

[25] Note the timing here assumes the agent cannot observe contemporaneous inflation and instead can only adjust expectations based on its forecast of $t-1$ inflation relative to realized inflation. We have explored how our results depend on this timing assumption by simulating dynamics assuming the agent adjusts expectations based on its forecast and observation of contemporaneous inflation. The main qualitative findings still hold.

the channel through which the central bank's forecasting history influences its forecast credibility and the latter where the household's biased perception of its own forecast precision can further distort credibility from the Bayesian benchmark. Households may exhibit recency bias when evaluating the central bank's forecast precision. Specifically, we assume the household only uses information from the central bank's N most recent forecasts and cares most about the most recent information. This gives

$$\tilde{\beta}_{\pi,t}^{-1} = \sum_{k=1}^{k=N} \lambda_k \big(\| \mathbb{E}_{CB,t-k}\{\pi_{t-k}\} - \pi_{t-k} \| \big) \tag{19}$$

where $\sum_{k=1}^{k=N} \lambda_k = 1$ and $\lambda_k > \lambda_{k-1} > \lambda_{k-2} > ... > \lambda_{k-N}$.

We use this simple model to consider the implications for inflation of an unanticipated cost-push shock under different assumptions about expectation formation and, particularly, the role that endogenous credibility and recency bias play. Endogenous credibility can matter for at least two reasons. First, the economic agent will not necessarily view the central bank's inflation projections as fully credible. Second, the agent's perception of the central bank's forecast credibility can change over time depending on the central bank's forecast performance.

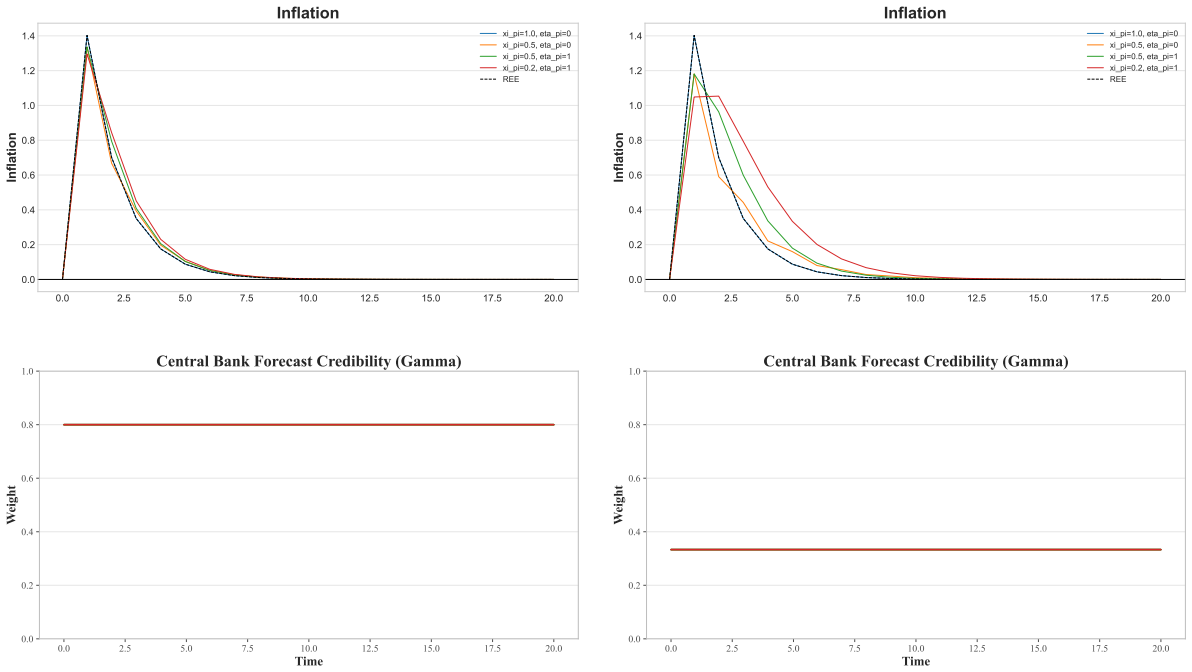
6.1.2 Fixed Credibility

We first consider the effects of the cost-push shock in an environment of fixed credibility. We begin with the assumption that the central bank forms more precise beliefs than that household. Specifically, we assume that $\alpha = 5$ and $\beta = 20$ which means that the optimal weight on the central bank forecast is $\Gamma_{\pi,t} = 0.8 \; \forall \; t$. We then explore the impact of four different forms of household expectation formation:

1. $\xi = 1$ so that the household forms noisy rational expectations too.
2. $\xi = 0.5$ so that there is a mass of boundedly-rational households with $\eta_x = 0$ such that these households' expectations are completely anchored.
3. $\xi = 0.5$ but $\eta_x = 1$ such that these household's expectations are completely backward-looking.
4. $\xi = 0.2$ and $\eta_x = 1$ such that these household's expectations are completely backward-looking.

We show the results of this exercise in panel (a) of Figure 14, where the top panel traces the response of aggregate inflation to this cost-push shock under the different assumptions about ξ, η. The lower panel shows the fixed values for $\Gamma_{\pi,t}$. The primary takeaway from this exercise is that with a fixed, and relatively high, weight on the noisy-rational central bank signal, the impact of different forms of non-rationality of the households matters less for inflation dynamics following an unanticipated inflation shock.

To illustrate the cost of low forecast credibility, we reconsider these same parameterizations of ξ, η but assume the household perceives the central bank's forecast credibility to



(a) $\Gamma = .33$ **(b)** $\Gamma = .33$

Figure 14: Responses to a cost-push shock under different levels of fixed forecast credibility. *Note*: This figure depicts the response of inflation and credibility to a cost-push shock when central bank forecast credibility is fixed at $\Gamma = .8$ and $\Gamma = .33$. For each fixed level of forecast credibility, we consider the combinations of ξ, η described in Section 6.1.2.

be lower than its own. Specifically, we set $\alpha = 10$ and $\beta = 5$ which yields $\Gamma_{\pi,t} = 0.33 \; \forall \; t$. We depict the results of this exercise in panel (b) of Figure 14. This decrease in Γ significantly erodes the central bank's ability to influence the inflation expectations of a boundedly rational agent, which leads to considerably more inflation persistence following the cost-push shock, relative to the world in which the central bank better maintains its forecast credibility.

6.1.3 Inflation shocks with variable credibility

We now explore the effect on inflation when the exogenous cost-push shock leads to forecast errors and the households update their perception of β, but α remains fixed at its true level. As a result of changes to $\tilde{\beta}$, the shock changes Γ. The question is whether the resulting dynamics of Γ lead to meaningful changes in inflation persistence.

To do this, we assume the agent bases her estimates of $\tilde{\beta}$ on the four most recently observed periods where $\lambda_k = [.63, .238, .09032, .03904]$ for $k = \{1, 2, 3, 4\}$, respectively. Note that these are the values produced by Equation (16) assuming λ=.622 and combined account for 99.73% of the total weighting. If $t < 4$, we use the same weighting scheme for all available λ_k and, at all horizons, normalize the weights so that they sum to one.

We show the results of this in Figure 15 where we again vary the assumptions on how the household updates its beliefs. Interestingly, allowing $\tilde{\beta}$ to adjust endogenously follow the initial credibility loss does little to speed up recovery relative to our exercise in panel

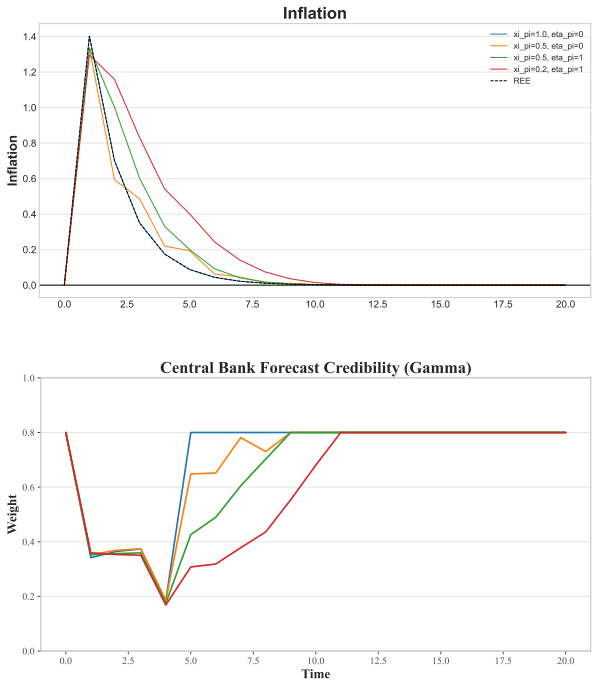


Figure 15: Response to cost-push shock with endogenous credibility.
Note: This figure shows the response of inflation and forecast credibility to a 1pp inflation shock for all combinations of $\xi = [0, .5]$ and $\eta = [0.5]$.

(b) of Figure 14, where $\tilde{\beta}$ never adjusts and Γ remains fixed at a relatively low-level fur the duration of the inflation shock. This is because the initial loss of credibility leads to large and persistent forecast errors, especially for small values of ξ, which translate into persistently low levels of Γ.

7 *Contextual Communication*

As a final exercise, we ask whether contextualizing communication that reinforces the central bank's inflation outlook and rationalizes its forecasting history can influence forecast credibility. From a practical perspective, this question is important because most central banks devote considerable resources to crafting and publishing contextualizing communication, often alongside projections. For example, alongside forecasts, inflation-targeting central banks typically publish a monetary policy report, and the information contained in these reports has been shown to influence market behavior (Hansen et al. 2019). Can this sort of communication enhance forecast credibility? If so, to what extent? And what sort of messaging is most effective in that it enables the central bank to talk its way out of a low-credibility position?

7.1 Contextual Communication Treatments

To answer these questions, we incorporate contextualizing communication into *T7* from our *Timing* treatments. Specifically, alongside the final history (*Late)*, we publish a written statement alongside the central bank's graphical forecast before allowing subjects

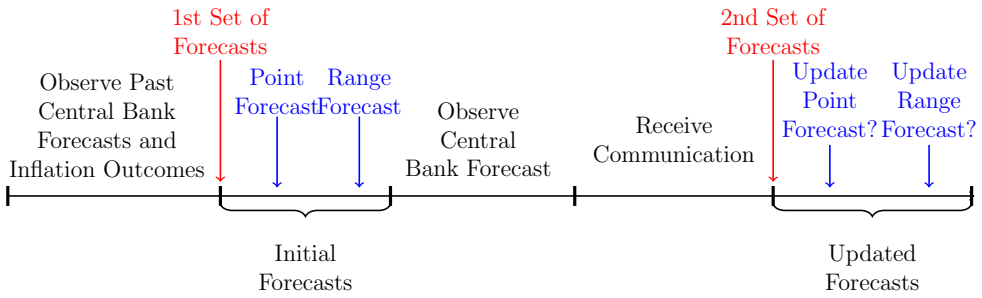


Figure 16: Timing of decision period in *Contextual Communication* treatments

Note: This figure depicts the flow of a single decision period in the *Contextual Communication* treatments of our experiment. The only difference here relative to any other treatments is that *Contextual Communication* treatments include a written statement from the central bank alongside the central bank's numerical inflation forecast.

to update their inflation expectations. We focus on *T7* where participants experience *Late* last, which allows them maximal time to learn the experimental environment before encountering written communication. We summarize the timing of decision periods in *Contextual Communication* in Figure 16.

Written statements in these treatments convey information about the central bank's inflation outlook, whether the source of poor historical forecast performance in *Late* is endogenous or exogenous, and about the central bank's forecast performance relative to peer forecasting institutions. Arguably, introducing written statements alongside projections increases information complexity from the perspective of participants. To address this, we also include a control text that provides generic information about the central bank. Because our control text is uninformative about the central bank's outlook or its forecasting history, we can use this control text as a baseline for our treatments that either reinforce the bank's outlook or rationalize its forecasting history. We provide additional details about each statement in Figure 17 and provide the full text of each statement in appendix A3.

Control We provide a general description of central banking.

Control + Outlook Repeats text from *Control* but also includes a written outlook on inflation that matches the graphical forecast and adds no new information for participants. This allows us to discern whether reinforcing graphical information via text can better convey important economic information.

Exogenous + Relative Performance *Control + Outlook* but includes an additional paragraph explaining that the decline in historical forecast performance resulted from exogenous forces and also says whether the bank performed better or worse than peer forecasting institutions.

Endogenous + Relative Performance As *Exogenous + Relative Performance* except that the central bank explains the decline in historical forecast performance resulted from endogenous forces.

Figure 17: *Contextual Communication* treatments

Note: This table provides general descriptions of the statements provided to subjects in our *Contextual Communication* treatments.

Additionally, we developed our contextual communication treatments so that complexity is roughly identical across the written statements where we measure complexity using the Flesch-Kincaid reading grade level.[26] We summarize these treatments in Table 7.

Treatment Summary - Communication				
	Name	**Sample Size**	**Flesch-Kincaid**	
			Score	*Reading Level*
T12	*Control*	160	8	10th-12th
T13	*Control + Outlook*	151	8.3	10th-12th
T14	*Exogenous + Better*	131	8.5	10th-12th
T15	*Exogenous + Worse*	152	8.5	10th-12th
T16	*Endogenous + Better*	157	8.4	10th-12th
T17	*Endogenous + Worse*	137	8.4	10th-12th

Table 7: Treatment summary for *Contextual Communication*

Note: The table summarizes our *Contextual Communication* treatments. Subjects completed a single treatment, described by rows *T12* through *T17*. Note that we include Flesch-Kincaid scores (*Score* column) and reading levels (*Reading Level* column) to demonstrate the similarity in the complexity of the written statement included in each of our six *Contextual Communication* treatments.

7.2 The Effects of Communication

We present estimates of central bank forecast credibility in *Contextual Communication* in Figure 18.[27] For ease of interpretation, we center forecast credibility in *Control* to zero and then normalize estimates of forecast credibility in the remaining treatments to *Control*. Thus, this figure reports estimates of effects relative to *Control*.[28]

Our primary result is that contextualizing communication can significantly increase the central bank's forecast credibility. Interestingly, this is true even with *Control + Outlook*, which reinforces the central bank's inflation outlook without providing any new information to our participants. This could be because participants are better at extracting qualitative or narrative information from text, if the process of reading text yields a better synthesis of information, or if simply seeing the information again but in text form somehow reinforces learning. Additionally, it could be that being seen to attempt to communicate helpfully is beneficial for the central bank's reputation as suggested in Haldane and McMahon (2018).

Taken together, these findings suggest that the design and the delivery of central bank communication are important. Narrative communication can provide important and useful context in the delivery of forecast performance, especially where that performance

[26]The Flesch-Kincaid reading grade level is a readability metric that gauges the complexity of a text and estimates the grade level at which an individual can comprehend a piece of writing. The formula takes into account the total words, sentences, and syllables in a text, producing a numerical score that corresponds to a U.S. grade level. McMahon and Naylor (2023) discuss alternative measures of conceptual complexity.

[27]Table A-4 presents the regression results.

[28]Note that updating in *T12* is lower than in its *Timing* counterparts that do not include any narrative communication, which is what rational inattention or other theories of costly information processing might predict. Because of this, Bayesian benchmark values are slightly higher here than in other instances where we

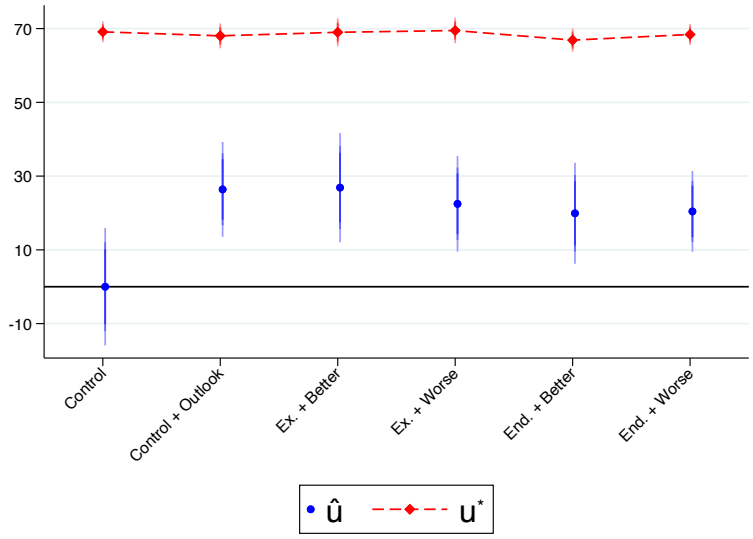


Figure 18: Estimated forecast credibility in *Contextual Communication* treatments.
Note: Forecast credibility (blue circles) in *Contextual Communication* treatments. Red diamonds indicate the optimal weighting of the central bank's forecasting in the updated inflation expectation of the average participant, assuming this participant is a rational Bayesian who equally weights all available historical information. Shading around both types of markers indicate 99%, 95%, and 90% confidence intervals.

may not be so strong for a period of time. For example, narrative communication may help safeguard against forecast credibility loss following exogenous shocks. This could ensure that the dynamics of the economy more closely resemble those in panel (a) of Figure 14, rather than those in Figure 15. Exploring this further is an important area for future research.

8 Conclusion

Central bank communication has emerged over the last few decades as a mainstay of central banking because it offers policymakers an effective way to manage expectations. A key component of communication is the central bank's economic outlook, which banks often publish as forecasts of key economic variables. Yet, this newly-established tool carries with it new concerns. Primarily, policymakers must now worry about how best to build and safeguard their forecast credibility so that publishing forecasts and communicating about their economic outlook remains potent. Though we know in practice that policymakers care deeply about their credibility (Blinder 2000), much less is known about the determinants and dynamics of this credibility. To address this shortcoming, we have used a novel experimental framework to study the causal relationship between features of historical forecast performance and forecast credibility.

We show that the link between historical forecast performance and forecast credibility is not as sharp as theory might predict, which is perhaps due to an inability of people to accurately reflect on their forecast precision when considering new signals from the central bank. Additionally, we show that it isn't just a central bank's historical forecast performance that matters. Instead, our subjects exhibit considerable recency bias when

evaluating forecast performance to form a perception of the central bank's forecast credibility. Taken together, this suggests that historical forecast performance can influence a central bank's forecast credibility, but that discrete changes in forecast performance can quickly shift perceived credibility.[29]

An implication of this is that forecast credibility is not static. Instead, credibility is an endogenous component of communication that central banks can both win and lose, a feature conspicuously absent in most theoretical work on the topic. Though this implies that banks can lose their ability to manage expectations via forecasting whenever unexpected economic shocks lead to poor forecast performance, it also implies that banks can rebuild forecast credibility. However, we show that these dynamics of credibility are asymmetric – building credibility is a much slower process than losing it. We also demonstrate that low-frequency communication can bolster a bank's forecast credibility even when it does not convey new information about the bank's economic outlook or the conditions underlying historical forecast performance.

These findings imply that central banks might tailor their communication strategies to mitigate the effects of recency bias by emphasizing long-term performance over recent short-term fluctuations, thereby enhancing overall credibility and effectiveness. Though we show that short-term forecast errors influence how people perceive central bank credibility, we also show that long-term performance and consistency play a crucial role in shaping these perceptions. Central banks can provide a more stable and reliable basis for perceived credibility by focusing on and communicating their long-term forecasting track record. This approach helps to buffer the negative impact of recent short-term forecasting inaccuracies and foster a more resilient form of credibility that withstands the inevitable fluctuations in short-term economic conditions.

We embed these ideas into an otherwise standard New Keynesian model to demonstrate how accounting for endogenous credibility can affect aggregate inflation dynamics. We show that if a central bank can maintain forecast credibility following an exogenous shock that leads to sharp but temporary forecast errors, then inflation dynamics are almost identical to those produced under rational expectations even if households are boundedly rational. However, if this same shock leads to a sharp loss in forecast credibility then the economy experiences persistent inflation dynamics.

On an optimistic note, our study shows that narrative communication can help a central bank that finds itself in a position of low forecast credibility regain some ability to influence inflation expectations with its publicized forecasts. An important area for future research is whether, and how, this sort of narrative communication might prevent credibility loss in a dynamic setting. Moreover, further research should focus on the micro-level behavior driving our results. While our experimental design rules out some potential biases, understanding the behavioral mechanism underlying our results could provide deeper insight into what sort of narrative communication will most effectively defend and restore forecast credibility following protracted periods of low forecast precision. An experimental design to understand the asymmetric recency bias we observe across our Early and Late histories may be particularly revealing of the mechanism.

[29] A5 provides empirical evidence that suggests these experimental results align with real-world behavior in financial markets in the United Kingdom.

References

Adam, K. (2007). Experimental evidence on the persistence of output and inflation. *The Economic Journal 117*(520), 603–636.

Ahrens, S., J. Lustenhouwer, and M. Tettamanzi (2019). The stabilizing effects of publishing strategic central bank projections. *Macroeconomic Dynamics*, 1–43.

Arifovic, J. and L. Petersen (2017). Stabilizing expectations at the zero lower bound: Experimental evidence. *Journal of Economic Dynamics and Control 82*, 21–43.

Armantier, O., S. Nelson, G. Topa, W. Van der Klaauw, and B. Zafar (2016). The price is right: Updating inflation expectations in a randomized price information experiment. *Review of Economics and Statistics 98*(3), 503–523.

Assenza, T., P. Heemeijer, C. H. Hommes, and D. Massaro (2013). Individual expectations and aggregate macro behavior.

Bao, T., C. Hommes, J. Sonnemans, and J. Tuinstra (2012). Individual expectations, limited rationality and aggregate outcomes. *Journal of Economic Dynamics and Control 36*(8), 1101–1120.

Blinder, A. S. (2000). Central-bank credibility: Why do we care? how do we build it? *American Economic Review 90*(5), 1421–1431.

Bulutay, M. (2023). Better than perceived? correcting misperceptions about central bank inflation forecasts. Technical report, Technische Universit at Berlin.

Candia, B., O. Coibion, and Y. Gorodnichenko (2020). Communication and the beliefs of economic agents. In *Navigating the Decade Ahead: Implications for Monetary Policy* (Aug. 27-28, 2020 ed.)., EConomic Policy Symposium.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023). Anchored inflation expectations. *American Economic Journal: Macroeconomics 15*(1), 1–47.

Chen, D. L., M. Schonger, and C. Wickens (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

Clarida, R., J. Galí, and M. Gertler (1999, December). The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature 37*(4), 1661–1707.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte (2020). Inflation expectations as a policy tool? *Journal of International Economics 124*, 103297.

Coibion, O., Y. Gorodnichenko, and M. Weber (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy 130*(6), 1537–1584.

Cornand, C. and C. K. M'baye (2018). Does inflation targeting matter? an experimental investigation. *Macroeconomic Dynamics 22*(2), 362–401.

Cornand, C. and C. K. M'baye (2018). Band or point inflation targeting? an experimental approach. *Journal of Economic Interaction and Coordination 13*(2), 283–309.

Eusepi, S. and B. Preston (2010). Central bank communication and expectations stabilization. *American Economic Journal: Macroeconomics 2*(3), 235–71.

Evans, G. W., S. Honkapohja, and R. Marimon (2001). Convergence in monetary inflation models with heterogeneous learning rules. *Macroeconomic Dynamics 5*(1), 1–31.

Flynn, J. P. and K. Sastry (2022). The macroeconomics of narratives. Mimeograph.

Fuster, A. and B. Zafar (2023). Survey experiments on economic expectations. In *Handbook of Economic Expectations*, pp. 107–130. Elsevier.

Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton, USA and Oxford, UK: Princeton University Press.

Goy, G., C. Hommes, and K. Mavromatis (2022). Forward guidance and the role of central bank credibility under heterogeneous beliefs. *Journal of Economic Behavior & Organization 200*, 1240–1274.

Greene, W. H. (2003). *Econometric analysis*. Pearson Education India.

Guido, A., A. Martinez-Marquina, and R. Rholes (2022). Reference dependence and the role of information frictions.

Haldane, A. and M. McMahon (2018, May). Central bank communication and the general public. *108*, 578–83.

Hansen, S., M. McMahon, and M. Tong (2019). The long-run information effect of central bank communication. *Journal of Monetary Economics 108*, 185–202.

Hommes, C. and J. Lustenhouwer (2019). Inflation targeting and liquidity traps under endogenous credibility. *Journal of Monetary Economics 107*, 48–62.

Hommes, C., D. Massaro, and I. Salle (2019). Monetary and fiscal policy design at the zero lower bound: Evidence from the lab. *Economic Inquiry 57*(2), 1120–1140.

Hommes, C., D. Massaro, and M. Weber (2019). Monetary policy under behavioral expectations: Theory and experiment. *European Economic Review 118*, 193–212.

Hubert, P. (2015). Do central bank forecasts influence private agents? forecasting performance versus signals. *Journal of Money, Credit and Banking 47*(4), 771–789.

King, R. G., Y. K. Lu, and E. S. Pastén (2008, December). Managing expectations. *Journal of Money, Credit and Banking 40*(8), 1625–1666.

Kryvtsov, O. and L. Petersen (2021). Central bank communication that works: Lessons from lab experiments. *Journal of Monetary Economics 117*, 760–780.

Macaulay, A. and W. Song (2023, May). News media, inflation, and sentiment. *AEA Papers and Proceedings 113*, 172–76.

Malmendier, U. and S. Nagel (2016). Learning from inflation experiences. *The Quarterly Journal of Economics 131*(1), 53–87.

McMahon, M. and M. Naylor (2023). Getting through: Communicating complex information. CEPR Discussion Papers 18537, C.E.P.R. Discussion Papers. CEPR Discussion Paper.

Mokhtarzadeh, F. and L. Petersen (2021). Coordinating expectations through central bank projections. *Experimental Economics 24*(3), 883–918.

Morris, S. and H. S. Shin (2002). Social value of public information. *American Economic Review 92*(5), 1521–1534.

Newey, W. K. and K. D. West (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, 777–787.

Petersen, L. and R. Rholes (2022). Macroeconomic expectations, central bank communication, and background uncertainty: A covid-19 laboratory experiment. *Journal of Economic Dynamics and Control*, 104460.

Pfajfar, D. and B. Žakelj (2014). Experimental evidence on inflation expectation formation. *Journal of Economic Dynamics and Control 44*, 147–168.

Pfajfar, D. and B. Žakelj (2016). Uncertainty in forecasting inflation and monetary policy design: Evidence from the laboratory. *International Journal of Forecasting 32*(3), 849–864.

Pfajfar, D. and B. Žakelj (2018). Inflation expectations and monetary policy design: Evidence from the laboratory. *Macroeconomic Dynamics 22*(4), 1035–1075.

Rholes, R. and L. Petersen (2021). Should central banks communicate uncertainty in their projections? *Journal of Economic Behavior & Organization 183*, 320–341.

Shiller, R. J. (2017, April). Narrative Economics. *American Economic Review 107*(4), 967–1004.

Thakral, N. and L. T. Tô (2021). Daily labor supply and adaptive reference points. *American Economic Review 111*(8), 2417–43.

Walsh, C. E. (2017). *Monetary theory and policy*. MIT press.

Weber, M., F. D'Acunto, Y. Gorodnichenko, and O. Coibion (2022, August). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives 36*(3), 157–84.

Woodford, M. (2003). *Interest and Prices*. Princeton University Press.

Woodford, M. (2005). Central bank communication and policy effectiveness. *Proceedings - Economic Policy Symposium - Jackson Hole* (Aug), 399–474.

Online Appendix for *"Building Central Bank Credibility: The Role of Forecast Performance"*

A1 Additional Tables and Figures

Summary of Bonus Payments					
		Early			
N	**Treatments**	**Initial - Point**	**Initial - Range**	**Updated - Point**	**Updated - Range**
516	*Timing*	.82	.40	.82	.40
348	*ForecastPerformance*	.83	.40	.83	.41
719	*Contextual Communication*	.81	.37	.82	.39
		Consistent			
	Timing	.33	.01	.47	.05
	ForecastPerformance	.34	.01	.55	.13
	Contextual Communication	.33	.01	.47	.06
		Late			
	Timing	.68	.23	.68	.23
	ForecastPerformance	.67	.22	.66	.22
	Contextual Communication	.69	.23	.70	.25

Table A-1: This table reports, in U.S. dollars, average earnings for Initial and Updated point and range forecasts of inflation provided by subjects in each history across our three primary treatments.

Summary of Bonus Payments - ***Forecast Performance***					
N	**Treatments**	**Initial - Point**	**Initial - Range**	**Updated - Point**	**Updated - Range**
90	*Great*	.37	.00	.72	.25
90	*Good*	.34	.01	.56	.14
72	*Moderate*	.34	.01	.54	.09
180	*Bad*	.32	.00	.48	.05
96	*Terrible*	.32	.02	.38	.02

Table A-2: This table reports, in U.S. dollars, average earnings for Initial and Updated point and range forecasts of inflation provided by subjects in our *Forecast Performance* treatments.

Summary of Bonus Payments - ***Contextual Communication***					
N	**Treatments**	**Initial - Point**	**Initial - Range**	**Updated - Point**	**Updated - Range**
130	*Control*	.70	.23	.69	.25
124	*Control+Outlook*	.71	.21	.75	.25
107	*Exogenous+Good*	.66	.23	.67	.24
121	*Exogenous+Bad*	.69	.24	.69	.22
122	*Endogenous+Good*	.68	.23	.73	.26
115	*Endogenous+Bad*	.67	.22	.69	.25

Table A-3: This table reports, in U.S. dollars, average earnings for Initial and Updated point and range forecasts of inflation provided by subjects in our *Contextual Communication* treatments.

Table A-4: Regression Table for Contextual Communication

	(1) $\hat{u}$	(2) $\hat{u}$	(3) $\hat{u}$	(4) $\hat{u}$
Control + Outlook	29.83****	29.82****	30.58****	30.00****
	(8.257)	(8.274)	(8.497)	(8.401)
Ex. + Better	27.84****	27.82***	27.71***	28.58****
	(8.387)	(8.463)	(8.511)	(8.425)
Ex. + Worse	23.43***	23.41***	23.74***	23.95***
	(7.920)	(8.039)	(8.234)	(7.884)
End. + Better	20.87**	20.88***	21.09**	20.64**
	(8.105)	(8.053)	(8.183)	(8.040)
End. + Worse	21.38***	21.38***	21.65***	21.84***
	(7.453)	(7.434)	(7.510)	(7.365)
Uncertainty		0.00162	-0.00167	0.00275
		(0.0351)	(0.0355)	(0.0359)
Demographics			✓	✓
Survey Responses				✓
Control	0.000	-0.222	7.099	-2.913
	(6.132)	(6.741)	(9.012)	(17.81)
N	679	679	674	674

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$, **** $p < .001$

Note: This table provides estimates of $\hat{u}$ for each of our *Contextual Communication* treatments obtained via OLS. Column 1 provides baseline estimates of forecast credibility, column 2 controls for forecast uncertainty, column 3 layers in demographic controls, and column 4 includes controls from our pre-experiment survey questions.

A2 Simulated Economic Histories

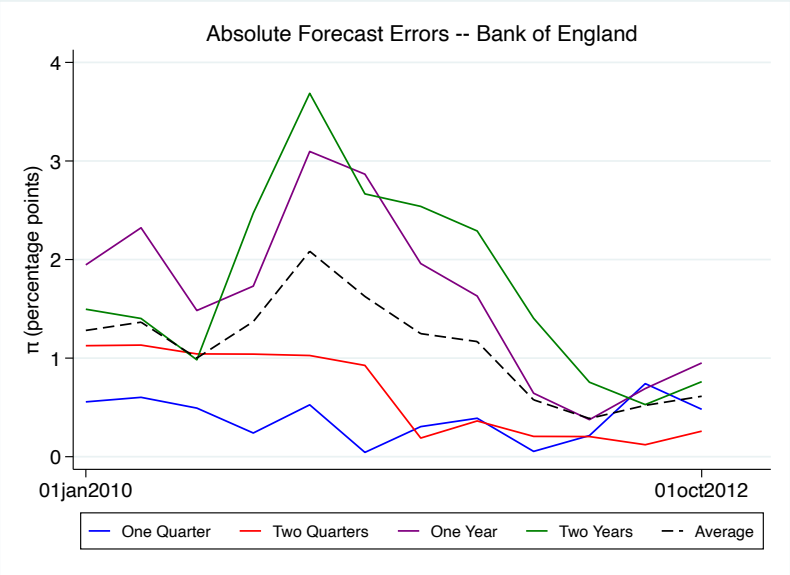


Figure A-1: Data used for calibrating economic histories

Note: This figure depicts inflation forecast errors at one quarter, two quarter, one year, and two year forecast horizons from the start of 2010 through the end of 2012.

To motivate the calibration of our simulated histories, we draw on the Bank of England absolute forecast error data shown in Figure A-1. At the one-quarter horizon, absolute errors averaged roughly 50–75 basis points at the start of the sample before declining below 25 basis points by the end. Errors at two quarters ahead were of a similar order of magnitude, averaging close to 100 basis points in 2010 and falling steadily thereafter. At longer horizons, however, forecast errors were markedly larger: one-year-ahead forecasts exhibited errors on the order of 150–300 basis points in the early years, while two-year-ahead forecasts were frequently above 200 basis points and at times approached 350 basis points before receding. Averaging across horizons, the Bank of England's forecast errors began at well above 100 basis points in early 2010 and converged toward roughly 40–60 basis points by late 2012.

Our simulated histories are calibrated to reproduce this qualitative pattern, while adopting somewhat conservative magnitudes that align with the forecast horizons most relevant to our experiment. In the *Early* history, forecast errors begin high (171 bps in Year 1) and fall to very low levels (13 bps in Year 3), capturing the dynamic of an economy moving out of a downturn as central bank forecasts regain accuracy. The *Late* history reverses this trajectory, preserving the full-sample average but shifting the deterioration to later periods, thereby representing an economy entering a downturn. Finally, the *Consistent* histories hold forecast performance steady at different levels—ranging from very accurate (*Great*, 13 bps) to persistently poor (*Terrible*, 171 bps)—thus bracketing the range of forecast precision observed in the historical data.

By anchoring our calibration in observed forecast errors across multiple horizons, while adopting values that are slightly conservative and centered on the short- to medium-term forecast horizons most relevant for central banks, we ensure that our simulated histories reflect real-world dynamics while remaining tractable for experimental subjects, thereby preserving the external validity of our findings.k

A3 Central Bank Messages in *Contextual Communication*

T12 - Control:
The Fed uses interest rate policy to stabilize prices and keep employment high. We base monetary policy on how healthy the economy is now and how healthy we think it will be in the future. We use forecasts to guide our decisions. We do our best when making forecasts but the world is uncertain and forecasts are never perfect.

T13 - Control+Outlook
The Fed uses interest rate policy to stabilize prices and keep employment high. We base monetary policy on how healthy the economy is now and how healthy we think it will be in the future. We use forecasts to guide our decisions. We do our best when making forecasts but the world is uncertain and forecasts are never perfect.

Over the last year, our forecasts under-predicted inflation. Our best guess is that inflation will decrease next quarter.

T14 - Exogenous + Better
The Fed uses interest rate policy to stabilize prices and keep employment high. We base monetary policy on how healthy the economy is now and how healthy we think it will be in the future. We use forecasts to guide our decisions. We do our best when making forecasts but the world is uncertain, and forecasts are never perfect.

Over the last year, our forecasts under-predicted inflation. This is because the pandemic lasted longer than initially expected and caused supply shortages.Our forecasts over this period were **more accurate** than private sector forecasts and other central banks. Our best guess is that inflation will decrease next quarter.

T15 - Exogenous + Worse
The Fed uses interest rate policy to stabilize prices and keep employment high. We base monetary policy on how healthy the economy is now and how healthy we think it will be in the future. We use forecasts to guide our decisions. We do our best when making forecasts but the world is uncertain and forecasts are never perfect.

Over the last year, our forecasts under-predicted inflation. This is because the pandemic lasted longer than initially expected and caused supply shortages. Our forecasts over this period were **less accurate** than private sector forecasts and other central banks. Our best guess is that inflation will decrease next quarter.

T16 - Endogenous + Better
The Fed uses interest rate policy to stabilize prices and keep employment high. We base monetary policy on how healthy the economy is now and how healthy we think it will be in the future. We use forecasts to guide our decisions. We do our best when making forecasts but the world is uncertain and forecasts are never perfect.

Over the last year, our forecasts under-predicted inflation. This resulted from interest rates being too low for too long. Our forecasts over this period were **more accurate** than private sector forecasters and other central banks. Our best guess is that inflation will decrease next quarter.

T17 - Endogenous + Worse
The Fed uses interest rate policy to stabilize prices and keep employment high. We base monetary policy on how healthy the economy is now and how healthy we think it will be in the future. We use forecasts to guide our decisions. We do our best when making forecasts but the world is uncertain and forecasts are never perfect.

Over the last year, our forecasts under-predicted inflation. his resulted from interest rates being too low for too long. Our forecasts over this period were **less accurate** than private sector forecasters and other central banks. Our best guess is that inflation will decrease next quarter.

A4 Salience of Real-World Inflation

We conducted a series of five treatment waves sequentially over the course of approximately one year, starting in February 2022 and ending in March 2023. We described the timing of our treatments in Table A-5. During this time, the United States – as many other economies – was experiencing considerable inflation, which constituted the first salient change in price dynamics in more than a decade. Though inducing preferences in an experimental setting ought to insulate results from real-world economic dynamics, Petersen and Rholes (2022) provides some evidence that central bank communication may be susceptible to real-world shocks. In the context of this experiment, one might be concerned about differences in headline inflation throughout our treatment waves, were these differences salient to participants, did it change inflation preferences, and should we be concerned that this matters for our results?

First, we note the last reported value for CPI inflation leading into each treatment wave Table A-5, was at its lowest in our final wave of treatments at approximately 6% and at its highest in our third treatment wave at about 8%. Though the two-percent variation in inflation is not trivial, it does constitute considerably less variation than what participants experienced in the time leading up to our experiment when inflation rose from approximately 2% to 7.5%.

Was this variation salient to participants? We show in Figure A-2 (main treatment waves, and Figure A-3 (robustness treatment waves), that the large majority of subjects were aware of the most recent measure of headline inflation. These same figures also show that preferences for real-world inflation were identical across treatment waves, indicating that the salient difference in actual inflation did not change inflation preferences. In all treatment waves, the majority of participants indicated a desire for low inflation ranging from 1% to 3%, with approximately 90% or more of our participants indicating a preference for inflation between -1% and 3%.

Treatments	**Name**	**Dates**	**Last Reported CPI Inflation**
T1 - T6	*Timing*	Feb. 25 - Feb. 28, 2022	7.5% in Jan.
T7 - T14	*ForecastPerformance*	Mar. 28 - Mar. 29, 2022	7.9% Feb.
T15 - T20	*Contextual Communication*	May 27 - May 28, 2022	8.2% in Apr.
T21 - T26	*Medium-Term*	Dec. 14 - Dec. 15, 2022	7.13% in Nov.
T27 - T28	*ReversedShock*	Mar. 24 - Mar. 25, 2023	5.98% in Feb. 2023

Table A-5: Actual inflation during each treatment wave.

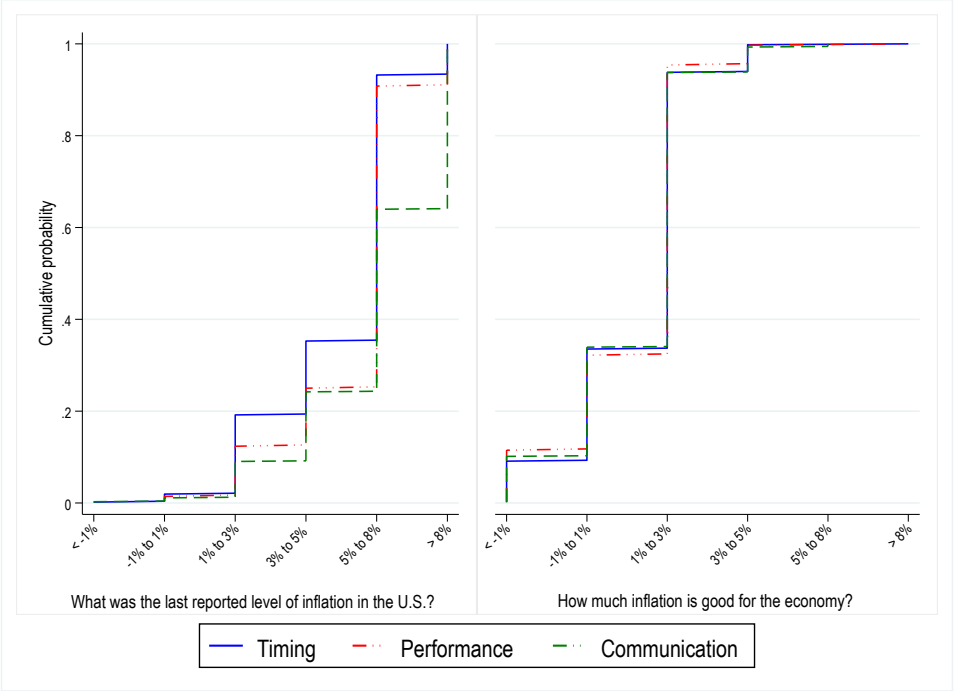


Figure A-2: This figure depicts kernel density functions of participants' inflation preferences alongside perceptions of prevailing inflation for *Timing ForecastPerformance*, and *Contextual Communication* treatments.

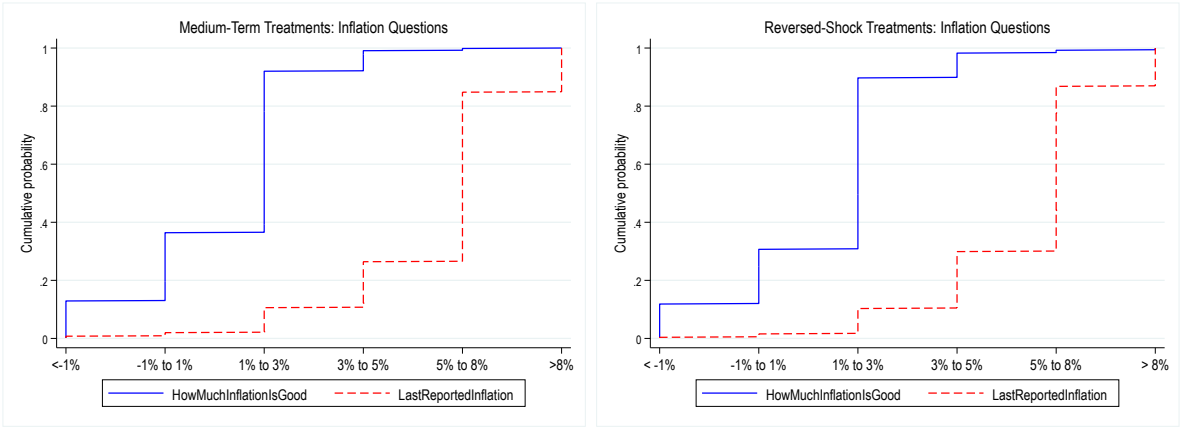


Figure A-3: This figure depicts cumulative distribution functions of participants' inflation preferences alongside perceptions of prevailing inflation for *Medium-Term* treatments.

A4.1 Hypothesis 2: Forecast credibility and stated forecast precision

Equation (4) also elucidates that a central bank's forecast credibility doesn't depend entirely on things it can control. Instead, forecast credibility depends both on the central bank's forecast precision and a participant's belief about her own forecasting credibility. This leads us to our second hypothesis:

Hypothesis 4. *For a given economic history, the central bank's forecast credibility increases in a participant's own forecast uncertainty.*

Intuitively, this hypothesis says that a participant who exhibits more forecast uncertainty in the Initial Forecast should update more toward the central bank's forecast.

We turn again to Table 4 and consider the coefficient estimates in the ForecastUncertainty row, which estimate the relationship, on average, between forecast uncertainty α^{-1} and perceived forecast credibility $\hat{u}_i$. Regardless of specification, participants exhibit no significant response to forecast uncertainty when updating their inflation forecasts. We depict this graphically in Figure A-4, which also shows that there is no relationship between forecast uncertainty and perceived central bank credibility. This result violates the logic of Equation (4) and suggests that, for a given forecast history, people would respond to new signals from the central bank in the same way regardless of their own uncertainty about future economic conditions.

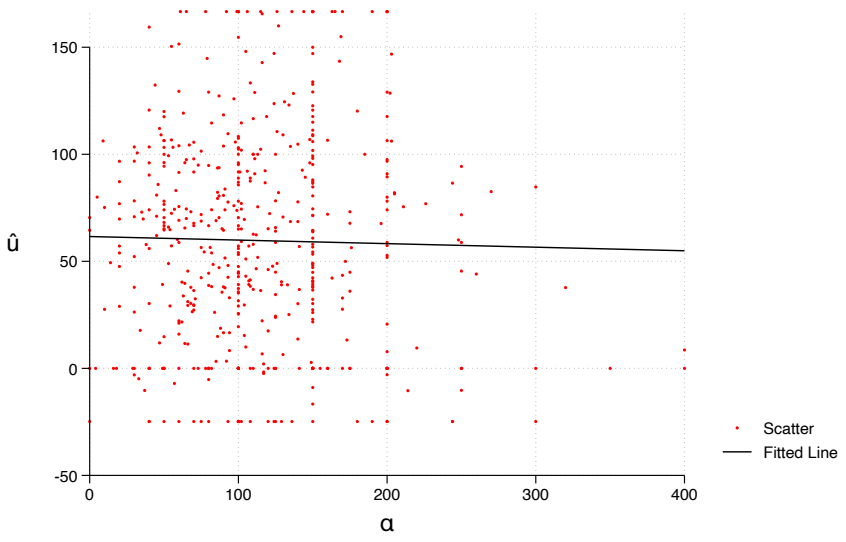


Figure A-4: Individual-level forecast uncertainty and forecast update weight.
Note: This figure presents a scatter plot of individual-level forecast uncertainty and the forecast update weight as a measure of perceived forecast credibility of the central bank.

This is surprising. Intuitively, a signal that conveys some clarifying information ought to be more valuable in instances of higher confusion, which is what Equation (4) says – an uncertain agent should more highly value new signals that help her better predict the evolution of important aggregates than a 'certain' agent who thinks she has a good grasp on how those aggregates will evolve.

A5 Evidence of recency bias in real-world data

Are our results an artifact of our experimental setting or do the results from our experiment generalize to a real-world setting? To explore this, we combine a high-frequency identification approach with the Bank of England's (BoE) quarterly Inflation Report (IR) (now called Monetary Policy Report). We ask "do markets react more strongly to the BoE's forecast information whenever the BoE's forecast credibility is high? Hubert (2015) explores forecast performance and market news for numerous central banks but focusing on the Bank of England is ideal since it releases the forecast information separately to the policy decision with a lag of about a week in our sample.

We project real yield market news of different maturities that occur in the 24-hour window surrounding IR releases onto conditioning information and variables capturing the BoE's recent forecasting performance. Based on our experimental results, better forecast performance should lead to stronger market reactions to information in the IR. We use data from between Q3 1997, after the Bank gained operational independence, through Q2 2015 (when the provision of information changed). We have 72 Inflation Report releases (observations).

To categorize the BoE's forecast performance over time, we have to recognize that central banks provide forecasts for many periods into the future and so forecast credibility in the real world is itself multi-dimensional. In order to try to be agnostic about the relationship between forecast errors, forecast horizon, and forecast credibility, we measure the central bank's forecast error for each forecast horizon during each quarter and collapse these horizon-specific measures of forecast performance into a single dimension using factor analysis. The result is a one-dimensional measure of forecast performance that accounts for forecast errors at each forecast horizon during each quarter, which we depict in A-5. Though the BoE has sometimes provided forecasts with as much as three-year horizons, this practice was not consistent during our time sample. Because of this, we focus on the BoE's forecasts up to eight quarters ahead.

Using this factor, we create a set of indicator variables denoting periods when the BoE's forecast performance has been above its sample average forecast for the previous one, two, or three quarters. We record 39, 34, and 29 instances where these indicators take on a positive value, respectively.

We use changes in one-year, three-year, and five-year gilts that occur within a 24-hour window surrounding an IR release to measure market reactions to the information contained in the IR.[1] Following from our experimental results, our hypothesis is:

Hypothesis 5. *Yields will respond more strongly to the information contained in the IR whenever the BoE's recent forecast credibility, proxied by its forecast performance, is above the sample average level of forecast credibility.*

[1] We use a 24-hour window following Hansen et al. (2019), who argue the longer window is necessary because of the volume of information contained in the BoE's IR. This is compared to isolated policy announcements, where market participants can quickly discern and react to information.

To test Hypothesis 5, we estimate the following equation:

$$|\Delta y_i| = \alpha_i + \zeta_l \mathbb{I}_{l,t} + \sum_{x,j} \psi_{x,j,i} \Delta PC_{x,j,t} + \eta_{1,i} FTSE_{t-1} + \chi_i X_{i,t} + \eta_{2,i} VIX_{t-1} + \epsilon_{i,t} \quad \text{(A.20)}$$

where $\Delta y_{i,t}$ captures changes in the i-year gilt that occur in the 24-hour window around the IR release, $\Delta PC_{x,j,t}$ is a set of six factors summarizing new information contained in the contemporaneous IR regarding the first three central moments ($x = \{1, 2, 3\}$) of the BoE's outlook on inflation and output ($j = \{\pi, Y\}$), and $\mathbb{I}_{l,t}$ is an indicator capturing whether the BoE's forecast performance has exceeded its sample average for the last $l = \{1, 2, 3\}$ quarters. As controls, $X_{i,t}$ contains controls that account for prevailing economic conditions (unemployment, output, and inflation), $FTSE_{t-1}$ is a daily, market-based measure of economic uncertainty, and VIX_{t-1} captures general economic uncertainty. We account for the possibility of autocorrelation and heteroskedasticity using Newey-West errors with 3 lags.[2]

Table A-6: Regression Table: Recency Bias in Markets

	(1) $\mathbb{I}_{1,t}$	(2) $\mathbb{I}_{2,t}$	(3) $\mathbb{I}_{3,t}$
1yr Gilt	0.036	0.050**	0.064***
	(0.027)	(0.021)	(0.02)
3yr Gilt	0.0307	0.0358*	0.0564****
	(0.025)	(0.020)	(0.012)
5yr Gilt	0.011	0.0199	0.0369***
	(0.023)	(0.019)	(0.013)

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$, **** $p < .001$

Note: This table depicts estimates of ζ_1 in Equation (A.20). $\mathbb{I}_{l,t}$ is an indicator capturing whether the BoE's forecast performance has exceeded its sample average for the last $l = \{1, 2, 3\}$ quarters.

We report $\hat{\zeta}_i$ in Table A-6, where columns correspond to different durations of above-average forecasting performance, $\mathbb{I}_{l,t}$, and rows to our outcomes of interest, $|\Delta y_{i,t}|$. Results from this estimation exercise indicate that markets respond more strongly to the information contained in the BoE's inflation report whenever the BoE's forecast credibility, proxied by its forecast accuracy, is above its historical average.

For each gilt maturity we consider, the estimated effect size of above-average forecast performance is strongest whenever this performance has persisted for at least three quarters and decreases monotonically in the duration of above-average performance. Additionally, the credibility premium garnered by above-average forecast performance has less impact as we extend deeper into the term structure.

These results align with our finding of recency bias in *Timing*. Further, the fact that this effect persists for longer-term maturities aligns with our finding that both short- and longer-term expectations respond in a qualitatively similar way to recent forecast performance. Perhaps reassuringly it is not just a single quarter of worse performance that matters for the result. Our findings align with Hubert (2015), who shows that

[2]Lag selection is based on Newey and West (1987) and Greene (2003).

central bank forecasts influence the expectations of private agents. However, that work also finds that forecast performance does not drive this result, which is different to what we find. Whereas that work considers whether survey-based forecasts respond to a gap in forecasting performance between private agents and the central bank, we consider if market participants respond differently to central bank communication depending upon the bank's recent forecast performance relative to its historical performance.

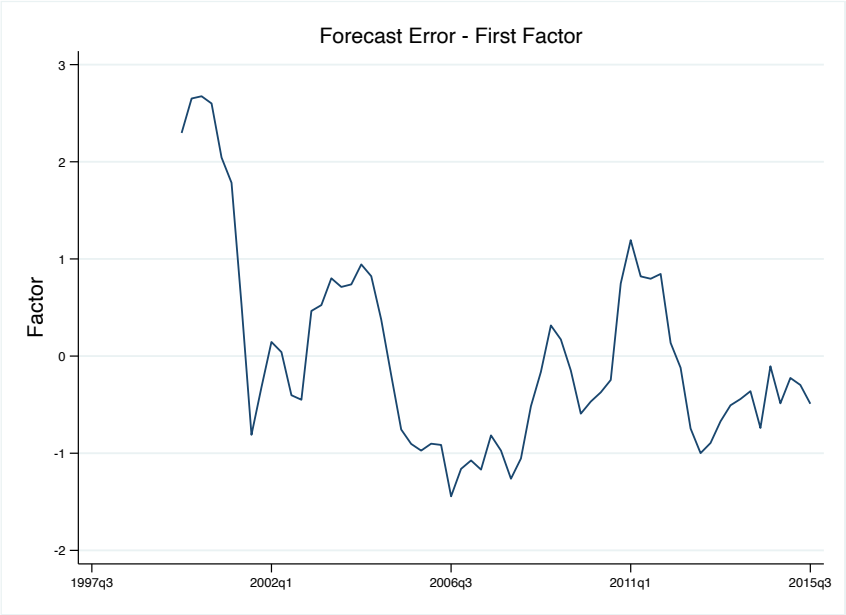


Figure A-5: Bank of England Forecast Performance: Factor Summary Variable
Note: This figure depicts our measure of the Bank of England's forecast error over time. The Bank of England forecasts the current and at least the next eight quarters of inflation in each quarter of our sample. To form our measure of the BoE's composite forecast error, we collapse these nine quarterly forecast errors via principal component factor analysis and use the first principal component that results from this exercise. This represents the linear combination of errors that captures the most common variation across individual forecast errors over time.

A6 The role of forecast bias

What if participants believe that the central bank's forecast error is biased? In theory, we typically model no systematic component to the central bank's forecast error. However, participants may perceive $\gamma \neq 0$ because our experimental histories contain only twelve quarters of data based on volatile, real-world time series (we provide details on how we create these histories in Section 3.3). If this is true, then not accounting for this bias can lead to systematically

To account for this, we can rewrite Equation (2) as:[3]

$$\pi_{cb} - \gamma = \pi + \epsilon, \quad \epsilon \sim \mathcal{N}\left(0, \frac{1}{\beta}\right). \tag{A.21}$$

Equation (A.21) says that once we adjust the central bank signal for its bias, we can apply the same logic as before. Intuitively, suppose that $\gamma < 0$ so that the central bank systematically under forecasts inflation. When the central banks signals its inflation forecast, the true signal from the central bank is adjusted upward and this new, higher, signal is used in the optimal update. That is, in Equation (4), we use $\pi_{cb} - \gamma > \pi_{cb}$ as the central bank's signal. Note that Figure 1 is unchanged once we make this bias correction since optimal updating scheme depends only on forecast precision (α_i^{-1}, β^{-1}).

Of course, our measure of the optimal update rate should also reflect the bias adjustment:

$$u_{\gamma,i} \equiv \frac{\mathbb{E}(\pi|\pi_{cb}) - \bar{\pi}_i}{(\pi_{cb} - \gamma - \bar{\pi}_i)} \tag{A.22}$$

Once this adjustment is done correctly, and assuming i updates according to the Bayesian optimal, $u^*_{\gamma,i} = \frac{\beta}{\alpha+\beta}$. Where $\gamma = 0$, $u^*_i = u^*_{\gamma,i}$ but if $\gamma \neq 0$, $\frac{u^*_{\gamma,i} - u^*_i}{u^*_i} = \frac{\gamma}{(\pi_{cb} - \gamma - \bar{\pi}_i)}$.

A6.1 Re-estimating our main results

We now account for the possibility of perceived bias ($\gamma \neq 0$) by adjusting our estimates of central bank forecast credibility according to Equation (A.22), where we assume that participants believe that the central bank's forecast bias is equal to the historical average forecast error so that

$$\gamma_{HistAvg} = \frac{1}{12} \sum_{k=t-1}^{k=t-12} (\mathbb{E}_{cb}\{i_k\} - i_k)$$

. We provide values of $\gamma_{HistAvg}$ for all economic histories in Table 2.

A6.1.1 *Forecast Performance*

Figure A-6 plots average treatment effects assuming our participants observe no systematic component in the central bank's forecast error (blue dots, baseline results) and

[3] We replace $\tilde{\epsilon} \sim \mathcal{N}\left(\gamma, \frac{1}{\beta}\right)$ in Equation (2) with $\gamma + \epsilon$ where $\epsilon \sim \mathcal{N}\left(0, \frac{1}{\beta}\right)$. If $\gamma = 0$, $\tilde{\epsilon} \equiv \epsilon$ trivially.

also assuming that participants account for a systematic error component of the central bank's inflation forecast (purple squares). For these biased estimates, we assume that subjects use the entire forecast history to discern the magnitude and direction of this systematic error component. Adjusting our estimates of forecast credibility to account for forecast bias preserves our results qualitatively and leads to little quantitative change (see also Table A-7).

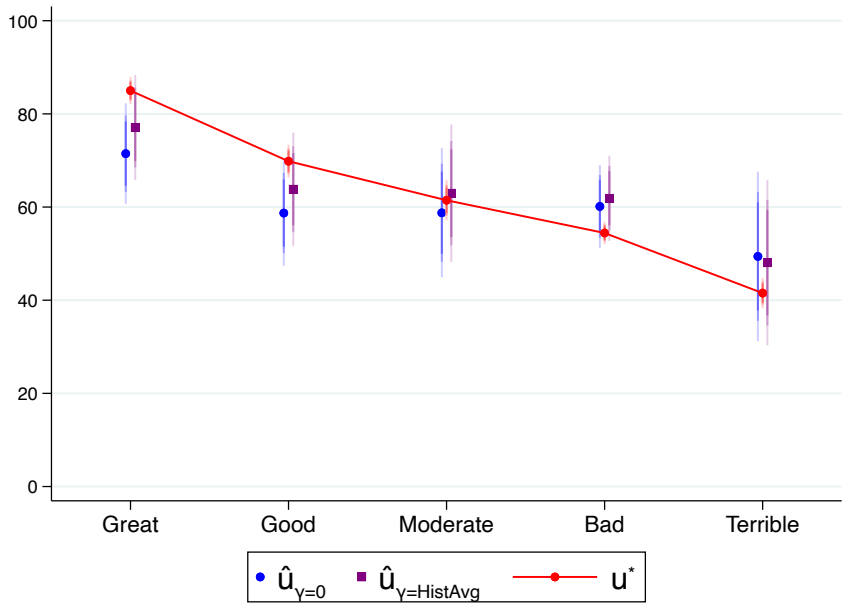


Figure A-6: *Forecast Performance* using $\hat{u}_{\gamma_{HistAvg}}$

Note: This figure presents estimates of central bank forecast credibility in the *Forecast Performance* treatments. It compares scenarios assuming that subjects do not perceive any systematic forecast errors (as in Equation (4), blue circles) with scenarios where subjects do perceive systematic forecast errors (as in Equation (A.22), purple squares). Shaded bands surrounding the point estimates depict 99% (lightest), 95%, and 90% (darkest) confidence intervals. The connected red triangles depict the optimal level of credibility, on average, from the perspective of a rational Bayesian agent.

Bias-adjusting our credibility estimates effectively leads to higher measures of forecast credibility in all but our *Terrible* treatments. This is unsurprising, given that $\gamma_{HistAvg} > 0$ for *Great*,...,*Bad* but not for *Terrible*. This is because adjusting Equation (A.22) for positive values of $\gamma_{HistAvg}$ shrinks the denominator of our estimation equation.

Intuitively, this resembles the assumption that the central bank consistently over-predicts inflation. Rectifying this over-prediction suggests that subjects align more closely with the central bank's signal than if we suppose participants perceive $\gamma = 0$. This is because the true signals participants received were lower than the central bank's published forecast.

Despite this mechanical increase in estimated forecast credibility, our finding that participants exhibit over-precision survives. We note this in Table A-7, which presents the output of a series of regressions capturing unconditional estimates of u^* and $u^*_{\gamma_{HistAvg}}$ (columns one and two) and their corresponding deviations from the equal-weighting Bayesian benchmark adopted throughout this paper (columns 3 and 4). Comparing columns 3 and 4 shows that, qualitatively, deviations from the equal-weighting Bayesian benchmark are robust to bias adjustment. However, the mechanical increase in estimates forecast credibility that results from bias adjustment decrease significance.

Table A-7: Regression Table for Forecast Performance: Bias Adjusted

	(1) $\hat{u}$	(2) $\hat{u}_{\gamma HistAvg}$	(3) $\hat{u} - u^*$	(4) $\hat{u}_{\gamma HistAvg} - u^*$
Great	71.16****	76.71****	-12.60***	-9.167**
	(4.177)	(4.273)	(4.186)	(4.326)
Good	58.37****	63.47****	-11.60***	-6.511
	(4.336)	(4.608)	(4.443)	(4.700)
Moderate	58.39****	62.67****	-3.061	1.211
	(5.434)	(5.622)	(5.698)	(5.867)
Bad	59.70****	61.52****	5.273	7.084*
	(3.421)	(3.479)	(3.575)	(3.633)
Terrible	47.31****	46.77****	6.282	5.733
	(6.967)	(6.614)	(7.044)	(6.733)
N	528	528	524	524

Robust standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$, **** $p < .001$

Note: This table presents estimates of forecast credibility for each *Forecast Performance* treatment. Column 1 is based on the assumption that there is no perception of a systematic forecast error (as in Equation (4)), whereas Column 2 is based on the assumption that subjects do perceive a systematic forecast error (as in Equation (A.22)). Columns 3 and 4 provide the corresponding estimates of deviations from optimal forecast credibility.

A6.1.2 *Timing*

We note in Table 2 that average forecast errors are larger in *Early* and *Late*. This subsection considers our *Timing* results assuming that our participants perceive a systematic bias in the central bank's forecasts so that we must adjust $\hat{u}$ to account for the values of $\gamma_{HistAvg}$ for *Early* and *Late* given Table 2. Note that we do not revisit *Consistent* results here since credibility estimates for this history remain essentially unchanged (note in Table 2 that the average forecast error for *Consistent* is only 2 basis points). Our interest is in whether our main finding – that the time profile of historical forecast errors – causes participants to more heavily weigh recent information when forming a perception of the central bank's forecast credibility.

Table A-8: Results of t-tests and Descriptive Statistics

	Early		*Late*	
	Mean	SE	Mean	SE
Early	28.25	1.38		
Late	$p < .001$		15.68	3.13

Comparing forecast credibility in *Early* and *Late* using $\hat{u}_{\gamma_{HistAvg}}$

Note: This table reports the mean and standard error (SE) of $u^*_{\gamma_{HistAvg}}$ for the *Early* and *Late* histories in the *Timing* treatments. It also includes the p-value from a two-sample, two-sided t-test comparing the means between *Early* and *Late* groups.

We present these updated estimates in Figure A-7, which shows that subjects are now estimated to underweight the central bank's inflation forecast for both *Early* and *Late*. This is consistent with the results in Figure 5 where participants tend to underweight very good performance; the net effect of over-weighting recent performance but under-

weighting great performance is not, ex-ante, obvious. Nonetheless, the finding that timing matters remains. Further, the qualitative result that participants exhibit more recency bias in *Late* than *Early* also survives. We confirm this in Table A-8, which shows that forecast credibility is almost two-fold larger in *Early* than in *Late*, and this difference in estimate forecast credibility across histories is statistically significant ($p < .001$).

Recency Bias for $\gamma \neq 0$

Table A-9: Estimated Values of λ

	γ_0	$\gamma_{HistAvg}$
Early	0.245	0.275
	(0.0170)	(0.0160)
Late	0.622	0.560
	(0.0198)	(0.0222)

Note: This table presents estimates of recency bias in the *Early* and *Late Timing* treatments. It distinguishes between conditions where we assume subjects perceive no systematic forecast error (γ_0 column) and where they do perceive a systematic forecast error ($\gamma_{HistAvg}$ column). Estimates are obtained via OLS regression, with robust standard errors reported in parentheses.

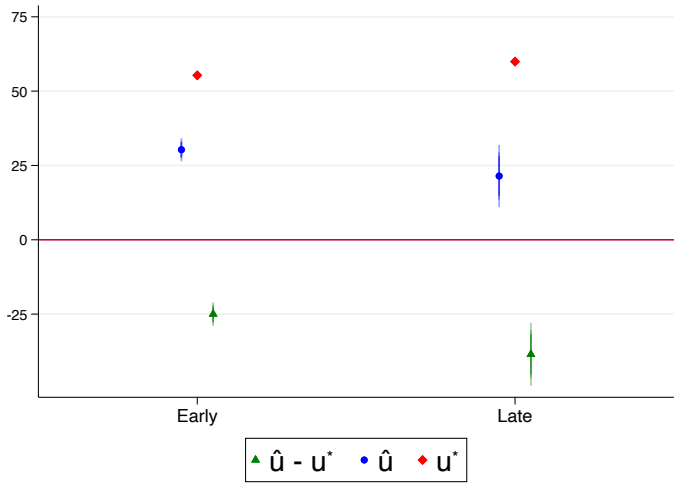


Figure A-7: Perceived forecast credibility in *Timing*: $\hat{u}_{\gamma HistAvg}$

Note: This figure depicts estimates of forecast credibility for the *Early* and *Late Timing* treatments (represented by blue circles), under the assumption that subjects perceive a systematic forecast error, as specified in Equation (A.22). It also illustrates the Bayesian optimal level of updating as a benchmark (red diamonds) and deviations from this benchmark (green triangles). Shaded bands around the point estimates indicate 99% (lightest), 95%, and 90% (darkest) confidence intervals.

Dynamics of Credibility for $\gamma \neq 0$

A6.1.3 Dynamics of forecast credibility

We further explore our bias-adjusted forecast credibility estimates to gain some insight into the dynamics of perceived forecast credibility. To do this, we use the fact that

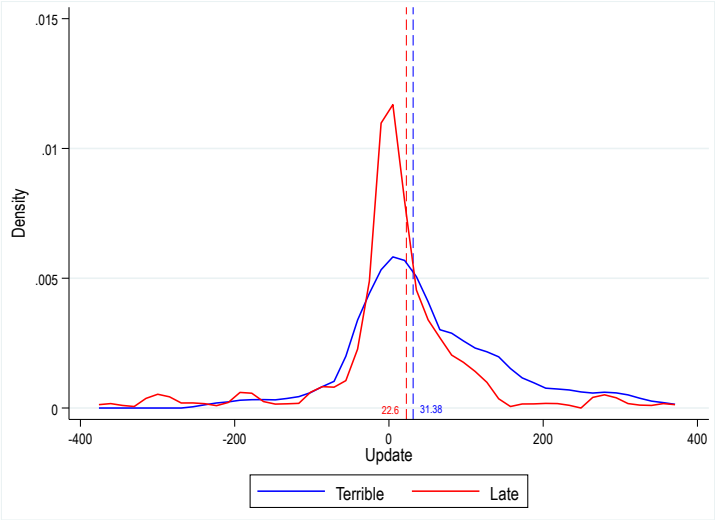


Figure A-8: *Consistent-Terrible* vs. *Late*

Note: This figure depicts kernel density estimates of forecast credibility assuming subjects perceive a systematic forecast error from the central bank, $\hat{u}_{\gamma HistAvg}$ as in Equation (A.22), from *Late* (red, dashed line) and *Terrible* (blue, solid line). The corresponding vertical dashed lines denote the mean of $u^*_{\gamma_{HistAvg}}$.

the central bank's historical forecast precision in *Terrible* from *ForecastPerformance* is identical to the central bank's forecast precision in the final year of *Late* from *Timing*. By comparing estimated forecast credibility economic histories, we can learn something about how quickly forecast credibility erodes. Similarly, we can also use the fact that the central bank's historical forecast precision in *Great* is identical to the bank's historical forecast precision in the final year of *Early*.

We first compare perceived credibility measures from *Terrible* and *Late*, which we depict as kernel density estimates in Figure A-8. First we note that the mean level of perceived credibility is not statistically different across treatments ($p = .704$, two-sample t-test).[4] Overall, results suggest that *Terrible* forecast precision for a single year leads to perceived forecast credibility that is as low, on average, as if subjects see *Terrible* forecast precision over the entire economic history.

However, this does not hold when comparing *Great* and *Early* in Figure A-9. Instead, we see that the mean level of perceived forecast credibility is significantly higher in *Great* than in *Early* ($p < .001$) and that the distributions are highly significantly different ($p < .01$, Kolmogorov-Smirnov (KS) test). This suggests that seeing *Great* forecast performance over the full sample history leads to significantly higher credibility than seeing it over only the last year.

These results align with our estimated weighting functions. In *Late*, the deterioration of the central bank's forecast performance induces a very strong recency bias. On average, participants in that experiment base the majority of their perception of the central bank's forecast credibility on the very last historical observation. In *Early*, an analogous improvement in forecast precision does not induce the same degree of recency bias. Though participants primarily focus on the final year of forecast performance following both histories, our estimated weighting function from *Early* exhibits a fatter right tail.

[4]Results from a Kolmogorov-Smirnov test indicate that the perceived credibility is slightly lower in *Late* than in *Terrible* ($p = .044$). This is driven by the slightly lower mass of positive updates in *Late*.

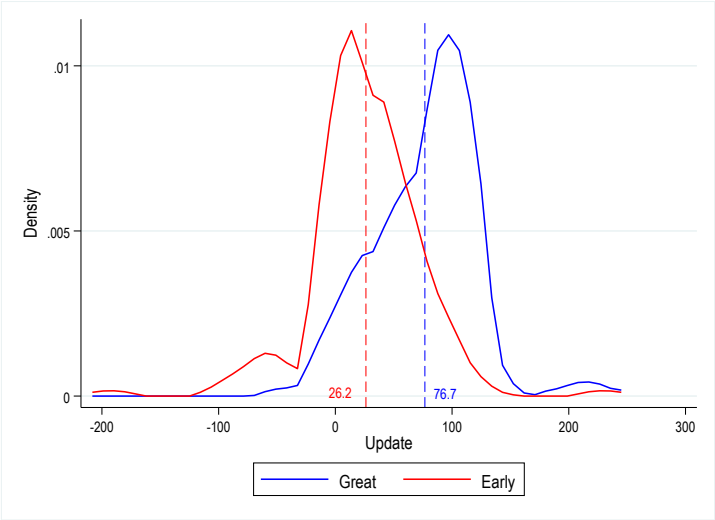


Figure A-9: *Consistent-Great* vs. *Early*

Note: This figure depicts kernel density estimates of forecast credibility assuming subjects perceive a systematic forecast error from the central bank, $\hat{u}_{\gamma HistAvg}$ as in Equation (A.22), from *Early* (red, dashed line) and *Great* (blue, solid line). The corresponding vertical dashed lines denote the mean of $u^*_{\gamma HistAvg}$.

Intuitively, this suggests that poor forecast performance lingers longer in people's minds when deciding how much faith to place in the central bank's ability to predict inflation accurately.

A6.1.4 *Contextual Communication*

This section reconsiders our primary result from *Contextual Communcation* assuming that $\gamma_{HistAvg}$. We depict these results in Figure A-10.

We note two main points. First, communication still increases perceptions of the central bank's forecast credibility. In fact, assuming $\gamma_{HistAvg}$ *strengthens* the estimated credibility gains in most treatments, and so much so in *Exogenous + Better* that the central bank recovers the Bayesian optimal level of forecast credibility via its contextualizing statement.

Second, acknowledging that the central bank's historical average forecast precision was lower than peer forecasting institutions is quite detrimental. This effect is most pronounced in *Exogenous + Worse*, though statistically the effects are similar for *Endogenous + Worse*. Announcing the bank outperformed peer forecasting institutions yields $\hat{u}_{\gamma_{HistAvg}} = 43.85$ while announcing relative under-performance yields $\hat{u}_{\gamma_{HistAvg}} = 10.37$. These differences are highly significant ($p < .001$, two-sided t-test).

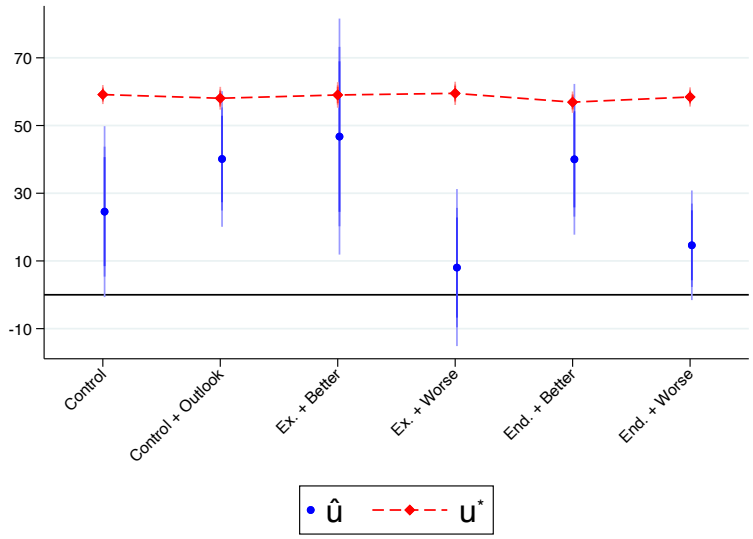


Figure A-10: Forecast credibility in *Contextual Communication* using $\hat{u}_{\gamma HistAvg}$
Note: This figure presents forecast credibility estimates for *Contextual Communication* treatments (blue circles), assuming subjects perceive a systematic forecast error ($\gamma_{HistAvg}$). The red diamonds represent the Bayesian optimal weighting of the central bank's forecast in the updated inflation expectation of a rational Bayesian participant, who equally considers all available historical information. Shaded bands around the markers denote 99% (lightest), 95%, and 90% (darkest) confidence intervals.

A7 Sensitivity Analysis

Our baseline analysis uses data Winsorized at 5th and 95th percentile so that extreme outliers do not drive our results. We chose these cutpoints because they were sufficient to eliminate extreme outliers in all instances for our data so that cutpoints remain consistent throughout. However, we understand that Winsorizing our data introduces a decision point into our analysis. Because of this, we explore in this section the sensitivity of results to the choice of cut points. To do this, we reproduce results from *Forecast Performance* (Figure A-11) and *Timing* (Figure A-12), which together comprise all eight histories that we use in this experiment. We show estimates of forecast credibility using cutpoints that retain data relative to our baseline (1st and 99th percentile) and trim additional data relative to the baseline (10th and 90th percentiles).

A7.1 *Forecast Performance*

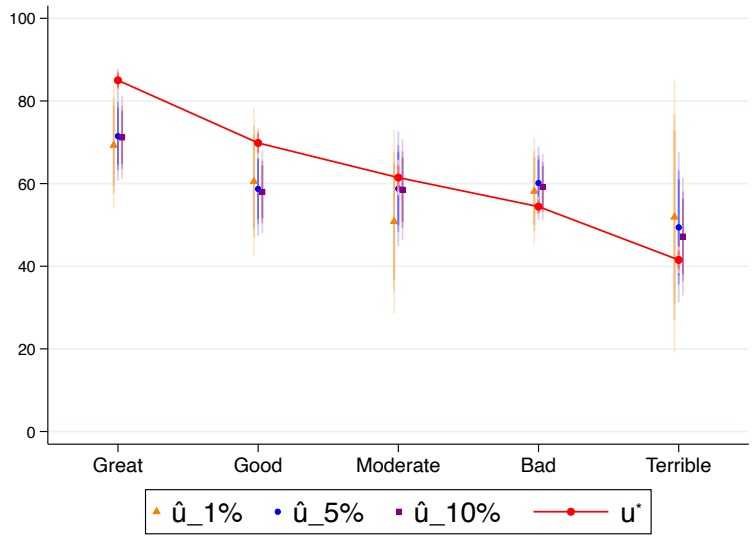


Figure A-11: Sensitivity Analysis for *Forecast Performance* results

Note: This figure provides estimates of forecast credibility, $\hat{u}$, in *Forecast Performance* using data that is Winsorized at the 1st and 99th percentiles (orange triangles), the 5th and 95th percentiles (blue circles), and the 10th and 90th percentiles (purple squares). Shaded bands around the point estimates indicate 99% (lightest), 95%, and 90% (darkest) confidence intervals.

A7.2 *Timing*

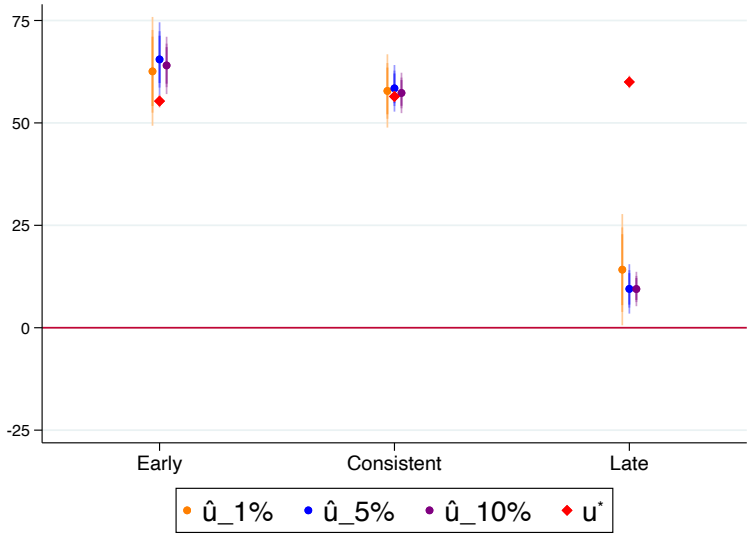


Figure A-12: Sensitivity Analysis for *Timing* results

Note: This figure presents estimates of forecast credibility, $\hat{u}$, for *Timing* treatments using Winsorized data. Orange triangles indicate data Winsorized at the 1st and 99th percentiles, blue circles at the 5th and 95th percentiles, and purple squares at the 10th and 90th percentiles. Shaded bands around the point estimates denote confidence intervals of 99% (lightest), 95%, and 90% (darkest).

A8 Alternative Approach to Estimating α_i

This section explores the robustness of our conclusions about whether and when subjects tend to over or under-update relative to a Bayesian benchmark, u^*. Note that the results here have no bearing on our primary finding that forecast credibility is imperfect, responds endogenously to short-term forecast errors, and that this response is asymmetric.

Whereas our baseline results assume $\alpha_i = \frac{1}{range_forecast}$, we might instead assume that the upper ($R_{l,i}$) and lower bounds ($R_{l,i}$) of i's range forecast represent percentiles for a distribution and then estimate α_i based on this assumption. In particular, we explore alternative estimates of α_i assuming forecast errors are normally distributed so that $R_u,\ R_l$ represent upper and lower percentiles in a normal distribution.

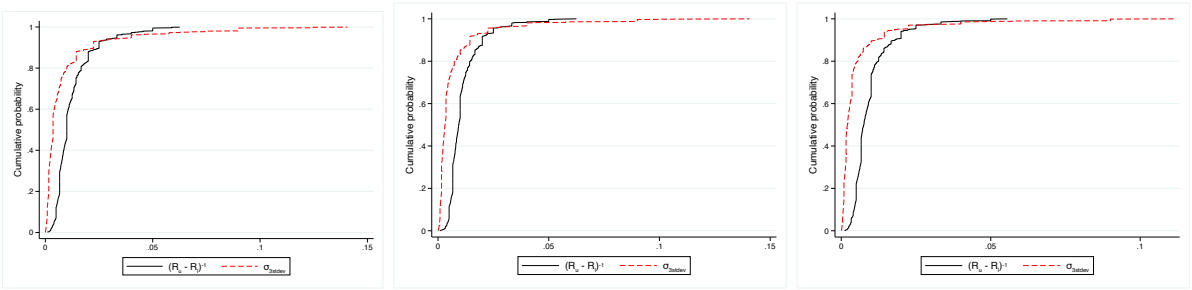
Given this assumption, we can estimate α_i as:

$$R_{u,i} - R_{l,i} = \mathbb{E}_{i,12}(\pi_{13}) + \sigma X - [\mathbb{E}_{i,12}(\pi_{13}) - \sigma X]$$
$$\implies \quad \alpha_i = \frac{1}{\sigma^2} = \frac{4X^2}{(R_u - R_l)^2}$$

where X represents the z-score corresponding to $R_{u,i},\ R_{l,i}$.

Suppose i's forecast bounds are symmetric and lie two (i.e. $X = 2$) or three ($X = 3$) standard deviations from i's point forecast. In the first case, we would have that $\alpha_i = \frac{16}{R_u - R_l}$ and in the second $\alpha_i = \frac{36}{R_u - R_l}$. Given that $u_i^* = \frac{\beta_T}{\alpha_i + \beta_T}$, this implies a higher u^* (i.e. i should place more weight on the central bank's forecast) in the first case than the second. Intuitively, this says that if the subject believes her forecast range incorporates less ($X = 2$) potential inflation realizations (i.e. less of her uncertainty) then she should place more weight on the central bank's signal than the world in which that same forecast range is meant to capture more potential future realized values of inflation.

We show the cumulative distribution functions of α_i using this alternative estimation strategy in A-13. Each subfigure displays the CDF of our baseline estimates of α_i alongside an alternative where we assume that upper and lower range bounds capture three standard deviations (σ_{3stdev}) of possible inflation values.



(a) *Early* **(b)** *Late* **(c)** *Consistent*

Figure A-13: Comparing baseline and alternative estimates of α_i

Additionally, we show our main results, using all three Winsorizing schemes discussed

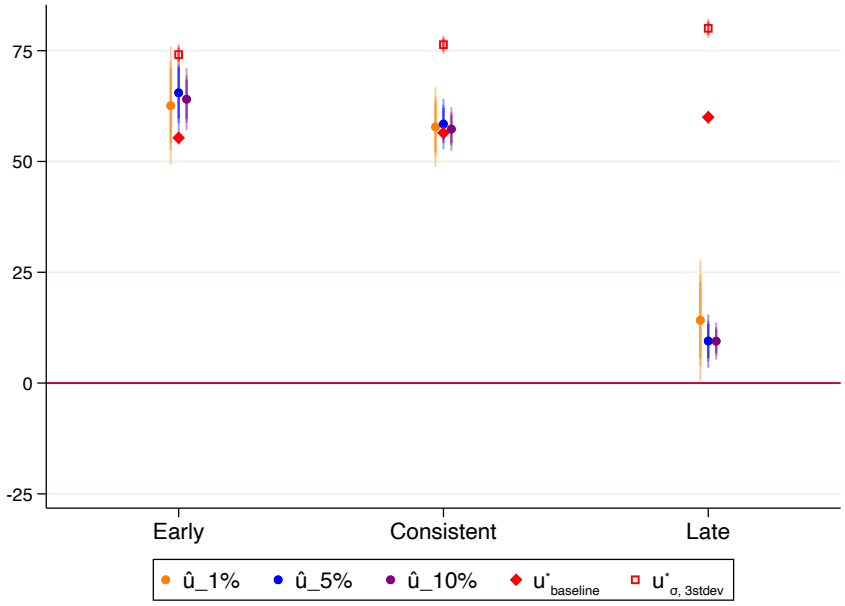


Figure A-14: Main results using both our baseline and alternative approach to estimating α_i.

in A7, using our baseline estimates of α_i to create our Bayesian benchmark ($u^*_{baseline}$) and this alternative approach to estimating α_i ($u^*_{\sigma_{3stdev}}$). We do this in A-14.

A9 Instructions

This section contains our experimental instructions for all treatments.

A9.1 *Contextual Communication*, *ForecastPerformance*, *Timing*, and *ReversedShock* Instructions

Experimental Instructions

You will now proceed to our experiment. If you read these instructions carefully and make appropriate decisions, you may earn a considerable bonus payment in addition to the participation payment. The bonus depends directly on the quality of your decisions.

You can access these instructions throughout the experiment. You may toggle the instructions on and off using the button labeled 'Instructions' below the 'Next' button on any page.

We will quiz you over these instructions on the following page. If you submit the quiz with at least one wrong answer more than three times then we will end the experiment early.

Your Objective in the Experiment

Your job in this experiment is to **forecast inflation**. Inflation is a measure of how prices change over an observed period of time. By 'inflation forecast' we mean your best guess of what inflation will be at a certain point in time. The more accurate your inflation forecasts, the more bonus money you earn!

You will provide two types of inflation forecasts:

- **Point Forecast:** Your 'Point Forecast' of inflation is your best guess of the exact value inflation will be at a certain point in time.
- **Range Forecast:** Your 'Range Forecast' of inflation allows for some uncertainty by letting you provide a range of possible values, defined by upper and lower inflation bounds, that you think will almost certainly contain the actual value of inflation.

Additional Definitions:

- **Central bank:** These national institutions provide banking services for the government, issue currency, and set interest rates to control inflation and maintain economic stability. Examples are the Federal Reserve in the United States and the Bank of England in the United Kingdom. An important part of a central bank's job is to provide economic forecasts to the public. Some examples of things central banks forecast are inflation and unemployment.

- **Forecast error:** A forecast error is the difference between an inflation forecast and inflation at a specific time. Your goal in this experiment is to have the smallest forecast error possible.
- **Quarter:** A quarter is a common unit of time for economic data. One quarter is equal to three months so that each year has four quarters. Central banks usually provide quarterly forecasts.

The experiment:

This experiment consists of **three decision periods**. In each decision period, you will form two sets of inflation forecasts. We call these your Initial Forecasts and your Updated Forecasts. The imagine below shows the flow of a decision period.

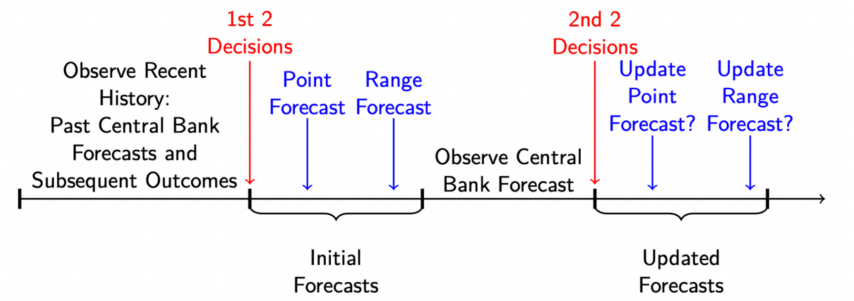


Figure: Experimental Timeline: A single decision period

1. We provide 12 quarters of history of inflation (blue line and dots) (*black line and dots* in *ReversedShock*) alongside the central bank's corresponding forecasts for those quarters (black line and dots) (*blue line and dots* in *ReversedShock* treatments). We also provide a summary of the central bank's historical forecast performance (text, next to the left side of the chart) that includes absolute forecast errors for each year and for the overall historical period.
    - Note that the difference between these two dots within the same quarter represents the central bank's forecast error for that quarter.
2. After viewing this history, you will provide your Initial Forecasts:
    - A point forecast (red dot) of inflation for the next quarter (Quarter 13)
    - Your corresponding range forecast of inflation.
3. After forming your Initial Forecasts, we will reveal to you the central bank's inflation forecast (on the next screen).
4. You will then provide your Updated Forecasts:
    - You will again provide a point forecast and a corresponding range forecast of inflation.
    - Your Updated Forecasts can be the same as your Initial Forecasts, use some of the same values, or use completely new values.

- We provide information about your Initial Forecasts both graphically and numerically whenever you are forming your Updated Forecasts.

5. After providing your Updated Forecasts, we will reveal the actual value of inflation for the forecasted period and inform you of your forecast performance.

6. **You will play through three decision periods with different economic data in each decision period.**

How our software scores your performance:

- *Point forecast:*
  - A perfect forecast earns exactly \$1.
  - The larger your forecast error (above or below), the less you earn.
- *Range forecast:*
  - If inflation does not fall inside your forecast range, you earn nothing for your range forecast.
  - The total range of your forecast is given by the gap between the upper bound and lower bound of your range forecast.
  - If actual inflation is inside your forecast range, you score $P = \frac{1}{1+totalrange}$.
  - The larger the range you create the less money you earn for your range forecast.

Suppose that actual inflation turns out to be 2.5%

- If you set your range from 1% to 3% then you would earn $P = \frac{1}{1+2} = \$.33$
- If you set your range from 1% to 5% then you would earn $P = \frac{1}{1+4} = \$.2$
- If you set your range from 3% to 5% then you would earn nothing since actual inflation is not within your range.
- you set your point forecast to 2.5% then you would earn \$1
- If you set your point forecast to 3.5% (or 1.5%) then you would earn \$0.50
- If you set your point forecast to 4.5% (or 0.5%) then you would earn \$0.25

You will get paid for your performance in one set of forecasts (Initial or Updated) in one of the 3 decision periods:

- Our software randomly chooses one of your three decision periods.
- For that decision period, the software randomly chooses either the initial forecasts or the updated forecasts.

- We pay you for this set of inflation forecasts as a bonus payment.

This means you need to take both the Initial Forecasts and the Updated Forecasts equally seriously when making your decisions.

Interacting with the data and inputting your forecasts:

The historical data:

- You may hover your mouse over any dot on the figure to see its exact value, which will appear in the upper left-hand corner of the graph.
- We remind you of your Initial Forecasts graphically (red dot and red shading) and numerically when forming your Updated Forecasts.

Providing your Point Forecast:

- You may submit positive values (prices are going up), negative values (prices are going down), or a value of zero.
- You can input your point forecast of inflation by clicking on the graph in the shaded 'Your Forecast' section and then dragging/dropping the dot that appears there.
- The dot will be red for your Initial Forecast and blue for your Updated Forecast.
- You may also type your forecast into the clearly labelled input text box.

Providing your Range Forecast:

- Our software will randomly generate upper and lower bounds for your range forecast (shaded region surrounding your point forecast).
- You may click on and drag these upper and lower bounds to whatever values you prefer.
- You can also drag the entire forecast range up and down.
- Your forecast range can be as big or small as you prefer.
- You may choose to have more or less range above your point forecast than below, and vice versa.
- Your upper (lower) bound must always be equal to or above (below) your point forecast - the software will prevent impossible range inputs.

A9.1.1 *Medium-Term* instructions

Your Objective in the Experiment

Your job in this experiment is to **forecast average inflation**. Inflation is a measure of how prices change over an observed period of time. By 'average inflation forecast' we mean your best guess of what inflation will be, on average, over a given time span. The more accurate your average inflation forecasts, the more bonus money you earn!

You will provide two types of average inflation forecasts:

- **Point Forecast:** Your 'Point Forecast' of inflation is your best guess of the exact value inflation will be, on average, over a given time span.
- **Range Forecast:** Your 'Range Forecast' of inflation allows for some uncertainty by letting you provide a range of possible values, defined by upper and lower inflation bounds, that you think will almost certainly contain the actual value of average inflation.

Additional Definitions:

- **Central bank:** These national institutions provide banking services for the government, issue currency, and set interest rates to control inflation and maintain economic stability. Examples are the Federal Reserve in the United States and the Bank of England in the United Kingdom. An important part of a central bank's job is to provide economic forecasts to the public. Some examples of things central banks forecast are inflation and unemployment.
- **Forecast error:** A forecast error is the difference between an inflation forecast and inflation at a specific time. Your goal in this experiment is to have the smallest forecast error possible.
- **Quarter:** A quarter is a common unit of time for economic data. One quarter is equal to three months so that each year has four quarters. Central banks usually provide quarterly forecasts.
- **Average Inflation:** Average inflation is what you think inflation will be, on average, for a given time span. For example, suppose inflation is 3% in one quarter and then 4% in the next. Average inflation for these two quarters is $\frac{3\%+4\%}{2} = \frac{7}{3}=$ 3.5%. Suppose instead, we want to know of average inflation for a year where inflation was 6% in the first quarter, 1% in the second, 3% in the third, and 2% in the fourth. Average inflation for the year would be $\frac{6\%+1\%+3\%+2\%}{4}$ $\frac{12}{4} = 3\%$. Remember, we are asking you to forecast average inflation.

The experiment:

This experiment consists of **three decision periods**. In each decision period, you will form two sets of inflation forecasts. We call these your Initial Forecasts and your Updated Forecasts. The imagine below shows the flow of a decision period.

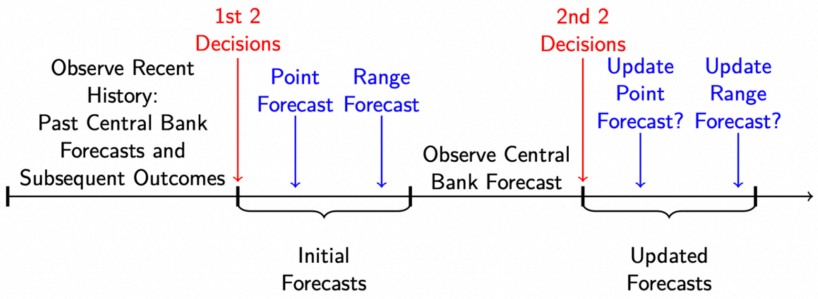


Figure: Experimental Timeline: A single decision period

1. We provide a 12-quarter history of inflation (blue line and dots) alongside the central bank's corresponding one-period-ahead forecasts for those quarters (black line and dots). For example, a one-period-ahead forecast would be if the central bank forecasts inflation for the fourth quarter of a year while in the third quarter of that same year. We also provide a summary of the central bank's historical forecast performance (text, next to left side of chart) that includes absolute forecast errors for each year and for the overall historical period. This historical data is **quarterly**. We provide **12 quarters** worth of historical data, which is equivalent to **three years** of data. We then ask you to forecast average inflation **for the next three years** .
    - Note that the difference between the blue dot and black dot within the same quarter represents the central bank's forecast error for that quarter.
2. After viewing this history, you will provide your Initial Forecasts:
    - A point forecast (red dot) of average inflation for the **next three years**.
    - Your corresponding range forecast of average inflation.
3. After forming your Initial Forecasts, we will reveal to you the central bank's forecast of average inflation for the next three years (on the next screen).
4. You will then provide your Updated Forecasts:
    - You will again provide a point forecast and a corresponding range forecast of average inflation.
    - Your Updated Forecasts can be the same as your Initial Forecasts, use some of the same values, or use completely new values.
    - We provide information about your Initial Forecasts both graphically and numerically whenever you are forming your Updated Forecasts.
5. After providing your Updated Forecasts, we will reveal the actual value of average inflation for the forecasted time span and inform you of your forecast performance.
6. **You will play through three decision periods with different economic data in each decision period.**

How our software scores your performance:

- *Point forecast:*
  - A perfect forecast earns exactly $1.
  - The larger your forecast error (above or below), the less you earn.
- *Range forecast:*
  - If actual average inflation does not fall inside your forecast range, you earn nothing for your range forecast.
  - The total range of your forecast is given by the gap between the upper bound of range forecast and the lower bound of range forecast.
  - If actual average inflation is inside your forecast range, you score $P = \frac{1}{1+totalrange}$.
  - The larger the range you create the less money you earn for your range forecast.

Suppose that actual average inflation turns out to be 2.5%

- If you set your range from 1% to 3% then you would earn $P = \frac{1}{1+2} = \$.33$
- If you set your range from 1% to 5% then you would earn $P = \frac{1}{1+4} = \$.2$
- If you set your range from 3% to 5% then you would earn nothing since actual average inflation is not within your range.
- you set your point forecast to 2.5% then you would earn $1
- If you set your point forecast to 3.5% (or 1.5%) then you would earn $0.50
- If you set your point forecast to 4.5% (or 0.5%) then you would earn $0.25

You will get paid for your performance in one set of forecasts (Initial or Updated) in one of the 3 decision periods:

- Our software randomly chooses one of your three decision periods.
- For that decision period, the software randomly chooses either the initial forecasts or the updated forecasts.
- We pay you for this set of inflation forecasts as a bonus payment.

This means you need to take both the Initial Forecasts and the Updated Forecasts equally seriously when making your decisions.

Interacting with the data and inputting your forecasts:

The historical data:

- You may hover your mouse over any dot on the figure to see its exact value, which will appear in the upper left-hand corner of the graph.

- We remind you of your Initial Forecasts graphically (red dot and red shading) and numerically when forming your Updated Forecasts.

Providing your Point Forecast:

- You may submit positive values (prices are going up), negative values (prices are going down), or a value of zero.
- You can input your point forecast of average inflation by clicking on the graph in the shaded 'Your Forecast' section and then dragging/dropping the dot that appears there.
- The dot will be red for your Initial Forecast and blue for your Updated Forecast.
- You may also type your forecast into the clearly labelled input text box.

Providing your Range Forecast:

- Our software will randomly generate upper and lower bounds for your range forecast (shaded region surrounding your point forecast).
- You may click on and drag these upper and lower bounds to whatever values you prefer.
- You can also drag the entire forecast range up and down.
- Your forecast range can be as big or small as you prefer.
- You may choose to have more or less range above your point forecast than below, and vice versa.
- Your upper (lower) bound must always be equal to or above (below) your point forecast - the software will prevent impossible range inputs.

A9.2 Comprehension Quizes, Survey Questions, & Economic Literacy Questions

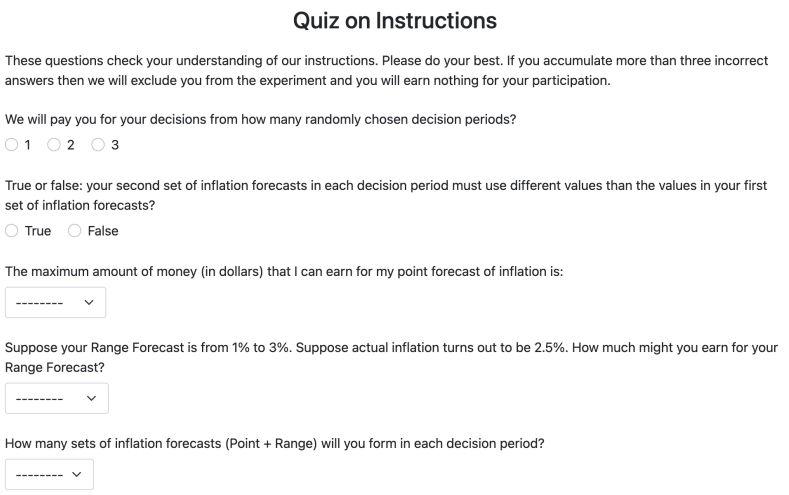


Figure A-15: This is a screenshot of the comprehension quiz faced by all subjects before beginning our experiment. Subjects who failed the quiz three times were excluded.

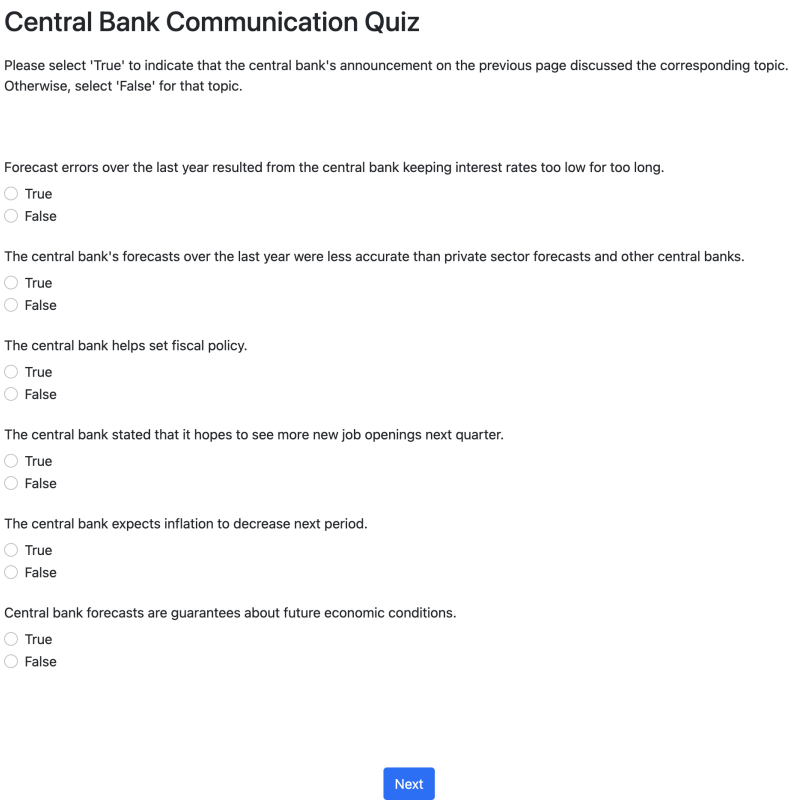


Figure A-16: This screenshot provides an example of the comprehension quiz we administered to subjects in our low-frequency communication treatments. The quiz came after subjects updated both their point and rage forecasts.

A9.2.1 Survey Questions

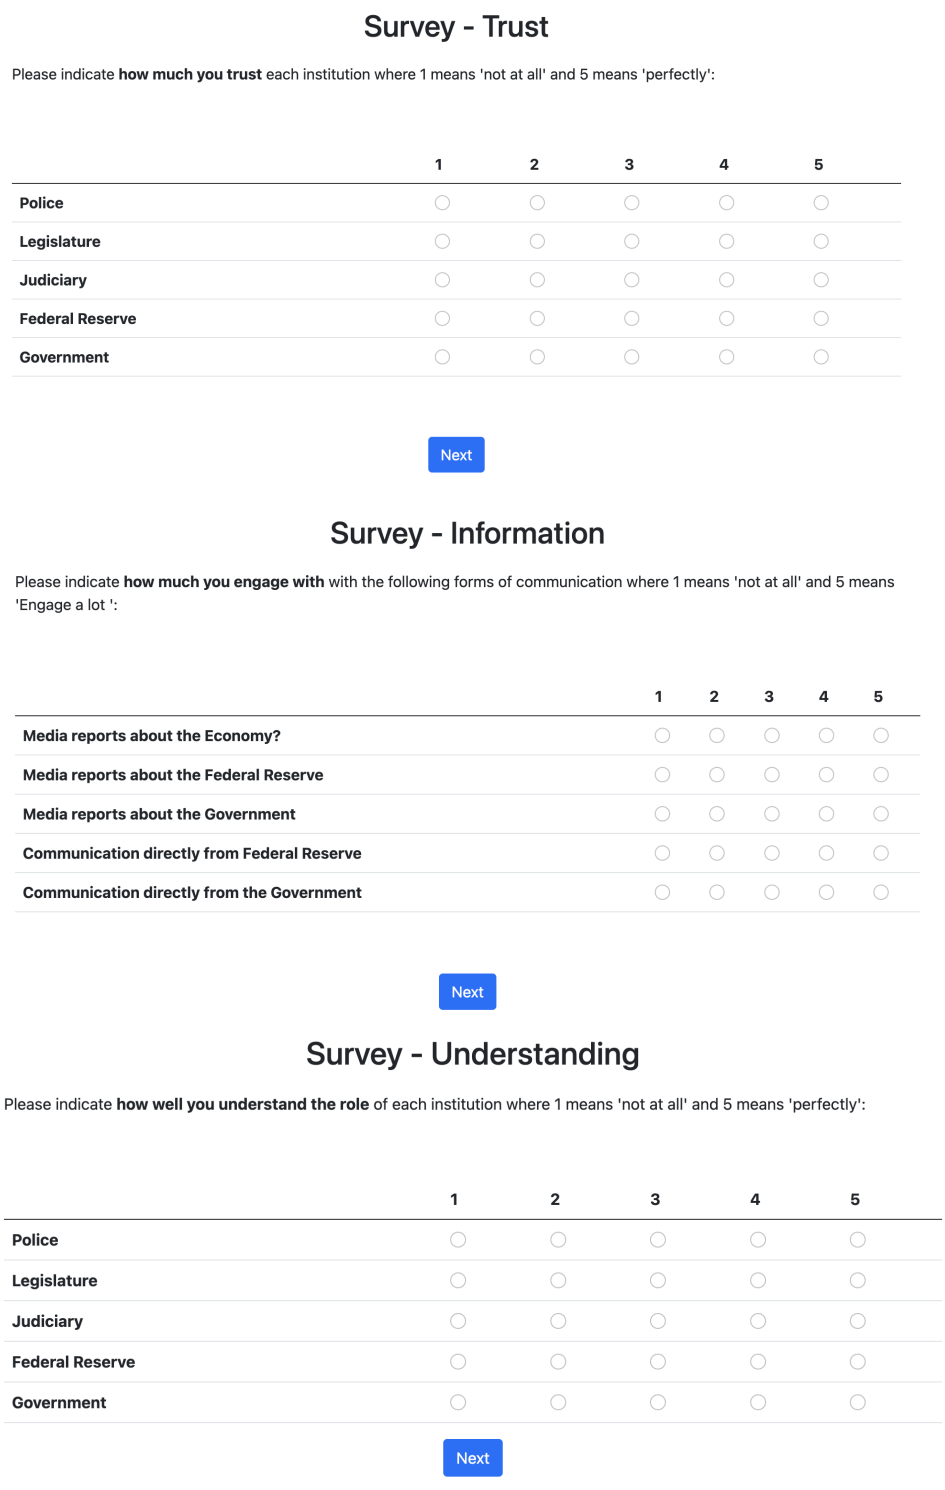


Figure A-17: Survey questions asking subjects about their level of trust and understanding for various U.S. institutions, and about preferences for information sources.

A9.2.2 Economics Literacy Questions

1. WagesRecession: What do you think happens to real wages (i.e. purchasing power) during a recession?
    - Fall on average
    - Unchanged on average
    - increase on average
    - Don't know
2. EmploymentRecession: What do you think happens to employment during a recession?
    - Goes up
    - Stays the same
    - Goes down
    - Don't know
3. InflationRecession: On average, what do you think happens to inflation during a recession?
    - It increases
    - It decreases
    - It remains unchanged
    - Don't know
4. BorrowRates: Suppose you need to borrow money. Which condition is best for you?
    - Interest rates are low
    - Interest rates are about average
    - Interest rates are high
    - Interest rates are irrelevant
    - Don't know
5. SavingRates: Suppose you are saving money. Which condition is best for you?
    - Interest rates are low
    - Interest rates are about average
    - Interest rates are high
    - Interest rates are irrelevant
    - Don't know
6. WhoSetsRates: Which U.S. institution sets the interest rate?
    - Congress

- President
- Federal Reserve
- Don't know

7. AvPi: What was the last reported level of inflation in the U.S.?
   - $< -1\%$
   - -1% to 1%
   - 1% to 3%
   - 3% to 5%
   - 5% to 8%
   - $> 8\%0$
8. InflationPreference: How much inflation do you think is good for the economy?
   - $< -1\%$
   - -1% to 1%
   - 1% to 3%
   - 3% to 5%
   - 5% to 8%
   - $> 8\%$