Why Do Central Banks Surprise the Public?

Paul Bousquet*

March 27, 2026
PRELIMINARY AND INCOMPLETE
[Click here for latest version]

Abstract

A typical interest rate target announcement by the Federal Reserve is not perfectly anticipated, suggesting they reveal information of some form. I argue policy surprises are mainly caused by unpredictable changes in the Fed's beliefs that are perceived as purely subjective judgment, rather than unique knowledge about the economy's trajectory. These "uninformative" belief surprises are nevertheless fundamentally different from traditional structural policy shocks and in practice are not spanned by the Fed's private point forecasts. A central bank acting as if its beliefs put increased probability on adverse outcomes rationalizes many features of the monetary policy status quo and gives surprises an interpretation of revealing the assessed balance of risks. When the central bank also faces a difficult learning problem, the finite sample persistence resulting from slow learning rates can generate many properties of the empirical distribution of policy surprises, including its state-dependence and estimated hump-shaped responses, even when these traits will vanish over time.

*University of Virginia, pbousquet@virginia.edu. This draft will be updated regularly; a few subsections of the Appendix and all but Section D of the Online Appendix are to be written.

1 Introduction

Structural monetary policy shocks in a macroeconomic model typically take the form of an unpredictable and arbitrary deviation from a policy rule. Central banks do not behave this way in practice yet often make decisions the public does not perfectly anticipate. For scheduled meetings outside the zero lower bound (ZLB) since 1994, the average market-implied forecast error of the Federal Reserve's overnight target rate announcement (Kuttner, 2001) is around 3 basis points, roughly translating to a 12% uncertainty level over choices 25 basis points apart. These monetary policy surprises, high-frequency movements in the yield curve around policy announcements, are commonly used to help describe the empirical effects of policy (Ramey, 2016). But any potential utility from their application can only be earned downstream of understanding their source; assumptions about policy surprises only make sense within the context of a model of policymaking (Sims, 1982).

Why are frequent surprises with modest but unignorable intensity the norm? Given a strong consensus that central bankers do not "flip coins" (Leeper et al., 1996; Nakamura and Steinsson, 2018b), policy decisions must reveal information of some kind, either about the central bank's reaction function form or its inputs (Campbell et al., 2012). The two leading explanations, the "central bank information effect" (CBI) and "response to news" (RTN), each correspond to a special case of one of these possibilities. CBI (Romer and Romer, 2000; Nakamura and Steinsson, 2018a; Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021) posits the public performs a signal extraction problem with policy decisions to learn about the economy, while RTN (Bauer and Swanson, 2023b; Cieslak et al., 2019) says the public is learning over time the central bank is more responsive to economic conditions than it thought. Less explored is an alternative offered in Caballero and Simsek (2022) and Sastry (2026): policy surprises have an "information effect" because they shed light on the central bank's beliefs, but they are regarded as useful only for forecasting future policy decisions (and not learning about the broader economy). In this paper, I argue these *uninformative* belief surprises (UBS) are the primary source of policy surprises in the data. The implications of this finding are further explored with a treatment of how belief surprises propagate in a broad class of models before restricting attention to a specific model that can rationalize many empirical properties of interest rates and their expectations.

Though the prevailing interpretations of policy surprises focus on different mechanisms, they arrive at a shared conclusion that policy surprises are primarily caused by the central bank having superior information of some form. If this is true, central banks should be superior forecasters of their future actions. Figures 1 & 8 show policy forecasts of Fed staff and voting members are not systematically better (Amodeo, 2026) and exhibit the same autocorrelated errors previously noted in private sector interest rate forecasts (Cieslak, 2018; Farmer et al., 2024). These graphs are instead consistent with a world that is intrinsically difficult to learn about. Specifically, when faced with uncertainty about a data generating process with components of varying persistence, as documented in

the term structure of interest rates (Atkeson and Kehoe, 2008), these forecast patterns are exactly what we should expect to see (Muth, 1961). This unobserved component structure provides the first building block for the ultimate model and is a helpful way to organize some thoughts and findings in the literature. First, this paradigm justifies why the central bank's beliefs would be perceived as uninformative. In a standard Bayesian benchmark, partial or full belief revelations from the central bank do act as a public signal, even if the information they act on is not necessarily more precise (Morris and Shin, 2002; Acosta, 2026).[1] But in a world hard for *everyone* to figure, the marginal information gain from one agent's subjective beliefs should be negligible, especially when there are many private belief aggregators available (e.g., through surveys and financial markets).[2] Also, much attention has been paid to the fact that information available prior to policy announcements is correlated with policy surprises (Bauer and Swanson, 2023b; Miranda-Agrippino and Williams, 2026).[3] But these estimated correlations are unstable, sensitive to outliers, and weak out of sample (Acosta, 2026; Sastry, 2026), strongly suggesting they only appear ex post (Eva and Winkler, 2023). When learning is slow, such finite sample artifacts will be abundant (Bauer and Swanson, 2023a) and are not reflective of any nuisance to be corrected.

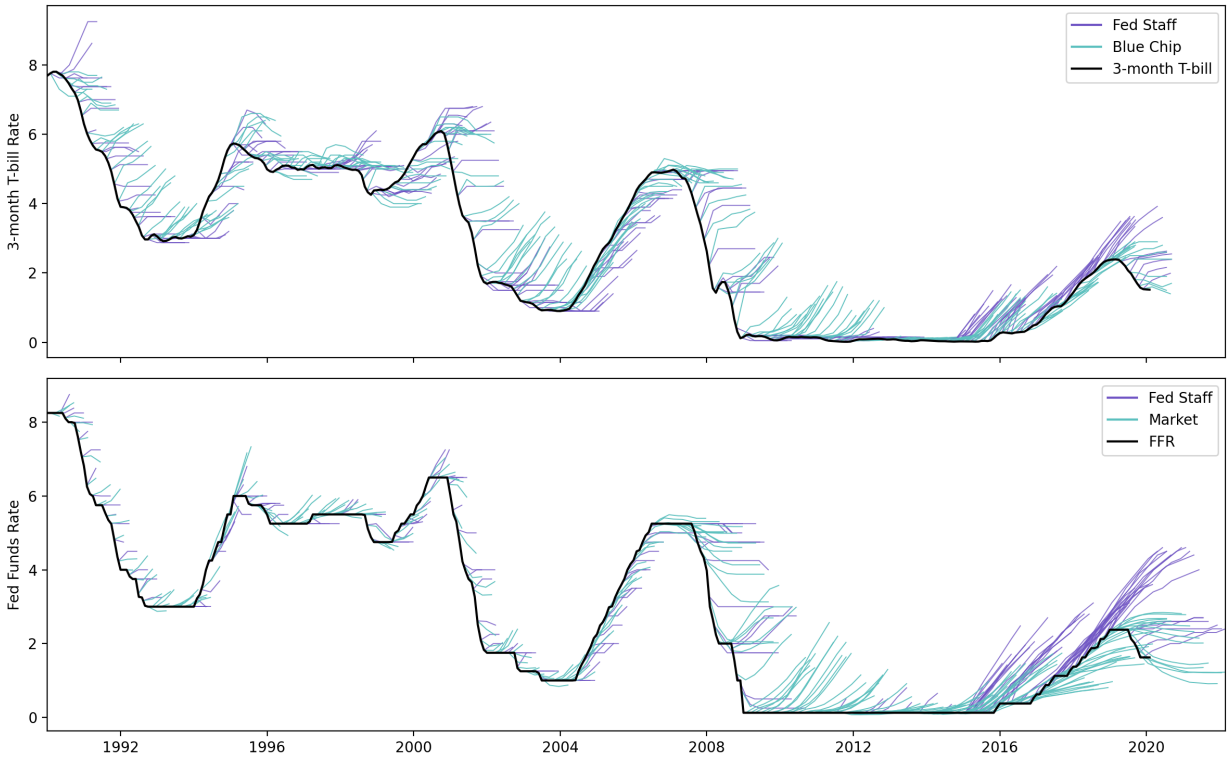


Figure 1: Forecasts of the 3-Month T-bill rate (DTB3) and the Federal Funds Rate. The top panel uses the Blue Chip Financial Forecast, the bottom panel uses Federal Funds Rate futures contracts, and the Fed staff forecasts in both panels are sourced from Greenbook/Tealbooks.

[1] Central banks arguably shape discussion on social and news media. Studying this distinct notion of an information effect first requires a formalization of how commentary, separate from data releases themselves, affects expectations (and whether it can under strict rationality).

[2] The formal statement of this idea: if private signals about an object with a linear law of motion are drawn from a shared distribution, observing the average signal makes the optimal gain on any individual signal equal to zero.

[3] Sastry (2026) points out that CBI cannot generate this fact on its own.

Giving agents a difficult learning problem goes a long way in explaining many important facts, but a complete environment needs a description of how central banks grapple with the pervasive uncertainty, the "defining feature" of monetary policy (Greenspan, 2003). Many representations of the Fed's reaction function take the form $g(\boldsymbol{X}_t)$, where $\boldsymbol{X}_t$ contains point forecasts of macroeconomic aggregates like inflation, output growth, and unemployment. With $\boldsymbol{X}_t$ in hand, policy surprises under these rules would be completely accounted for, a testable implication; the usual intuition that it's noteworthy to explain even a small amount of forecast error variation inverts when including information the forecaster did not have. Literally demanding $R^2 \approx 1$ is perhaps excessive, but we should see something meaningful if $g(\boldsymbol{X}_t)$ is a good approximation of the Fed.[4] I find $R^2 \approx 0$ during the Great Moderation period (excluding the ZLB), rejecting a broad class of reaction functions as descriptions. Since the Fed's beliefs about macroeconomic aggregates still presumably determine policy decisions, point forecasts must be far from a sufficient statistic.[5] This finding informs the second core component of the model: a central bank that places extra weight on the non-central parts of its belief distribution to hedge against the uncertainty it faces. I interpret this structure as "risk management" preferences (Kilian and Manganelli, 2008; Cieslak et al., 2025a), where the central bank acts *as if* it believes adverse models or scenarios are more likely to take out insurance against making decisions that are particularly costly. I present a special case of these preferences that produces periods of inaction and gradualism, a defining feature of the monetary policy status quo.[6]

Finally, returning to the inseparability of using and understanding policy surprises, a primary motivation of the CBI vs. RTN debate is to establish consensus on how best to incorporate surprises in empirical policy analysis. The identification holy grail is recovering structural policy shocks, but it remains unclear how such an object can arise if the central bank acts optimally from its perspective. Common informal arguments suggest it would be welcome news if policy surprises are a consequence of UBS (Nakamura and Steinsson, 2018a, fn. 20; Bauer and Swanson, 2023a, p. 119). The idea is the central bank picks from discrete options, and though the deliberations are involved, the reason for choosing one or the other is ultimately arbitrary from the public's perspective. And in the spirit of Sims (1998) or the Cochrane (2004) discussion of Romer and Romer (2004), meeting valid instrument requirements can substitute for a bona fide structural shock. Results from Caballero and Simsek (2022), where policy surprises occur because the public views the Fed's belief as a noisy version of its own, suggest uninformative belief surprises do not fit these criteria in general.[7] Though surprises and structural shocks yield similar high-frequency asset price dynamics, I show in Online Appendix D the correlation between leads of macroeconomic aggregates and policy surprises is essentially unrestricted in their model. The implication is that UBS do not have inherent bearing on monetary policy transmission. Section 3 builds on this theme by developing some results in a general class of models.

[4]The baseline RTN rule makes a similar prediction by imposing linearity and persistent underestimation of the response parameters.

[5]This aligns with themes from work linking policy announcements and risk premia, see Knox and Vissing-Jorgensen (2025) and Cieslak et al. (2025a) for a review. An alternative explanation is central bank accommodation of public opinion (Caballero and Simsek, 2022).

[6]Gradualism, incremental movement to a target, is subtly distinct from smoothing, where central banks always hit their (smoothed) target. Caballero and Simsek (2022) show belief disagreement provides a motive for gradualism but not enough to explain its scale in practice.

[7]If the public's dogmatic belief about the central bank turns out to be true, Sims (1987) shows the "as good as random" logic holds.

To generate policy surprises, there is a timing constraint placed on the central bank that they must announce the policy function for $t+1$ at time t. They can condition $t+1$ policy on information known at time t and a private signal they will receive between periods. In this setup, optimal policy will typically lead to policy surprises because the public understands the unpredictable component will be used efficiently. This foreshadows a conclusion that is perhaps surprising ex ante but is a very natural consequence of rational expectations – in a broad class of linear models, there will be no population impulse response to policy surprises. For example, suppose a central bank commits to setting i_{t+1} equal to its optimal forecast of the natural rate of interest r^*_{t+1}. In a simple "divine coincidence" model (Benigno and Woodford, 2004), aggregates will be proportional to $r^*_t - i_t$, the central bank's forecast error. Tomorrow's forecast error is expected to be zero both before and after the central bank receives its between-period signal. This means the policy surprise resulting from making use of their signal will be uncorrelated with aggregates (i.e., whether the forecast ended up too high or too low). A reasonable response to these results is to point to the wealth of previously estimated non-zero responses, but the "no response" statement only holds in population. For example, standard models must put in considerable effort to generate the "hump-shaped" responses that are often observed (Nakamura et al., 2025).[8] The aforementioned correlations inherent to slow learning environments can naturally lead to these shapes, amplified even further when the central bank practices gradualism. Extending this logic, it will also be possible to create the types of correlations of surprises with aggregates which have been described as "puzzling". These can occur even in linear models depending on what exactly the received signal is about or when the central bank effectively follows a conservative (attenuated or inertial) rule.[9] The findings all hold regardless of whether the central bank's signal is unique or is drawn from a distribution common to a continuum of agents and rendered "uninformative" by a simultaneous revelation of the average signal.

Before developing a framework for policy surprises, Section 2 lays out a large set of their empirical traits, including those pertinent to where positions are currently staked out in the literature. The RTN and CBI explanations are commonly referred to as passing a litmus test of rationalizing why surprises have been shown to be correlated with (i) information available prior to the policy announcement and (ii) professional forecasters revising their outlook in an expansionary direction. I add to and extend other comprehensive probes of (i) and (ii) (Bauer and Swanson, 2023a,b; Acosta, 2026; Sastry, 2026) showing they are unstable, rendering them unhelpful for distinguishing competing theories. More broadly, CBI and RTN are built around the public learning about the economy through the central bank and the importance of perceived reaction function shifts around policy announcements. I present evidence that these channels are not of systematic importance. In particular, while there may be some time variation in the central bank's rule or preferences (Bauer and Swanson, 2023b) and the process

[8] Traditional news shocks can generate hump-shaped responses (Cochrane, 1998; Milani and Treadwell, 2012; Adams and Barrett, 2025), but central banks do not perform forward guidance by orthogonal announcements of state-invariant future adjustments of random size and sign (Bianchi et al., 2024). Humps in my model occur because surprises imply a broader trajectory of interest rates on average and ex post.

[9] The latter connects to an earlier idea that partially hesitant policy action might cause a price puzzle (Sims, 1992; Balke and Emery, 1994).

of learning the correct solution to their decision problem (Sargent, 1999; Ilut and Valchev, 2025), in the Great Moderation period, we would expect such changes to mostly occur infrequently (Romer and Romer, 2023; Bianchi et al., 2025).[10] I also show policy surprises are conditionally skewed – surprises are more likely to share the same sign as the actual direction rates were adjusted – something that emerges when central banks practice gradualism.

The state-dependence of policy surprises further highlights a disconnect between surprises and typical structural policy shocks, the primary focus of Section 3. As demonstrated in McKay and Wolf (2023); Barnichon and Mesters (2023), there is no such thing as a "structure-free" approach to policy analysis because, at the least, theory is needed to make sense of the assumptions being made. Questions about identification in these environments are beyond the scope of this paper but are nonetheless significant. During the heat of the original battles in the rational expectations revolution, Chris Sims argued for thinking of policy decisions as realizations from a state-dependent stochastic process (Sims, 1982, 1987), implying "there is no possibility of taking policy actions that change the stochastic process followed by the economy" (Sims, 2008). In a companion paper (Bousquet, 2026a), I explore how this perspective can inform how to do empirics (see also Angrist et al. (2018); Kocherlakota (2019)) and further discuss why many common background assumptions are not innocuous.

Section 4 delivers the preceding body of facts through a model where the central bank has "risk management" preferences, all agents have uncertainty about the parameters governing the latent state, and policy surprises are not useful for learning about the economy. Policy surprises in this model reveal unpredictable changes in the central bank's balance of risks. The stance that private agents have a good understanding of the central bank's current reaction function but are unsure how their risk assessments have evolved is meant to reflect expectations formed immediately before a typical policy announcement. This simplification provides a unifying way to think about what high-frequency surprises tell us but leaves out other notable considerations, such as public uncertainty over whether discretionary reoptimizing has occurred (Debortoli and Lakdawala, 2016; King and Lu, 2025). An additional scenario to be considered is a private sector trying to learn about the central bank's beliefs and conduct, and in turn the central bank needs to learn about these higher-order beliefs in order to set optimal policy (Nimark, 2008; Stein and Sunderam, 2018; Caballero and Simsek, 2022; Rungcharoenkitkul and Winkler, 2026). Such infinite regress should seemingly imply more of a role for point forecasts to explain policy surprises, but there are limitations in our current understanding of these settings, and the two-sided learning concept undoubtedly is linked to why central bank communication is seen as important (Bernanke, 2004, 2013). Also, when there are low-frequency changes in preferences or general strategy, it's likely there will be a period where policy announcements are informative for learning about these changes (Cieslak et al., 2025b). Section 5 outlines these two criticisms further and discusses how this framework can be extended.

[10] U.S. central bankers largely share a strong sense of the properties of the overnight lending rate target as a policy tool (Williams, 2019).

2 Empirical Evidence on Policy Surprises

The purpose of this section is to present evidence on properties of U.S. monetary policy surprises, high-frequency movements in the yield curve around Federal Reserve policy announcements. Section 2.1 reviews the features of the monetary policy surprise data that have received the most attention in prior literature and shows they can be explained without appeals to a central bank with superior information. Section 2.2 tests basic implications of central banks having private information or only using point forecasts in their policy rule and finds negative results. Section 2.3 focuses on the potential quantitative importance of the central bank having private information about their policy function or the economy. Section 2.4 documents a new facts about policy surprises, including their conditional skewness and counter-cyclical dispersion. Policy surprise data (1993-2026) are sourced from Jarociński and Karadi (2025) and Acosta et al. (2025). Three different surprise series are used to give results across the short-end of the term structure: movement in the interest rate futures for the current period (MP1, reflecting the implied "forecast error" of the policy announcement), three months ahead (FF4), and three quarters ahead (ED4) (Kuttner, 2001; Acosta et al., 2025). Professional forecast data at the panel level is sourced from the Blue Chip Financial Forecasts (BCFF, 1993-2024) and Bloomberg's Economic Forecasts (ECFC, 2005-2025). The Blue Chip Economic Indicators (BCEI, 1993-2025) provides an added source of consensus forecast data. Federal Reserve staff forecasts are made available with a five-year lag by the Federal Reserve Bank of Philadelphia.

2.1 Overview of Existing Facts

There are two common explanations for why there is public uncertainty about monetary policy decisions. The "central bank information effect" (**CBI**) framing (Romer and Romer, 2000; Nakamura and Steinsson, 2018a; Jarociński and Karadi, 2020) argues the central bank has private knowledge about the economy, while the "response to news" (**RTN**) perspective Bauer and Swanson (2023a) says the central bank has private knowledge about its policy function. The discussion around these ideas typically revolves around special cases – the (superior) "information effect" and the (unexpectedly aggressive) "response to news".[11] These special cases are used primarily to explain surprises have been shown to be correlated with (i) information available prior to the policy announcement and (ii) professional forecasters revising their outlook in an expansionary direction. Comprehensive investigation of these facts in Bauer and Swanson (2023a,b); Acosta (2026); Sastry (2026) provide many caveats on these facts. CBI cannot explain (i) on its own and the correlations themselves are unstable. (ii) is not a robust feature of the data and the implied magnitudes are negligible. Table 1 summarizes the specific statements supporting these characterizations, which are also replicated in the Online Appendix. Before proceeding, I offer some additional contextualization for why past findings of (i) and (ii) appear to be at least somewhat spurious.

[11] Sastry (2026) and Knox and Vissing-Jorgensen (2025) provide comprehensive overviews and a general setting nesting all possibilities.

Table 1: Clarifications on Policy Surprise Correlations

Fact	Source	Details
Correlations between policy surprises and news are unstable, weak out of sample, and explain only a small amount of policy surprise variation	Sastry (2026), Acosta (2026), Bauer and Swanson (2023b)	Online Appendix (B.1)
Central bank private information cannot produce correlations with news on its own (i.e., absent other elements like different priors over news or misspecified learning)	Sastry (2026)	Online Appendix (B.2)
Correlations between policy surprises and professional forecast revisions, even when significant, have negligible magnitudes	Bauer and Swanson (2023b)	Online Appendix (B.3)
Puzzling correlations between policy surprises and revisions are sensitive to sample and often don't jointly exist for growth and unemployment forecasts	Bauer and Swanson (2023b)	Online Appendix (B.3)
Correlations between policy surprises and revisions are sensitive to revision timing	Bauer and Swanson (2023b) Acosta (2026)	Online Appendix (B.3)
In many of the strongest instances of comovement between revisions and surprises, stock indices and surprises did not comove.	Bauer and Swanson (2023b)	Online Appendix (B.3)
If news available prior to the FOMC meeting is included as controls, the effect of surprises on revision is generally a precisely estimated null.	Bauer and Swanson (2023b) Sastry (2026)	Online Appendix (B.3)

(i) at first blush seems troubling because it ostensibly violates the prediction of rational expectations that forecast errors should be uncorrelated with anything in the forecaster's information set. But given these correlations drift around considerably and have virtually no out of sample predictive power, we certainly cannot reject the hypothesis that they are finite sample artifacts in a world with slow learning (Bauer and Swanson, 2023a; Sastry, 2026).[12] It may be helpful to consider an analogy to econometrics developed further in Bousquet (2026a). If we think of private expectations more generally as nonparametric estimation of a conditional mean function where the data generating process is unknown, we should expect these estimates to share similar features as "purely statistical" (e.g., machine learning) estimation in small samples with high-dimensional data (Bianchi et al., 2022). In Appendix A.4, I show that conditional mean forecast errors from statistical methods in these environments can exhibit purported "predictability". This will also be the case for the model in Section 4.

Table 1 mentions a confounding factor of (ii): revisions are measured at monthly frequency, while monetary policy surprises are measured at an intraday frequency. This is especially problematic given the cyclical nature of monetary policy, detailed further in Section 2.4. In Appendix A.1, Figure 6 demonstrates that in some sense we should observe low-frequency comovement of surprises and revision to growth forecasts because of the raw correlation of surprises and growth. When for example contractionary surprises occur during times when all else equal growth will be higher than expected, then we would typically observe concurrent surprises and revisions even if the Fed is not the one "breaking the news". If surprises are not strong enough to cancel out the business

[12] Miranda-Agrippino and Williams (2026) use a high-frequency window around major data releases and find the same "predictability drift".

cycle movements they occur in, this is consistent with the small and sensitive results documented in Table 1.

Another difficulty of linking monthly revisions specifically to one Fed announcement is professional forecasters tend to update their beliefs at low frequencies, with the highest frequencies being in response to actual news releases. Martinez (2025) shows that among the forecasters submitting to the separate Blue Chip surveys (roughly two weeks apart on average), more than half of quarterly growth and inflation forecasts are not revised, with the lack of revisions increasing in the forecast horizon. Moderate rates of infrequent updating exist even in the Survey of Professional Forecasters, which is released quarterly. Bauer and Swanson (2023b) provide convincing evidence that professional forecasters do not regard policy announcements as times when any information about economic fundamentals is made public. It would therefore be surprising if (ii) did not go away with higher frequency data, and it seems to. High-frequency measures of revisions to expectations, such as correlations with stock indices, prediction market forecasts, or inflation expectations implied from breakeven rates on inflation-protected treasuries, reject systematic information effects (Bianchi et al., 2025; Acosta et al., 2025; Swanson et al., 2025).[13]

In terms of actual professional forecast revisions, I use two approaches to address the timing concern. One uses data from Bloomberg's ECFC survey, available at a daily frequency. Because of limitations on this data (many forecasters default to monthly updates), I look for the most recent update within the two weeks before and after policy surprises, delivering a weakly tighter window. Online Appendix B.3 shows no evidence of information effects, with most point estimates having the expected sign.

Second, these results are typically found using one particular Blue Chip consensus (BCEI), but using both (BCEI & BCFF) can also be used to better align the timing of Fed staff and consensus forecasts as well as better control for the documented sensitivity to unscheduled meetings. We can also control for the fact that these revision correlations may exist even when it goes *against* the Fed's beliefs. If the Fed has a more optimistic growth outlook but surprises in an expansionary direction, professional forecasters revising their forecasts down in the following month should not be attributed to learning from the Fed the economy is faltering.[14] Concretely, an information channel is opened only when the difference between Fed and professional forecasts shares the same sign as the surprise.[15] Though we might expect expansionary surprises are synonymous with a stronger belief of a weakening economy (and vice versa for negative surprises), only around 60% of scheduled Fed policy announcements feature this alignment of the sign of surprise and disagreement (average across inflation and growth forecasts, see Table 2). This also includes periods where the measured surprise was zero to account for announcements where the Fed potentially communicated new information about the economy but not policy.

Table 2 shows two different tests of (ii) with respect to a one standard deviation change in the policy surprise

[13] Acosta (2026) and Gormsen and Lazarus (2026) caveat that a meaningful information effect may exist but with a magnitude small enough to still be dominated by the conventional channels of monetary policy.

[14] The outlook may not be captured by Fed staff forecasts if they are modal rather than mean forecasts (Aruoba and Drechsel, 2026).

[15] This refers only to inflation and growth forecasts. For unemployment forecasts, the sign reverses. Also, this definition of surprise-disagreement alignment ignores that the difference in the base forecast could have the opposite sign as differences in perceived tail risks.

measures from 1993-2019 over the subsample away from the zero lower bound (ZLB). Both compare periods with and without surprise-disagreement alignment, arguably a necessary condition for an information effect. The first tests differences in revisions overall and the second tests only the revision's sign. Specifically, we test $\beta_1 - \beta_2 = 0$ and $\gamma_1 \neq 0$ in the following specifications

$$\textbf{Test A:} \quad \overline{\mathbb{E}}_{t+1}[y_{q+h}] - \overline{\mathbb{E}}_t[y_{q+h}] = \beta_0 + \beta_1 \cdot A_t^{y,h} z_t + \beta_2 \cdot (1 - A_t^{y,h}) z_t + \beta_A \cdot A_t^{y,h} + e_{t+1}^{y,h}$$

$$\textbf{Test B:} \quad \mathbb{1}_{\overline{\mathbb{E}}_{t+1}[y_{q+h}] - \overline{\mathbb{E}}_t[y_{q+h}] > 0} = \gamma_0 + \gamma_1 A_t^{y,h} + u_{t+1}^{y,h},$$

where $\overline{\mathbb{E}}_t[y_{q+h}]$ is the relevant time t Blue Chip consensus h-quarters ahead forecast of y, z_t is a given monetary policy surprise measure, and $A_t^{y,h}$ is an indicator variable of surprise-disagreement alignment (whether the surprise shares the same sign as the difference between the Fed and Blue Chip forecast). These results show that while there is generally greater surprise-revision comovement when there is surprise-disagreement alignment, there is scant evidence supporting an information effect story (further checks and details are in Online Appendix C.1). Most of the estimates are statistically insignificant, and even rarer is significant results across all surprise measures. The fluctuations across horizons for the same surprise series indicate a delicacy to sample size, as it's for instance hard to imagine an information effect not existing for one and four quarter ahead Real GDP forecasts but existing for two and three. The effect sizes are also almost always negligible, never amounting to more than a rounding error for Test A. Consistent with this, while Wald tests do reject that $\beta_1 - \beta_2$ is the same across h (for a given variable y), they do not reject that β_1 is the same across h. In Appendix A.1, Figure 7 provides a new illustration of how these correlation results are sensitive to outliers. A more general discussion on the possibility of the "information effect" is presented in Section 2.3.2. Summarizing, while information effects certainly can exist, there is no direct evidence they are a prevailing characteristic.

Another motivation for the focus on (i) and (ii) is the potential implications for regression exercises with policy surprises. That is – will the presence of the central bank's superior information or private knowledge of aggressiveness obscure the "true" policy shock the surprise measures contain? Given a strong consensus that central banks do not "flip coins" (Leeper et al., 1996; Ramey, 2016), there is no reason to think such an object exists. One common, but never formalized, argument draws an analogy to taste shocks in the industrial organization literature: the central bank largely follows some guiding principles, summarized reasonably well in a Taylor Rule, and the exact way it deviates from high-level representations is ultimately random from the perspective of everyone else. This motivates thinking about some persistent, structural policy shock process as a reasonable representation of observed policy surprises. Results in Section 3 show such an argument is untenable, generalizing a finding of Caballero and Simsek (2022). The source of the surprise determines the response path and it need not be aligned with the response to a hypothetical arbitrary disturbance.

Table 2

	CPI Inflation $t+h$ forecast					Real GDP Growth $t+h$ forecast					FFR
	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$	$h=4$
Test A: Difference in Forecast Revision Following Surprise Based on Type of Disagreement											
MP1	-0.162	0.092**	0.098***	0.027	0.069**	0.194	-0.075	0.105***	0.058**	0.033	0.069
	(0.214)	(0.039)	(0.037)	(0.041)	(0.028)	(0.127)	(0.102)	(0.039)	(0.022)	(0.036)	(0.048)
FF4	-0.215	0.053	0.080**	0.024	0.070***	0.182	-0.158	0.085**	0.041**	0.037	0.063
	(0.231)	(0.039)	(0.035)	(0.032)	(0.022)	(0.126)	(0.119)	(0.034)	(0.019)	(0.025)	(0.041)
ED4	-0.071	0.008	0.044**	0.044	0.028	0.087	-0.090	0.049*	-0.008	0.016	0.030
	(0.126)	(0.037)	(0.022)	(0.042)	(0.018)	(0.078)	(0.095)	(0.029)	(0.021)	(0.016)	(0.034)
Test B: Differential Probability of Surprise-Revision Comovement Based on Type of Disagreement											
MP1	0.033	-0.095	0.210***	0.065	0.085	0.107	0.005	-0.028	0.083	0.082	0.042
	(0.081)	(0.071)	(0.059)	(0.067)	(0.068)	(0.078)	(0.076)	(0.073)	(0.069)	(0.064)	(0.083)
FF4	0.058	-0.128*	0.101	0.079	0.057	0.088	-0.008	-0.035	0.066	0.094	-0.083
	(0.080)	(0.071)	(0.064)	(0.067)	(0.069)	(0.077)	(0.079)	(0.071)	(0.069)	(0.063)	(0.083)
ED4	0.005	-0.055	0.075	0.131*	0.093	0.043	-0.111	0.037	0.112*	0.000	0.015
	(0.079)	(0.070)	(0.065)	(0.067)	(0.069)	(0.079)	(0.077)	(0.070)	(0.067)	(0.066)	(0.079)

Columns index the dependent variable (forecast revisions) and each entry is a separate regression. Robust standard errors for each test in parentheses, significance levels * $p<0.10$ ** $p<0.05$, *** $p<0.01$. See Online Appendix C.1 for additional information/exercises.

It has already been noted that many justifications of the presence of an i.i.d wedge in policy decisions, such as discretionary reoptimizing, private information, and changes in taste, are fundamentally not an "exogenous" shock in the spirit that term is typically employed, even if the onset of such changes is truly random (Miranda-Agrippino and Ricco, 2021; Hack et al., 2024). But while individual explanations have been pointed out as insufficient, they have never been taken together, and in the background most frameworks will assume there is some arbitrary variation to be recovered. I explore how we might think about empirical exercises with policy surprises without such an assumption in Bousquet (2026a), and here I relay some takeaways relevant to the literature on the cause of policy surprises. Efforts in Jarociński and Karadi (2020, 2025) seek to partition "information shocks" from policy shocks. But unless the central bank has some sort of a trembling hand, all policy surprises are information revelations of a certain form (Campbell et al., 2012). This means such partitioning can be useful to separate information revelations we think are more proximate to a textbook policy shock (e.g., an unexpectedly bullish reading of public data) from information the central bank might be revealing about its institutional traits, like acts of discretion (Debortoli and Lakdawala, 2016; King and Lu, 2025) or insufficient commitment to its goals (Cieslak et al., 2025b; Ai and Hu, 2025). But still, even what we might call "relevant" information shocks will not be isomorphic to policy rule innovations in a standard model. In Appendix A.4, I show standard VAR approaches (recursive, sign, or zero restrictions) fit with a "small" VAR system consisting just of output, inflation, and interest rates on data simulated from a version of Smets and Wouters (2007) without monetary policy shocks will nonetheless recover "masquerading" policy shocks (Wolf, 2020). The responses can be similar to those estimated from US data and the VARs themselves pass diagnostic stability checks. As far as correlation with lagged news being an obstacle for the econometrician, in Appendix A.4, I show unnecessary controls in specifications where there is a "true policy shock" generically introduce

finite sample bias. Including controls, for instance to account for a "response to news" (Bauer and Swanson, 2023a), implicitly redefines the policy surprise as a new object that may not even be statistically associated with changes in the target policy instrument (Kolesár and Plagborg-Møller, 2025; Bousquet, 2026c), with the direction and severity of bias depending on the structure of the data generating process (Bousquet, 2026a,b).

2.2 Current Explanations Do Not Explain Surprises

If policymaking is intentional, the central bank's information set, properly controlled for, should account for all the variation in policy (Leeper et al., 1996; Cochrane, 2004; Aruoba and Drechsel, 2026). What about the variation in policy surprises? The CBI and RTN frameworks both suggest policy surprise regressions of a particular kind should have $R^2 \approx 1$. Specifically, the following interest rate rules are the baseline point of illustration for both mechanisms in the literature (Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023b; Sastry, 2026)

$$\textbf{CBI}: \quad i_t = \boldsymbol{\phi}' \cdot \widehat{\boldsymbol{X}}_t^{\boldsymbol{CB}}$$

$$\textbf{RTN}: \quad i_t = \boldsymbol{\phi}_t' \cdot \boldsymbol{X}_t.$$

In both models, the central bank follows a Taylor Rule with respect to some variables $\boldsymbol{X}_t$. In CBI, the central bank has private information about $\boldsymbol{X}_t$, whereas in RTN the central bank has private knowledge that the time-varying response parameter $\boldsymbol{\phi}_t'$ is larger than the public thinks. Both environments imply policy surprise variation should be completely explainable, either by private forecasts or by public data. Though these particular linear rules are not treated as beyond reproach, they are meant to be good first order approximations of the truth and also directly inform recommendations of how to use policy surprises. Also, these rules typically include a "policy shock", but even if such an object exists, there is now little collective doubt its magnitude is small (Sims, 1980; Nakamura and Steinsson, 2018b).[16] It's also worth noting that the CBI specification also applies to the "agree to disagree" channel proposed by Sastry (2026) – if policy surprises come from the Fed having a different mean forecast based on their subjective interpretation, belief differences should explain surprises.

Little, if any, variation of policy surprises is explained by private central bank forecasts and news. This is an extremely robust finding throughout a battery of specifications, detailed in Online Appendix C.2. Below, example results are presented for the 3 aforementioned surprise series. For CBI, we look at the difference between the Federal Reserve staff internal forecasts and the Bluechip (BCFF) consensus forecast (for inflation and output at horizons $h = 0, \ldots, 4$), giving a monthly sample from 1993-2019. For RTN, we use the "news" controls advocated in Bauer and Swanson (2023b), which include payroll surprises, employment growth, changes in the S&P 500, and changes in the term structure of interest rates, giving a sample at daily frequency from 1993-2025. Because the objective is describ-

[16] R^2 can also attenuate from measurement error. For example, Aruoba and Drechsel (2026) point out the Fed Staff's forecasts are better characterized as modal (rather than mean). The R^2 will also be smaller under RTN if the Fed is sometimes less aggressive than expected.

ing the "typical" conduct of monetary policy, periods at the zero lower bound as well as months with a significant unscheduled meeting are excluded.[17] Table 3 presents the results. The last column shows a specification combining the CBI and RTN variables with unscheduled meetings included, constituting a de facto ceiling on the R^2.

Table 3: Policy Surprise Variation Explained by CBI and RTN

Adjusted R^2

	CBI	RTN	CBI+RTN	CBI+RTN
MP1	<0	<0	<0	.03
FF4	<0	.02	0	.05
ED4	.06**	.07**	.17***	.19***
N	152	179	152	156
w/ Unsched.				✓

HC1 Significance levels * $p < 0.10$ ** $p < 0.05$, *** $p < 0.01$ for Wald test (all $\beta = 0$). See Online Appendix C.2 for additional information/exercises.

Online Appendix C.2 details the many other ways these exercises could be run. Other definitions of "news", such as current period forecasts or forecast revisions, do not improve performance beyond the Bauer and Swanson (2023b) controls, nor do inclusions of different variable forecasts. Though Fed staff forecasts have been repeatedly established as largely representative (Romer and Romer, 2004), it's possible at the margin having the forecasts of the FOMC members themselves would make a difference, but the same qualitative findings recur when instead using the FOMC's summary of economic projections (SEP). Albeit with a short sample (from Bloomberg's ECFF since 2018), incorporating survey data on the *expected* median FOMC forecast also has virtually no explanatory power. Lastly, if we take the point estimates from these regressions at face value, Online Appendix C.2 shows some are puzzlingly signed, though statistically insignificant (e.g., the Fed having higher one quarter ahead inflation expectations is associated with negative policy surprises).

The results constitute a resounding rejection of the Federal Reserve approximately following a Taylor Rule at the meeting-level frequency, or a time-varying Taylor Rule with private information on aggressiveness. In some sense, this is not surprising. Some might argue the meeting-by-meeting decisions carry excessive noise, though at lower frequencies the Taylor Rule also seems to be a poor description (Nakamura et al., 2025). Regardless, there must exist a strong link between a meeting-level rule and a lower-frequency rule, so providing an explanation is paramount. Assumptions about how policy is made are also a prerequisite to identifying policy surprises (Sims, 1982, 1987), meaning this question needs to be thoroughly addressed before considering how policy surprises should be used.

[17] The ZLB periods typically have an attenuating effect. Fed staff forecasts are only available for months with a scheduled meeting, limiting the number of possible exclusions.

2.3 Current Explanations Do Not Explain the Status Quo

2.3.1 Expectations of Future Policy and the Reaction Function

Though some specific recommendations of Bauer and Swanson (2023a) are based on the specific interest rate rule described above, it's carefully framed as motivation for something that should be "helpful" but not necessarily a panacea.[18] The bigger idea underlying Bauer and Swanson (2023a,b) is the Fed has private knowledge on its responsiveness. While a large part of the motivation was to develop a framework to explain facts (i) and (ii) discussed previously, we can take this explanation at face value for an explanation of why policy surprises persist. It's certainly plausible that the Fed has private knowledge on its reaction function, particularly if we think of the two zero lower bound episodes as periods where new parts of the state space were being explored. But in general, the Fed is not a better forecaster of future policy than private agents; Amodeo (2026) shows they are actually slightly worse across their full sample. More importantly for the RTN argument, Figure 1 shows the Fed and the private sector make the same type of forecast errors – with a notable exception of the ZLB liftoff, they both underestimate the extent to which rates will be adjusted (Cieslak, 2018; Farmer et al., 2024). This suggests the movement of interest rates is driven by a data generating process with a substantial persistent component, consistent with Atkeson and Kehoe (2008). Learning in such environments is difficult (Muth, 1961), meaning this should naturally lead to a policy rule with ample incorporation of current uncertainty (Greenspan, 2003; Orphanides and Williams, 2007).

Bauer and Swanson (2023b)'s "response to news" is one of many papers to argue that perceptions of the Fed's policy rule have drifted over time (e.g., Bauer et al., 2024). But such results are more narrowly showing linear policy rule fits exhibit drift, subject to the well-established difficulties of policy inference from data generated in a particular equilibrium (Cochrane, 2011; Occhino, 2019; Beaudry et al., 2024). This leads to a question of whether there has actually been (perceived) reaction function drift, or merely drift in empirical representations. Anecdotally, the stability of long-run inflation expectations suggests stability along these fronts, though arguably this is still consistent with drift if the changes will average out along some more stable reaction function over time (Cieslak et al., 2024). Though reaction function drift can occur at low frequencies, the Bauer and Swanson (2023b) framework for understanding policy surprises suggests policy announcements are the main driver of reaction function revelation. We would therefore expect the revisions to expectations of future policy that are not driven by changes in expectations of macroeconomic fundamentals to be persistent. One simple test of this hypothesis is how surprises at the shortest end of the yield curve propagate to nearby maturities. Unless changes to the reaction function or shocks in the economy are expected to quickly mean-revert, there should be what looks like a level shift in near-term policy expectations. More formally, an exact level shift occurs under the RTN rule if $\overline{\mathbb{E}}_t[\boldsymbol{\phi}_{t+h}] = \overline{\mathbb{E}}_t[\boldsymbol{\phi}_t]$

[18] Econometrically, this is unclear without further assumptions (Bousquet, 2026a,b).

and $\overline{\mathbb{E}}_t[\boldsymbol{X}_{t+h}] = \boldsymbol{X}_t$ (and $R^2 = 1$). Figure 2 shows that is not the case – nearby maturities do not move 1 for 1 with surprise changes to today's interest rate target and are decreasing in the futures contract horizon (along with the R^2). Though this is relative to an extremely persistent benchmark, Appendix A.2 discusses a more general case where $\boldsymbol{X}_{t+h}$ is a function of several, less persistent structural shocks and shows Figure 2 is consistent with RTN only if the most persistent shock accounts for a tiny amount of the variance of i_t. These results are more straightforwardly consistent with beliefs, rather than policy rule revisions, being the predominant driver of surprises.

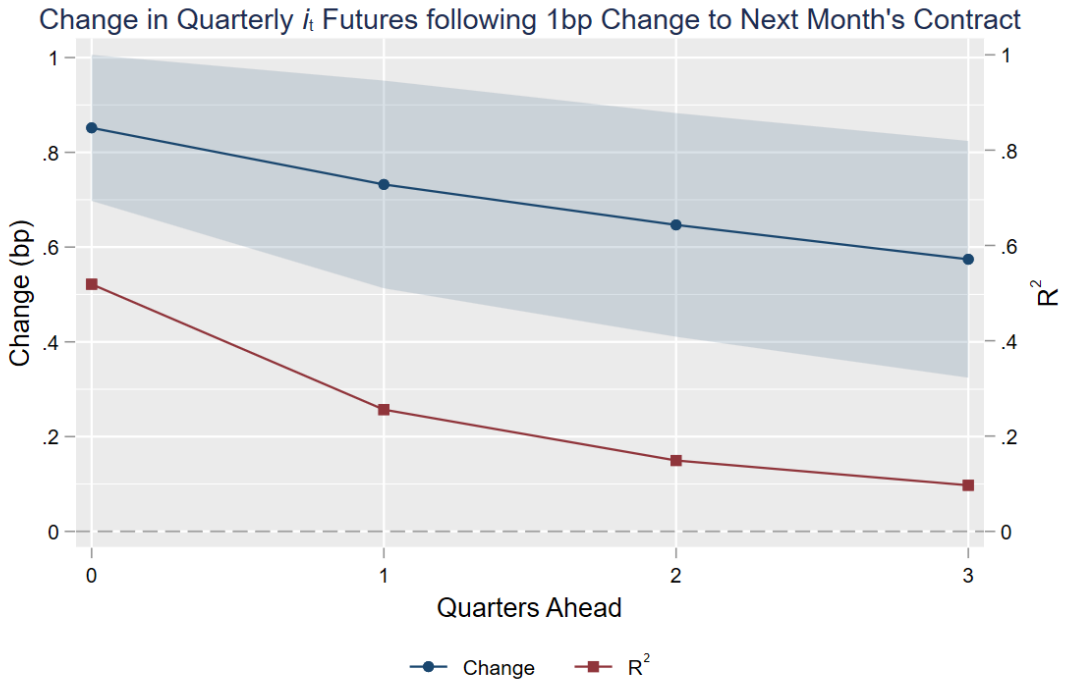


Figure 2: ED*h* futures on FF2 futures in narrow window around scheduled policy announcements away from ZLB, 1993-2025. Wald tests reject null hypothesis that coefficients are not decreasing in *h*. See Online Appendix C.4 for additional information/exercises.

Another place to look for persistence of policy expectation revisions not driven by macro news is to look at the autocorrelation of residuals $u_{t,i}$ from the following specification

$$\Delta x_{t,i} = \boldsymbol{\beta}' \cdot \widehat{\text{macro}}_{t,i} + \alpha_i + \tau_t + u_{t,i},$$

where α_i and τ_t are time and individual fixed effects. Here, this is run at the panel level (with forecasters from BCFF) and with two choices of $x_{t,i}$ – one year ahead interest rate forecasts and forecast disagreement relative to the Fed. The panel setting is meant to flexibly control for belief heterogeneity (i.e., other parts of the "aggregate belief distribution") and the additional inclusion of the disagreement variable tries to further remove the influence of revisions stemming from disagreement about the outlook. The $\widehat{\text{macro}}_{t,i}$ variable contains relevant factors to the macroeconomic outlook plausibly unrelated to contemporaneous changes in policy and is clearly the sticking point of these exercises. $u_{t,i}$ is negatively autocorrelated across a range of specifications for $\widehat{\text{macro}}_{t,i}$, suggestive

evidence inconsistent with an environment where agents are constantly learning about a drifting policy rule.[19]

In a "response to news" world, we would also expect to see belief convergence over time. Even if the public is perpetually underestimating the Fed's responsiveness, we should see forecasts of future policy becoming more aligned if the primary wedge is the central bank has private reaction function knowledge. Across several measures of future policy disagreement, I find it does not decline over time. As a primary specification, we can look at whether disagreement about interest rates one year ahead of the forecast date is expected to be smaller k periods after a monetary shock (or more generally a monetary policy announcement). The RTN theory predicts positive monetary surprises occur when the Fed has a higher forecast for future rates, so the change in disagreement should be negative (and vice versa for negative surprises). Figure 3 shows there is no such negative covariance across various surprise measures. One objection could be that slow learning and limited sample size could be consistent with persistent disagreement even in expectation. Table 2 shows revisions to interest rate forecasts following surprises are not materially stronger when the difference between Fed staff and professional forecasts indicates a potential "information effect" (when the difference shares the same sign as the surprise). While it's possible the Fed action can be a surprise even to the staff (Aruoba and Drechsel, 2026), Figure 8 shows in Appendix A.2 this trend and broader disagreement persists when using the forecasts of the Fed governors and presidents. This is suggestive that the disagreement is a consequence of pervasive uncertainty about what the future actually holds for the economy, rather than private reaction function knowledge.

2.3.2 Information Effect

It's not hard to imagine situations where monetary policy announcements also communicate information about the economy. During the Great Financial Crisis, it's plausible the public sometimes inferred the Fed's private knowledge on the day to day activity of the financial system informed some of their decisions. Event study decompositions of specific announcements, in the vein of Del Negro et al. (2023) and Bianchi et al. (2025), are definitely appropriate to assess whether meaningful information about the economy was revealed. But the thesis these occurrences describe dynamics at monetary policy events more generally is in conflict with well-documented behaviors of forecasters discussed in Section 2.1. Similar to response to news for policy forecasts, if the public is learning from the Fed we would expect disagreement about economic trajectory to get smaller over time. I find this is not the case – Figure 3 is a reflective indicator of the persistence of disagreement for other variables. As mentioned in the discussion for Table 2, it's also frequently the case that contractionary policy surprises do not occur when the Fed sees stronger output or inflation forecasts (and likewise for expansionary surprises and weaker outlooks). Some of the point estimates in isolation in Table 2 do support the presence of a marginal information effect, but also raise

[19] $\widehat{\text{macro}}_{t,i}$ includes combinations of individual-level forecast revisions that predate policy announcements, macroeconomic data that predates both policy announcements and forecasts (e.g., Bauer and Swanson, 2023b), and the predicted value of the disagreement with Fed growth and inflation forecasts using only these factors. More details can be found in Online Appendix C.3.

the question of how professional forecasters can successfully differentiate when one is present, and what the cause of surprises is otherwise perceived to be.

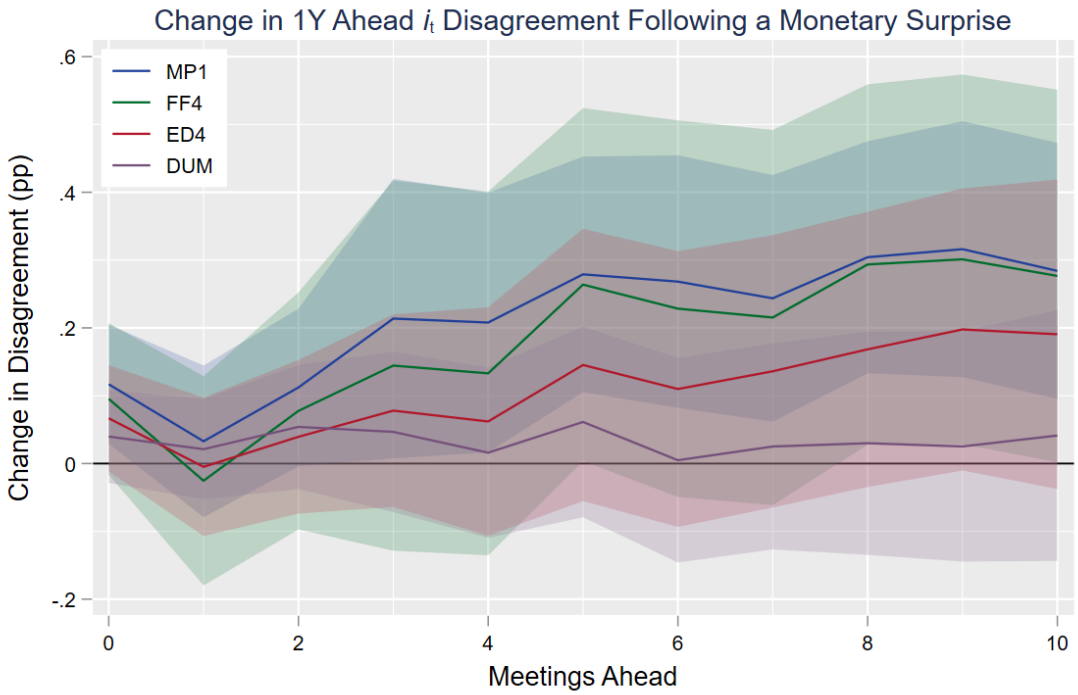


Figure 3: The change in disagreement $t+k$ meetings ahead (for Fed vs. BCFF consensus forecast, 1993-2019). "DUM" is a dummy to indicate a positive (+1) or negative (-1) surprise to handle the ZLB. More details and variations, including at the panel level, are in Online Appendix C.5.

Acosta (2026) puts forward two important explanations for an "information effect" different from the simplest interpretation that the Fed has superior knowledge: agents still should learn from each other in standard Bayesian environments (Morris and Shin, 2002) and policy announcements could emphasize different pieces of public information, helping people know what to pay attention to. The Bayesian learning analogy does not hold together when considering the central bank as just one agent among a continuum of private agents and various measures of the average private signal are available through surveys and private markets. In Appendix A.5, I document that forecasters of specific data releases (e.g., the latest CPI print) do not seem to update their beliefs based on professional forecasters with impressive records who submit a public forecast well in advance. More generally, the universe of signals is so large that the weight on any measure-zero subset will be negligible under optimal learning, especially when there are several signal aggregators available. This is consistent with Bianchi et al. (2022) who use a machine learning proxy for individual forecasters and find that incorporating the forecasts of other individuals does not improve performance. For the idea that central banks teach others how to learn, Sastry (2026) shows forecast errors in months without a policy announcement are statistically indistinguishable from months with policy announcements. To resurrect the world of Morris and Shin (2002), one might argue the Fed has a *unique*

private signal. Nakamura and Steinsson (2018a) point out that the Fed produces Industrial Production data, and in months where timing aligns they will have access to this data. I show in Appendix A.1 there is no evidence of distinguished information effects in these instances. Without these considerations of unique data access, it's hard to argue the Fed has a unique signal in a way e.g., Goldman Sachs does not.

2.4 Distribution of Surprises

Monetary policy shocks in DSGE models are typically white noise innovations to a Taylor Rule. One could reinterpret the finding that empirical policy surprises are correlated with lagged news as implying that policy surprises have a state-dependent distribution. While Sections 2.1, 2.2 cast doubt on the relevance of some of the specific correlations, it's clear that if central banks do not "flip coins", then the distribution of policy surprises must have a distribution that depends on something about the central bank, either its reaction function (e.g., discretion, taste shifts, or generally being unknown) or something it's reacting to. Figure 4 presents clear evidence of state-dependence: MP1 surprises are conditionally skewed towards the direction policy was actually adjusted. There also appears to be a counter-cyclical uncertainty effect; even controlling for some outliers from unscheduled meetings, the distribution of surprises occurring as a result of interest rate cuts is flatter with a much longer tail. While the CBI and RTN paradigms can be made consistent with this distribution, it's an indication that something fundamental about both monetary policy and the data generating process informing its conduct are missing from the benchmark models.

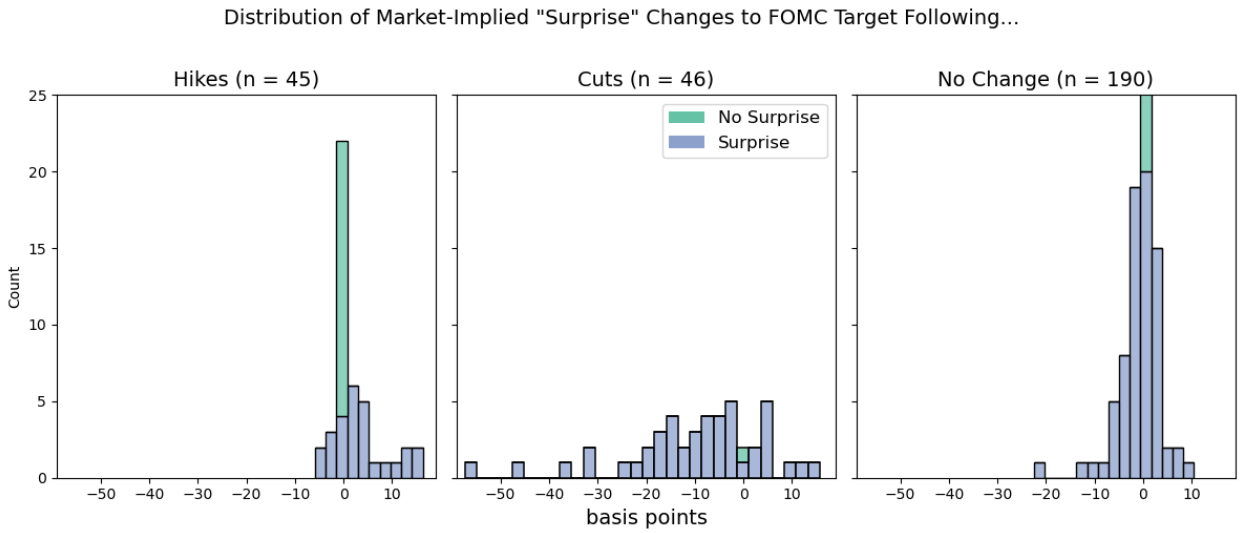


Figure 4: High-frequency changes (MP1) in futures from 1990-2024, cumulated by month (Jarociński, 2024). Note: the y-axis is truncated to have uniform scaling across panels; there are over 100 "no surprise" events when the interest rate is unchanged.

The conditionally skewed distribution of surprises can also be used to say something about the Fed's gradualism tendencies. Though there is often second-guessing at the time, there are few or no episodes in US Monetary policy

beginning with Volker where the ex post consensus is the Fed acted too aggressively.[20] Figure 4 represents a propensity of the Fed to "catch up". In Appendix A.2, Figure 9 demonstrates this more formally; large change in the target makes it substantially less likely monetary policy will reverse course at some point over the next year. This is very different behavior from standard DSGE models, where a move of a given sign or size does not alter the likelihood of future policy changes. This motivates a model of policymaking developed in Section 4 where the central bank does not act aggressively unless it's certain there is substantial room for future adjustments.

3 Belief Surprises vs. Structural Shocks

3.1 A General Framework for Policy Surprises

There are several ways to introduce unpredictable movements in policy other than the central bank introducing arbitrary disturbances to a policy rule. Fundamentally, we can still group all such approaches under a banner of coming from the public's lack of access to some form of information. For example, in a Markov switching setup (e.g., Davig and Leeper, 2007), there would be no policy surprises if the latent states were known in advance. If the central bank has private information (e.g., a private signal or private knowledge of its reaction function), closing the information asymmetry eliminates policy surprises. It's been suggested that one reason why there might be recoverable structural policy shocks is because even though the public understands the central bank is trying to act optimally according to its objectives, there remains a fundamental opacity for the exact inclination that led them to make a particular decision. The structure for policy surprises in this paper is meant to emulate this description as closely as possible.

To generate surprises, we consider a class of models with a particular constraint on the central bank and general information structure. The central bank must announce a policy function at the start of period t that will determine the realization of policy instruments at $t+1$. The policy function is measurable with respect to all common knowledge at t and a signal the central bank receives between periods. This structure is purposefully broad – this can nest cases where there are several policy instruments, signals, or different assumptions about what information is common to all agents. Without loss of generality, we will represent a scalar case with policy function

$$i_{t+1}(s) : \overline{\mathcal{I}}_t \times \mathbb{R} \to \mathbb{R} \tag{1}$$

where $\overline{\mathcal{I}}_t$ is the common time t information set implicit in the time $t+1$ subscript of the function. For expositional purposes, consider this policy function the mapping that determines tomorrow's interest rate. If agents knew ahead of time the central bank would receive signal s between periods, they would know tomorrow's interest rate exactly. The central bank can decide to make tomorrow's interest rate known regardless of the signal's realization

[20] Cieslak et al. (2025b) offers the full extent of the 2017-2019 tightening as a possible exception.

by making the function constant in s. More traditional uses of the signal would for instance include committing to use the signal as prescribed by a standard Kalman filter.

3.2 Irrelevance Results for Belief Surprises

Let $\mathbb{E}[\cdot]$ denote the objective expectations operator with respect to a given data generating process.

Proposition 1: From (1), let $z_t = i_t - \mathbb{E}[i_t \mid \overline{\mathcal{I}}_{t-1}]$ and let $s \notin \overline{\mathcal{I}}_{t-1}$ denote the additional random variable i_t can be conditioned on. If $\mathbb{E}[y_{t+h} \mid \overline{\mathcal{I}}_{t-1}, s] = \mathbb{E}[y_{t+h} \mid \overline{\mathcal{I}}_{t-1}]$, then $\text{Cov}(y_{t+h}, z_t) = 0$.

This is a simple correlation of rational expectations, but applied to policy surprises it is arguably harder to reconcile with standard intuition. First, we can consider a purely statistical explanation by the law of iterated expectations

$$\text{Cov}(y_{t+h}, z_t) = \mathbb{E}\Big[z_t \Big(\mathbb{E}[y_{t+h} \mid \overline{\mathcal{I}}_{t-1}, s] - \mathbb{E}[y_{t+h} \mid \overline{\mathcal{I}}_{t-1}]\Big)\Big] = 0,$$

where the last step follows from the Proposition's sufficient condition.[21] Note that z_t is a known object given the central bank's between-period information set. Optimal policy for the central bank, because of its timing constraint, involves a heavy forecasting component. Because z_t cannot be correlated with the central bank's forecast error, this property will extend to many of the aggregates y_t in a broad class of models. For example, consider the simplest "divine coincidence" equilibrium in a New Keynesian model (Benigno and Woodford, 2004) where inflation π_t and the output gap $\widetilde{y}_t$ are directly proportional to $i_t - r_t^*$, the gap between the interest rate and the natural rate of interest. If the central bank follows a rule where it seeks to track the natural rate, it will always announce a policy function of $i_{t+1} = \mathbb{E}[r_{t+1}^* \mid \overline{\mathcal{I}}_{t-1}, s]$ under the constraints in this setup. This means future inflation and output gap are always expected to be zero, both before and after the central bank receives its signal. So policy surprises z_t must be uncorrelated with the realized output gap and inflation at all horizons. This result generalizes broadly in linear settings. Notice the logic is not altered even if there exists a "central bank information effect" because the commitment to the policy rule already tells private agents all they need to know about the future. The result can also apply under more complicated policy functions and models (e.g., if we include cost-push shocks and break the "divine coincidence"). The sole requirement is the information structure is such that expectations about aggregates do not change once the central bank receives its signal.

Proposition 1 is an "irrelevance result" because it shows a broad class of environments where policy surprises have no propagation. But it may be useful to consider a different kind of irrelevance result: one that illustrates the propagation of belief surprises can go directly against the underlying monetary transmission mechanism. Again, we consider the aforementioned New Keynesian model but explicitly add cost push shocks

$$\widetilde{y}_t = \mathbb{E}_t[\widetilde{y}_{t+1}] - i_t + r_t^* + \mathbb{E}_t[\pi_{t+1}] \tag{2}$$

[21] This shows the sufficient condition of Proposition 1 can be weakened slightly.

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \widetilde{y}_t + u_t. \tag{3}$$

We will consider two ways to get a "puzzling" comovement between policy surprises and inflation. The first will involve a hypothetical where the central bank follows a known, non-optimal forecast rule. As discussed further in Online Appendix D, this first result has a precedent in Caballero and Simsek (2022) where the central bank follows a non-optimal forecast rule because of its dogmatic beliefs, making the correlation between policy surprises and aggregates essentially unrestricted. Suppose $\text{Var}(u_t) = 0$, r^*_{t+1} follows a (known) linear law of motion with Gaussian innovations, and the central bank receives a signal about r^*_{t+1} between t and $t+1$. Optimal policy in this model is for the central bank to commit to a rule of $i_{t+1} = b_t(s)$, where b_t is the objective forecast for r^*_{t+1} prescribed by the Kalman filter after observing signal s.

Proposition 2: Suppose the policy rule is $i_{t+1} = \alpha b_t(s)$ with $\alpha < 1$. In the equilibrium in (2), (3) generated by $\text{Var}(u_t) = 0$, inflation and the output gap are directly proportional to $i_t - r^*_t$ and $\text{Cov}(\pi_t, z_t), \text{Cov}(\widetilde{y}_t, z_t) > 0$.

Proposition 2 inverts the logic of Proposition 1: when the private sector forecast of interest rates is too high, it on average will imply that interest rates were realized as lower than the optimal level. This is because the central bank is following an "attenuated" forecast rule, so the powerful corollaries of rational expectations no longer hold. Similar results hold if the central bank instead follows an inertial forecast rule. There is also a connection to optimal policy in an unconstrained environment. If a central bank follows the real rate rule in Nakamura and Steinsson (2018a) $i_t - \mathbb{E}_t[\pi_{t+1}] = \psi r^*_t + \phi \pi_t$ with $\psi < 1$, then rates adjust less than one-for-one with the natural rate. This means an impulse to r^*_t will generate positive responses of inflation and the output gap in textbook New Keynesian models.

Proposition 2 considers an environment with suboptimal policy, but a similar irrelevance result can still hold in this environment under optimal policy by changing what the central bank receives a signal about. For simplicity, assume $r^*_t \ \forall t \geq 0$ is known at $t = 0$. Instead, the central bank faces the same timing constraint but receives a between-period signal s about u_{t+1}. Appendix A.6 shows optimal policy leads to the following equilibrium expressions, largely determined by the central bank's forecast $m_t = \mathbb{E}[u_t \mid \overline{\mathcal{I}}_{t-1}, s]$

$$i_t = r^*_t + c m_t \qquad \widetilde{y}_t = -c m_t \qquad \pi_t = u_t - \kappa c m_t, \tag{4}$$

for a specific $c \in \mathbb{R}$ pinning down the optimal response parameter. If u_t is i.i.d normal with zero mean, the policy surprise is simply $z_t = i_t - r^*_t = c m_t$. This z_t shares a much different relationship with inflation than a hypothetical structural policy shock in this model would.

Proposition 3: In the equilibrium described by (4), $\text{Cov}(\pi_t, z_t) > 0$.

4 A Model of Monetary Policy and Policy Surprises

We will now build a model motivated by the discussed facts about policy surprises

- Policy surprises not explained by baseline forecasts
- Conditional skewness of policy surprises (i.e., comovement of policy action and surprises)
- Finite-sample correlations with lagged news and "hump-shaped" correlations with leads of macroeconomic aggregates
- Lack of evidence for learning from the central bank about the economy or reaction function drift during policy announcements

and interest rates

- Autocorrelated errors of future policy action
- Secular component in the term structure of interest rates
- Inaction and gradualism in the setting of monetary policy

Section 4.1 establishes the first core component of the model: the natural rate of interest is governed by two unobserved components, and agents in the model do not know the parameters of the transition equation of the latent states. Section 4.2 discusses the information structure of the model and restates the constraints the central bank faces. Section 4.3 overviews a broad class of preferences where a decision maker behaves *as if* they have a different probability distribution. Section 4.4 interprets these preferences in the context of monetary policy as central banks performing "risk management" and justifies why these preferences would (or possibly should) arise. Section 4.5 presents results from the model and discusses the findings.

4.1 A Natural Rate with Unobserved Components

We begin with the aforementioned simplest log-linear New Keynesian equilibrium conditions under log utility

$$\widetilde{y}_t = \mathbb{E}_t[\widetilde{y}_{t+1}] - i_t + r_t^* + \mathbb{E}_t[\pi_{t+1}] \tag{2}$$

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \widetilde{y}_t + u_t \tag{3}$$

The discussion in Section 2 motivates a limited information environment of some form, but these environments can substantially alter the equilibrium (and make solving for it intractable). To introduce a tractable structure (i.e., one that will not break the standard (2),(3) system), I assume all agents observe r_t^* contemporaneously and know the structure of the law of motion, but do not observe the latent states or know the parameters governing their evolution. The system is

$$r_t^* = a_t + d_t \tag{5}$$

$$a_t = \phi_1 a_{t-1} + \phi_2 a_{t-2} + \varepsilon_t^a \tag{6}$$

$$d_t = d_{t-1} - \varepsilon_t^z \tag{7}$$

$$\varepsilon_t^j \overset{\text{i.i.d}}{\sim} \mathcal{N}(0, \sigma_j^2) \quad j = \{a, d\}. \tag{8}$$

Microfoundations for this environment are provided in Appendix A.6 as well as in Caballero and Simsek (2022) and Rungcharoenkitkul and Winkler (2026). Agents perform optimal learning given their prior over the parameter vector $\boldsymbol{\theta} = \{\phi_1, \phi_2, \sigma_a, \sigma_d\}$ and their information set, which includes the history of natural rate realizations but not its components. Treating beliefs as traditional state variables in this case would lead to an intractable state evolution process, so an anticipated learning assumption is imposed (Kreps, 1998). This means agents update their beliefs about $\boldsymbol{\theta}$ at the start of each period but believe they will run their Kalman filter at each point in the future using their current posterior distribution. Each vector $\boldsymbol{\theta}$ households and firms place positive probability on can be thought of as indexing the class of possible models they might live in. I assume private agents aggregate over models using their probability distribution for $\boldsymbol{\theta}$, the analog of expected utility maximization for model uncertainty. To formally nest all structure, consider a Bellman equation formulated under smooth ambiguity aversion preferences (Klibanoff et al., 2009)

$$V_t(\boldsymbol{x}_t) = \max_{\mathbf{y}_t} \left[\int_{\boldsymbol{\theta}} f_{\boldsymbol{\theta},t} \left(u(\mathbf{y}_t) + \beta \int_{\boldsymbol{\varepsilon}} f_{\boldsymbol{\varepsilon},t} \cdot V_t(\boldsymbol{x}_{t+1}(\mathbf{y}_t, \boldsymbol{\varepsilon}; \boldsymbol{\theta})) \, \mathrm{d}\boldsymbol{\varepsilon} \right)^{1-\eta} \mathrm{d}\boldsymbol{\theta} \right]^{\frac{1}{1-\eta}}, \tag{9}$$

where $\boldsymbol{x}_t$ is the state vector, $\mathbf{y}_t$ is the control vector, $f_{\boldsymbol{\theta},t}$ and $f_{\boldsymbol{\varepsilon},t}$ are the time t probability distributions over $\boldsymbol{\theta}$ and the vector of structural shocks $\boldsymbol{\varepsilon}$, and η is a preference parameter. The t subscripts on the value function refer to the anticipated utility assumption – agents believe they will use today's probability distributions ad infinitum, but tomorrow they will update their beliefs and effectively solve an entirely new dynamic program. Also, while today's state vector is fixed, the evolution equation for $\boldsymbol{x}_{t+1}$ depends on what model nature has actually selected. Finally, agents do not have any inherent distaste for model uncertainty, meaning their preferences correspond to $\eta = 0$.

As shown in Figure 5, both the structure of the natural rate and the assumptions about parameter uncertainty will reproduce dynamics observed both in the term structure of interest rates (Atkeson and Kehoe, 2008) and interest rate forecasts (Cieslak, 2018; Farmer et al., 2024).

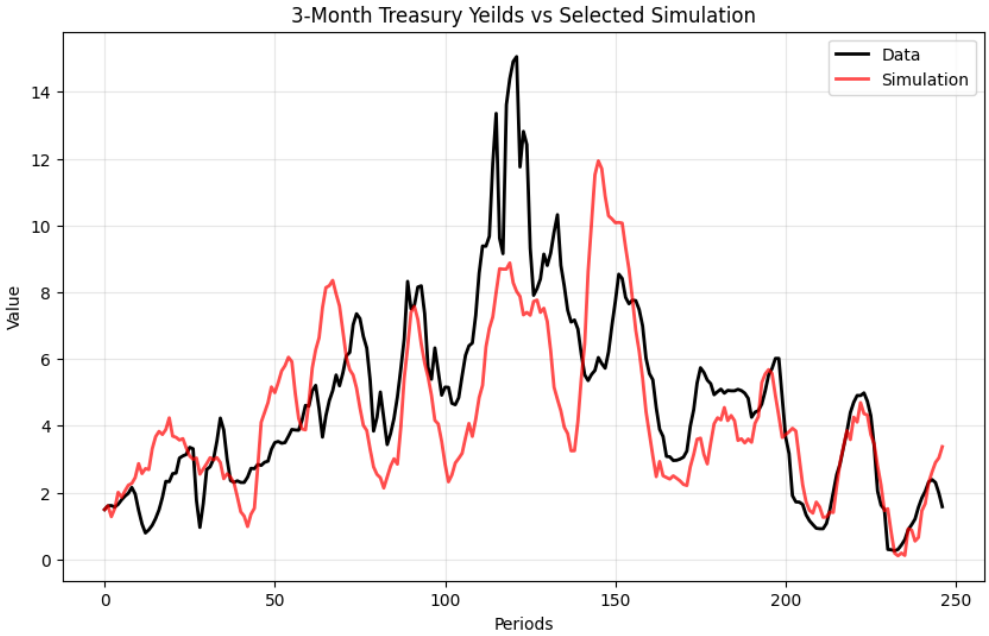


Figure 5

4.2 Information Structure and the Central Bank's Constraints

Under the assumption that r_t^* is common knowledge, a fully informed and unconstrained central bank for essentially any reasonable set of preferences would choose $i_t = r_t^*$ and achieve the "divine coincidence" (Benigno and Woodford, 2004). In kind, we impose an implementation lag assumption detailed in Section 3 that the central bank must announce its policy function for i_{t+1} at t and cannot renege when period $t+1$ arrives. Given a primary focus on policy surprises, we will introduce an information structure where information is received by agents between periods. The central bank can condition its $t+1$ policy function on information it will receive, but because what will be received is not known ahead of time, there still will be policy surprises. To make this precise, we first need to introduce the exact information structure.

The baseline assumption in the model is the central bank does not possess a unique signal structure for its private information. We can think about indexing all agents in the economy using $i \in [0,1]$ and without loss of generality letting the central bank correspond to agent $i = 1$. All agents receive a base signal between t and $t+1$ of the cyclical component a_t, and this signal potentially is (conditionally) biased as in Caballero and Simsek (2022). The idiosyncratic signal noise and conditional bias are drawn from common, known distributions

$$s_{i,t} = a_t + \mu_{i,t} + \varepsilon_{i,t}^s \tag{10}$$

$$\mu_{i,t} = \overline{\mu}_t + \varepsilon_{i,t}^{\mu} \tag{11}$$

$$\overline{\mu}_t = \varepsilon_t^{\overline{\mu}} \tag{12}$$

$$\varepsilon^{s}_{i,t}, \varepsilon^{\mu}_{i,t}, \varepsilon^{\overline{\mu}}_{t} \overset{\text{i.i.d}}{\sim} \mathcal{N}\left(0, \text{diag}(\sigma_s^2, \sigma_\mu^2, \sigma_{\overline{\mu}}^2)\right). \tag{13}$$

Connecting with discussion in Section 2, at $t+1$ a noisy aggregation of all public signals is revealed

$$p_{t+1} \sim \mathcal{N}(a_t + \overline{\mu}_t, \sigma_p^2) \tag{14}$$

where σ_p is known to all agents. The timing convention can be thought of as similar to data releases or surveys about the current period being released with a lag.

We consider a special case where bias is common ($\sigma_\mu = 0$) but unknown ($\sigma_{\overline{\mu}} > 0$) and there is no measurement noise ($\sigma_p = 0$). These simplifications are merely to reduce the dimensionality of the problem and are not material to any conclusions. Under these conditions, the time t state vector $\boldsymbol{x}_t$ contains the beliefs about the latent states households and firms enter the period with $\hat{b}_{t-1}$, the natural and nominal interest rate histories through time t, and the history of the aggregated signal through t. The objects in $\boldsymbol{x}_t$ dated time t are revealed to all agents at the start of the period, after which agents update their beliefs about the parameter vector $\boldsymbol{\theta}$ as well as the histories of latent states. After this process, beliefs are common to all agents under these assumptions.

4.3 Generalized "*as if*" Preferences

An important piece of the problem yet to be specified is the central bank's preferences. It will be instructive to build a formal foundation for a general class of preferences common in environments where agents do not know the underlying data generating process. These preferences can be broadly described as agents placing weights on specific scenarios (adverse models and/or outcomes) beyond what is implied by their objective probability distribution. This can be thought of as some form of ambiguity aversion or alternatively an extension of α-minmax preferences (Arrow and Hurwicz, 1972). Sims (2001) argues such preferences can be thought of as a special case of Bayesian decision making with a different probability distribution, or as "risk-sensitive" loss that acts as if nature has twisted the probability distribution in some way (Jacobson, 1973; Del Negro and Schorfheide, 2009). This way of thinking can be expanded further by allowing each scenario to potentially be evaluated under its own conditional probability distribution. A general form follows

$$\max_{\{\boldsymbol{y}_t\}_{t\geq 0}} -\sum_{t=0}^{\infty} \beta^t \left(\overline{\theta} \mathbb{E}_0 \left[\ell(\boldsymbol{y}_t; \boldsymbol{x}_t) \right] + \sum_{s \in \mathcal{S}_t} \int \theta_s(\boldsymbol{x}_t) \ell(\boldsymbol{y}_t; \boldsymbol{x}_t) \, \mathrm{d}\boldsymbol{x}_t \right).$$

For example, maxmin expected utility preferences correspond to $\overline{\theta} = 0$ and a single scenario s with $\theta_s = 1$ if $\boldsymbol{x}_t = \text{argmin} -\ell(\boldsymbol{y}_t; \boldsymbol{x}_t)$ and 0 otherwise. However, this structure is general enough to nest other possibilities not typically considered in this class of preferences. One component represents standard expected utility maximization, and the other is added weight which is weighted both by some intrinsic distaste θ_s and also how likely that scenario

is under a possibly scenario-specific probability measure. The latter feature carries a distinction between forecasts made under the current objective posterior and forecasts made under a conditional realization of some event. To make this distinction clear, let the sole state variable be a three-state Markov process $x_t \in \{G, B, D\}$ with an unlikely disaster state and the set of scenarios be singleton with $x_t^s = D$. Looking at tomorrow's individual loss and recognizing $\theta_s(x_t)$ as effectively a probability kernel, this looks like

$$\begin{aligned} L_1 &\equiv \mathbb{E}_0\left[\ell(\boldsymbol{y_1}; x_1)\right] + \theta_s \mathbb{E}_0^s\left[\ell(\boldsymbol{y_1}; x_1^s)\right] \\ &= \mathbb{P}(x_1 = G \mid x_0)\ell(\boldsymbol{y_1}; G) + \mathbb{P}(x_1 = B \mid x_0)\ell(\boldsymbol{y_1}; B) + \mathbb{P}(x_1 = D \mid x_0)\ell(\boldsymbol{y_1}; D) + \theta_s \mathbb{E}_0^s\left[\ell(\boldsymbol{y_1}; x_1^s)\right] \\ &\equiv C_1 + \mathbb{P}(x_1 = D \mid x_0)\ell(\boldsymbol{y_1}; D) + \theta_s \mathbb{E}_0^s\left[\ell(\boldsymbol{y_1}; x_1^s)\right] \end{aligned}$$

$\mathbb{E}_0^s$ allows for two possibilities in forming the policymaker's "as if" probability distribution for period 1 loss: the "scenario" component of loss may condition on a history $\{x_0, D\}$, i.e.,

$$\mathbb{E}_0^s\left[\ell(\boldsymbol{y_1}; x_1^s)\right] = \mathbb{E}_0\left[\ell(\boldsymbol{y_1}; x_1) \mid x_1 = D\right] = \ell(\boldsymbol{y_1}; D),$$

or $\ell(\boldsymbol{y_1}; x_1^s)$ can be treated as a random variable realized with respect to some probability kernel, e.g.,

$$\mathbb{E}_0^s\left[\ell(\boldsymbol{y_1}; x_1^s)\right] = \mathbb{P}(x_1 = D \mid x_0)\ell(\boldsymbol{y_1}; D),$$

leading to two possibilities

$$L_1 = C_1 + \begin{cases} \left[\mathbb{P}(x_1 = D \mid x_0) + \theta_s\right]\ell(\boldsymbol{y_1}; D) \\ (1 + \theta_s)\,\mathbb{P}(x_1 = D \mid x_0)\,\ell(\boldsymbol{y_1}; D) \end{cases}.$$

Returning to the Sims (2001) analogy, these preferences can either add fixed mass to the probability distribution or stretch it at a set of points by a fixed proportion. The importance of this distinction in the context of the model can be seen by applying these preferences to monetary policymaking.

4.4 Risk Management: Why Do Central Banks Have these Preferences?

Section 2 finds evidence against the usual presentation of the "information effect" (e.g., Nakamura and Steinsson, 2018a) and "response to news" (Bauer and Swanson, 2023b) ideas, but each maps into a more fundamental decomposition of the policymaking process. Any given decision is a product of the perceived optimal mapping and the beliefs informing how that policy rule is carried out. For example, if the Fed is a discretionary policymaker mimicking a committed policymaker (King and Lu, 2025) or a dove masquerading as a hawk (Caravello et al., 2026), this constitutes private knowledge of its reaction function *form* in the spirit of "response to news". Whereas if the Fed has unique beliefs about the economy, this in effect is private knowledge of reaction function *inputs*, a particular "information effect". Given that the Fed forecasts of future policy have comparable accuracy and make

similar, autocorrelated mistakes (outside of the zero lower bound episodes), this suggests the second category is more relevant to the central bank's decision problem and their relationship with the public. The central bank must navigate through an incredibly uncertain environment and its subjective assessments, thorough but ultimately imprecise, are subject to unpredictable shifts as conditions evolve (Greenspan, 2003).

To better articulate the policymaking environment considered here, especially with respect to unpredictable changes in outlook, it may be useful to consider an illustrative anecdote. Suppose an all-powerful oracle says to a central banker that they can reveal any piece of data (including financial market and sentiment indicators) that will be available at the next month's meeting, and furthermore they can even "freeze time" so they can take as long as they want to process the information and ask for as much data as they want. In exchange for this information (and time), they would have to commit to their policy choice privately to the oracle before time restarts. Would the central banker be indifferent about this arrangement, or perhaps even prefer it given the quantitative information processing advantage? The answer is almost surely no; the world's data generating process is too complicated to not exercise subjective judgment in forecasting and decision-making. As discussed in Section 4.3, these considerations are inevitable when a standard forecast is insufficient for making an optimal decision.

Virtually all central bankers invoke the concept of "risk management" when explaining their rationale (Kilian and Manganelli, 2008; Cieslak et al., 2025a). This framing is critical to understanding policy surprises. When the Fed is deciding between cutting 25bp or holding firm, what are they deciding between? Their decision boils down to what they deem the most pressing concerns to the economy, and the announcement reveals their thinking. Though there are confounding factors like acts of discretion (Debortoli and Lakdawala, 2016), changes in tastes (Cieslak et al., 2024), and intra-committee politics (Howes et al., 2026), evidence from deliberations suggests if the public could perfectly anticipate how the central bank was going to balance risk, they would be able to know the policy decision with certainty. The confounding factors listed are ways where policy is intentional but nonetheless still has a stochastic component, and this idea is explored in more detail in Bousquet (2026a). Cieslak et al. (2025a) provides an exhaustive overview of the descriptive and narrative evidence for the Federal Reserve following a risk management approach. Because this philosophy is so embedded in the world of central banking, it is not hard to find quotes from each of the past several Fed chairs emphasizing its importance, and a selection is found below.

> "We are not—at least I am not—proposing that we go to a purely mechanical reserve targeting approach. [...] There are all these elements of judgment that enter into the process."
>
> — Paul Volcker, October 1979

> "In implementing a risk-management approach to policy, we must confront the fact that only a limited number of risks can be quantified with any confidence... Other risks are essentially unquantifiable...As

a result, risk management often involves significant judgment on the part of policymakers."

— Alan Greenspan, August 2003

"In such circumstances, the most likely outcome—referred to as the modal forecast—for the economy may be fairly benign, but there may be a significant risk of more severe adverse outcomes. In such circumstances, the central bank may prefer to take out insurance."

— Ben Bernanke, January 2008

"We necessarily operate in an environment in which there's a great deal of uncertainty. In such an environment, it makes sense to use a risk-management approach to identify and avoid the big mistakes. That's one reason I favor a cautious approach."

— Janet Yellen, April 2016

"As is often the case, we are navigating by the stars under cloudy skies. In such circumstances, risk-management considerations are critical."

— Jerome Powell, August 2023

Now we have sufficient background to apply the formalization in Section 4.3. A classical presentation of environments where agents do not know the true data generating process (e.g., Hansen and Sargent, 2008) specifies preferences over a neighborhood of models that are considered "reasonable", or sufficiently hard to reject. Instead, it's possible a distrust of one's objective distribution leads to a probability floor: a decision maker putting some fixed weight on a scenario even if the perceived probability is zero. Paraphrasing Evans (2019), one way central bankers frame risk management is "what would be the costs if our forecasts are wrong"? This is a humility condition about the possibility of having a fundamentally off-centered belief about where the economy is heading and wanting to hedge against that scenario. Section 4.5 will show how this framing can help explain many behaviors of central banks and generate particular features of the data.

4.5 Equilibrium, Results, and Discussion

We begin with a baseline where the central bank has a standard loss function, cost-push shocks are eliminated ($\text{Var}(u_t) = 0$), and the preferences of all agents follow $\eta = 0$ from (9).

Definition. The discretionary equilibrium is a sequence of aggregates $\{\widetilde{y}_t, \pi_t, i_{t+1}\}_{t=0}^{\infty}$ and beliefs $\{f_{\boldsymbol{\theta},t}, \hat{b}_t\}_{t=0}^{\infty}$ such that:

1. Initial conditions are given as $(r_0^*, i_0^*, f_{\boldsymbol{\theta},-1}, \hat{b}_{-1}, p_0)$ and state evolution equations follow (5)–(8) and (10)–(14)

2. $(\widetilde{y}_t, \pi_t, i_t)$ and public signal p_t defined in (14) are realized simultaneously at the start of period t

3. Upon realization of aggregates, agents share a common filtration

$$\mathcal{F}_t = \{\widetilde{y}_k, \pi_k, i_k, p_k, f_{\boldsymbol{\theta},k-1}, \hat{b}_{k-1}\}_{k=0}^{t} \tag{15}$$

and form updated beliefs $f_{\boldsymbol{\theta},t}$ about $\boldsymbol{\theta}$ using Bayes rule subject to $\mathcal{F}_t$. All agents believe $f_{\boldsymbol{\theta},t+h} = f_{\boldsymbol{\theta},t}\ \forall h \geq 0$.

4. The common expectations operator $\overline{\mathbb{E}}_t$ generates $\hat{b}_t = (\overline{\mathbb{E}}_t[a_t], \overline{\mathbb{E}}_t[d_t])$ and satisfies

$$\overline{\mathbb{E}}_t[\cdot] = \int_{\boldsymbol{\theta}} \mathbb{E}[\cdot \mid \boldsymbol{\theta}, \mathcal{F}_t] f_{\boldsymbol{\theta},t}\, \mathrm{d}\boldsymbol{\theta} \tag{16}$$

5. $(\widetilde{y}_t, \pi_t)$ is a measurable function of $\boldsymbol{x}_t \subset \mathcal{F}_t$ and satisfies (2)–(3) with respect to expectation operator $\overline{\mathbb{E}}_t[\cdot]$

6. $i_{t+1} : \boldsymbol{x}_t \times \mathbb{R} \to \mathbb{R}$ is a measurable function of $\boldsymbol{x}_t \subset \mathcal{F}_t$ and its signal (10) and solves dynamic program (9) with flow utility $\widetilde{y}_t^2 + \lambda \pi_t^2$ taking as given the action of future policymakers.

7. $(\widetilde{y}_t, \pi_t)$ are the unique solution to min $\widetilde{y}_t^2 + \lambda \pi_t^2$ across the set of all possible solutions to (2)–(3)

Despite the additional layers, the equilibrium is standard.

Proposition 4. The equilibrium output gap and inflation are proportional to the interest rate gap $(\widetilde{y}_t, \pi_t \propto i_t - r_t^*)$ and optimal monetary policy features policy surprises (i_{t+1} is a function of its between period signal).

Proof: See Appendix A.6.

A monetary policy surprise in this model $z_t = i_t - \overline{\mathbb{E}}_{t-1}[i_t]$ is defined with respect to the common beliefs of all agents prior to the between-period individual signals. This model falls under the class of environments in Proposition 1, which means there is no population response to z_t. As discussed, this does not rule out observing non-zero responses in small samples. Suppose an econometrician in this environment has a time series of aggregates $\{\widetilde{y}_t, \pi_t, i_t\}_{t=1}^{T}$ and high-frequency policy surprises $\{z_t\}_{t=1}^{T}$. They naturally do not know the economy's structure, where a hypothetical one-time random movement in the interest rate has no impact beyond the initial period. As detailed in Appendix A.6 and visualized in Figure 5, I select true parameters for the data generating process that produce dynamics similar to the 3-month treasury yield data for the US from 1960-2008. I find that in roughly half of simulations, an econometrician will estimate responses that peak from 2 to 6 quarters out, qualitatively matching commonly estimated "hump-shaped" responses Ramey (2016) and "long and variable lags" (Friedman, 1960). The initial period, the only point where the true impulse response is nonzero, is almost always underestimated. Given Proposition 1, any inference about the true impulse response will be inherently mistaken.

Why does this happen? What these responses are uncovering is a consequence of the same conditional skewness of surprises observed in Figure 4. Agents rationally underestimate the extent policy will be adjusted in the future,

making a "hump-shaped" correlation appear ex post. Because of the persistent forecast errors in the natural rate, a positive surprise on average occurs in a situation where policy will tighten progressively further in the future more than was expected (likewise for negative surprises). These correlations will shrink to zero as the sample grows, but the persistent forecast errors can initially have a dominating effect.[22] Surprises are genuinely unpredictable, so adding controls and lag-augmentation (Montiel Olea and Plagborg-Møller, 2021) does not change the qualitative results. However, it does make some superficial quantitative difference because surprises do exhibit ex post correlations with lagged aggregates, leading to some changes in point estimates but preserving the general shape. More importantly, an econometrician in this environment would find the same "predictability" of empirical surprises, even though this is purely a finite sample artifact of slow learning.

There is a major shortfall of the model with standard preferences relative to the data: it predicts the central bank's point forecasts account for all policy surprise variation, as the baseline "information effect" and "response to news" formulations. To address this discrepancy, we introduce different preferences for the central bank which cause other parts of their belief distribution to matter even though the underlying model is linear and (conditionally) Gaussian. This approach will also deliver "hump-shaped" responses nearly universally in small samples, rather than just frequently. These patterns will be amplified because in the special case considered, the central bank will let its target drift away and take small steps towards it (gradualism), as well as have periods of inaction. Surprises will now be conditionally skewed even for longer sample sizes by reflecting a propensity to "catch up", accentuating the finite sample effect.

Returning to the idea of the central bank having "risk management" preferences, these emerge naturally when there is significant uncertainty that extends beyond baseline risk and instead comes from a lack of knowledge about the overall data generating process. A decision maker in this environment may be concerned that listening to their "objective" probability distribution will lead them down a costly path. At a high level, we can capture this idea of the central bank being led astray by the available data and forecasts with a notion of a biased signal. This can be accomplished by extending the preferences of Section 4.3 with a scenario directly on the between-period signal the central bank receives. For simplicity, the optimal policy problem is recast as a problem of forecasting the natural rate since $i_t = r_t^*$ would be the unconstrained outcome. The central bank's decision is to pick a function $g_t(s)$, which takes as input the information set available at time t and outputs a rule for policy to be set as a function of the signal they will receive between periods. Their problem is

$$\max_{g_t(s)} \left\{ -\overline{\theta}\mathbb{E}_t\left[(g_t(s) - r_{t+1}^*)^2\right] + \theta \min_{b \in \mathcal{F}_t} -\mathbb{E}_t\left[(g_t(s+b) - r_{t+1})^2\right] \right\}$$

where $b \in \mathcal{F}_t$ represents a feasible set of signal bias. Here, we can extend this idea further to generate several features of the monetary policy status quo. Rewriting $g_t(s) = \alpha_t(s)(b_t(s) - i_t) - i_t$, where $b_t(s)$ is the objective

[22] There is a similar intuition in Caballero and Simsek (2022) discussed further in Online Appendix D.

Kalman update after seeing a signal s given the information set at t, the problem is isomorphic to choosing $\alpha_t(s)$ (i.e., since $b_t(s)$ and i_t are deterministic functions of the time t information set, any desired policy can be achieved by picking $\alpha_t(s)$). A constraint of $\alpha_t(s) \geq 0$ can be interpreted as being sure enough to quell any "paranoia" (Sargent, 2015) – policy will not adjust until the (preference-weighted) worst-case bias that would be entertained wouldn't push us to move in the *opposite* direction as what our belief is. This effectively acts as a fixed cost of sufficient confidence the central bank must pay before taking action, which also produces periods of inaction, gradualism, and through specifications of $\mathcal{F}_t$ can also lead to the counter-cyclical dispersion of policy surprises shown in Figure 4. This framework is extremely flexible, and can accommodate a continuous or multi-dimensional scenario set to motivate more comprehensive risk management tasks. These sparse setups are nevertheless able to demonstrate features of the data describing the essence of central banking, suggesting rich implications from branching out from this starting point.

5 Conclusion

This paper argues that the unpredictable component of monetary policy can be attributed to "uninformative" belief surprises – unpredictable shifts in the central bank's subjective risk assessments that reveal little about the economy or the reaction function but are nonetheless consequential for understanding observed interest rate dynamics. The empirical evidence points against the two leading explanations: perceptions about the economy and the reaction function do not seem to vary meaningfully following a typical policy announcement. Instead, a model is presented where the focus is the data generating process agents are grappling with. Pervasive uncertainty can create many features of the data, including finite-sample relationships that will disappear eventually but are overwhelmed initially by persistence induced from slow learning rates. Surprises in this model reveal the subjective balance of risks of the central bank, in line with how central bankers describe their own job and how financial media describes what tips the balance of a particular decision. The unobserved component structure can on its own replicate key empirical features: conditionally skewed surprises, persistent forecast errors, and hump-shaped impulse responses that intrinsically emerge from optimal policy under uncertainty rather than from transmission lags. Risk management preferences—where policymakers weight adverse scenarios beyond objective probabilities—rationalize gradualism and inaction while offering a coherent interpretation of what announcements actually communicate. Accounting for these sources of misspecification leads to dramatically different recommendations for the macroeconomic models we should write down and how empirical policy analysis should be conducted.

References

Acosta, Miguel. 2026. "The Perceived Causes of Monetary Policy Surprises." Working Paper.

Acosta, Miguel, Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino. 2025. "Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database." Working Paper 2025-30, Federal Reserve Bank of San Francisco.

Adams, Jonathan J., and Philip Barrett. 2025. "What Are Empirical Monetary Policy Shocks? Estimating the Term Structure of Policy News." June, Working Paper.

Ai, Hengjie, and Xinxin Hu. 2025. "A Model of Fed Information Effect." Working paper.

Amodeo, Francesco. 2026. "Mind the Gap: Disagreement and Credible Monetary Policy." Working Paper.

Angrist, Joshua D., Òscar Jordà, and Guido M. Kuersteiner. 2018. "Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited." *Journal of Business & Economic Statistics*, 36(3): 371–387.

Arrow, Kenneth J., and Leonid Hurwicz. 1972. "An Optimality Criterion for Decision Making Under Ignorance." In *Uncertainty and Expectations in Economics*. 1–11.

Aruoba, S. Boragan, and Thomas Drechsel. 2026. "Identifying Monetary Policy Shocks: A Natural Language Approach." *American Economic Journal: Macroeconomics*, forthcoming.

Atkeson, Andrew, and Patrick J Kehoe. 2008. "On the need for a new approach to analyzing monetary policy." *NBER Macroeconomics Annual*, 23(1): 389–426.

Balke, Nathan S., and Kenneth M. Emery. 1994. "Understanding the price puzzle." *Economic and Financial Policy Review* 15–26.

Barnichon, Régis, and Geert Mesters. 2023. "A Sufficient Statistics Approach for Macro Policy." *American Economic Review*, 113(11): 2809–2845.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam. 2024. "Perceptions About Monetary Policy." *The Quarterly Journal of Economics*, 139(4): 2227–2278.

Bauer, Michael D., and Eric T. Swanson. 2023a. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual*, 37.

——— 2023b. "An Alternative Explanation for the "Fed Information Effect"." *American Economic Review*, 113(3): 664–700.

Beaudry, Paul, Franck Portier, and Andrew Preston. 2024. "Some Inference Perils of Imposing a Taylor Rule." Working Paper.

Benigno, Pierpaolo, and Michael Woodford. 2004. "Optimal Monetary and Fiscal Policy: A Linear-Quadratic Approach." In *NBER Macroeconomics Annual*. 18.

Bernanke, Ben S. 2004. "Gradualism." Speech, 5.

——— 2013. "Communication and Monetary Policy." Speech.

Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma. 2022. "Belief Distortions and Macroeconomic Fluctuations." *American Economic Review*, 112(7): 2269–2315.

——— 2025. "A Structural Approach to High-Frequency Event Studies: The Fed and Markets as Case History." Working Paper.

Bianchi, Francesco, Leonardo Melosi, and Giovanni Nicolò. 2024. "Is There a Forward Guidance Puzzle?" 7, Working Paper.

Bousquet, Paul. 2026a. "Policy Identification Without Arbitrary Policy." Working Paper.

——— 2026b. "Minimal Assumptions to Uncover Macro Dynamics with "Shocks"." Working Paper.

——— 2026c. "Where Do Price and Output Puzzles Come From?", Working Paper.

Brave, Scott A., R. Andrew Butters, and David Kelley. 2019. "A New "Big Data" Index of U.S. Economic Activity."

Caballero, Ricardo J., and Alp Simsek. 2022. "Monetary Policy with Opinionated Markets." *American Economic Review*, 112(7): 2353–2392.

Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity* 1–80, See Chris Sims' discussion comments.

Caravello, Tomás E., Alex Carrasco, and Pedro Martínez-Bruera. 2026. "Reassessing Central Bank Reputation: Beyond Long-Run Expectations." Working Paper.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James M. Robins. 2018. "Double/Debiased Machine Learning for Treatment and Structural Parameters." *The Econometrics Journal*, 21(1): C1–C68.

Cieslak, Anna. 2018. "Short-Rate Expectations and Unexpected Returns in Treasury Bonds." *Review of Financial Studies*, 31(9): 3265–3306.

Cieslak, Anna, Stephen Hansen, Michael McMahon, and Song Xiao. 2024. "Policymakers' Uncertainty." Working Paper.

Cieslak, Anna, Stephen Hansen, and Hao Pang. 2025a. "Risk Management in Monetary Policy: Review with Asset Pricing Implications." 28th Annual Conference of the Central Bank of Chile.

Cieslak, Anna, Michael McMahon, and Hao Pang. 2025b. "Did I Make Myself Clear? The Fed and the Market under the 2020 Monetary Policy Framework." Working paper.

Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen. 2019. "Stock Returns over the FOMC Cycle." *The Journal of Finance*, 74(5): 2201–2248.

Cochrane, John. 2011. "Determinacy and Identification with Taylor Rules." *Journal of Political Economy*, 119(3): 565–615.

Cochrane, John H. 1998. "What do the VARs mean? Measuring the output effects of monetary policy." *Journal of Monetary Economics*, 41(2): 277–300.

——— 2004. "Comments on "A new measure of monetary shocks: Derivation and implications"." July, NBER EFG meeting.

Davig, Troy, and Eric M. Leeper. 2007. "Generalizing the Taylor Principle." *American Economic Review*, 97(3): 607–635.

Debortoli, Davide, and Aeimit Lakdawala. 2016. "How Credible Is the Federal Reserve? A Structural Estimation of Policy Re-optimizations." *American Economic Journal: Macroeconomics*, 8(3): 42–76.

Del Negro, Marco, Marc Giannoni, and Christina Patterson. 2023. "The Forward Guidance Puzzle." *Journal of Political Economy Macroeconomics*, 1(1): .

Del Negro, Marco, and Frank Schorfheide. 2009. "Monetary Policy Analysis with Potentially Misspecified Models." *American Economic Review*, 99(4): 1415–1450.

Eva, Kenneth, and Fabian Winkler. 2023. "A Comprehensive Empirical Evaluation of Biases in Expectation Formation." Working Paper.

Evans, Charles L. 2019. "On Risk Management in Monetary Policy." Federal Reserve Bank of Chicago.

Farmer, Leland E. 2021. "The discretization filter: A simple way to estimate nonlinear state space models." *Quantitative Economics*, 12(1): 1–45.

Farmer, Leland E., Emi Nakamura, and Jón Steinsson. 2024. "Learning About the Long Run." *Journal of Political Economy*, 132(10): 3334–3377.

Frank, Murray Z., Jing Gao, and Keer Yang. 2025. "Behavioral Machine Learning? Regularization and Forecast Bias." Working Paper.

Friedman, Milton. 1960. *A Program for Monetary Stability*.: Fordham University Press.

Gormsen, Niels Joachim, and Eben Lazarus. 2026. "Interest Rates and Equity Valuations." NBER Working Paper.

Goulet Coulombe, Philippe, and Karin Klieber. 2025. "Opening the Black Box of Local Projections." arXiv:2505.12422.

Greenspan, Alan. 2003. "Remarks at Jackson Hole Symposium."

Hack, Lukas, Klodiana Istrefi, and Matthias Meier. 2024. "Identification of Systematic Monetary Policy." Working Paper.

Hansen, Lars Peter, and Thomas J. Sargent. 2008. *Robustness*. Princeton, NJ: Princeton University Press.

Howes, Cooper, Marc Dordal i Carreras, Olivier Coibion, and Yuriy Gorodnichenko. 2026. "How Monetary Policy Is Made: Lessons from Historical FOMC Discussions." Working Paper 34632, NBER.

Ilut, Cosmin, and Rosen Valchev. 2025. "Learning Optimal Behavior Through Reasoning and Experiences." Working Paper.

Jacobson, David H. 1973. "Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games." *IEEE Transactions on Automatic Control*, AC-18 124–131.

Jarociński, Marek. 2024. "Estimating the Fed's unconventional policy shocks." *Journal of Monetary Economics*, 144, p. 103548.

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

——— 2025. "Disentangling Monetary Policy, Central Bank Information, and Fed Response to News Shocks." April, Working Paper.

Kilian, Lutz, and Simone Manganelli. 2008. "The Central Banker as a Risk Manager: Estimating the Federal Reserve's Preferences under Greenspan." *Journal of Money, Credit and Banking*, 40(6): 1103–1326.

King, Robert G., and Yang K. Lu. 2025. "Evolving Reputation for Commitment: Understanding Inflation and Inflation Expectations." Working Paper.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2009. "Recursive Smooth Ambiguity Preferences." *Journal of Economic Theory*, 144(3): 930–976.

Knox, Benjamin, and Annette Vissing-Jorgensen. 2025. "The Effect of the Federal Reserve on the Stock Market: Magnitudes, Channels, and Shocks." *Annual Review of Financial Economics*, 17.

Kocherlakota, Narayana R. 2019. "Practical Policy Evaluation." *Journal of Monetary Economics*, 102 29–45.

Kolesár, Michal, and Mikkel Plagborg-Møller. 2025. "Dynamic Causal Effects in a Nonlinear World: The Good, the Bad, and the Ugly." *Journal of Business & Economic Statistics*, 43(4): 737–754.

Kreps, David M. 1998. "Anticipated Utility and Dynamic Choice." In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*. eds. by Donald P. Jacobs, Ehud Kalai, Morton I. Kamien, and Nancy L. Schwartz: Cambridge University Press, 242–274.

Kuttner, Kenneth N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics*, 47(3): 523–544.

Leeper, E. M., C. A. Sims, and T. Zha. 1996. "What does monetary policy do?" *Brookings papers on economic activity* 1–78.

Martinez, Andrew B. 2025. "How Do Macroeconomic Expectations React to Extreme Weather Shocks?", Working Paper.

McKay, Alisdair, and Christian K. Wolf. 2023. "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica*, 91(5): 1695–1725.

Milani, Fabio, and John Treadwell. 2012. "The Effects of Monetary Policy "News" and "Surprises"." *Journal of Money, Credit and Banking*, 44(8): 1667–1692.

Miranda-Agrippino, Silvia, and Giovanni Ricco. 2021. "The Transmission of Monetary Policy Shocks." *American Economic Journal: Macroeconomics*, 13(3): .

Miranda-Agrippino, Silvia, and John C. Williams. 2026. "Interest Rate Surprises When the Fed Doesn't Speak." Working Paper.

Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2021. "Local Projection Inference is Simpler and More Robust Than You Think." *Econometrica*, 89(4): 1789–1823.

Morris, Stephen, and Hyun Song Shin. 2002. "The Social Value of Public Information." *American Economic Review*, 92 1521–1534.

Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica*, 29(3): 315–335.

Nakamura, Emi, Venance Riblier, and Jón Steinsson. 2025. "Beyond the Taylor Rule." Working Paper.

Nakamura, Emi, and Jón Steinsson. 2018a. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Nakamura, Emi, and Jón Steinsson. 2018b. "Identification in Macroeconomics." *Journal of Economic Perspectives*, 32(3): 59–86.

Nimark, Kristoffer. 2008. "Monetary policy with signal extraction from the bond market." *Journal of Monetary Economics*, 55(8): 1389–1400.

Occhino, Filippo. 2019. "The Flattening of the Phillips Curve: Policy Implications Depend on the Cause."

Orphanides, Athanasios, and John C. Williams. 2007. "Robust monetary policy with imperfect knowledge." *Journal of Monetary Economics*, 54(5): 1406–1435.

Ramey, Valerie. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*. Chap. 2.

Romer, Christina D., and David H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review*, 90 429–457.

——— 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94(4): 1055–1084.

——— 2023. "Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years." *American Economic Review*, 113(6): 1395–1423.

Rungcharoenkitkul, Phurichai, and Fabian Winkler. 2026. "The Natural Rate of Interest Through a Hall of Mirrors." *Journal of Monetary Economics*, 157.

Sargent, Thomas. 1999. "Comment on 'Policy Rules for Open Economies' by Laurence Ball." In *Monetary Policy Rules*. ed. by John Taylor, 31 of National Bureau of Economic Research Studies in Business Cycles 144–154.

Sargent, Thomas J. 2015. "Sets of Models and Prices of Uncertainty." Keynote address at the Conference for Expectations in Dynamic Macroeconomic Models.

Sastry, Karthik. 2026. "Disagreement About Monetary Policy." *American Economic Journal: Macroeconomics*, forthcoming.

Sims, Christopher A. 1980. "Macroeconomics and Reality." *Econometrica*, 48(1): 1–48.

——— 1982. "Policy Analysis with Econometric Models." *Brookings Papers on Economic Activity*, 13(1): 107–164.

——— 1987. "A Rational Expectations Framework for Short-Run Policy Analysis." In *New Approaches to Monetary Economics*. eds. by William A. Barnett, and Kenneth J. Singleton, Cambridge, UK: Cambridge University Press, 293–308.

——— 1992. "Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy." *European Economic Review*, 36(5): 975–1000.

——— 1998. "Comment on Glenn Rudebusch's "Do Measures of Monetary Policy in a VAR Make Sense?"." *International Economic Review*, 39(4): 933–941.

——— 2001. "Pitfalls of a Minimax Approach to Model Uncertainty." *American Economic Review*, 91(2): 51–54.

——— 2008. "Improving Monetary Policy Models." *Journal of Economic Dynamics and Control*, 32(8): 2460–2475.

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

Stein, Jeremy C., and Adi Sunderam. 2018. "The Fed, the Bond Market, and Gradualism in Monetary Policy." *The Journal of Finance*, 73(3): .

Swanson, Eric T., Renxuan Wang, and Yanbin Wu. 2025. "The Effects of Monetary Policy on Macroeconomic Expectations: High-Frequency Evidence from Traded Event Contracts." Working Paper.

Williams, John C. 2019. "Meeting of the Federal Open Market Committee April 30–May 1, 2019." p. 57.

Wolf, Christian K. 2020. "SVAR (Mis-)Identification and the Real Effects of Monetary Policy Shocks." *American Economic Journal: Macroeconomics*, 12(4): 1–32.

Appendix

A.1 Professional Forecast Revisions and Surprises

A.1.1 Surprises Are Associated With Growth

Figure 6 shows monetary policy surprises and leads of industrial production and real GDP growth (Brave et al., 2019). This is possibly to be expected from Section 2.4, as tightening surprises are associated with tightening cycles, which take place during times when economic growth is stronger (and vice versa for cutting cycles). Table 1 demonstrates that the professional forecaster revision results are sensitive to the inclusion of additional controls. The results in the literature are just plotting the raw correlations, as demonstrated here. Including control variables, as is common in regression analysis with surprises, would likewise also change these results. Results also change when unscheduled meetings are dropped, and the sign even reverses for the MP1 series (Bousquet, 2026c).

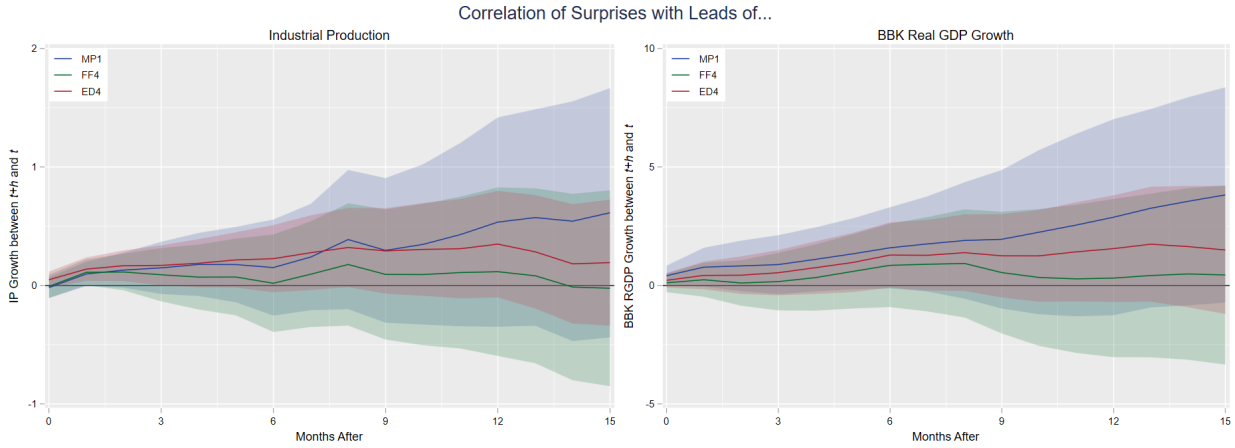


Figure 6

A.1.2 Revision Regressions are Driven by Outliers

Section 2.1 discusses regressions of professional forecaster revisions on surprises. Figure 7 demonstrates a common theme for these results: they are sensitive to outliers. Here, we take an example from Table 2 which looks at the correlation on the subsample where a surprise change in monetary policy occurs in a way directionally consistent with an information effect. This surprise-disagreement alignment occurs when the Fed tightens more than is expected and also is expecting a hotter economy (and vice versa for expansionary policy surprises). Goulet Coulombe and Klieber (2025) shows a simple decomposition to represent regression estimates as a cumulating weighted average over time. In other words, each time period receives weight ω_t and the sum of all weights will equal the point estimate. To appeal to the law of large number, we should expect to see the partial summation converging before the sample ends. Figure 7 looks at one of the more significant correlations in Table 2 (for 2

quarters ahead inflation forecast and a one standard deviation MP1 surprise) and shows it's driven by a handful of dates. This is a common pattern across all results for the professional forecaster revision correlations. It's important to keep in mind this can still be consistent with an information effect, but clearly not one occurring frequently or dominantly.

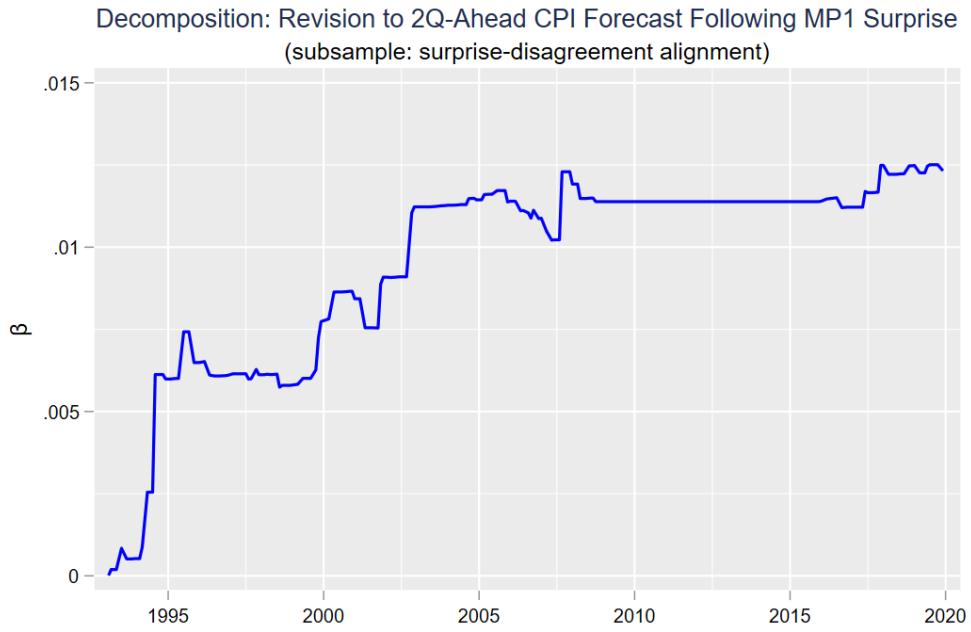


Figure 7

A.1.3 Industrial Production Data and Information Effects

Table 4 tests a hypothesis made in Nakamura and Steinsson (2018a): are stock returns on days where the Fed potentially had access to industrial production data statistically distinct? In line with Jarociński and Karadi (2020), it's possible that in a day where there was a contractionary monetary policy surprise could be coupled with a positive stock market response if the FOMC had access to better than expected industrial production data. Though it's possible there are a few isolated dates where this could have occurred, there is no evidence for a systematic or quantitatively relevant effect. Under the hood of Table 4 is the fact that the S&P 500 still has the conventional negative return even on potential days of early data access.[23] This is consistent with the fact that the observed magnitudes are negligible. Bousquet (2026c) and Figure 7 also point out that these exercises are intrinsically sensitive to outliers. In particular, several of the potential days for an early access information effect (stock-surprise comovement) occurred when the actual surprise component of industrial production pointed in the opposite direction of the surprise.

[23] As discussed in Section 2.1 and Acosta (2026), an information effect presumably could still include the return being *less* negative.

Table 4: Difference in S&P 500 Return Based on Early Access Industrial Production Data

	Scheduled	w/ Unscheduled
MP1	0.023	−0.032
FF4	−0.091*	−0.004
ED4	0.027*	0.027
N	184	196

HC1 Significance levels * $p < 0.10$ ** $p < 0.05$, *** $p < 0.01$ for Wald test. An early access day is defined as when industrial production was released between one and seven days before an FOMC press release (1993-2025). The coefficients are the high-frequency difference in return (pp) for a 1 basis point surprise.

A.2 Forecasting Future Policy

A.2.1 FOMC Policy Forecasts

Figure 1 shows Fed Staff are not superior forecasters of interest rates. But what about members of the Federal Open Market Committee (FOMC)? These comparisons are a bit more difficult because there is a much more limited forecast history (beginning during the first zero lower bound episode), set of variable (yearly averages), and frequency (quarterly). Figure 8 performs the best possible comparison using the Summary of Economic Projections and the Survey of Professional Forecasters, with quadratic interpolation to approximate a path of rates implied by the annual averages. Over the sample period (2015-2025) it's clear the FOMC does not have a superior forecasting performance. In particular, the initial zero lower bound liftoff ended up relatively in the middle of the FOMC and professional forecaster trajectories.

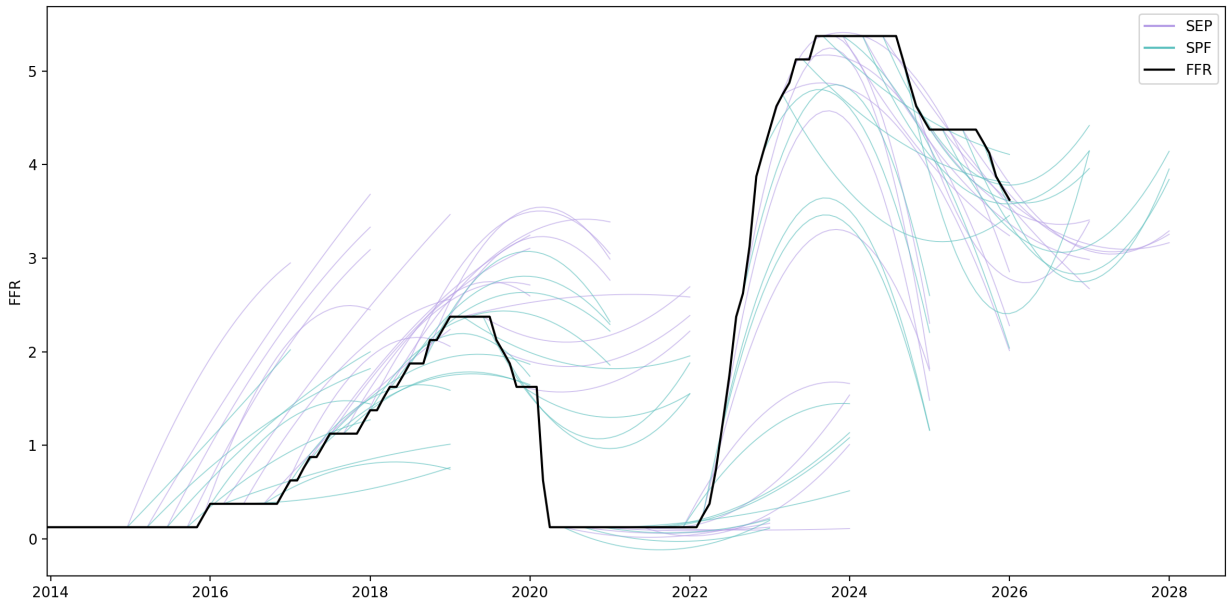


Figure 8: FOMC vs. Professional Forecast of the Fed Funds Rate

A.2.2 Large Moves and Future Reversals

Figure 9 looks at the subsample where interest rate target changes occurred from 1987-2025 and tests whether a large target change predicts a decreased probability of reversing course. Specifically, whether any reversal occurs from t to $t+h$, defined as $R_{t+h} = \mathbb{1}_{\min_{k \le h}\{\Delta i_{t+k} \cdot \Delta i_t\} \ge 0}$ and fitting a linear probability model

$$= \alpha + \beta \, \mathbb{1}_{|\Delta i_t| > .25} + e_{t+h}.$$

The results indicate large changes tend to occur early in a monetary policy cycle when central bankers believe they have ample room to adjust.

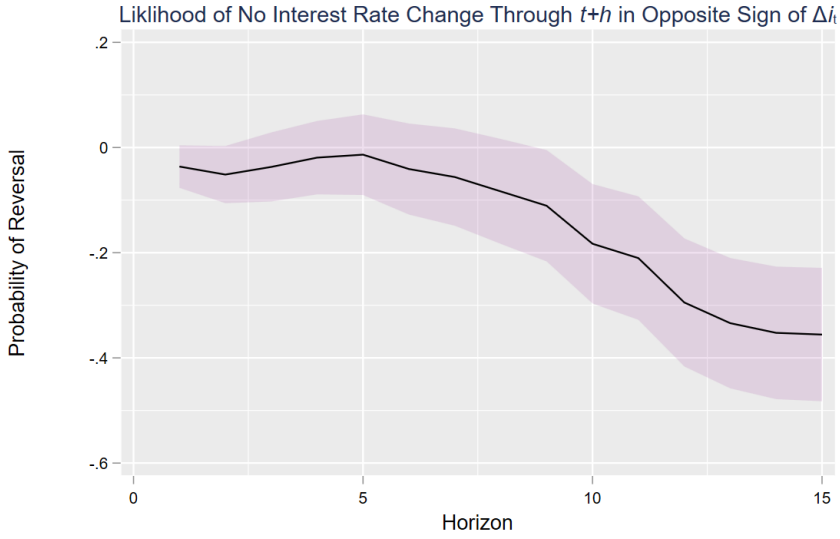


Figure 9

A.3 Policy Forecast Revisions under Existing Frameworks

Finally, an exposition on how interest rate forecasts should shift following a policy surprise under "response to news" (RTN). RTN posits $i_t = \boldsymbol{\phi}_t' \cdot \boldsymbol{X}_t$. In a standard, log-linear DSGE model, we can write

$$\boldsymbol{\phi}_t' \cdot \mathbb{E}_t[\boldsymbol{X}_{t+h}] = \sum \psi_i \rho_i^h \varepsilon_{t,i},$$

where $\{\varepsilon_{t,i}\}_{i=1}^N$ are independent structural shocks, ρ_i is the persistence of a given structural shock, and ψ_i represents sum of contributions of a structural shock to the elements of $\boldsymbol{X}_t$ (weighted by $\boldsymbol{\phi}_t$). Assuming the best guess of the future Taylor Rule is today's (perceived) Taylor Rule, regressing a post-announcement revision to an interest rate futures contract dated $t+h$ and the forecast error of the interest rate target announcement announcement yields

$$\beta(h) = \frac{\sum \omega_i \rho_i^h}{\sum \omega_i} \qquad R^2(h) = \beta(h)^2 \cdot \frac{\sum \omega_i}{\sum \omega_i \rho_i^{2h}}$$

$$R^2(h) \to \frac{\omega_m}{\sum \omega_i},$$

where $\omega_i = \psi_i^2 \sigma_i^2$, σ_i^2 is the variance of structural shock i, and $m = \text{argmax}_i \, \rho_i$ (assumed to be unique). $R^2(h)$ is bounded away from zero by the variance contribution of the most persistent structural shock to interest rates. Evidence presented in this paper and elsewhere (Atkeson and Kehoe, 2008; Farmer, 2021) would indicate that to be a large number. It's worth noting the "information effect" Taylor Rule does not make the same prediction. Analogous expressions can be derived where the weights are instead the variance in revisions about $\boldsymbol{X}_t$ that occur around policy announcements, and while the limit of $R^2(h)$ still does not depend on h, it's scaled by the Kalman gain, presumably pushing it close to 0.

A.4 Misidentification of Structural Shocks and Forecast Error Predictability[24]

A.4.1 Inherency of Ex Post Predictability

When forecast errors are correlated with information in the forecaster's information set, this "predictability" is often described as a rejection of a full information rational expectations benchmark. Eva and Winkler (2023); Bauer and Swanson (2023a); Sastry (2026) note that many documented predictability results are based on correlations that, while sometimes statistically significant ex post, are unstable and with little to no predictive power out of sample. In other words, the estimated relationships do not offer much insight into how to improve forecasts. This suggests the correlations are *ex post* artifacts of an environment where it is difficult to learn. To demonstrate this further, we can extend an analogy made in Bianchi et al. (2022); Frank et al. (2025) – do the forecast errors of machine learning algorithms have similar empirical traits? The answer depends on the application and candidate data generating process (DGP). In general, "predictability" seems to be an inherent trait of high-dimensional estimation in small samples (Frank et al., 2025).

As an example, consider a classic environment in the machine learning literature (Chernozhukov et al., 2018) where the underlying DGP is sparse but unknown and the data set is high dimensional $K >> T$. If the underlying DGP is indeed sparse, a lasso procedure will produce consistent estimation. Suppose we are trying to find the causal effect of a treatment D on an outcome Y, with time series data generated from the following system

$$Y_t = \theta D_t + g(\mathbf{X}_t) + v_t$$

$$D_t = q(\mathbf{X}_t) + u_t.$$

I assume $\dim(\mathbf{X}_t) = K$ is 2 orders of magnitude larger than the length of the time series T, but the underlying DGP is linear and sparse (i.e., $g(\cdot), q(\cdot)$ are linear functions of the same few components of $\mathbf{X}_t$). To estimate θ, lasso

[24] Code for these exercises is available at this paper's GitHub repository, https://github.com/paulbousquet/WhySurprise

estimation is performed on Y_t and D_t, then the residuals of these first stages $(\widetilde{Y}_t, \widetilde{D}_t)$ are regressed on one another. While the convergence of $\hat{\theta} \to \theta$ is remarkably fast in this environment, $\{(\widetilde{Y}_t, \widetilde{D}_t)\}_{t=1}^{T}$ are ex post predictable. Specifically, if we regress these residuals on the "relevant" components of $\mathbf{X}_t$ as well as a few other randomly selected variables, 100% of these sets of residuals show several statistically significant correlations. Just as in standard rational expectations, these correlations would indeed vanish as $T \to \infty$. These results extend naturally to the job of professional forecasters and financial market participants; they must sort through reams of data to produce some conditional mean function. This exercise shows even "good" estimation procedures will generate some seeming idiosyncrasies in finite sample.

A.4.2 Masquerading Shocks

Suppose the economy follows a data generating process consistent with a VAR but an econometrician only includes a subset of the variables in estimation. Wolf (2020) shows that the recovered structural shocks will necessarily be functions of several *true* structural shocks, i.e., the recovered structuring shocks are "masquerading" as true primitive objects. This thinking can be extended to the case where there are actually no structural shocks present. Figure 10 shows estimated impulse responses on a small VAR system using data generated by the Smets and Wouters (2007) model at their estimated posterior modes, with the exception that the monetary policy shock variance is set to zero. The results show that the recovered figures will mechanically produce reasonable-looking responses, with the dotted black line represented the actual impulse response Smets and Wouters (2007) estimate. Standard diagnostic tests to check for identification issues do not indicate any issues.

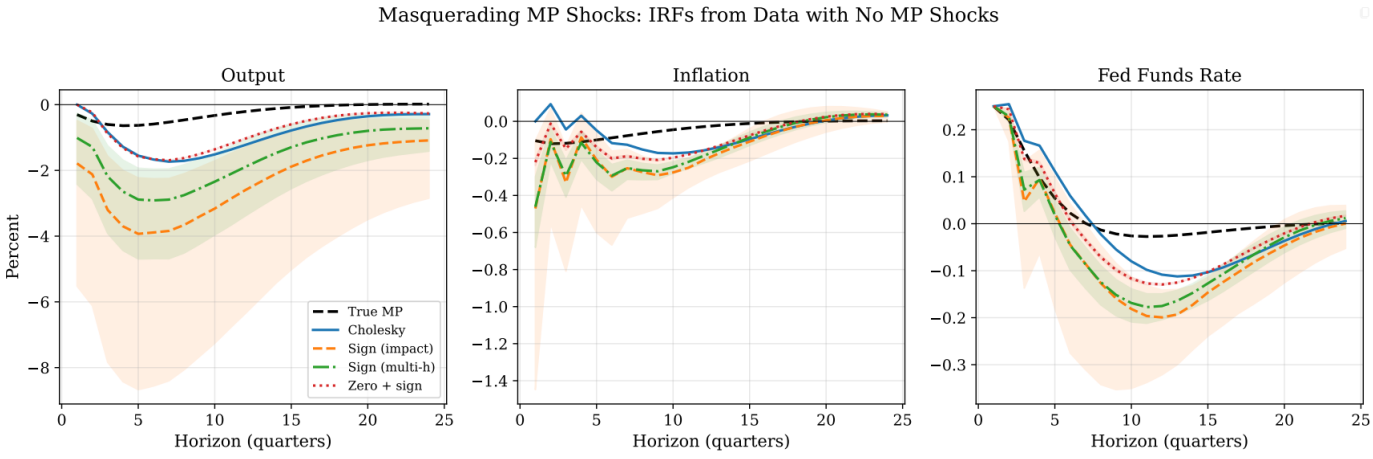


Figure 10

A.5 Do Professional Forecasters Learn from Other Individual Forecasts?

To be written.

A.6 Model Details

A.6.1 Deriving Equilibrium Conditions and the Natural Rate

To be written.

A.6.2 Optimal Policy under $\eta = 0$

To be written.

A.6.3 Proof of Proposition 4

Let the union of the time t information set $\mathcal{I}_t$ and the central bank's between period signal be $\overline{\mathcal{I}}_t$. We have

$$i_t = \mathbb{E}[r_t^* \mid \overline{\mathcal{I}}_t] - \gamma_{t-1}\widetilde{\imath}_{t-1},$$

where $\widetilde{\imath}_t = i_t - r_t^*$ is the interest rate gap. Define $z_t = i_t - \mathbb{E}[i_t \mid \mathcal{I}_{t-1}]$. By Lemma 1, we want to show that $\text{Cov}(\widetilde{\imath}_{t+h}, z_t) = 0 \ \forall h \geq 0$. We will start with the $h = 0$ case. Note that z_t and $\gamma_{t-1}\widetilde{\imath}_{t-1}$ are known to the central bank (measurable function of $\overline{\mathcal{I}}_{t-1}$). Further, the difference in one step ahead forecasts of tomorrow's natural rate (r_{t+1}^*) with respect to $\overline{\mathcal{I}}_t$ and $\mathcal{I}_t$ is mean zero. So it suffices to show that

$$\mathbb{E}\Big[(r_t^* - \mathbb{E}[r_t^* \mid \overline{\mathcal{I}}_{t-1})(\mathbb{E}[r_t^* \mid \overline{\mathcal{I}}_{t-1} - \mathbb{E}[r_t^* \mid \mathcal{I}_{t-1}])\Big] = 0.$$

Because beliefs about the structure about the economy are correct, this population moment is zero by the law of iterated expectations. The $h \geq 1$ cases follow identically.

A.6.4 Optimal Policy under $\eta \neq 0$

To be written.