

What Policy Surprises Reveal About the Policy Rule

Paul Bousquet*

July 31, 2026

preliminary

[\[Click here for latest version\]](#)

Abstract

I document U.S. monetary policy surprises are not explained by the Federal Reserve's point forecasts, are skewed towards the direction of actual policy adjustment, and are counter-cyclically dispersed. These facts are consistent with a policy rule depending on the central bank's entire belief distribution and a strong preference for hedging against uncertainty. I model monetary policy under ambiguity aversion and show a special case where the central bank exhibits both of these behaviors by placing additional weight on a scenario where the available data leads it astray. 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.

1 Introduction

Policy announcements by central banks are often not perfectly anticipated. 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. 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). Because forecast errors of the announcement are a consequence of what information must have been revealed, specific assumptions about the central bank's policy rule map into clear predictions about monetary policy surprises. This paper shows features of policy surprises that cannot be understood through standard policy rules and presents an alternative framework to generate them. The common theme is the importance of *uncertainty* for central bankers not captured by these rules.

Setting good monetary policy relies heavily on good economic forecasting (and nowcasting). Naturally, most policy rules often only take as input realizations and point forecasts of macroeconomic aggregates like inflation, output growth, and unemployment. Policy surprises under these rules should be completely accounted for once these inputs are known, 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 a "point forecast only" rule is a good approximation of the Fed. I find $R^2 \approx 0$ during the Great Moderation period (excluding the ZLB), rejecting a broad class of reaction functions as descriptions. These tests suggest had the Fed counterfactually made its internal forecasts public before its policy meetings, policy surprises would not have been any smaller on average. Since the Fed's beliefs about macroeconomic aggregates still presumably determine policy decisions, point forecasts must be far from a sufficient statistic.¹

In what environments would it make sense for this to be true? Most immediately, a world where following point forecasts can lead to undesirable outcomes. Caution is often paired with inertia, a feature of the policy rule observable even in raw time series data. Central banks frequently opt for inaction ("wait and see") and practice gradualism: even when there is consensus current policy is far from target, changes are small and broadly predictable.² These regularities show up clearly in the distribution of policy surprises: they are likely to share the same sign as the actual direction of rate adjustment (conditionally skewed). In other words, in the status quo, an announcement is the central bank's answer to either "have conditions changed materially to the point that a

¹This aligns with themes from work linking policy announcements and risk premia; see Knox and Vissing-Jorgensen (2025) and Cieslak et al. (2025) for a review.

²Gradualism 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.

target adjustment is needed?" or "how aggressive should the adjustment be?", where the latter takes some form of intervention as a given. The distribution of policy surprises shows most announcements are answering the first question and reflect a propensity for the central bank to "catch up" to its target. The distribution of surprises also reveals an asymmetry: it is considerably more dispersed during cutting cycles (counter-cyclical monetary policy uncertainty). This suggests a fundamental asymmetry in the policy rule.

Uncertainty, the combination of risk and ambiguity, has been described as the "defining feature" of monetary policy (Greenspan, 2003). This creates a natural motive for the central bank to follow a rule that depends on higher-order moments of its belief distribution. I formalize this idea within a generalized framework of ambiguity aversion where the central bank directly 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 Manganello, 2008; Cieslak et al., 2025), where, to take out insurance against making particularly costly decisions, the central bank acts *as if* it believes adverse models or scenarios are more likely. However, ambiguity aversion in general does not necessarily produce a policy rule that is a good descriptive fit of modern central banking. While ambiguity aversion can lead to this flavor of cautiousness (Brainard, 1967), it can also lead to *more* aggressive adjustments (Sargent, 1999; Barlevy, 2011). I show that a generic form of smooth ambiguity aversion (Klibanoff et al., 2009) in general leads to a more aggressive central bank, inconsistent with the behavior we observe. I instead present a case that can generate inaction and gradualism: the central bank is ambiguity averse about possible bias in its private signal.

A final feature of interest in the distribution of monetary policy surprises is its relationship to lags and leads of macroeconomic aggregates. A large literature uses monetary policy surprises to assess the effects of monetary policy (Ramey, 2016) and documents "hump-shaped" correlations with leads of macroeconomic aggregates, with magnitudes peaking several periods after the surprise's initial impulse. Standard models must put in considerable effort to match this fact (Nakamura et al., 2025). A different literature on policy surprises has noted they are correlated with information available prior to the policy announcement (Bauer and Swanson, 2023b; Miranda-Agrippino and Williams, 2026). I add a difficult learning environment to my model which can help explain both facts. Specifically, agents in the model must disentangle permanent and transitory shocks, a structure consistent with documented facts on the term structure of interest rates (Atkeson and Kehoe, 2008) and interest rate expectations (Muth, 1961; Farmer et al., 2024). Adding model uncertainty through the parameters of the natural rate's law of motion can create slow learning rates generating persistence that manifests itself in the types of correlations surprises have with aggregates. The "hump-shaped" correlations occur in the model because of the cyclical nature of the natural rate, amplified by the central bank practicing gradualism. The correlations with lagged data are simply finite sample, "ex post" artifacts that will disappear with sufficient data (Bauer and Swanson, 2023a; Eva and Winkler, 2023). These facts support a policy rule with sufficient incorporation of adjustments to

uncertainty (Orphanides and Williams, 2007).




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

2 Do We Learn From Policy Surprises?

The purpose of the next two sections is to present evidence on properties of U.S. monetary policy surprises, high-frequency movements in the yield curve around Federal Reserve policy announcements. Policy surprise data (1993-2026) are sourced from Jarociński (2024) and Acosta et al. (2026). 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., 2026). 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.

Given the large literature on what causes policy surprises, it's worth exploring in depth the existing perspectives and where they come from. This is particularly important for ensuring interpretations of data do not ignore

other factors. 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 typically revolves around the special cases of a (superior) "information effect" and an (unexpectedly aggressive) "response to news".³ The existing literature supports the view that both mechanisms may sometimes be present but are not habitually present. To support this conclusion, this section overviews the existing body of evidence and provides some new evidence directly targeted to the mechanisms. Section 2.1 reviews some particular correlations of policy surprises that have received much attention. Section 2.2 reviews some other arguments for CBI and Section 2.3 assesses more broadly if the policy rule of the Fed is perceived to drift. The evidence presented in Section 3 about the form of the policy rule also provides a test of the specific forms of CBI and RTN usually considered. Finally, parallel work (Bousquet, 2026a) argues that the extent of disagreement between the Fed and private sector suggests the learning channel must have a limited role.

2.1 Correlations of Interest

The CBI and RTN special cases are used primarily to explain 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. Bauer and Swanson (2023a,b); Acosta (2026); Sastry (2026) comprehensively investigate these facts and provide many caveats. 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.

(i) at first glance 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).⁴ It may be helpful to consider an analogy to econometrics developed further in ?. 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 with "purely statistical" (e.g., machine learning) estimation in small samples with high-dimensional data (Bianchi et al., 2022). In Appendix

³Sastry (2026) and Knox and Vissing-Jorgensen (2025) provide comprehensive overviews and a general setting nesting all possibilities.

⁴Miranda-Agrippino and Williams (2026) use a high-frequency window around major data releases and find the same "predictability drift".

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 statistically 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)

A.4, I show that conditional mean forecast errors from statistical methods in these environments can exhibit purported "predictability". This will be true also for the model in Section 4.⁵ Section 3 discusses the idea that these correlations could be a nuisance for an econometrician seeking to estimate the effects of policy.

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 3.2. In Appendix A.1, Figure 5 demonstrates that in some sense we should observe low-frequency comovement of surprises and revisions to growth forecasts because of the raw correlation of surprises and growth. For example, when contractionary surprises occur during times when, all else equal, growth will be higher than expected, 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 cycle movements they occur in, this is consistent with the small and sensitive results documented in Table 1.

Linking monthly revisions specifically to one Fed announcement is also difficult because professional forecasters tend to update their beliefs at low frequencies, with updates most frequent 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

⁵Ex post predictability can be inherent even in standard rational expectations Markov switching models (Hajdini and Kurmann, 2026).

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., 2026; Acosta et al., 2026; Swanson et al., 2026).⁶ A more general discussion on the possibility of the "information effect" is presented in Section 2.2. Summarizing, while information effects certainly can exist, there is no direct evidence they are a prevailing characteristic.

In terms of actual professional forecast revisions, (ii) is typically based on one particular Blue Chip consensus (BCEI), but both surveys (BCEI & BCFF) can also be used to better align the timing of Fed staff and consensus forecasts and to better control for the documented sensitivity to unscheduled meetings.⁷ A previously unnoted potential confounder of these revision correlations is they may exist even when revisions move in the opposite direction of the Fed's beliefs. If the Fed has a more optimistic growth outlook but surprises in an expansionary direction, downward revisions by professional forecasters in the following month should not be attributed to learning from the Fed the economy is faltering.⁸ Though we might expect expansionary surprises are synonymous with the Fed having a relatively weaker outlook (and vice versa for contractionary surprises), only around 60% of scheduled Fed policy announcements feature this alignment between the signs of the surprise and the disagreement (averaged across inflation and growth forecasts; see Table 3). With these considerations in mind, Table 3 in Appendix A.1 shows no consistent evidence for an information effect: the few effects distinct from zero are not economically significant and are driven by outliers (Figure 6).

Another motivation for the focus on (i) and (ii) is the potential implications for regression exercises with policy surprises.⁹ 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 central banks do not "flip coins" (Leeper et al., 1996; Nakamura and Steinsson, 2018b), it is unclear why such an object would exist. One common, informal 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. In parallel work (Bousquet, 2026a), I show this conjecture is untrue in general.

⁶Acosta (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.

⁷Online Appendix B.3 shows results with Bloomberg's ECFC survey to provide another test of timing sensitivity.

⁸The outlook may not be captured by Fed staff forecasts if they are modal rather than mean forecasts (Aruoba and Drechsel, 2026).

⁹Though some specific recommendations of Bauer and Swanson (2023a) are based on the specific interest rate rule described above, they are carefully framed as being directionally helpful but not necessarily a panacea. Econometrically, this is unclear without further assumptions (Bousquet, 2026c,b).

2.2 Information Effect

It is not hard to imagine situations where monetary policy announcements also communicate information about the economy. During the Great Financial Crisis, it is plausible the public sometimes inferred that some of the Fed's decisions were informed by its private knowledge of the day-to-day activity of the financial system. Event study decompositions of specific announcements, in the vein of [Del Negro et al. \(2023\)](#) and [Bianchi et al. \(2026\)](#), are definitely appropriate to assess whether meaningful information about the economy was revealed. But treating these occurrences as representative of 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 the economy's trajectory to get smaller over time, but I show in Online Appendix C this does not occur. As mentioned in the discussion for [Table 3](#), it is 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). Taken in isolation, some of the point estimates in [Table 3](#) do support the presence of a marginal information effect, but they also raise the questions of how professional forecasters can successfully tell when one is present and what the cause of surprises is otherwise perceived to be.

[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 the central bank is viewed 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, consistent with evidence in [Klein \(2022\)](#). 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. [Faust et al. \(2004\)](#); [Nakamura and Steinsson \(2018a\)](#) point out that the Fed produces Industrial Production data, and in months where timing aligns they will have early

access to this data. I show in Appendix A.1 there is no evidence of distinct information effects in these instances.¹⁰ Without these considerations of unique data access, it is hard to argue the Fed has a unique signal in a way e.g., Goldman Sachs does not.

2.3 Perceived Drift in the Policy Rule

The main idea underlying Bauer and Swanson (2023a,b) is the Fed has private knowledge of 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 as an account of why policy surprises persist. It is certainly plausible that the Fed has private knowledge of 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 the Fed is 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 the 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), which 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 implies significant reaction function updates occur around policy announcements. 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 $\bar{\mathbb{E}}_t[\phi_{t+h}] = \bar{\mathbb{E}}_t[\phi_t]$ and $\bar{\mathbb{E}}_t[X_{t+h}] = X_t$ (and $R^2 = 1$). Figure 2 shows that is not the case — nearby maturities do not move one-for-one with surprise changes to today's interest rate target, and

¹⁰The Bank of England publishes inflation reports a few days after policy announcements. Hubert (2025) finds interest rate futures and asset prices have a stronger contractionary relationship during these releases (relative to announcements), incompatible with superior information.

their responses (along with the R^2) are decreasing in the futures contract horizon. Though this is relative to an extremely persistent benchmark, Appendix A.2 discusses a more general case where 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 .¹¹ 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 the autocorrelation of residuals $u_{i,t}$ from the following specification

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

where α_i and τ_t are individual and time fixed effects. Here, this is run at the panel level (with forecasters from BCFF) and with two choices of $x_{i,t}$ — 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 $\widehat{\text{macro}}_{i,t}$ variable contains factors relevant to the macroeconomic outlook but plausibly unrelated to contemporaneous changes in policy. $u_{i,t}$ is negatively

¹¹When shocks to the rule affect X_{t+h} , Ricco and Savini (2025) show RTN makes different but still counterfactual predictions. For example, a contractionary surprise can cause longer yields to fall because of the effect on the path of inflation.

autocorrelated across a range of specifications for $\widehat{\text{macro}}_{i,t}$, suggestive evidence inconsistent with an environment where agents are constantly learning about a drifting policy rule.¹²

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, with details provided in [Bousquet \(2026a\)](#). One objection could be that slow learning and limited sample size could be consistent with persistent disagreement even in expectation. [Table 3](#) 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 is possible the Fed's action can be a surprise even to the staff ([Aruoba and Drechsel, 2026](#)), [Figure 7](#) (Appendix A.2) shows this trend and broader disagreement persist when using the forecasts of the Fed governors and presidents. This suggests that the disagreement is a consequence of pervasive uncertainty about what the future actually holds for the economy, rather than private reaction function knowledge.

3 Policy Surprises and the Policy Rule

This section moves beyond the existing debate about informational asymmetries and looks for evidence of what can be gleaned about the central bank's systematic behavior from policy surprises. [Section 3.1](#) tests basic implications of central banks only using point forecasts in their policy rule and finds negative results. [Section 3.2](#) documents new facts about policy surprises, including their conditional skewness and counter-cyclical dispersion.

3.1 The Policy Rule is Not Spanned by Point Forecasts

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? One common form for the central bank's rule is to assume policy rules take only as input the central bank's point forecasts about macroeconomic aggregates ([Romer and Romer, 2004](#); [Miranda-Agrippino and Ricco, 2021](#)). We can think of this as a central bank information effect (CBI) rule — because the central bank has private knowledge of its forecasts, it surprises the public. So these rules predict that if counterfactually the forecasts had been known, there would be no surprise movements in policy. [Bauer and Swanson \(2023b\)](#) instead argue the central bank only "responds to news" (RTN) that is available to everyone and the private information

¹² $\widehat{\text{macro}}_{i,t}$ 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 central bank has its monetary policy rule. The special case they present (the RTN rule) carries a testable implication: if the central bank follows a time-varying linear rule taking only public data as input and markets always underestimate how aggressively it will respond, public data should completely explain surprises. Together, the CBI and RTN rules cover the vast majority of specifications of the central bank's reaction function. Though these rules are never treated as beyond reproach, they are meant to serve as sufficient approximations. So while rejecting $R^2 \approx 1$ is consequential only under some assumptions that should perhaps not be taken literally, explaining a non-trivial share of variance is arguably a necessary condition of a good approximation.¹³

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.1. 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 Blue Chip (BCFF) consensus forecast (for inflation and output at horizons $h = 0, \dots, 4$), 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 to describe the "typical" conduct of monetary policy, periods at the zero lower bound as well as months with a significant unscheduled meeting are excluded.¹⁴ Table 2 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 2: Policy Surprise Variation Explained by CBI and RTN

<i>Adjusted R²</i>				
	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.1 for additional information/exercises.

Online Appendix C.1 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 does including forecasts of other variables. 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

¹³For example, R^2 will be smaller under RTN if the Fed is sometimes less aggressive than expected, and using belief differences as a proxy for errors in the forecast of the Fed's forecast relies on the law of iterated expectations. A more traditional source of R^2 attenuation is measurement error of the true forecasts ([Aruoba and Drechsel, 2026](#)). 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](#)).

¹⁴The 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.

using the FOMC's Summary of Economic Projections (SEP). Albeit with a short sample (from Bloomberg's ECFC since 2018), incorporating survey data on the *expected* median FOMC forecast also has virtually no explanatory power.¹⁵ Lastly, if we take the point estimates from these regressions at face value, Online Appendix C.1 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 its 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). More generally, these results show the Fed makes decisions that are not captured by simple point forecasts.

3.2 The Distribution of Surprises as Evidence of Cautious Policymaking

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. If central banks do not "flip coins", then the distribution of policy surprises must depend on something about the central bank, either its reaction function (e.g., discretion, taste shifts, or generally being unknown) or something it is reacting to. Figure 3 presents clear evidence of state-dependence: MP1 surprises are conditionally skewed toward 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 indicates that benchmark models are missing something fundamental about both monetary policy and the data generating process informing its conduct.

¹⁵See Martinez and Sinclair (2026) for more information on this particular questionnaire.




Figure 3: 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 overnight rate target 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 recent US monetary policy history where the ex post consensus is the Fed acted too aggressively.¹⁶ Figure 3 represents a propensity of the Fed to "catch up". In Appendix A.2, Figure 8 demonstrates this more formally; large changes in the target make 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. Additionally, Portier (2026) shows longer-term yields move more in response to policy surprises when the Fed is in a "plateau" phase, extended periods of no changes to the target with intermittent uncertainty about the direction of the next policy action. During these phases, the Fed's assessment of the economy tips the balance towards the next move, and from the historical record, the public understands that the initial step out of a plateau will be the first of many in that direction. These observations motivate a model of policymaking developed in Section 4 where the central bank does not act aggressively unless it is certain there is substantial room for future adjustments.

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)

¹⁶Cieslak et al. (2026) offers the full extent of the 2017-2019 tightening as a possible exception.

- Counter-cyclical dispersion of policy surprises
- Finite-sample correlations with lagged news and "hump-shaped" with leads of macroeconomic aggregates

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

discusses the information structure of the model and restates the constraints the central bank faces. [Section 4.1](#) overviews a broad class of preferences where a decision maker behaves *as if* they have a different probability distribution. It then interprets these preferences in the context of monetary policy as central banks performing "risk management" and justifies why they would (or possibly should) arise. [Section 4.2](#) establishes the second 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. This is necessary to generate more persistence in the model to match some of the empirical facts. [Section 4.3](#) presents model results on optimal policy, outcomes, and what relationships an econometrician would uncover with data simulated from this model.

4.1 Generalized "as if" Preferences and Risk Management

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 -maxmin 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

$$\max_{\{y_t\}_{t \geq 0}} - \sum_{t=0}^{\infty} \beta^t \left(\bar{\theta} \mathbb{E}_0[\ell(y_t; \mathbf{x}_t)] + \sum_{s \in S_t} \int \theta_s(\mathbf{x}_t) \ell(y_t; \mathbf{x}_t) d\mathbf{x}_t \right).$$

For example, maxmin expected utility preferences correspond to $\bar{\theta} = 0$ and a single scenario s with $\theta_s = 1$ if $\mathbf{x}_t = \text{argmin} -\ell(y_t; \mathbf{x}_t)$ 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 that potentially depends both on some intrinsic distaste for a scenario and on how

likely that scenario is under a possibly scenario-specific probability measure. Returning to the [Sims \(2001\)](#) analogy, these preferences can either add fixed mass to a given probability distribution or stretch it at a set of points by a fixed proportion. In other words, we can represent the kernel $\theta_s(s)$ as

$$\theta_s(\cdot) = \begin{cases} f_s(\cdot) + \omega_s(\cdot) \\ f_s(\cdot) \times \omega_s(\cdot), \end{cases}$$

where $f_s(\cdot)$ are scenario and state-specific probability densities and weights. The second type is more common, and in fact some forms of ambiguity and risk aversion are isomorphic in linear-quadratic Gaussian environments ([Hansen and Sargent, 2008](#)). The first type corresponds to weighting scenarios even when we put no probability on them, which is outside of the scope of standard presentations of ambiguity aversion. The importance of this distinction in the context of the model can be seen by applying these preferences to monetary policymaking.

Why might the central bank have preferences in this class? 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](#)). Virtually all central bankers invoke the concept of "risk management" when explaining their rationale ([Kilian and Manganelli, 2008](#); [Cieslak et al., 2025](#)). This framing is critical to understanding policy surprises. When the Fed is deciding between cutting 25 bp and 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.¹⁷ [Cieslak et al. \(2025\)](#) 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 in [Appendix A.2](#).

Now we have sufficient background to apply the formal preference class. A standard 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 is 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:

¹⁷The confounding factors listed are ways in which policy is intentional but nonetheless still has a stochastic component, and this idea is explored in more detail in ?.

one's belief about where the economy is heading could be fundamentally off-centered, and one wants to hedge against that scenario.

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. Suppose the central bank must set policy with a lag but can condition it on a private signal it receives between periods. For simplicity, the optimal policy problem is recast as a forecasting problem, which is similar to what would arise in a "divine coincidence" world (Benigno and Woodford, 2004). 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

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

where $b \in \mathcal{B}_t$ represents a feasible set of signal biases. Here, we can extend this idea further to generate several features of the monetary policy status quo. Define $\mu_t(s)$ as the objective Kalman update after seeing a signal s given the information set at t and rewrite $g_t(s + x)$ for $x = 0, b$ as $\alpha_t(s)(\mu_t(s + x) - i_t) + i_t$.¹⁸ 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 no longer leads to adjusting rates in the *opposite* direction from the one our belief implies. This effectively acts as a fixed cost of sufficient confidence the central bank must pay before taking action, which produces periods of inaction and gradualism and, through specifications of $\mathcal{B}_t$, can also lead to the counter-cyclical dispersion of policy surprises shown in Figure 3. To see this, let $\Delta_t(s) = \mu_t(s) - i_t$ be tomorrow's expected distance from the natural rate if a signal s is received but policy is held constant at i_t . Appendix A.6 shows inaction occurs when

$$|\theta (\mu_t(s + b) - i_t)| > |\Delta_t(s)|$$

If $\Delta_t(s) \approx 0$, then current policy is relatively close to optimal and the magnitude of the insurance term might dominate, making it not worthwhile to pay the de facto fixed cost. When policy moves out of the inaction region, the fixed probability private agents put on staying there leads to the conditional skewness of surprises. Counter-cyclical dispersion simply comes from a specification of $b \in \mathcal{B}_t$, specifically when the central bank is asymmetrically worried about its current stance being over-optimistic.

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.

¹⁸Since $\mu_t(s)$ and i_t are deterministic functions of the time t information set, the problem is isomorphic to choosing $\alpha_t(s)$.

4.2 Adding a Learning Problem

The previous subsection demonstrated a preference class that will generate the qualitative patterns observed in the data, but it falls short on the quantitative patterns. To capture the strength of the documented correlations in policy and policy surprises, the underlying data generating process needs a force with substantial persistence.

First, we present the exogenous learning environment faced by agents in the macroeconomic model, keeping their roles and optimization problems abstract for now. Suppose there is a continuum of agents $i \in [0, 1]$ seeking to make forecasts of r_t^* , which is the sum of two unobserved components:

$$r_t^* = a_t + d_t, \quad (1)$$

$$a_t = \rho a_{t-1} + \varepsilon_t^a, \quad (2)$$

$$d_t = d_{t-1} + \varepsilon_t^z, \quad (3)$$

$$\varepsilon_t^j \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_j^2), \quad j \in \{a, d\}. \quad (4)$$

We assume the structure of this system of equations is known, but the parameter vector $\theta = (\rho, \sigma_a, \sigma_d)$ that indexes the data generating process is unknown.

Agents observe a private signal and an aggregated signal. Between time t and $t + 1$ each agent receives a private signal of the transitory component:

$$s_{i,t} = a_t + \bar{\mu}_t + \varepsilon_{i,t}^s, \quad \bar{\mu}_t = \varepsilon_t^{\bar{\mu}}, \quad (5)$$

where $\bar{\mu}_t$ is a common conditional bias and $\varepsilon_{i,t}^s$ is independent across i . A public aggregate signal is realized at $t + 1$:

$$p_{t+1} = a_t + \bar{\mu}_t \quad (6)$$

All variances ($\sigma_s, \sigma_{\bar{\mu}}, \sigma_p$) are known and $\sigma_{\bar{\mu}} > 0$ prevents the latent a_t from being revealed exactly. For expositional simplicity, we assume agents begin at $t = 0$ with the same priors over θ and the latent states. This means that at the start of each period, agents share common beliefs and filtration.

Each agent enters period t with a prior $f_{\theta,t-1}$ over θ and posterior beliefs about the latent components. After observing r_t^* at the beginning of the period, agents update by Bayes rule to produce $f_{\theta,t}$. For tractability, we therefore impose anticipated learning (Kreps, 1998): today the agent updates beliefs, but for the purpose of forming forward-looking decisions she behaves as if her current posterior $f_{\theta,t}$ will remain her belief at every future date.

Again, we can think of θ indexing a world of possible models agents may live in. An infinite-lived agent solving

a decision problem in this environment solves a dynamic program of the following form (Klibanoff et al., 2009)

$$V_t(\mathbf{x}_t) = \max_{\mathbf{y}_t} \left[\int_{\boldsymbol{\theta}} f_{\boldsymbol{\theta},t} \left(U(\mathbf{t}) + \beta \int_{\boldsymbol{\varepsilon}} f_{\boldsymbol{\varepsilon},t} \cdot V_t(\mathbf{x}_{t+1}(\mathbf{t}, \boldsymbol{\varepsilon}; \boldsymbol{\theta})) d\boldsymbol{\varepsilon} \right)^{1-\eta} d\boldsymbol{\theta} \right]^{\frac{1}{1-\eta}}, \quad (7)$$

where $\mathbf{x}_t$ is the state vector, $\mathbf{y}_t$ is the control vector, $\mathbf{t} = (\mathbf{x}_t, \mathbf{y}_t)$, $f_{\boldsymbol{\theta},t}$ 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 $\mathbf{x}_{t+1}$ depends on what model nature has actually selected.

4.3 A New Keynesian Model where r_t^* Has Unobserved Components

We now embed this learning mechanism into a model with a central bank solving their optimal policy problem with discretion under ambiguity aversion preferences. We restrict other agents in the model to $\eta = 0$, which essentially is the analog of expected utility preferences under model uncertainty. The learning mechanism has been carefully defined so the private sector block of the textbook New Keynesian model (Benigno and Woodford, 2004) can remain perfectly intact. Households have separable, isoelastic preferences over consumption and leisure, and firms choose prices to maximize profits subject to (Rotemberg, 1982) adjustment costs in prices. Under the listed assumptions, the model admits an aggregate representation structurally identical to a classic presentation of the log-linear New Keynesian model. The learning in the model is about r_t^* , the interest rate that would prevail in a flexible price economy and is a function of only stochastic processes in baseline setups. Microfoundations for this environment are provided in Appendix A.6 as well as in Caballero and Simsek (2022) and Rungcharoenkitkul and Winkler (2026). The structure described in the preceding subsection has been shown to perform well at fitting data on interest rates and interest rate forecasts (Muth, 1961; Atkeson and Kehoe, 2008; Farmer et al., 2024), as Figure 4 further illustrates.




Figure 4

A monetary policy process needs to be determined to close the model. The previous signal structure will be extended to create policy surprises. Without loss of generality, we consider the central bank to be agent $i = 1$. We impose an implementation lag: at t the bank announces a measurable function $i_{t+1} : \bar{\mathcal{I}}_t \rightarrow \mathbb{R}$, where $\bar{\mathcal{I}}_t$ contains everything known at t and the between-period signals revealed before $t + 1$. This commitment cannot be reneged on. This restriction will create surprise movements in policy if the central bank decides to make use of their between-period signal. Because the central bank is just one agent in the continuum, its signal does not provide additional information about the latent states once the aggregated signal is observed. Under these conditions, the discretionary equilibrium is given by

Definition. The discretionary equilibrium is a sequence of aggregates $\{\tilde{y}_t, \pi_t, i_t, r_t^*\}_{t=0}^\infty$ and beliefs $\{f_{\theta,t}, \hat{b}_t, \hat{p}_t\}_{t=0}^\infty$ such that:

1. States evolve according to (1)—(4) and (5)—(6) with initial conditions $(i_0, r_0^*, f_{\theta,-1}, \hat{b}_{-1}, \hat{p}_{-1}, p_0)$
2. $(\tilde{y}_t, \pi_t, i_t, r_t^*)$ and public signal p_t defined in (6) are realized simultaneously at the start of period t
3. Upon realization of aggregates, agents share a common filtration

$$\mathcal{F}_t = \{\tilde{y}_k, \pi_k, i_k, r_k^*, p_k, f_{\theta,k-1}, \hat{b}_{k-1}, \hat{p}_{k-1}\}_{k=0}^t \quad (8)$$

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

4. The common expectations operator $\bar{\mathbb{E}}_t$ satisfies

$$\bar{\mathbb{E}}_t[\cdot] = \int_{\theta} \mathbb{E}[\cdot | \theta, \mathcal{F}_t] f_{\theta,t} d\theta \quad (9)$$

and generates expectations $\hat{b}_t = \bar{\mathbb{E}}_t[b_t]$ and $\hat{P}_t = \bar{\mathbb{E}}_t[(b_t - \hat{b}_t)(b_t - \hat{b}_t)']$ over latent state vector b_t .

5. $(\tilde{y}_t, \pi_t)$ is a measurable function of $\mathcal{X}_t \subset \mathcal{F}_t$ satisfying

$$\tilde{y}_t = \bar{\mathbb{E}}_t[\tilde{y}_{t+1}] - i_t + r_t^* + \bar{\mathbb{E}}_t[\pi_{t+1}] \quad (10)$$

$$\pi_t = \beta \bar{\mathbb{E}}_t[\pi_{t+1}] + \kappa \tilde{y}_t. \quad (11)$$

6. $i_{t+1} : \mathcal{X}_t \times \mathbb{R} \rightarrow \mathbb{R}$ is a measurable function of $\mathcal{X}_t \in \mathcal{F}_t$ and its signal (5) and solves dynamic program (7) with flow utility $-(\tilde{y}_t^2 + \lambda \pi_t^2)$ taking as given the action of future policymakers. i_{t+1} is converted into an equivalent feedback rule to ensure $(\tilde{y}_t, \pi_t)$ are the unique solution to $\min \tilde{y}_t^2 + \lambda \pi_t^2$ across the set of all possible solutions to (10)–(11).

The following result demonstrates the equilibrium is standard if the central bank is not ambiguity averse.

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

Proof: See Appendix A.6.

While ambiguity aversion is capable of generating patterns consistent with cautious policymaking (Brainard, 1967), it can also lead to *more* aggressive adjustments (Sargent, 1999; Barlevy, 2011). In this model, standard ambiguity aversion has exactly this effect.

Proposition 2. Comparing the discretionary equilibrium for $\eta = 0$ and $\eta \neq 0$, the unconditional variance of i_t is strictly lower when $\eta = 0$.

Proof: See Appendix A.6.

As discussed, ambiguity aversion still has strong anecdotal support for central banking, motivating the alternative form considered earlier in this section. The learning environment and these preferences combine to match the patterns observed in the data. In this model, a monetary policy surprise, $z_t = i_t - \bar{\mathbb{E}}_{t-1}[i_t]$, is defined with respect to the common beliefs of all agents prior to the between-period individual signals. Suppose an econometrician in this environment has a time series of aggregates $\{\tilde{y}_t, \pi_t, i_t\}_{t=1}^T$ 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 4, 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. With standard sample sizes, 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 true impulse response function is only nonzero on impact, meaning any inferences 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 3. Agents rationally underestimate the extent to which 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 when policy will subsequently tighten by progressively 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.¹⁹ Surprises are genuinely unpredictable, so adding controls and lag-augmentation (Montiel Olea and Plagborg-Møller, 2021) does not change the qualitative results. However, because surprises do exhibit ex post correlations with lagged aggregates, it does make some superficial quantitative difference: point estimates change somewhat, but the general shape is preserved. 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. These patterns can exist even when the central bank has $\eta = 0$ preferences but will be amplified in the special case considered. The aforementioned conditional skewness of surprises means that hikes imply the interest rate gap will get smaller over time as the central bank "catches up" (vice versa for cuts), leading to estimated peak effects farther out.

5 Conclusion

The existing debate over the cause of policy surprises supposes that the central bank has private information about the economy or its policy decision rule. This paper finds evidence against the models used to illustrate these views but finds that policy surprises are nonetheless informative about the central bank's systematic behavior. Because point forecasts provide no explanatory power for surprises, this suggests the meeting-by-meeting decisions of policymakers are heavily influenced by higher-order moments of the economy's joint distribution. This is further evidenced by the skewed distribution of policy surprises, which is caused by gradualism that is presumably caused by risks away from forecasts of the economy's central tendencies. To these ends, I develop a model where the central bank takes out insurance against data leading it astray. This model can successfully replicate the key features of the policy surprise distribution.

¹⁹There is a similar intuition in Caballero and Simsek (2022) discussed further in Online Appendix D.

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.** 2026. “Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database.” working paper.
- Amodeo, Francesco.** 2026. “Mind the Gap: Disagreement and Credible Monetary Policy.” Working Paper.
- 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.
- Barlevy, Gadi.** 2011. “Robustness and Macroeconomic Policy.” *Annual Review of Economics*, 3 1–24.
- 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.
- Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma.** 2022. “Belief Distortions and Macroeconomic Fluctuations.” *American Economic Review*, 112(7): 2269–2315.
- 2026. “A Structural Approach to High-Frequency Event Studies: The Fed and Markets as Case History.” Working Paper.
- Bousquet, Paul.** 2026a. “How Central Bank Announcements Affect Expectations.” Working Paper.

- . 2026b. “Minimal Assumptions to Uncover Macro Dynamics with "Shocks".” Working Paper.
- . 2026c. “Where Do Price and Output Puzzles Come From?”, Working Paper.
- Brainard, William.** 1967. “Uncertainty and the Effectiveness of Policy.” *American Economic Review*, 57(2): 411–425.
- 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.
- 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.** 2025. “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.** 2026. “Did I Make Myself Clear? The Fed and the Market under the 2020 Monetary Policy Framework.” Working paper.
- Cochrane, John.** 2011. “Determinacy and Identification with Taylor Rules.” *Journal of Political Economy*, 119(3): 565–615.
- Cochrane, John H.** 2004. “Comments on "A new measure of monetary shocks: Derivation and implications".” July, NBER EFG meeting.
- 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., Emi Nakamura, and Jón Steinsson.** 2024. “Learning About the Long Run.” *Journal of Political Economy*, 132(10): 3334–3377.
- Faust, Jon, Eric T. Swanson, and Jonathan H. Wright.** 2004. “Do Federal Reserve Policy Surprises Reveal Superior Information About the Economy?” *Contributions to Macroeconomics*, 4(1): .
- 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.”
- Hajdini, Ina, and André Kurmann.** 2026. “Predictable Forecast Errors in Full-Information Rational Expectations Models with Regime Shifts.” 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.
- Hubert, Paul.** 2025. “Information Effects and Central Bank Macroeconomic Forecasts.” 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.
- 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.
- Klein, Tony.** 2022. “Agree to disagree? Predictions of U.S. nonfarm payroll changes between 2008 and 2020 and the impact of the COVID-19 labor shock.” *Journal of Economic Behavior & Organization*, 194 264–286.
- 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.
- 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?” *Econometrics Journal*, forthcoming.
- Martinez, Andrew B., and Tara M. Sinclair.** 2026. “When the Fed Reveals its Hand: The SEP and Monetary Policy Surprises.” Working Paper.
- 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.
- 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.
- Portier, Rose.** 2026. “The Phase-Dependent Effects of Monetary Policy: Plateau vs Cycle.” Working Paper.
- Ramey, Valerie.** 2016. “Macroeconomic Shocks and Their Propagation.” In *Handbook of Macroeconomics*. Chap. 2.
- Ricco, Giovanni, and Emanuele Savini.** 2025. “Decomposing Monetary Policy Surprises: Shock, Information, and Policy Rule Revision.” Working paper.
- 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.
- Rotemberg, Julio J.** 1982. “Monopolistic Price Adjustment and Aggregate Output.” *Review of Economic Studies*, 49 517–531.
- Rungcharoenkitkul, Phurichai, and Fabian Winkler.** 2026. “The Natural Rate of Interest Through a Hall of Mirrors.” *Journal of Monetary Economics*, 157.
- Sargent, Thomas J.** 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.
- 2015. “Sets of Models and Prices of Uncertainty.” Keynote address at the Conference for Expectations in Dynamic Macroeconomic Models.

- Sastry, Karthik A.** 2026. “Disagreement about Monetary Policy.” *American Economic Journal: Macroeconomics*, 18(2): 486–522.
- Sims, Christopher A.** 1980. “Macroeconomics and Reality.” *Econometrica*, 48(1): 1–48.
- 2001. “Pitfalls of a Minimax Approach to Model Uncertainty.” *American Economic Review*, 91(2): 51–54.
- Smets, Frank, and Rafael Wouters.** 2007. “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review*, 97(3): 586–606.
- Swanson, Eric T., Renxuan Wang, and Yanbin Wu.** 2026. “The Effects of Monetary Policy on Macroeconomic Expectations: High-Frequency Evidence from Traded Event Contracts.” Working Paper.
- 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 5 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 3.2, 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 5

A.1.2 A More Precise Information Effect Test with Forecast Revisions

?? discusses a common test of information effects: looking at the correlations of interest rate surprises and forecast revisions. Table 3 makes some adjustments to the default specification of these exercises, and so the description is self-contained, I restate some of the exposition from the main body of the paper. 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. 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 3). 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 3 shows two different tests of (ii) with respect to a one standard deviation change in the policy surprise 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

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

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

where $\bar{\mathbb{E}}_t[y_{q+h}]$ 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.2). 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 .

A.1.3 Revision Regressions are Driven by Outliers

Section 2.1 discusses regressions of professional forecaster revisions on surprises. Figure 6 demonstrates a common theme for these results: they are sensitive to outliers. Here, we take an example from Table 3 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 6 looks at one of the more significant correlations in Table 3 (for 2

Table 3

	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 (0.214)	0.092** (0.039)	0.098*** (0.037)	0.027 (0.041)	0.069** (0.028)	0.194 (0.127)	-0.075 (0.102)	0.105*** (0.039)	0.058** (0.022)	0.033 (0.036)	0.069 (0.048)
FF4	-0.215 (0.231)	0.053 (0.039)	0.080** (0.035)	0.024 (0.032)	0.070*** (0.022)	0.182 (0.126)	-0.158 (0.119)	0.085** (0.034)	0.041** (0.019)	0.037 (0.025)	0.063 (0.041)
ED4	-0.071 (0.126)	0.008 (0.037)	0.044** (0.022)	0.044 (0.042)	0.028 (0.018)	0.087 (0.078)	-0.090 (0.095)	0.049* (0.029)	-0.008 (0.021)	0.016 (0.016)	0.030 (0.034)
Test B: Differential Probability of Surprise-Revision Comovement Based on Type of Disagreement											
MP1	0.033 (0.081)	-0.095 (0.071)	0.210*** (0.059)	0.065 (0.067)	0.085 (0.068)	0.107 (0.078)	0.005 (0.076)	-0.028 (0.073)	0.083 (0.069)	0.082 (0.064)	0.042 (0.083)
FF4	0.058 (0.080)	-0.128* (0.071)	0.101 (0.064)	0.079 (0.067)	0.057 (0.069)	0.088 (0.077)	-0.008 (0.079)	-0.035 (0.071)	0.066 (0.069)	0.094 (0.063)	-0.083 (0.083)
ED4	0.005 (0.079)	-0.055 (0.070)	0.075 (0.065)	0.131* (0.067)	0.093 (0.069)	0.043 (0.079)	-0.111 (0.077)	0.037 (0.070)	0.112* (0.067)	0.000 (0.066)	0.015 (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.2 for additional information/exercises.

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.²⁰




Figure 6

²⁰The magnitudes are small enough that negligible noise infusions to the surprises decrease effect sizes sometimes over 50

A.1.4 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.²¹ This is consistent with the fact that the observed magnitudes are negligible. Bousquet (2026c) and Figure 6 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.

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 7 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.

²¹As discussed in Section 2.1 and Acosta (2026), an information effect presumably could still include the return being less negative.




Figure 7: FOMC vs. Professional Forecast of the Fed Funds Rate

A.2.2 Large Moves and Future Reversals

Figure 8 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 $R_{t+h} = \mathbb{1}_{\min_{k \leq h} \{\Delta i_{t+k} \cdot \Delta i_t\} \geq 0}$ 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 8

A.2.3 Quotes About Risk Management

“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**

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 $i_t = \phi'_t \cdot X_t$. In a standard, log-linear DSGE model, we can write

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

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 X_t (weighted by ϕ_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

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

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

where $\omega_i = \psi_i^2 \sigma_i^2$, σ_i^2 is the variance of structural shock i , and $m = \operatorname{argmax}_i \rho_i$ (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 et al., 2024) 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 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 Predictability²²

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 $K \gg T$. 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

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

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

²²Code for these exercises is available at this paper's GitHub repository, <https://github.com/paulbousquet/WhySurprise>

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 estimation is performed on Y_t and D_t , then the residuals of these first stages $(\tilde{Y}_t, \tilde{D}_t)$ are regressed on one another. While the convergence of $\hat{\theta} \rightarrow \theta$ is remarkably fast in this environment, $\{(\tilde{Y}_t, \tilde{D}_t)\}_{t=1}^T$ 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 $T \rightarrow \infty$. 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 9](#) 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 9

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 $\eta = 0$ (Proof of Proposition 2) and Discussion of Policy Surprises

Define the interest-rate gap $g_t \equiv i_t - r_t^*$ and conjecture that the bank's announced policy admits a stationary Markov representation

$$i_{t+1} = \hat{r}_{t+1}^*(s) + \gamma(r_t^* - i_t) \iff \bar{\mathbb{E}}_t[g_{t+1}] = -\gamma g_t, \quad (12)$$

for some $\gamma \in [0, 1)$, where $\hat{r}_{t+1}^*(s) = \bar{\mathbb{E}}_t[r_{t+1}^* | s]$ is the expectation conditional on the central bank's between-period signal. Under (12) the gap evolves as

$$g_{t+1} = -\gamma g_t + \varepsilon_{t+1}, \quad \mathbb{E}_t[\varepsilon_{t+1}] = 0, \quad (13)$$

and beliefs/signal noise enter only through the conditional law of ε_{t+1} . In particular, $\mathbb{E}_t[g_{t+1}]$ is independent of those nuisance elements. Substituting the conjecture (12) into (10)–(11), write $\tilde{y}_t = -\phi(\gamma)g_t$ and $\pi_t = -\theta(\gamma)g_t$. Matching coefficients yields, with

$$D(\gamma) \equiv 1 + (1 + \beta + \kappa)\gamma + \beta\gamma^2, \quad (14)$$

the closed forms

$$\phi(\gamma) = \frac{1 + \beta\gamma}{D(\gamma)}, \quad \theta(\gamma) = \frac{\kappa}{D(\gamma)}. \quad (15)$$

The period loss coefficient is then

$$\tilde{y}_t^2 + \psi \pi_t^2 = \Lambda(\gamma) g_t^2, \quad \Lambda(\gamma) \equiv \frac{(1 + \beta\gamma)^2 + \psi \kappa^2}{D(\gamma)^2}. \quad (16)$$

Because g_t^2 is a sufficient statistic for the per-period objective, the value function admits a stationary quadratic form $V(g) = P g^2 + \alpha_0$. A guess-and-verify approach leads to an equation pinning down γ^* that solves the central bank's problem.

We will now demonstrate that this economy satisfies the conditions of Proposition 1. Let the union of the time t information set $\mathcal{I}_t$ and the central bank's between period signal be $\bar{\mathcal{I}}_t$. We have

$$i_t = \mathbb{E}[r_t^* | \bar{\mathcal{I}}_t] - \gamma \tilde{y}_{t-1},$$

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

$$\mathbb{E}\left[(r_t^* - \mathbb{E}[r_t^* | \bar{\mathcal{I}}_{t-1}]) (\mathbb{E}[r_t^* | \bar{\mathcal{I}}_{t-1}] - \mathbb{E}[r_t^* | \mathcal{I}_{t-1}])\right] = 0.$$

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.6.3 Optimal Policy under $\eta \neq 0$

To be written.