Mind the Gap:
Disagreement and Credible Monetary Policy

Francesco Amodeo
UC San Diego

May 2026

Abstract

This paper studies the impact of interest rate disagreement between the Federal Reserve and financial markets on the transmission of U.S. monetary policy. First, I construct a simple real-time measure of disagreement and show that it improves short- and medium-horizon forecasts in standard macroeconomic VARs. Second, using lag-augmented local projections, I show that disagreement dampens monetary policy effectiveness, attenuating its impact on output and offering an explanation for apparent anomalies such as the "activity puzzle." Third, I demonstrate that market expectations, as reflected in tradable asset prices, outperform Fed staff forecasts, and that information-advantage tests reveal no systematic Fed edge, challenging prevailing narratives. I show that standard New Keynesian frameworks with incomplete information contain a latent disagreement structure, which, when formalized, yields a parsimonious explanation of the documented evidence. The model generates disagreement from three distinct sources — fundamentals, policy rule, and credibility. A general state-space model identifies the importance of each and sheds light on key normative implications.

Keywords: Expectations, Monetary Policy, Disagreement, Credibility

Email: famodeo@ucsd.edu Webpage: francescoamodeo.com Dashboard: mindthegap.com

I am indebted to James Hamilton for his invaluable guidance throughout my Ph.D. journey. I am grateful for inspiring discussions with Allan Timmermann, Nir Jaimovich, Johannes Wieland and Giacomo Rondina for insightful comments by Edoardo Briganti, Jake Chaquiça, Paula Donaldson, Domenico Giannone, Francesco Giovanardi, Juan Herreño, Anna Lipinska, Michael Melvin, Thomas Mertens, Paolo Surico, Ali Uppal, Chris Wolf, Johnathan Wright and seminar participants at UC San Diego, Rady School of Management, UCSD Global Policy School, Federal Reserve Board, Johns Hopkins University, Bank of Italy, Bank of Spain, CEMFI, Midwest Macro Meetings 2025, 2025 Annual EAYE Conference, 2025 South Cal International Macro-Fin Conference, 2025 IMIM Rising Stars Series, and the joint ECB-FRB IRFMP Conference. All errors my own.

1 Introduction

In the wake of the Great Financial Crisis, on January 27–28, 2009, the Federal Reserve held its first scheduled Federal Open Market Committee (FOMC) meeting after the policy rate hit the zero lower bound. In an attempt to reassure distressed markets, the Fed announced that interest rates would remain near zero, emphasizing that "*[…] the Committee continues to anticipate that economic conditions are likely to warrant exceptionally low levels of the federal funds rate for some time.*"[1] Nevertheless, at the end of the trading day, market expectations reflected in federal funds futures appeared to anticipate a steady but immediate tightening, revealing a sharp gap with the Fed's outlook. This paper refers to this gap as "Fed–Market interest rate disagreement" and studies how it systematically shapes the transmission of monetary policy.

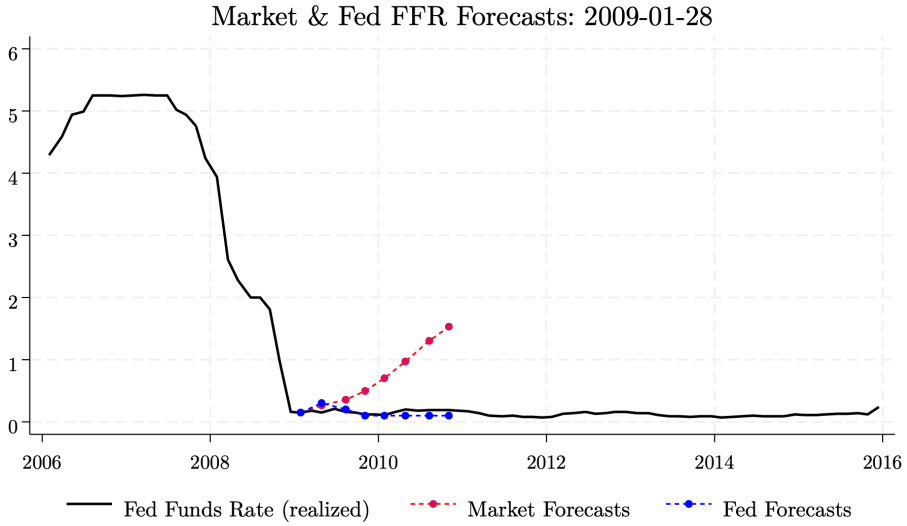


Figure 1: Black solid line —effective federal funds rate. Red dashed — meeting-day market-implied expectations from federal funds futures. Blue dashed — Fed policy-rate expectations from the Tealbook. Dots mark quarterly horizons.

Monetary transmission is central to modern macroeconomics. Over the last few decades, economists and policymakers have increasingly focused on the role of expectations in generating monetary non-neutrality. Workhorse macroeconomic models operate under the assumption that financial markets and the Fed share information and form expectation similarly, implying a common policy outlook. I show that this is rarely true. Fed and market interest rate forecasts disagree throughout 1990–2019, the gaps are serially correlated, and absolute differences at times exceed 250 basis points. These are some of the facts that motivate the

[1] Excerpt from the FOMC Statement, dated January 28, 2009.

analysis in this study.

The paper makes six contributions. First, I construct a novel real-time measure of Fed–Market disagreement at FOMC-meeting frequency, which spans all available forecasting horizons. To my knowledge, this is the first real-time Fed–Market disagreement series at FOMC-meeting frequency which combines Fed staff's projections with traded futures aligned to the Tealbook cutoff date, with the intent to measure the forward-path state variable relevant for monetary transmission. Second, I show that incorporating this novel measure into standard VARs improves the forecasts of macroeconomic variables at short and medium horizons, indicating valuable predictive content for policy decisions and macro aggregates. Third, using state of the art estimation, I demonstrate that these large, persistent expectation differentials attenuate – and sometimes reverse – the conventional effects of monetary shocks, generating state-dependent policy transmission with attenuation increasing in disagreement. Fourth, forecast-comparison and information advantage tests find no systematic Fed edge, undermining a prominent narrative that postulates an information channel as primary explanation for some puzzling evidence. Fifth, I provide a microfounded mechanism: in a generalized New Keynesian framework, the expectation gap enters the Euler equation as a wedge, which can explain the empirical evidence and offers a parsimonious theory of monetary transmission grounded in conventional intertemporal substitution and disagreement-dependent learning. Sixth, I endogenize disagreement and model learning through an Extended Kalman Filter that decomposes its sources into disagreement about fundamentals, rule misperception, and forward guidance, yielding a quantitative attribution assessment of their relative contribution and providing a real-time index of Fed's policy effectiveness.

The core empirical finding estimates the state-dependent effect of Fed-Market disagreement on the transmission of monetary shocks. Using lag-augmented local projections, I find that when disagreement is low, tightening shocks reduce activity with familiar hump-shaped dynamics, consistently with standard theoretical predictions. As disagreement rises, the response of industrial production attenuates, and at high levels it can reverse sign, offering a beliefs-based rationale for the so called "real activity puzzle.", i.e., the observed positive response of activity to high-frequency hiking shocks (Campbell et al., 2017). Quantitatively, a 25-bp increase in average disagreement – a shift from the median to the 75^{th} percentile – reduces the magnitude of the output response at a horizon of 12 months by approximately 60%. In summary, disagreement induces a pronounced state-dependent attenuation, which is pervasive over the last 30 years.

These patterns naturally raise the question of whether monetary announcements reflect information effects. The prevailing view holds that markets' information sets are a subset of the Fed's, due to access to privileged data and superior knowledge of the reaction function. This underpins the "information channel" of monetary policy, the conjectured mechanism under which markets read a policy surprise as news about fundamentals rather than an unambiguous deviation from the rule

(Nakamura and Steinsson, 2018). Under this view, an unexpected rate hike may reflect either a true policy shock or the Fed's expectation of stronger future fundamentals, invalidating empirical specifications that do not distinguish between the two, and yielding puzzling findings like the counterintuitive response of real activity.

A corollary of this view is that the Fed's larger information set should translate into superior forecasting performance. I test this by assembling projections from Tealbooks[2] and federal funds future contracts. In contrast, I find that markets forecast the policy path more accurately than the Fed at horizons of 1–8 quarters. The edge mildly declines with the forecasting horizon but remains statistically significant and robust to variations in the measuring cutoff date. To ensure a valid comparison, it is paramount to ensure that expectations are captured at a common point in time: in fact, using market expectations after the FOMC announcements and the press releases would grant traders extra days of news, beyond the Fed's officially and unofficially released data (Vissing-Jorgensen, 2020). Therefore, for this exercise, I measure market expectations at each Tealbook internal cutoff date, which on average precedes the FOMC meeting by 6.6 days.

Moreover, I apply state of the art testing to this newly constructed dataset to assess the presence (and the evolution) of any Fed informational advantage. Previous studies find substantial evidence of Fed's edge, although decreasing over time. Leveraging the timeliness and unambiguity of my market-implied data, I find no evidence of a systematic information advantage, with minor exceptions around specific episodes concentrated around the zero lower bound era. These results align with, yet are distinct from, the finding of market superiority in policy-rate forecasting, and together they cast doubt on the plausibility of the information channel of monetary policy, thereby motivating the introduction of a more general theoretical framework.

I show that disagreement is a latent state in New Keynesian frameworks, one that can be formally accounted for by relaxing highly restrictive assumptions like full information. As a result, the gap between Fed and market expected policy paths enters the generalized Euler equation as a wedge. Monetary transmission depends on the policy rate path that private agents expect after the announcement, not only on the surprise observed on impact. For instance, if markets expect a smaller increase in future rates than the Fed intends, the rise in expected real rates relevant for spending/investment is smaller, so the contractionary effect on output is attenuated.
I then study the model's impulse responses in low- and high-disagreement states, closely mirroring the empirical design, to show how this mechanism maps into the observed state dependence. The transparency of the framework translates the empirical restrictions into testable implications, which I verify with model-implied regressions: the data exhibit the theory's predicted differential expectation behavior

[2]Informal name for the *Current Economic and Financial Conditions* report, prepared by the Federal Reserve Board's staff before each meeting.

across disagreement states. In summary, treating disagreement as an endogenous wedge allows to track its own dynamic responses, providing a parsimonious and intuitive mechanism confirmed by the data.

Finally, the theory isolates three sources of disagreement: fundamentals, the policy rule coefficients, and the limited commitment of forward guidance. These map, respectively, to the information channel, the response-to-news channel, and the credibility channel, which to my knowledge have not been studied jointly. To operationalize learning, I estimate a state-space model treating all three as latent states through an Extended Kalman filter, mapping observed rate surprises and wedge dynamics into expectation updates. The filter accommodates the policy rule's nonlinearity and delivers time-varying Kalman gains that depend on the prevailing level of disagreement. Intuitively, higher disagreement induce markets to place a larger discount on Fed's actions and communication, so identical surprises shift expectations less and transmission weak. This yields quantitative attribution shares for each source of disagreement and regime-dependent impulse responses that match the documented attenuation. In other words, the same monetary surprise moves the latent states (and, as a consequence, the Euler wedge) by different amounts in low- versus high-disagreement regimes, which can ultimately be considered a real-time proxy of Fed's credibility. Methodologically, the contribution is a rigorous yet flexible learning framework that illustrates markets' state-contingent expectation formation process, unifying shocks, beliefs, and transmission in one system.

Related Literature Woodford (2005) famously argues that *"not only do expectations about policy matter, but… very little else matters,"* echoing the rational expectations idea that only unanticipated monetary changes affect output (Lucas, 1972). This study maintains rational expectations but relaxes full information (Blinder et al., 2008), focusing on the pervasive expectation heterogeneity between financial markets and the Federal Reserve and on how this misalignment shapes monetary transmission.

Theoretically, intertemporal substitution is a core channel of monetary transmission: expectation management and credible commitment have been central since the early New Keynesian models (Clarida et al., 1999, Woodford, 2001), and more recent behavioral frameworks reinforce this view, highlighting communication as a key lever of policy (Sims, 2003, Lorenzoni, 2009, Eusepi and Preston, 2010, Gabaix, 2020). A contemporaneous and complementary contribution by Jimenez-Gimpel (2026) studies optimal central bank communication when markets are ambiguity averse over future policy rates. Closest to this paper, Caballero and Simsek (2022) model Fed–market disagreement and study policy implications when markets are opinionated and information effects are muted. In contrast, I first establish the empirical prevalence and state-dependent relevance of disagreement for transmission; then show it enters a generalized New Keynesian framework as a wedge that explains attenuation; finally, I decompose its sources, nesting leading theories.

Empirically, extensive work identifies monetary shocks and their effects (Ramey, 2016), in particular using narrative instruments and high-frequency surprises around announcements (Kuttner, 2001, Romer and Romer, 2004, Bernanke and Kuttner, 2005, Gurkaynak et al., 2005, Ramey, 2016). Measurement has sought to separate policy from information effects. Recent work decomposes market reactions into monetary and information shocks (Melosi, 2017, Nakamura and Steinsson, 2018, Jarociński and Karadi, 2020), helping explain "puzzles" common in high-frequency specifications. Under imperfect information, announcements correlate with forecast revisions, generating positive activity responses to tightening that reflect news rather than causal policy effects. Relatedly, studies analyze the heterogeneous perceptions of the policy rule and their implications for belief dispersion and outcomes (Hamilton et al., 2011, Couture, 2021, Bauer et al., 2024, Sastry, 2026).

Predictability of external instruments ("monetary surprises") has motivated orthogonalization using surveys and Fed projections (Cieslak, 2018, Miranda-Agrippino and Ricco, 2021, Bauer and Swanson, 2023a, Amodeo, 2025). This paper complements that agenda by studying the dynamics of disagreement — the gap between Fed and market expectations around the FOMC statement and the press conference — providing a clean, price-based proxy for residual differences in beliefs that can be studied in relation to monetary transmission and efficacy.

Ex post measurement also incorporates Fed communication, now a core policy instrument. By steering beliefs about the future policy path, communication shapes intertemporal choices of firms and households (Eggertsson et al., 2003). Markets may read guidance as signals about fundamentals ("Delphic") rather than pure commitments ("Odyssean") (Campbell et al., 2012, Melosi, 2017). In this setting, disagreement represents limited credibility, constraining the Fed's ability to influence policy expectations (Moscarini, 2007).

The information channel posits that policy announcements reveal the Fed's private knowledge about fundamentals, so markets interpret surprises as information rather than pure unexpected policy (Nakamura and Steinsson, 2018, Andrade et al., 2019, Cieslak and Schrimpf, 2019). More broadly, the Fed may possess an information advantage regarding the future path of its policy rule or planned deviations from it, which manifests as superior policy-rate forecasts. Early studies report a Fed edge (Romer and Romer, 2000, Gavin and Mandal, 2003), while later work finds erosion and instability by sample and horizon (D'Agostino and Whelan, 2008, Hoesch et al., 2023).

My results differ sharply. Markets consistently outperform the Fed in the short and medium run; moreover, state of the art tests show no systematic Fed information advantage, with only episodic exceptions. Two key reasons justify this divergence with previous results. First, I compare Tealbook projections against market-implied expectations (prices), avoiding notable survey issues, like cheap talk (Crawford and Sobel, 1982), behavioral frictions(Bordalo et al., 2020, Kohlhas and Walther, 2021, Bianchi et al., 2022), and strategic bias (Gemmi and Valchev, 2023, Amodeo, 2024). Second, traded futures prices allow precise alignment of mar-

ket expectations to the Tealbook's internal cutoff, eliminating vintage mismatches that plague survey comparisons (Vissing-Jorgensen, 2020). Together, these findings challenge the information-channel premise as the dominant explanation for announcement responses.

More broadly, the paper relates to the literature on disagreement in measured expectations (Mankiw et al., 2003, Andrade and Le Bihan, 2013, Andrade et al., 2019). In a similar vein, Dong et al. (2025) show that household inflation disagreement weakens forward guidance and dampens the effects of monetary shocks, consistent with uncertainty about the inflation target. On the household side, Jeong et al. (2026) offer a complementary perspective, finding that the pass-through of policy news to consumers' inflation expectations is mediated by endogenous attention. Angeletos and Lian (2018) argue that without common knowledge, forward-guidance effects are mitigated through higher-order beliefs, providing a complementary mechanism for the attenuation documented here. On transmission, Zohar (2024) shows that disagreement and uncertainty behave differently across the business cycle, illustrating why disagreement may capture different forces beyond uncertainty; relatedly, Benchimol et al. (2023) use a high-frequency approach to show that uncertainty amplifies stock-market reactions to monetary surprises. Both papers treat uncertainty as a distinct state variable shaping macro and financial responses; I follow this distinction, keeping disagreement conceptually separate and results empirically robust by controlling for standard macro and financial uncertainty measures. The rest of the paper is organized as follows. Section 2 describes the data used to compute shocks and disagreement. Section 3 outlines all key empirical results, first describing the main properties of my disagreement measure and then showing its effects on monetary transmission, with a rich discussion of the findings. Section 4 introduces a theoretical framework to model disagreement within a generalized New Keynesian framework with heterogeneous expectations, and leverages the core empirical restriction to test the model. Finally, Section 5 proceeds to microfound the sources of disagreement in relation to the learning implied by the model, and outlines a methodology to quantitatively assess the expectation formation process of financial markets in an environment with private information, forward guidance, unobserved fundamentals and policy rule coefficients. Section 6 provides concluding remarks and directions for future research.

2 Data

Throughout the paper, the analysis draws on three data blocks. First, the standard macroeconomic series used as outcomes and conditioning variables. Second, the instruments for the monetary shocks, primarily based on high-frequency surprises around FOMC announcements. Third, the expectations for the Fed, financial markets, and—where relevant—professional forecasters. Variable definitions, transformations, and sources are documented with each application and consoli-

dated in Appendix A. The next subsections describe the core inputs and the novel construction choices that underlie the empirical analysis.

2.1 Measuring Fed's Expectations

The main sources for the Federal Reserve's expectations are the Tealbooks (previously Greenbooks), distributed to the Federal Open Market Committee prior to each scheduled FOMC meeting, and containing in-depth analysis of current economic and financial conditions and projections of the Federal Reserve Board staff. They are released to the public with a five-year lag; as of this writing, the most recent release covers the December 2019 meeting. I construct a time series of Fed forecasts produced for each FOMC meeting based on digitized data from the Federal Reserve Bank of Philadelphia, which provides expectations for interest rates from 1981 to 2008, and on hand collected data for the predictions made between 2008 and 2019. Figure 2 showcases an exemplificative snapshot of a recent Tealbook release.

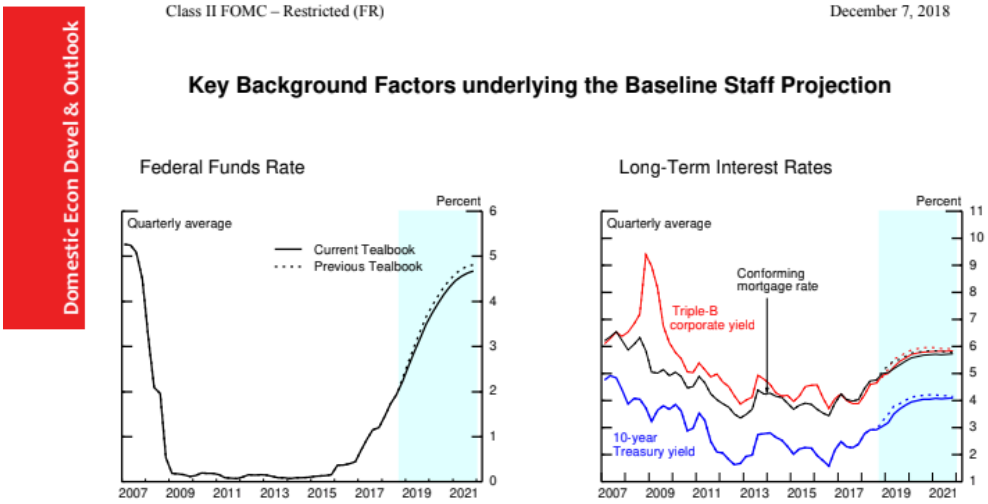


Figure 2: Example Greenbook/Tealbook projection of the Federal Funds Rate (FFR) path used for an FOMC meeting (vintage indicated); Source: Federal Reserve Board, Tealbook.

The precise dates used are public, obtained from Nakamura and Steinsson (2018), recently extended by Acosta et al. (2024). The key variable under study, the Federal Funds Rate (FFR), is forecasted up to several years in the future, and I follow Caballero and Simsek (2022)'s methodology in interpolating the annual Tealbooks forecast to quarterly horizons up to ten quarters ahead.[3] Importantly, the Tealbooks projections are prepared a few days in advance of the scheduled meetings. To compare the forecast accuracy between the Fed and financial markets,

[3]See Appendix A.2 for a more detailed account.

the precise temporal alignment of information sets is of the essence. Therefore, I manually collect each exact day differential between the internal release of the Tealbooks and the FOMC Meeting it refers to; across the 233 scheduled monetary decisions in the sample, the average gap is 6.6 calendar days.[4]

For auxiliary analysis, I collect additional data from the Fed's Summary of Economic Projections (SEP), which started including FFR projections since 2012. However, several reasons limit the SEP's scope for the purpose of this study: first, the much shorter sample; second, the reduced frequency (four releases yearly); moreover, its being a collection of all FOMC participants' projections has raised questions regarding their nature of *predictions* or *intentions*, and other reservations of political type (Gerlach and Stuart, 2019, Bordo and Istrefi, 2023); relatedly, the Committee's published projections are better understood as scenarios guarding against model misspecification than as its central expectation (Ellison and Sargent, 2012).[5]

More fundamentally, the Tealbook is the forecast on which the policy decision is conditioned, rather than an estimate that ambiguously targets the most likely or the desired path. The staff projection circulated to the Committee before each meeting is the real-time information set the FOMC brings to the decision (Romer and Romer, 2004), and the forecast the literature uses to represent the Fed's reaction function (Orphanides, 2001, Bauer and Swanson, 2023b). Its release with a five-year lag does not preclude its use as the Fed's conditioning measure, since the empirical design conditions on disagreement predetermined at the meeting.[6]

2.2 Measuring Market's Expectations

To measure markets' expectations, I use daily federal funds rate (FFR) futures prices for maturities up to 36 months obtained from proprietary Bloomberg data.[7] These contracts are highly liquid, widely used contracts: in recent years, daily trading has typically amounted to several hundred thousand contracts, corresponding to current average notional volumes on the order of $1–2 trillion per day across maturities. Available data from traded futures spans from 1988 to the present, although with remarkable heterogeneity in trading horizon, as early in the sample futures priced only 2 to 3 quarters ahead FFR realizations. I collect all available

[4] Further details in A.2.

[5] Appendix D.7 re-estimates the headline local-projection result with SEP-implied disagreement in place of the Tealbook-based measure and shows that the key documented channel is qualitatively preserved both on the 2012:01–2019:12 overlap and on the extended SEP sample through 2023:12; a pooled-overlap Wald test corroborated the conclusion.

[6] Two features reinforce this choice. Romer and Romer (2008) find that the optimal combination of the staff and the Committee's projections places almost all weight on the former. Operationally, the Tealbook is produced for every scheduled meeting on a dated cutoff before the decision, whereas the SEP appears only quarterly, is released after the meeting, and reports year-end values.

[7] Bloomberg tickers: `"FF1 Comdty"` to `"FF36 Comdty"`.

closing prices from 1988 to 2024, quoted as 100 minus the average monthly FFR at settlement, and convert to implied forward rates following the literature via

$$\text{Forward Rate}_{t,m} = 100 - \text{Futures Price}_{t,m},$$

where t indexes trading days and $m \in \{0, \ldots, 36\}$ the monthly horizon. Forward rates are aligned with the format of Federal Reserve's Tealbook projections, e.g., they are aggregated to quarterly horizons (henceforth denoted by h) by averaging. Figure 3 displays an illustrative snapshot of a future contract as shown in Bloomberg Terminal, where column `Last` collects the Futures Price$_{t,m}$ in the formula above.

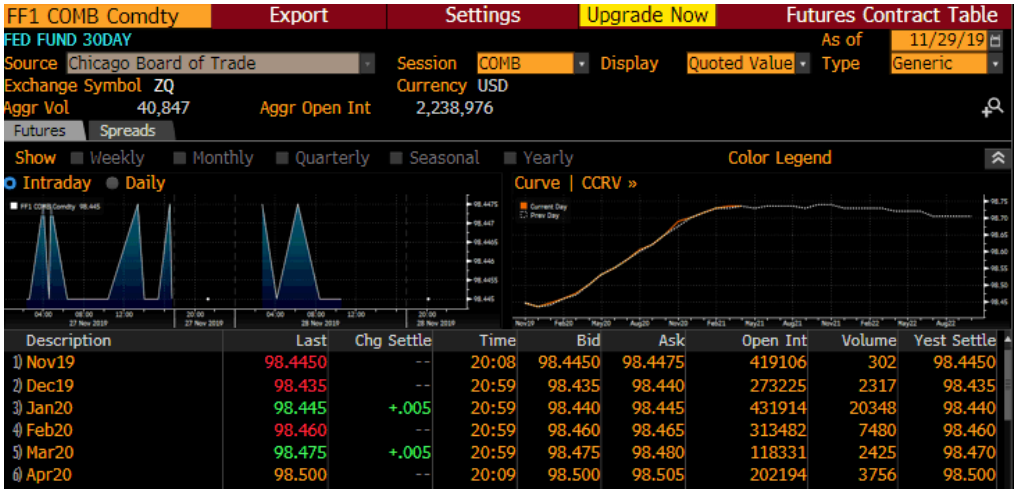


Figure 3: Exemplificative Bloomberg Terminal view of 30-Day Federal Funds futures — Generic 1st continuous contract (`FF1 COMB Comdty`).

Extracting expectations from traded securities insures against several concerns regarding survey forecasts, from *cheap talk* critiques (Crawford and Sobel (1982), Ottaviani and Sørensen (2006)) to the many documented biases professional forecasters and household exhibit (Laster et al. (1999), Malmendier and Nagel (2016), Bordalo et al. (2020), Gemmi and Valchev (2023), Amodeo (2024)). However, financial prices embed a risk premium potentially causing them to diverge from actual expectations (Piazzesi and Swanson (2008)).[8] Estimates of such premia are very small compared to the disagreements I document and are even more negligible for short-term interest rate futures (Schmeling et al. (2022)). To allow for this small discrepancy, in robustness I adopt commonly employed adjustments (Piazzesi and Swanson (2008), Diercks and Carl (2019)) – also in use at the Board of Governors – to adjust the extracted expectations.

Finally, although unsuited for some of the purposes in this study, I also analyze

[8] In Appendix E.1, I show a theoretical microfoundation of prices as (rational) measures of expectations for agents endowed with constant absolute risk aversion preferences.

the expectations proxied by the consensus estimates (simple average) of the Survey of Professional Forecasters. Although the SPF does not include forecasts of the overnight interest rate among its variables, I proxy for it by using the 3-month US Treasury Bill rate.

2.3 Monetary Surprises

To proxy for monetary shocks, I use several influential high-frequency surprise series based on traded securities' price variation around monetary decisions. The baseline shock series is an evolution of the high-frequency orthogonalized monetary policy surprises of Bauer and Swanson (2023b), as proposed in the parallel study Amodeo (2025). The data cover all FOMC announcements from 1989 to 2023, and calculate surprises ($\texttt{MPS}_t$) as the first principal component of shocks to the one, two, three, and four-quarter Eurodollar contracts in 30-minute windows around monetary announcements, consistently with evidence highlighting the relevance of longer-term expectations in financial markets (Gurkaynak et al., 2005).

To address recent concerns of predictability in high-frequency surprises, I orthogonalize each $\texttt{MPS}_t$ with respect to latest available release of six macrofinancial variables on the same day of the monetary announcement.[9] Following Amodeo (2025), orthogonalization takes place at the observation (meeting) level and then the purged series is aggregated to monthly frequency, so that shocks are ex ante orthogonal to information available at the time of the decision (i.e., they are unpredictable). By contrast, Bauer and Swanson (2023a) aggregate first and orthogonalize afterwards. The resulting series, $\texttt{MPS_ORTH}_t$, constitutes the baseline instrument for the monetary shocks throughout, although a range of alternative popular series – including the original Bauer and Swanson (2023a) – is used for robustness at every stage.

[9] These being the variables described in Bauer and Swanson (2023b): the surprise in nonfarm payrolls, employment growth, a measure of S&P 500 growth, the change in the slope of the yield curve, the log change in the Bloomberg Commodity Spot Price index, and the implied skewness of the 10-year Treasury yield. Details in Appendix A.3

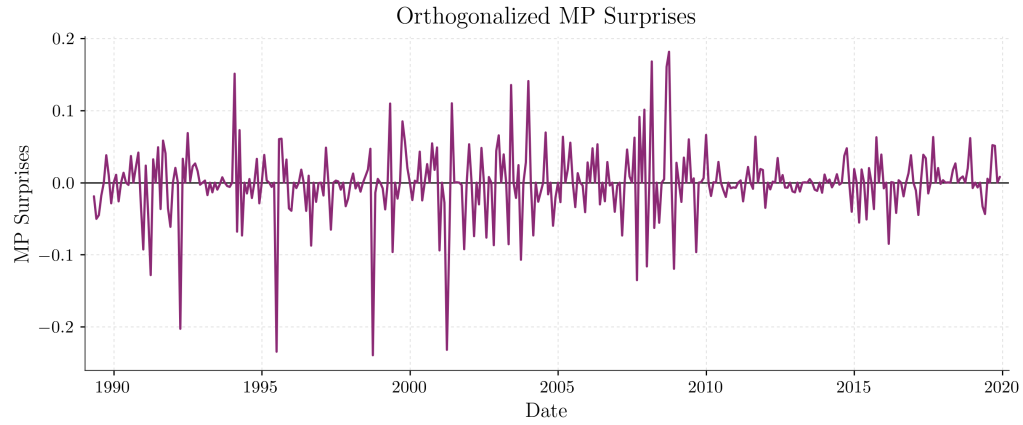


Figure 4: Orthogonalized high-frequency monetary surprises (`MPS_ORTH`) from Bauer and Swanson (2023b), Amodeo (2025).

3 Empirical Results

The main empirical findings are organized in three parts. First, I provide descriptive evidence of my measure of disagreement, defined as the difference between the Fed's and the market's expectations of the future path of the Federal Funds Rate, and argue that adding it to standard VARs improves the forecasting of macroeconomic aggregates and the policy rate.

Second, and crucially, I study the consequences of disagreement for monetary policy transmission. I find that disagreement attenuates the effects of monetary shocks on the economy, as I show for the response of output (industrial production) as a function of disagreement levels. Numerous variations on the empirical specification corroborate the relevance of this novel measure in explaining state-dependent effects of monetary policy and known outstanding puzzles.

Finally, I compare statistically the forecasting accuracy of Fed, Markets, and (in appendix) the consensus of the Survey of Professional Forecasters in predicting the future monetary policy. In contrast to the prevailing narrative and to previous results, I find that the market overperforms the Fed at all forecasting horizons, and, distinctly but consistently, that the Fed does not possess systematic information advantage throughout the sample, with episodic exceptions concentrated around the zero lower bound era.

3.1 Disagreement About Monetary Policy

For each forecast horizon h, define disagreement

$$D_t^h := \mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^M(i_{t+h}) \tag{1}$$

where i_{t+h} denotes the federal funds rate h quarters ahead, $\mathbb{E}_t^F$ is the Fed's expectation (from the Greenbook/Tealbook) and $\mathbb{E}_t^M$ is the market's expectation (implied from traded futures). Note that by construction (1) is signed and order-sensitive — swapping Fed and Market flips the sign; in most of the empirical analysis, I use the average across forecast horizons of its absolute value.

As a summary of differences between Fed and market expectations, I compute average disagreement simply as the arithmetic mean over the available horizons of D_t^h. Table 1 summarizes moments and regime means. Disagreement is slow-moving and highly persistent. Its mean is close to zero and statistically indistinguishable from no bias; the sign is roughly balanced over the sample, with a slightly positive tilt — the Fed's path exceeds markets in about 47% of months. Mean absolute disagreement is higher in recessions than in expansions, but the recession–expansion gap is imprecisely estimated and not statistically significant. Overall, disagreement is at most weakly countercyclical in magnitude and shows negligible cyclicality in sign.

Panel A. Unconditional moments				
	Mean	SD	AR(1)	$Pr > 0$
$\bar{D}_t$	0.058	0.490	0.925	0.465
$\lvert\bar{D}_t\rvert$	0.383	0.310	0.866	1
Panel B. Means by business-cycle regime				
	Exp. ($N = 333$)	Rec. ($N = 37$)	Rec−Exp	p-val
$\bar{D}_t$	0.065	-0.005	-0.069	0.734
$\lvert\bar{D}_t\rvert$	0.375	0.454	0.079	0.275

Table 1: Disagreement: Unconditional Moments and Business-Cycle Means.
Notes: Monthly data, 1988:10–2019:12. Units: percentage points. "Rec−Exp" and p-values from Newey–West regressions with 12 lags. $\bar{D}_t$ is the unweighted mean of $D_t^h = \mathbb{E}^F(i_{t+h}) - \mathbb{E}^M(i_{t+h})$; $|\bar{D}_t|$ is its absolute value. NBER recession dates.

Figure 5 plots the time series of the average D_t^h over all available horizons throughout 1988:10–2019:12, with NBER-dated recessions shaded. The series spends substantial time in both the top and bottom quintiles of its distribution, in episodes that map cleanly to recognizable moments in monetary-policy history. A top-quintile episode (in blue) is one in which the Federal Reserve's expected path of policy rates sits above what financial markets are pricing. The defining example is the long stretch covering the post-liftoff normalization cycle, when the Fed economists' projections persistently exceeded the path implied by futures: the Fed believed it would raise rates faster than markets believed it would. Symmetrically, a bottom-quintile episode (in red) is one in which markets price a steeper path than the Fed. The 1994 episode — discussed at length in Blinder (2023) — is the canonical case: bond markets, skeptical of Greenspan's success in achieving a "soft landing", priced continuing hikes while the Fed welcomed the promising

inflation data and stabilized its rate forecasts. The same logic applies through the post-2008 zero lower bound, where the Bernanke "lower-for-longer" commitment to keep rates near zero was repeatedly contested by markets pricing earlier exits. Table D1 in Appendix D.5 lists the full set of ten such episodes—four Fed-hawkish and six Market-hawkish—each lasting at least six months, together with the specific macroeconomic or policy source of the Fed–market gap in each case.

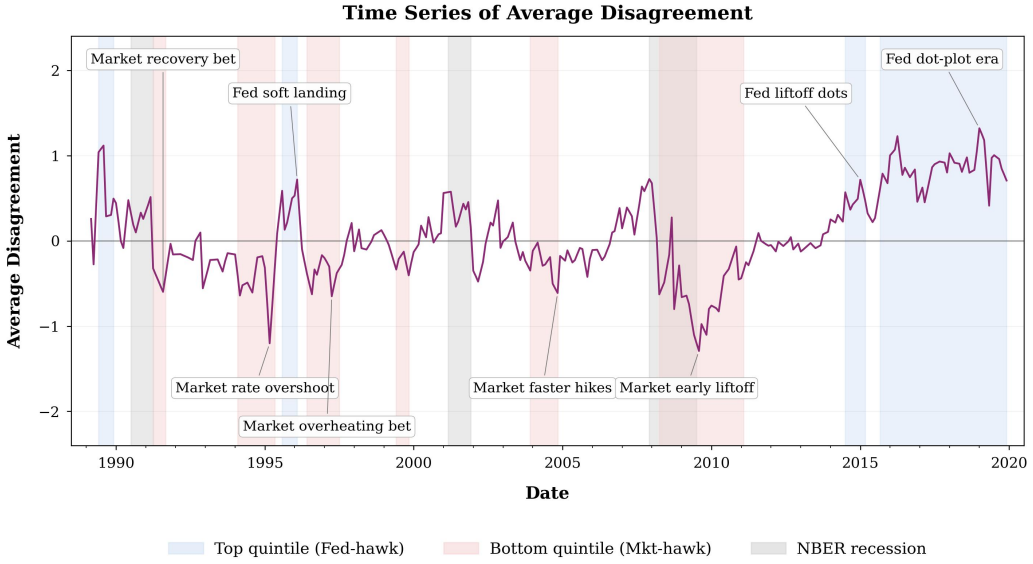


Figure 5: Time Series of Average Disagreement, 1988:10–2019:12. Pale blue (pale red) bands flag contiguous runs of at least 6 months in the top (bottom) signed quintile of $\bar{D}_t$, allowing 4-month gaps to bridge brief reversions. Gray bands mark NBER recession dates.

3.2 Predictive Power of Disagreement

As previously argued, disagreement between markets' and Fed's expectations is pervasive and likely reflect differential views on fundamentals, policy rules, or future behavior of the FOMC. Then, these frequent divergences allow me to construct a measure of disagreement that may carry information for future macroeconomic outcomes. This section shows that such measure is indeed helpful in predictive exercises in commonly specified macroeconomic vector autoregressions.

As a baseline, I estimate VARs with three variables – industrial production, consumer prices and the policy rate — and perform block Granger causality tests to establish if the newly obtained disagreement variable statistically improves the forecasting of any of the specified variables. Then, the null hypothesis H_0 is that all lag coefficients on $D_t^h, \ldots, D_{t-p}^h$ are zero. Rejection implies that p lags of D_t^h improves conditional forecasts of $[\mathrm{IP}_t, \pi_t, i_t]$.

Table 2 reports block F-statistics for horizons of one to eight quarters. The null

that disagreement (and its lags) does not Granger-cause macroeconomic aggregates is rejected at the 5% significance level across these horizons. At horizons seven and eight quarters, the F-statistics fall and the null cannot be rejected, indicating that the predictive power of disagreement weakly fades beyond about a year and a half.

More precisely, adding D_t^h yields non-trivial relative gains. Measured as the percent drop in unexplained variance, for i_t it peaks near 15–16% at $h = 5$–6; for π_t about 9–13% at $h = 5$–8; for IP_t about 4–8% at $h = 3$–6. This aligns with the block-Granger rejections.[10]

Null: H_0	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude $\{D^1_{t,\ldots,t-p}\}$	F(21,335)	2.468	0.000	Reject
Exclude $\{D^2_{t,\ldots,t-p}\}$	F(18,340)	2.261	0.003	Reject
Exclude $\{D^3_{t,\ldots,t-p}\}$	F(24,305)	2.224	0.001	Reject
Exclude $\{D^4_{t,\ldots,t-p}\}$	F(18,255)	2.852	0.000	Reject
Exclude $\{D^5_{t,\ldots,t-p}\}$	F(30,148)	2.065	0.002	Reject
Exclude $\{D^6_{t,\ldots,t-p}\}$	F(30,141)	1.960	0.005	Reject
Exclude $\{D^7_{t,\ldots,t-p}\}$	F(30,141)	1.417	0.092	Fail to Reject
Exclude $\{D^8_{t,\ldots,t-p}\}$	F(27,139)	1.315	0.156	Fail to Reject
Exclude $\{\bar{D}_{t,\ldots,t-p}\}$	F(18,340)	2.308	0.002	Reject

Table 2: Block Granger Causality: 1-8 quarters, and Average; Null H_0: $\Phi_{n1,p} = 0\ \forall p$ (no predictive power) at $\alpha = 0.05$ significance level. Based on an estimated VAR with a constant including $[FFR, D_t, CPI, IP]$, where D_t is either a vintage D_t^h or the summary measure $\bar{D}_t$. Number of lags is picked by the most conservative between the Akaike and Schwartz Criteria ($p = \max(AIC, SBC)$). See Appendix C for specifics on the details of the estimation.

Consistently with the rest of the evidence, Block Granger causality tests using $\bar{D}_t$ as a regressor reject the null that past values of average disagreement do not help predict macro variables. The average measure too, therefore, provides a useful addition in a reduced form forecasting exercise. Thus, I conclude that measures of disagreement between Fed and market expectations carries predictive content for monetary policy and real activity over short to medium horizons.

While Block-Granger causality can be a useful tool to establish the general informativeness of the various measures of disagreement, univariate testing raises the bar and can shed light on the specific variables that disagreement relates to. Appendix C provides univariate test results, details the mechanics of the estimation and includes a battery of robustness checks, including the estimation of

[10] To quantify gains, I use the partial R^2 for nested models:

$$\text{Gain} = \frac{R^2_{\text{aug}} - R^2_{\text{base}}}{1 - R^2_{\text{base}}} = 1 - \frac{\text{SSE}_{\text{aug}}}{\text{SSE}_{\text{base}}}.$$

This measures the percent drop in unexplained variance.

richer VARs like Gertler and Karadi (2015), which augments the system using Gilchrist and Zakrajšek (2012)'s measure of Excess Bond Premium, or Bauer and Swanson (2023b)'s even richer 6-variables system including indexes of the price of commodities and unemployment. Results remain qualitatively consistent.

3.3 Disagreement and the Transmission of Monetary Policy

After demonstrating the reduced-form usefulness of a measure of disagreement between the Fed and financial markets for macroeconomic forecasting, in this section I present the central empirical result of the paper. Does interest rate disagreement affect the transmission of monetary policy to the economy?

To gauge whether differences in expectations affect the pass-through of the Fed's policy, I estimate lag-augmented local projections in the spirit of Montiel Olea and Plagborg-Møller (2021). For each horizon h, the response of a macro variable y to a monetary policy shock s_t is obtained from the following specification:[11]

$$\mathbf{y}_{t+h} = \mathbf{c}_h + \boldsymbol{\alpha}_h s_t + \boldsymbol{\beta}_h s_t |\bar{D}_t| + \boldsymbol{\theta}_h \text{LAGS}_{t-1}^{L} + \boldsymbol{\varepsilon}_{t+h} \tag{2}$$

where $\text{LAGS}_{t-1}^{L} \equiv \big[\, \mathbf{y}_{t-1:t-L},\ s_{t-1:t-L},\ |\bar{D}_{t:t-L}|,\ s_{t-1:t-L}|\bar{D}_{t-1:t-L}| \,\big]'$, with respective coefficients $\boldsymbol{\theta}_h \equiv [\, \phi_h^{1:L},\ \gamma_h^{1:L},\ \delta_h^{0:L},\ \zeta_h^{1:L} \,]$; $\mathbf{y}_{t+h}$ collects the macroeconomic variables under analysis, s_t is a monetary policy shock, and $|\bar{D}_t|$ is the average across forecast horizons of the absolute Fed–market disagreement, which I use to capture *symmetric* effects of the magnitude of disagreement, below discussed and relaxed. I include L lags of both the independent and control variables, the shock, and the measure of disagreement, according to the latest guidance in Plagborg-Møller and Wolf (2022). Notice that in this exercise disagreement is measured before the FOMC announcement, using expectations formed *prior* to the meeting so that the state of beliefs is predetermined with respect to the surprise component of the decision. This allows me to interpret the interaction coefficients as state dependence of monetary transmission with respect to ex-ante Fed–market misalignment, without post-treatment contamination from information released at the FOMC itself.[12]

The main coefficients of interest are $\boldsymbol{\alpha}_h$ and $\boldsymbol{\beta}_h$. The former captures the response of the dependent variable to a positive shock s_t in the absence of disagreement ($|\bar{D}_t| = 0$); the latter measures the effect of the interaction term $s_t|\bar{D}_t|$, that is, it pins down how the impact of monetary policy depends on the state of alignment between the Fed and financial markets. In other words, $\boldsymbol{\beta}_h$ can be thought of as capturing the disagreement state dependence of monetary transmission to the economy.

[11] For brevity, I use the notation LAGS_{t-1}^{L} to denote the collection of lagged regressors: lags of the dependent and control variables; lags of the shock, disagreement, and interaction term.

[12] Alternative specifications where disagreement is measured *after* the FOMC announcement are included in Appendix D.2.

As a baseline, I estimate the h regressions in equation (2) using a simple monthly 3-variable system including the overnight interest rate, the consumer price index, and industrial production. As for the shocks, I select my own update (Amodeo, 2025) to the orthogonalized high-frequency monetary surprises introduced by (Bauer and Swanson, 2023b), although – as I discuss below (§ 3.3.2) – results are robust to the choice of surprise series. Finally, I select $L = 12$ lags, which is a customary and conservative choice, exceeding the selection of both the Akaike and the Schwartz Information Criteria, and an estimation horizon of three years ($h = 36$). Estimating (2) across horizons allows us to evaluate the same impulse response functions we would obtain from a structural vector autoregression (Plagborg-Møller and Wolf, 2021).

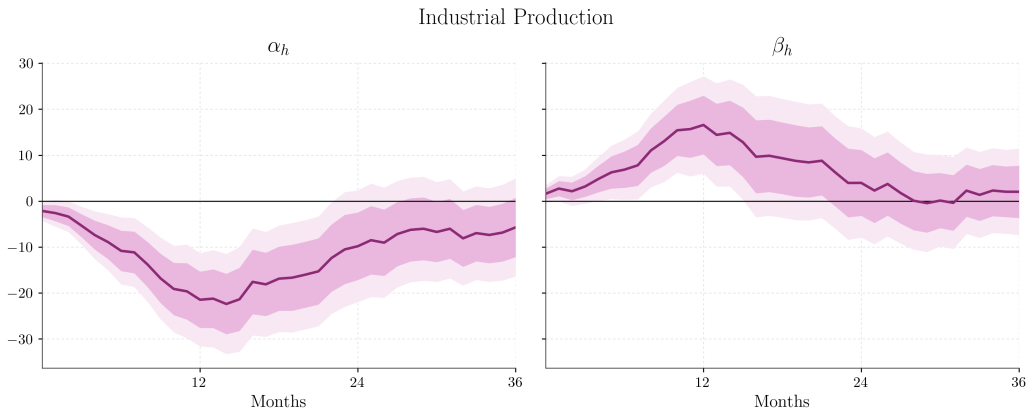


Figure 6: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Figure 6 displays the estimates of $\boldsymbol{\alpha}_h$ and $\boldsymbol{\beta}_h$ over all available observations ($N = 358$) in the sample (1990-2019).

The left panel ($\boldsymbol{\alpha}_h$) shows a highly significant, marked decrease in industrial production starting after a quarter from the impact of the shock. In states of no disagreement between financial markets and the Fed, the response of output has the sign predicted by standard economic theory and its dynamics follow a path that is slightly delayed and reversely hump-shaped, reaching a trough effect at around five to seven quarters after the impulse, mirroring well-known previous results.

The right panel ($\boldsymbol{\beta}_h$) displays the opposite tendency, with a moderately sluggish but statistically significant positive response that reaches its peak after about a year from the impact of the shock. The interpretation of this finding is key: in states of non-null disagreement, the (negative) response of industrial production to a monetary policy shock is *weakened* in a way that is proportional to the degree of misalignment between the interest rate expectations of financial markets and the

Fed. In other words, we observe an *attenuation* of the transmission of monetary innovations to the real economy that is increasing in the level of disagreement.

In this specification, I normalize $|\bar{D}_t|$ to have mean one, so as to be able to immediately compare the magnitudes of the left and right panels. Factoring out the shock s_t in equation (2), this normalization aids interpretability, allowing to obtain the *total average effect* estimated by OLS simply as the sum $\boldsymbol{\alpha}_h + \boldsymbol{\beta}_h$.[13] As is evident at first glance, and as I show graphically in Figure D1 in Appendix D, the total effect is heavily diminished and barely significant after five quarters, consistently with the low power of orthogonalized surprises documented in Amodeo (2025).

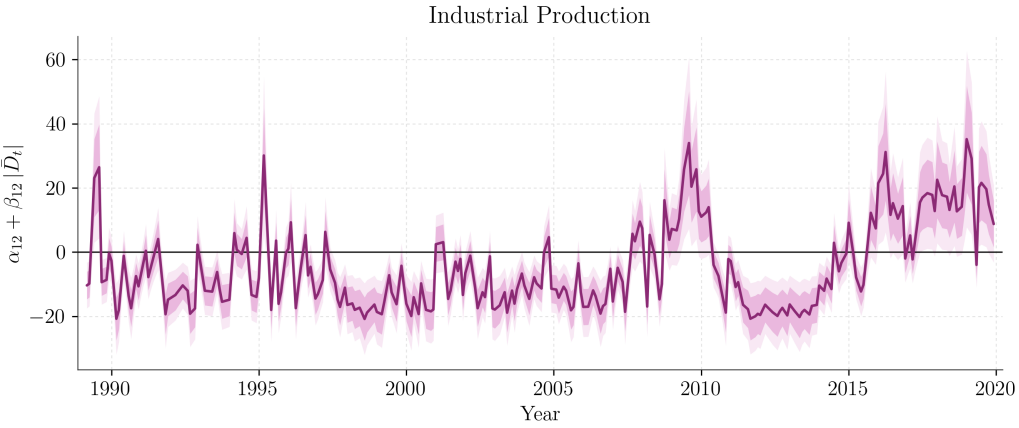


Figure 7: Impulse Response Functions coefficients of Industrial Production to an orthogonalized monetary policy shock evaluated at each $|\bar{D}_t|$ at horizon at $h = 12$, estimated from lag-augmented local projections, as in equation (2). Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Another useful way to visualize the findings emerging from the regressions in (2) is to fix the estimated coefficients at a given horizon h and allow the disagreement level vary in the time series of the sample. I implement this at the natural threshold of $h = 12$, around which both measured effects reach their maximum magnitude.[14] Figure 7 depicts the total average response of industrial production a year after a monetary shock. The occurrences of total *positive* effects are frequent and often statistically significant. In fact, across the sample, they constitute about 31% of all point estimates, with more than 2/3 (1/3) statistically significant at the 68% (90%) confidence level.

Previous unintuitive results capturing expansionary effects of monetary tightenings have been famously dubbed a "real activity puzzle" by the literature. In turn, the state-dependence founded on the degree of Fed-markets disagreement documented

[13] The estimation implements Eicker–Huber–White standard errors to allow valid inference under unknown heteroskedasticity.

[14] I include the results for $h = 6, 18, 24$ in Appendix D.

in Figure 6, can be further explored to assess if the degree of attenuation can offset and even revert the baseline contraction in output in specific episodes.

Figure 8 points to the comparison between the *total* effects on the date of minimum occurrence of disagreement, October 1997, and its maximum counterpart, on August 2009.[15] As expected, states of high disagreement coincide with a reversal of the expected sign of the response, with an expansion of industrial production starting after two to three quarters from impact. These findings suggest an explanation to the activity puzzle based on the differential outlooks on future policy held by financial markets and Fed.

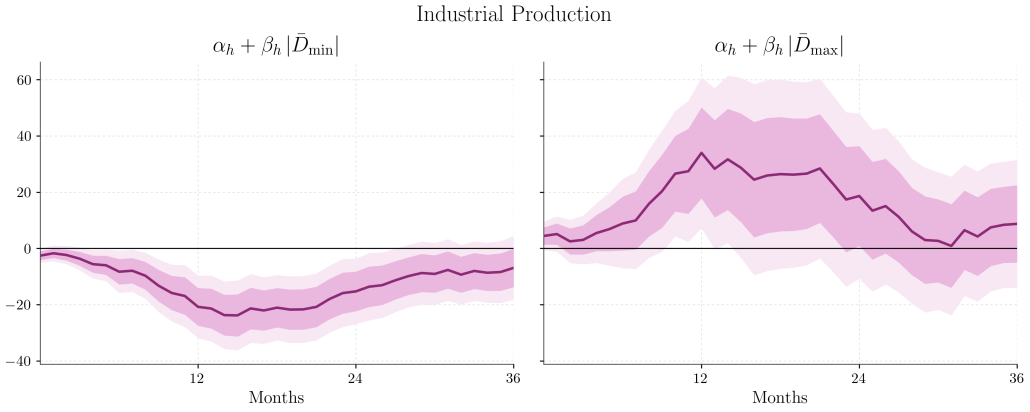


Figure 8: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Left panel: total marginal effect at (second) minimum absolute average disagreement level, $\hat{\boldsymbol{\alpha}}_h + \hat{\boldsymbol{\beta}}_h |\bar{D}_{1997:10}|$; right panel: total marginal effect at (second) maximum absolute average disagreement level, $\hat{\boldsymbol{\alpha}}_h + \hat{\boldsymbol{\beta}}_h |\bar{D}_{2009:08}|$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Another clarifying episode is December 16, 2015 — the first lift-off from the ZLB. The monetary surprise was positive (tightening), yet output rose in response. On that occasion, disagreement about the rate path was positive: the Fed projected a steeper FFR path than markets. In other words, markets priced a less severe incoming hiking cycle, i.e. projecting more *dovish* financial conditions and supporting demand — so the tightening shock was neutralized. Similar sign reversals appear along the sample in both directions of policy move sign (hike/cut) and realized surprise sign (tightening/easing).

In summary, I find evidence of a significant, protracted dampening of the transmission of monetary shocks to the economy depending on the degree of disagreement between financial markets and the Fed. This attenuation reduces the effectiveness of policy in a way that is proportional to the *beliefs gap* on future policy. The pre-

[15] Minimum and maximum computed after excluding ZLB and restricting to disagreement with at least five quarters of data. See Appendix D for more episodes.

dominant interpretation of the mechanism underpinning monetary non-neutrality is founded on different forms of intertemporal substitution by a variety of agents (households, firms, financial intermediaries, etc.).[16] In light of this, it is intuitive to view disagreement as a friction weakening monetary pass-through, at it acts like a wedge between the policy path the Fed intends to communicate and the path private agents use in their spending and investment decisions. Consistent with this view, the same surprises that move output leave the realized policy rate path little changed on average (Amodeo (2025)), suggesting they indeed transmit through revised beliefs about future policy. In Section 4, I will expand on this theme by leveraging those same fundamental structures governing intertemporal substitution in the New Keynesian model. Before doing that, however, I explore the robustness of my findings and control for relevant alternative mechanisms that may confound this narrative.[17]

3.3.1 Disagreement and Uncertainty

Two identification threats are key: cyclical comovement and uncertainty. Disagreement may rise in weak times or during risk spikes, and it may respond to the volatility of the economic environment. In fact, disagreement's endogeneity might invalidate the state-dependent interpretation of the estimates from the regressions in (2).

While the conservative lag-augmentation allows for a robust control of the state of the economy, disagreement seems inherently related to the uncertainty in the economic environment: a more uncertain or volatile outlook may widen the gap between Fed and market forecasts. I therefore augment equation (2) one proxy at a time, interacting the monetary policy shock with (*i*) the VIX, (*ii*) the monthly Baker et al. (2016) Economic Policy Uncertainty (EPU) index, and (*iii*) the cross-sectional dispersion of responses in the Survey of Professional Forecasters (SPF).[18] Let υ_t denote one such uncertainty proxy:

$$\mathbf{y}_{t+h} = \mathbf{c}_h + \boldsymbol{\alpha}_h s_t + \boldsymbol{\beta}_h s_t |\bar{D}_t| + \boldsymbol{\delta}_h s_t \upsilon_t + \boldsymbol{\theta}_h \text{LAGS}_{t-1}^L + \boldsymbol{\varepsilon}_{t+h} \tag{3}$$

where $\text{LAGS}_{t-1}^L \equiv \left[\mathbf{y}_{t-1:t-L},\ s_{t-1:t-L},\ |\bar{D}_{t:t-L}|,\ \upsilon_{t-1:t-L},\ s_{t-1:t-L}|\bar{D}_{t-1:t-L}|\right]'$, with respective coefficients $\boldsymbol{\theta}_h \equiv [\,\phi_h^{1:L},\ \gamma_h^{1:L},\ \delta_h^{0:L},\ \mu_h^{1:L},\ \zeta_h^{1:L}\,]$.

As Figure 9 shows for $\upsilon_t = VIX_t$, the marginal effect of disagreement (right panel) is left practically intact, while α_h – the response of industrial production to a monetary shock in the absence of disagreement *and* uncertainty – is accentuated. Appendix D.8 shows that the same conclusion survives when VIX_t is replaced by EPU or SPF dispersion: the right-panel coefficient $\boldsymbol{\beta}_h$ remains positive. This

[16] Appendix D.16 decomposes the effects across GDP components: (*i*) most of the output response — baseline and attenuation — runs through consumption and investment; (*ii*) contributions are proportional between the baseline and attenuation terms, implying uniform attenuation.

[17] Further robustness analysis is in Appendix D. For instance, D.6 excludes from the estimation sample the ZLB era (2008:12-2015:12).

[18] Appendix D.8 compares the three monthly proxies directly.

finding suggests that the attenuation associated with the level of disagreement in the economy is robust to the inclusion of standard uncertainty controls. This aligns with previous related questions on macrofinancial variable responses, as in Zohar (2024) and Benchimol et al. (2023).

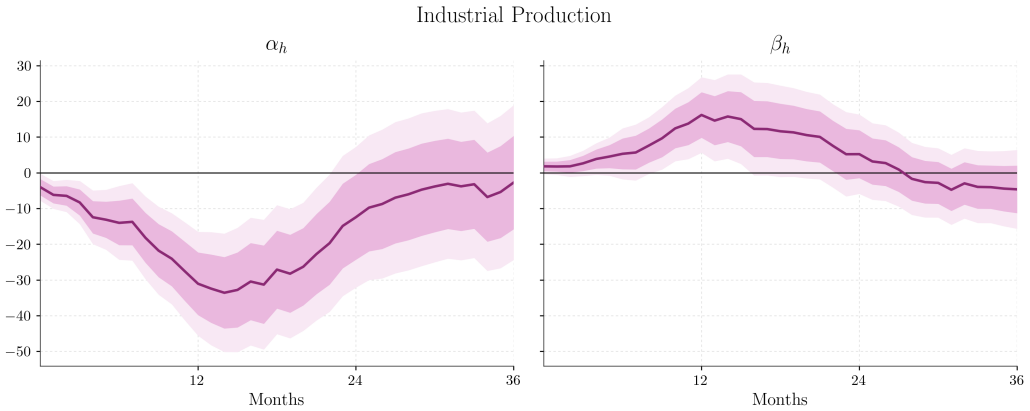


Figure 9: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections controlling for VIX_t and its lags, equation (3). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1991–2019, $N = 346$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

3.3.2 Disagreement and Monetary Surprises

Another valid concern is the potential causal interplay between the shocks proxy, s_t and the measure of disagreement adopted, $|\bar{D}_t|$ (or $\bar{D}_t$ for an asymmetric specification, see § 3.3.4 below). In fact, even if monetary shocks are exogenous (as suggested by Bauer and Swanson (2023b), Amodeo (2025)), the potential feedback between the policy surprises and the expected FFR paths of Fed and financial markets could hinder the inference drawn on the estimated coefficients (Angrist and Pischke (2009) and Gonçalves et al. (2024)). Intuitively, identification requires the "treatment" – the shock – to be orthogonal to the conditioning state (mathematically, the interacting variable) – disagreement.

First, I address this by investigating the correlation between the two objects of interest. In particular, a lead-lag correlation map is useful to understand the dynamic relationships between variables over time, allowing for delay or anticipation in the response of one to the other. Figure 10 displays two maps: on the left, the correlation between the baseline orthogonalized s_t and the absolute value of average disagreement $|\bar{D}_{t+h}|$, shifted ± 12 months; on the right, an alternative specification of the interacting variable without the absolute value, allowing then the full range of motion of the conditioning state in the regressions in (2).

All across the two-year range, the correlation between monetary policy shocks and disagreement is close to zero, and never significant. Appendix D.11 provides a

number of alternative comparisons, using a variety of influential surprise series, and details the construction of the 95% confidence bands.

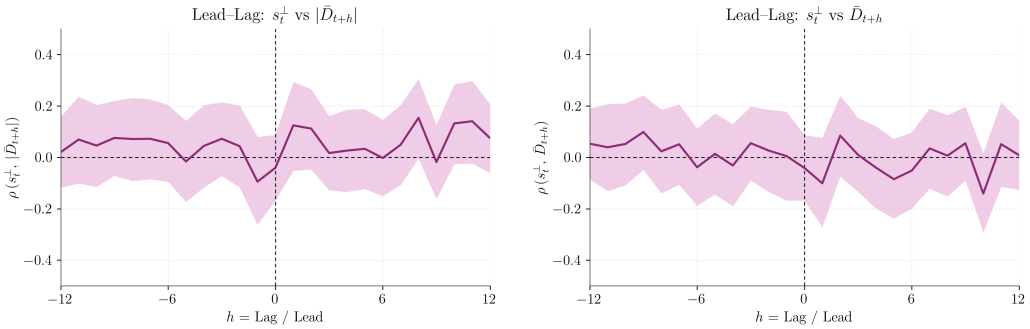


Figure 10: Left panel: Lead-Lag correlation map between orthogonalized monetary policy shocks, and absolute average disagreement. Right panel: Lead-Lag correlation map between orthogonalized monetary policy shocks, and average disagreement. 95% confidence bands computed with Fisher's z-transform (exact method). Estimation sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Moreover, I replace the orthogonalized surprises with leading alternative shock series (Gertler and Karadi (2015), Miranda-Agrippino and Ricco (2021), Jarociński and Karadi (2020)), and re-estimate all core specifications. I also estimate the same local projections using the original Bauer and Swanson (2023a) shock series, removing the adjustment in Amodeo (2025). The impulse responses and interaction patterns remain qualitatively unchanged: disagreement still attenuates conventional transmission on output. This stability reduces dependence on any one high-frequency identification technique, mitigates concerns about the orthogonalization treatment, and increases external validity. In sum, the mechanism is robust and economically significant across datasets, samples, and controls.

Finally, another relevant concern is that the attenuation I document might be spurious if the shock itself differs across disagreement states. I address this issue – treatment homogeneity – in Appendix D.12, proceeding in two steps: first, I estimate the same local projections with the interest rates (FFR, 1-year and 2-year Treasury yields) as outcome variables, including baseline and other influential shock series (D.12.1; D.12.2); second, I explicitly test that the identified shock is distributionally invariant across disagreement quartiles (D.12.3).
Although the FFR shows slightly differential effects across disagreement states, these go in the opposite direction as they would to explain the evidence of attenuation (Figure D16), strengthening, if anything, the role of disagreement in transmission; at longer maturities – more relevant and widely used – shocks are comparable across states. Crucially, graphical inspection of histograms and densities (Figures D24, D25) and nonparametric tests (Table D5) detect no differences in the conditional distribution of shocks across disagreement quartiles.

Overall, the evidence indicates that shocks are statistically indistinguishable across

disagreement states (i.e., treatment homogeneity is verified), securing inference and corroborating heterogeneous transmission as the source of attenuation.

3.3.3 Bauer and Swanson (2023b) specification

A broader type of consideration one might entertain in the evaluation of the evidence of attenuation emerging from equation (2), is the selection of the variables. While widely employed and heuristically insightful, a simple three-variable system including output, inflation, and interest rates might raise questions about its inability to capture the complexity of economic relationships in modern macroeconomic contexts. Similar considerations motivated a number of influential studies in recent years, spurring richer VAR specifications like in Gertler and Karadi (2015) or Bauer and Swanson (2023b).[19]

Below, I replicate my key finding using the latter's system of equations, which encompasses and expands on Gertler and Karadi (2015)'s specification. In particular, in addition to the existing variables, I include in $\mathbf{y}_t$ the excess bond premium (EBP_t) studied by Gilchrist and Zakrajšek (2012), the unemployment rate ($unemp_t$), and a measure of commodity prices ($Pcomm_t$).[20] The excess bond premium is widely considered an important indicator of the credit conditions in the economy, and Caldara and Herbst (2019) argue about its relevance in the estimation of monetary policy effects. The unemployment rate is another indicator of real economic activity on top of the already included industrial production index, and I include it for comparability with some of the aforementioned studies and other influential work by Miranda-Agrippino and Ricco (2021). Finally, Swanson and Williams (2014) advocate for the use of the 2-year Treasury yield instead of the overnight interest rate, due to it not being constrained by the zero lower bound period (2009-15), and it is then included in Bauer and Swanson (2023a)'s six variables system. Hence, I follow the same strategy and perform the analysis for the 2-year and 1-year Treasury yields as well as for the Fed Funds Rate.

In summary, I present in Figure 11 the estimates relative to industrial production estimated with the same methodology as in (2), but specifying a more comprehensive $\mathbf{y}_t = (IP_t, CPI_t, EBP_t, UNEMP_t, PCOMM_t, 2Y_t)'$. The results align very closely with the key findings of Figure 6, with the same pronounced and persistent evidence of attenuation, and a similar pattern for the left panel, in this instance featuring marginally wider confidence bands. It is worth noting that the slightly weaker dampening effects are mostly due to the inclusion of the 2-year Treasury yield – unsurprisingly less affected by the monetary policy surprises. Appendix D

[19] Notice, however, that this exercise serves purposes of *robustness*, not identification. Identification of the state-dependent effect comes from the interaction design under an exogenous policy shock, and as such it is not threatened by larger systems (an application of the Frisch-Waugh-Lovell theorem). A larger outcome vector only tests stability in richer environments.

[20] Since Sims (1980), commodity prices have been often included in common specifications of VARs and local projections, more as an object of analysis itself than as a controlling factor; in fact, alternative specification excluding such variable report almost indistinguishable results.

showcases several variations of the specification above, including the alternative estimates using the 1-year Treasury yield and the FFR, and systems à la Gertler and Karadi (2015), featuring EBP_t but *not* $UNEMP_t, PCOMM_t$. Results are virtually unchanged. Overall, I can conclusively infer that the evidence supporting the role of disagreement in attenuating monetary policy's effectiveness is robust to a number of checks, confirming the substantially state-dependent nature of its transmission to the economy.

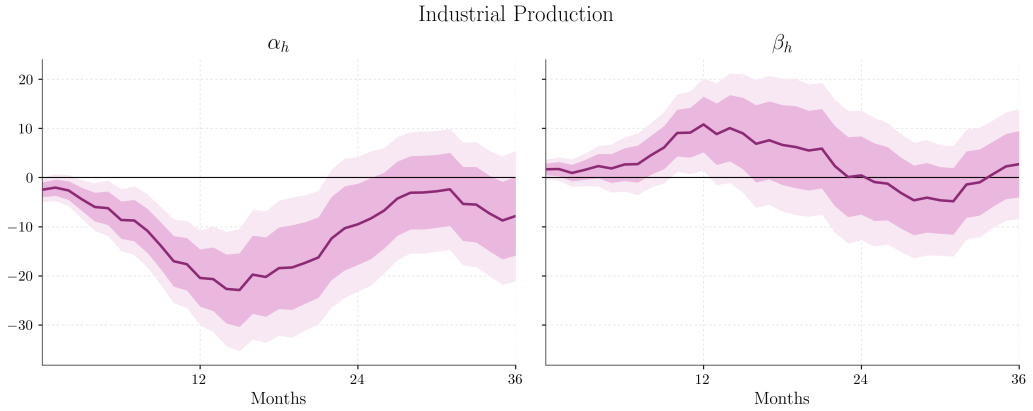


Figure 11: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections on augmented system à la Bauer and Swanson (2023b). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Misspecification sensitivity. Appendix D.9 implements the sensitivity analysis of Imbens (2003) to address *general* concerns about omitted variables. Interpreting a generic misspecification concern as an unobserved confounder U_t correlated with both the interaction term $s_t|\bar{D}_t|$ and the outcome $\mathbf{y}_{t+h}$, the method asks how much of the *residual* variation in each object U_t must explain to overturn the baseline estimate. For industrial production at $h = 12$, the symmetric robustness value is $r^2_{\text{min}} = 12.81\%$: an omitted variable would need to account for at least 12.81% of the residual variance in *both* the interaction term and the outcome (after partialling out all controls) to reverse the sign of $\hat{\beta}_{12}$. As a benchmark, VIX – a major measure of volatility and potential confounder – explains only 2.62% of the residual variance in $s_t|\bar{D}_t|$, implying that overturning the result would require a confounder substantially more predictive of the interaction term than standard uncertainty proxies.[21]

[21] Moreover, this is only a *necessary* condition: additionally, unobserved confounder U should explain a comparable share of residual variation in the outcome, making it an extremely high bar.

3.3.4 Asymmetric Disagreement

The local projection specification in (2) treats disagreement's effect symmetrically, i.e., it does not distinguish between cases where the Fed's projections for the FFR are higher than the market's and vice versa. What the local projection identifies is not the effect of disagreement on the *level* of outcome variable of interest, but how the transmission of a monetary shock changes with the prevailing state of disagreement. A belief gap that tilts demand in one direction can still dampen the pass-through of policy by the same amount whether markets price above or below the Fed: what weakens transmission, acting as the state variable, is how far the two paths stand apart, which corresponds to the magnitude of disagreement.

This can be tested letting the shock's effect depend separately on the signed and the absolute disagreement, the signed component is indistinguishable from zero at every horizon, while the absolute component carries the attenuation. I implement this formally with two complementary designs. The first splits the gap into its positive and negative parts and lets each load the shock separately, recovering attenuation of near-symmetric magnitude on both sides as postulated in the baseline design. The second includes the signed interaction $s_t \bar{D}_t$ and the absolute interaction $s_t|\bar{D}_t|$ together in a single specification, decomposing transmission into a sign-dependent and a magnitude-dependent loading, confirming the robustness of the baseline results. Appendix D.13 provides the full detail on these findings, and further shows they hold across disagreement timings, full and ZLB-excluded samples, and larger systems augmented with the excess bond premium, unemployment, and commodity prices.

Lastly, the disagreement measure that matters is specific to the rate path rather than to the underlying outlook: measuring Fed–Market inflation disagreement directly and conditioning the baseline on it leaves the attenuation essentially unchanged, while the inflation interaction is indistinguishable from zero at every horizon – as Appendix D.14 shows empirically.

3.4 Comparing Forecast Performance

The attenuation evidence raises a natural question. In standard models featuring the "information channel", a policy surprise partly reveals the Fed's private information about fundamentals. Does the Fed possess an informational edge over markets? If so, this narrative would imply superior forecasts of the policy path. I test this corollary. In order to gauge any evidence of superior reliability for Federal Funds Rate forecasts, I compare accuracy across forecasters—Fed staff (Tealbook) and market-implied expectations from federal funds futures. Importantly, market prices are aligned to the Tealbook internal cutoffs, removing any ex post advantage derived from news and data released in the days preceding the FOMC meeting and, obviously, from meeting-day policy decisions and announcements.

Accuracy varies markedly across the sample; in some periods the Fed appears to anticipate the path of the policy rate better than financial markets, whereas in others the market forecast is closer to the realized federal funds rate. Figure 12 depicts the root–mean-square forecast errors (RMSE) over all available quarterly forecasting horizons (excluding nowcasts), averaged by non-overlapping five-year bins.[22] As is immediately evident, rankings of ordinal accuracy vary throughout the sample and no forecaster systematically dominates the others.[23]

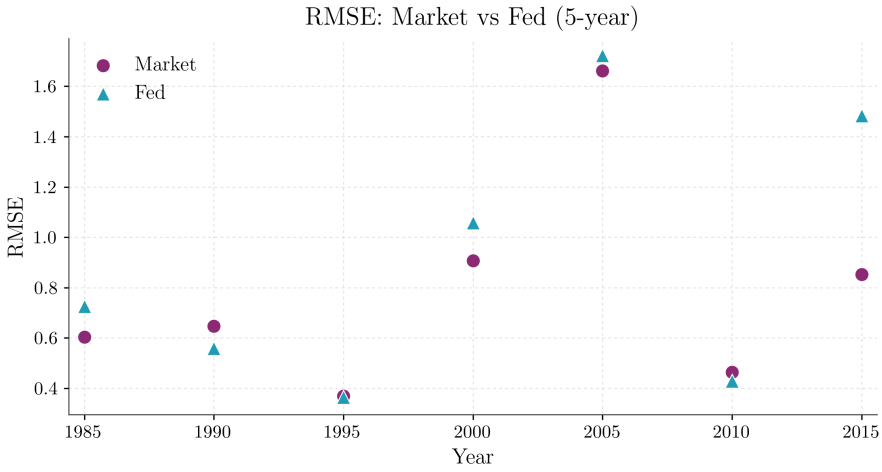


Figure 12: Five-year root mean squared forecast errors for the federal funds rate by forecaster — Market and Fed — averaged over non-overlapping bins.

3.4.1 Accuracy Comparison

To formally evaluate differences in predictive accuracy, I conduct Diebold and Mariano (2002) tests on all three pairs. I present below the main comparison, between financial markets and the Fed, relegating the others comparison and the technical details of the test to the Appendix.

Let d denote the mean loss differential between two forecasts (e.g., market minus Fed mean–squared error) and define the DM statistic

$$DM_h = \frac{\bar{d}_h}{\sqrt{\mathrm{Var}(\bar{d}_h)}} \sim \mathcal{N}(0,1)$$

where, under MSE loss function, $e^i_{t,h} = (y_t - \mathbb{F}^i_{t-h}(y_t))^2$, $\quad d_{t,h} = e^M_{t,h} - e^F_{t,h}$, $\quad \bar{d}_h = \frac{1}{T}\sum_{t=1}^{T} d_{t,h}$. Under the null hypothesis, both forecasts are equally accurate; values

[22] Single calendar year averages are depicted in Figure B1 Appendix B, and horizon-specific comparisons display similarly inconclusive insights.

[23] Appendix B also includes comparisons with the Survey of Professional Forecasters' consensus—whose timing is inherently coarser—yield similar rank-fluctuation patterns

statistically different from zero indicate that one party's forecasts have a lower average squared error.

Table 3 presents DM statistics for horizons from one to eight quarters. The test rejects equal predictive accuracy at every horizon. Across all horizons the market exhibits lower MSFE than the Fed, with gaps attenuating as the horizon lengthens.

Null: H_0	**N**	**DM**	**PValue**	**Winner**
Equal MSE at $h = 1$	243	-4.895	0.000	M
Equal MSE at $h = 2$	234	-3.943	0.000	M
Equal MSE at $h = 3$	186	-3.971	0.000	M
Equal MSE at $h = 4$	139	-4.367	0.000	M
Equal MSE at $h = 5$	127	-4.003	0.000	M
Equal MSE at $h = 6$	126	-3.855	0.000	M
Equal MSE at $h = 7$	115	-3.796	0.000	M
Equal MSE at $h = 8$	94	-3.996	0.000	M

Table 3: Diebold–Mariano tests of equal predictive accuracy (Market vs Fed) across horizons with Market's forecast on the day of the Tealbook (internal) release. Reported are sample size N, test statistic, p-value, and the lower-MSE winner.

These results suggest that the Fed does not possess superior information relative to markets, at least over horizons up to one year, and that the accuracy gap narrows for longer horizons. Two relevant remarks: (*i*) time t represents *the same* day, i.e., information sets for financial markets and the Fed are (manually) aligned to the date of internal release of each Tealbook; (*ii*) markets' expectations are risk-adjusted according to a rule of thumb proposed in Piazzesi and Swanson (2008) and used by the Federal Reserve Board. For robustness, results for different timings of the forecasts by the market (FOMC meeting day, day before) and risk-unadjusted are provided in Appendix B, with insubstantial variation in the findings.

Across horizons, the predictions embedded in financial markets' prices consistently outperform the Fed's staff. An influential narrative in the monetary policy debate is founded upon the notion of the Fed possessing some form of superior information with respect to market participants and the market as a whole. Although the existence of privileged information only available to the Fed's economists who contribute to the Tealbook projections is undeniable, this first result suggests that such privileged information does not constitute a significant edge in forecasting the future realizations of the Federal Funds Rate, despite it being a *de facto* control variable of the Fed.

3.4.2 Revisiting Fed's Informational Advantage

This subsection revisits explicitly whether the Fed holds an informational edge over private expectations, a precondition for information effects that confound monetary shock identification. Hoesch et al. (2023) proposes a test using survey forecasts as the benchmark, but surveys are infrequent and not synchronized with FOMC meetings or Fed forecasts, necessitating timing assumptions. I leverage the high-frequency market expectations derived from federal funds futures to re-estimate their test. As already discussed, traded expectations can be precisely aligned with the Tealbook information cutoff and continuously updated, ensuring comparable information sets and flexibility in testing the robustness of the approach. This improved timing enhances the credibility of the informational advantage test, extending its domain and offering a more stringent framework for assessing the conjectured Fed's forecasting edge.[24]
I restate the test's logic, replicate it in a refined version replacing survey data with the precisely timed, priced-in market expectations, and discuss the remarkably different findings. This preserves the original design while yielding cleaner inference and corroborating the evidence emerged from the Diebold-Mariano tests.

As in Hoesch et al. (2023), I estimate the following regression for $x_t = FFR_t$:

$$x_{t+h} - x^{\mathrm{M}}_{t+h|t} = \delta + \beta^F x^F_{t+h|t} + \beta^M x^M_{t+h|t} + \eta_{t+h}, \tag{4}$$

where $x^F_{t+h|t}$ denotes the Federal Reserve's Greenbook/Tealbook forecast and $x^M_{t+h|t}$ is the private sector's priced-in forecast at the same horizon. The null hypothesis $H_0 : \beta^F = 0$ tests whether the Fed's forecast adds no information beyond the markets' prediction. A significant β^F indicates that the Fed possesses an informational advantage, as forecasters would optimally weight both forecasts rather than rely on the private forecast alone.

Following closely the methodology in Hoesch et al. (2023), I allow for time variation in by estimating (4) in rolling windows of size m (meetings), obtaining $\hat{\beta}^F_t$ at each mid-window t. Then, the t-statistic below is calculated:

$$\tau^F_t = \hat{\beta}^F_t \Big/ \sqrt{\hat{\sigma}^2_F / m} \tag{5}$$

and the "Information-Advantage Fluctuation test statistic" corresponds to

$$\mathcal{I}_F = \max_t \left| \tau^F_t \right|. \tag{6}$$

Here $\hat{\sigma}^2_F$ is a heteroskedasticity and autocorrelation consistent estimate (Newey–West) of the diagonal element of the variance matrix (i.e., $\hat{\beta}^F_t$'s variance) within

[24] For instance, I replicate the test not only timing markets' forecast to the Tealbook cutoff date, but I also experiment with one and three days in advance futures prices, to allow a buffer in case the Tealbook predictions were formulated before their release date. Appendix D.17.1 reports the additional results.

the window.[25] The null $H_0 : \beta_t^F = 0 \quad \forall\ t$ (no informational advantage at any time) is rejected when $\mathcal{I}_F$ exceeds its simulated critical value;[26] rejection implies an information advantage "*at some point in time*".[27]

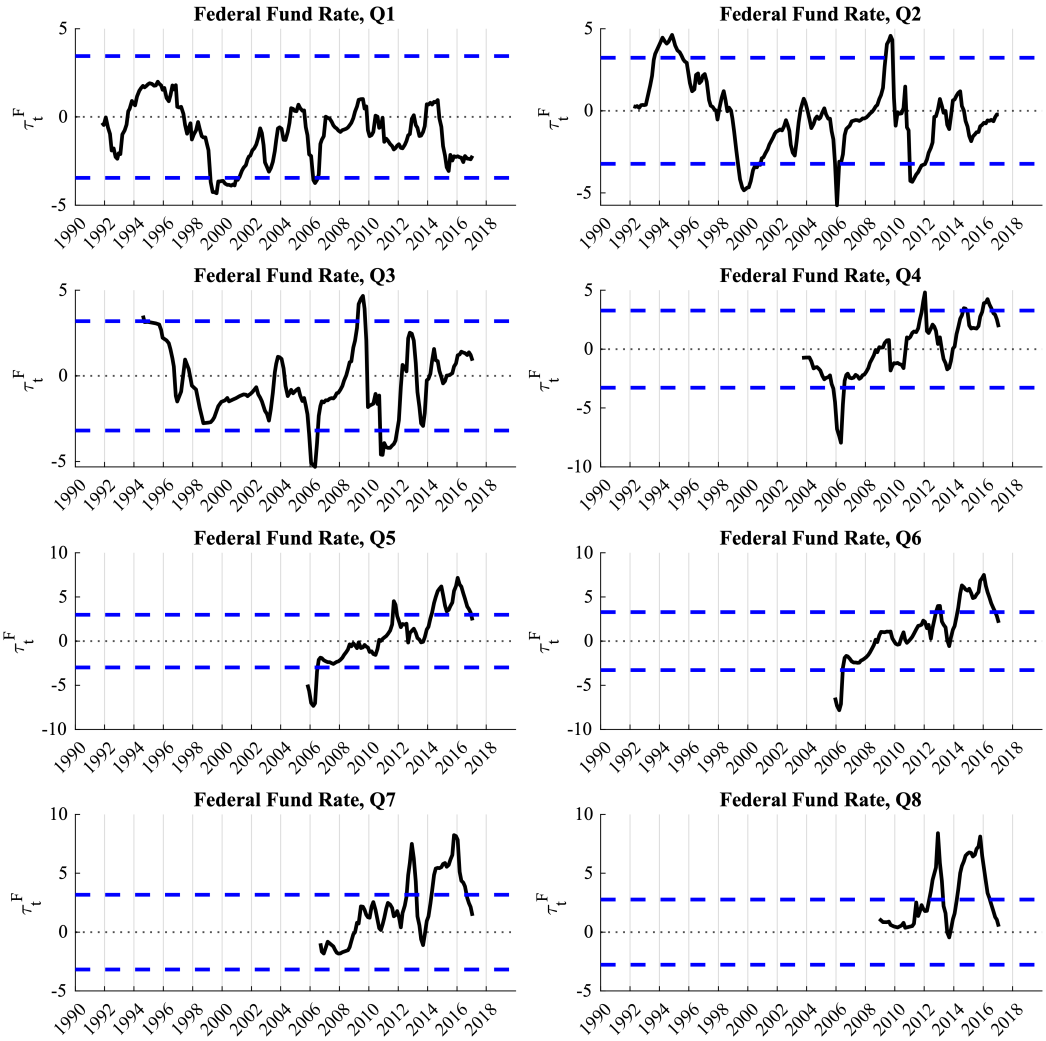


Figure 13: Information advantage—FFR. Rolling t-statistics τ_t^F from (4), estimated in $m = 30$–meeting windows (dated at the midpoint). Dashed red lines are the two-sided 5% Rossi and Sekhposyan (2016) critical values. HAC s.e. are Newey–West with bandwidth $m^{1/4}$. Market forecasts are extracted from Fed Funds Futures aligned to the Tealbook cutoff.

[25] Let $\widehat{V}_t$ be the Newey–West estimate of the long-run variance matrix V in the m-observation window at time t. Define

$$\hat{\sigma}_F^2 := e_F' \widehat{V}_t e_F = [\widehat{V}_t]_{FF}, \qquad \widehat{\mathrm{Var}}(\hat{\beta}_t) = \widehat{V}_t / m.$$

[26] Critical values are the two-sided t-statistics analog to the Wald test values for the survey and model-free forecasts reported in Table II, Panel C of Rossi and Sekhposyan (2016), resimulated to closely match the sample size and rolling window size.

[27] A rejection rules out the hypothesis that the central bank never had an information advantage. Previous studies using surveys report an evident advantage; by contrast, my results contradict this, as the statistic crosses the critical bands only episodically.

Figure 13 displays the test results at $h = 1, \ldots, 8$ quarters horizons. Comparing with the results reported in Figure 2 of Hoesch et al. (2023), the difference is striking.[28] The rolling statistics τ_t^F are overwhelmingly below the two–sided 5% bands at short and medium horizons, and they remain well within the bands for most of the sample. The only material excursions occur in the post-ZLB period at longer horizons (Q5–Q8), with brief spikes earlier on. These exceedances are narrow, clustered in the forward-guidance era, and short-lasted. The pattern clearly does not support a general informational edge for the Fed in forecasting the future stance of monetary policy.

3.4.3 Revisiting Time-Varying Accuracy

Complementing the information-advantage exercise, Hoesch et al. (2023) evaluates the *relative predictive accuracy* of Fed and markets forecasts with a rolling Giacomini and Rossi (2010), again employing survey data in the comparison.[29] I repropose this analysis, for each horizon h and window of m meetings, with the futures-based market data, timed at the closing price of the date of the Tealbook's internal release. Consistently with the Diebold-Mariano tests in section 3.4, as accuracy criterion I consider the mean–squared forecast errors (MSFE) of the market and the Fed.

Let the h-ahead forecast errors and their loss differential be

$$e_{s,h}^F = x_{s+h} - x_{s+h|s}^F, \qquad e_{s,h}^M = x_{s+h} - x_{s+h|s}^M, \qquad d_{s,h} = (e_{s,h}^M)^2 - (e_{s,h}^F)^2. \tag{7}$$

Within each window $s \in \{t - m + 1, \ldots, t\}$, I run the constant-only regression

$$d_{s,h} = \mu_{t,h} + u_{s,h}. \tag{8}$$

Similarly to (5), the t–statistic takes the following form:

$$\tau_{t,h}^{\mathrm{GR}} = \bar{d}_{t,h} \Big/ \sqrt{\hat{\sigma}^2{}_{t,h}/m}, \qquad \bar{d}_{t,h} = \frac{1}{m} \sum_{s=t-m+1}^{t} d_{s,h}. \tag{9}$$

I summarize time variation via the *Forecast-Accuracy Fluctuation* statistic

$$\mathcal{A}_h = \max_t \tau_{t,h}^{\mathrm{GR}}, \tag{10}$$

the null being $H_0 : \mu_{t,h} = 0 \; \forall \; t$ (i.e., equal MSFE at all times) is tested against the

[28] In Appendix D.17, I report the results for a variety of window lengths and window/sample ratios, including the $m = 60$ Hoesch et al. (2023) implement to facilitate the comparison.

[29] Distinct, conceptually, from information advantage tests, which ask whether one forecast adds information conditional on the other: it can be shown that such tests do not rank forecast accuracy unless both forecasts are unbiased and their errors uncorrelated. With correlation or bias, the mapping breaks. Giacomini and Rossi (2010), instead, targets accuracy directly by tracking the mean loss differential, regardless of the correlation structure.

one-sided alternative $H_1 : \mu_{t,h} > 0$ (i.e., Fed more accurate for at least one window). Intuitively, the null hypothesis is that the Tealbook and market's forecasts have the same predictive accuracy; under the alternative, positive values of the test statistic indicate that the Tealbook predictive performance is more accurate. Global rejection occurs when $\mathcal{A}_h$ exceeds the Giacomini–Rossi critical value.[30]

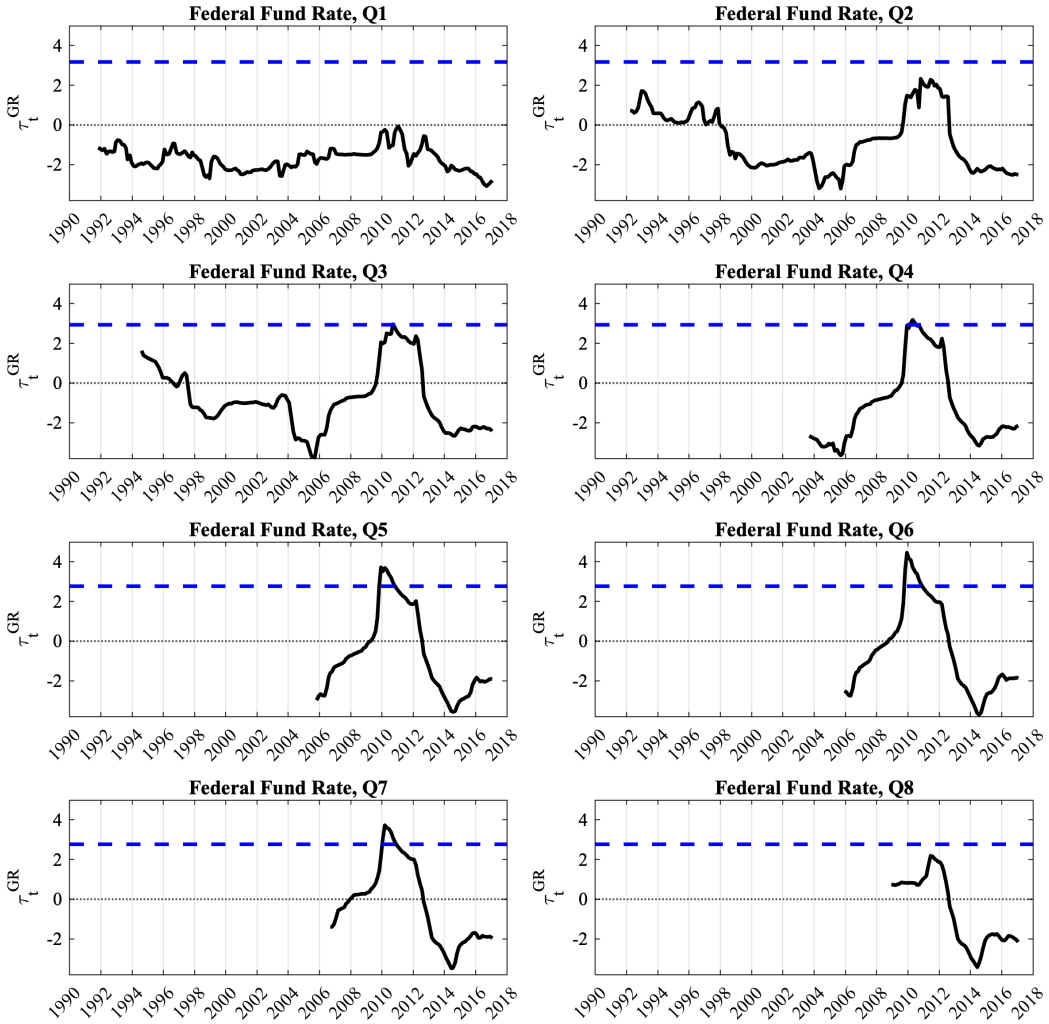


Figure 14: Forecast accuracy—FFR ($h = 1:8$). Rolling $\tau_{t,h}^{\mathrm{GR}}$ from regression 8 (positive $\approx$ Fed more accurate), estimated in $m = 30$–meeting windows and dated at the midpoint. Dashed red lines mark the one-sided 5% Giacomini and Rossi (2010) critical values. HAC s.e.: Newey–West with bandwidth $m^{1/4}$. Market forecasts are extracted from Fed Funds Futures aligned to the Tealbook cutoff.

Figure 14 displays the accuracy comparison results at $h = 1, \ldots, 8$ quarters horizons. The statistic τ_t^{GR} almost never sustains values above the one-sided 5% line. On the contrary, for most horizons, values below the threshold and below zero

[30] More technical details are reported in Appendix D.17.

are exceedingly more common, indicating that the market's MSFE is *at least* as low as the Fed's. In the medium run (Q6-Q7) there are short-lived peaks during 2009–2010, coinciding with the (maintained) predictions of the Fed to keep an accommodative stance of policy; however, these quickly reverse and the statistics turn negative thereafter. Unsurprisingly, the evidence points in the opposite direction of a Fed's supremacy in forecasting: the market is weakly—and often plainly—more accurate than the Fed in predicting the funds rate, consistent with the findings reported in 3.4.

4 Theoretical Framework

This section maps the evidence to theory. I first show that relaxing the full-information assumption tractably introduces disagreement in a baseline New Keynesian model. I then derive the model's impulse responses to characterize policy transmission in a disagreement economy, and use the empirically documented state-dependent attenuation to obtain model restrictions. These restrictions discipline the framework's microfoundations so that it reproduces the empirical patterns and yields sharp, testable predictions, which the model satisfies. The next section microfounds disagreement and market learning in a model where markets update around FOMC announcements under heterogeneous information, different policy beliefs and Fed's limited commitment.

4.1 Disagreement in the New Keynesian Model

Strong assumptions like full information prevent the occurrence of disagreement in the equilibrium of standard macro models. In fact, under full information and rationality, expectations of all economic agents coincide. To microfound the pervasive disagreement documented in the empirical section, I assume a general belief structure with differential information on fundamentals, heterogeneous priors on the monetary policy rule and limited commitment ("credibility") by the Fed. This Section uses the core structure of the New Keynesian model to accommodate the presence of disagreement and study its effects on monetary transmission.

Intuitively, the objective is to let the Fed announce one path for policy, while markets *act* on a potentially different perceived path, and to trace how this belief gap feeds back into aggregate demand through intertemporal substitution – a key channel of monetary propagation – in its simplest configuration: the Euler equation.

I start from a log-linearized consumption Euler equation, which is derived from intertemporal household optimization, and that therefore features *market* expec-

tations of future real rates and consumption:

$$c_t \;=\; -\frac{1}{\sigma}\,\mathbb{E}_t^M\Big[i_t - \pi_{t+1} - r_t^n\Big] + \mathbb{E}_t^M(c_{t+1}) \tag{11}$$

where σ is the intertemporal elasticity of substitution and r_t^n the natural rate of interest. Economically, this condition says that households choose consumption so that today's marginal utility equals the expected discounted marginal utility tomorrow: if they anticipate a lower path of future real rates, it becomes optimal to bring spending forward and increase current demand (consumption smoothing). Iterating (11) forward gives

$$c_t \;=\; -\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big[\mathbb{E}_t^M(i_{t+h}) - \mathbb{E}_t^M(\pi_{t+h+1}) - \mathbb{E}_t^M(r_{t+h}^n)\Big] \tag{12}$$

Define $D_{t+h|t} \equiv \mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^M(i_{t+h})$. Assume, for the time being, that inflation expectations are common, i.e. $\mathbb{E}_t^M(\pi_{t+h+1}) = \mathbb{E}_t^F(\pi_{t+h+1})$, and that r^n is a publicly known constant denoted by ρ. Adding and subtracting $\mathbb{E}_t^F(i_{t+h})$ yields[31]

$$c_t = -\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big[\mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^F(\pi_{t+h+1}) - \rho\Big] + \frac{1}{\sigma}\sum_{h=1}^{\infty} D_{t+h|t} \tag{13}$$

where the first term represents the iterated Euler Equation (12) if the Market shared the Fed's beliefs. That is, if the Market and the Fed *agreed*. I rewrite (13) as

$$c_t \equiv c_t^{\text{Agreement}} + W_t, \tag{14}$$

where $c_t^{\text{Agreement}}$ is the conventional New Keynesian consumption implied by the Euler Equation under full information rational expectations (FIRE) and

$$W_t \;\equiv\; \frac{1}{\sigma}\sum_{h=1}^{\infty} D_{t+h|t} \tag{15}$$

represents an *expectations wedge*. Because W_t enters additively, it behaves like a belief-driven demand shock: if, on average, markets expect a more accommodative future path than the Fed ($D_{t+h|t} > 0$), then $W_t > 0$ and consumption is stronger than under agreement; if they expect a tighter path ($D_{t+h|t} < 0$), $W_t < 0$ and demand is weaker than what the Fed intends. The wedge therefore summarizes how much of the announced tightening or easing is actually internalized in households' Euler equation. In other words, disagreement distorts the market's demand schedule through W_t, reflecting market's different beliefs on the future interest

[31] Relaxing $\mathbb{E}_t^M(\pi_{t+h+1}) = \mathbb{E}_t^F(\pi_{t+h+1})$ is immediate and leaves all qualitative results intact, only adding a parallel wedge proportional to inflation-expectations disagreement. See the derivation in E.2 for a formalization of this argument.

rate path.

4.2 Transmission of Monetary Policy Shocks

Monetary policy follows a rule of the form $i_t = \mathbf{\Phi}_t' \mathbb{E}_t^F(\mathbf{x}_t) + u_t^S$, where $\mathbf{\Phi}_t$ is an $n \times 1$ time-varying vector of coefficients, $\mathbf{x}_t$ is a $n \times 1$ vector representing (unobserved) economic fundamentals, and u_t^S is a zero-mean monetary policy shock endowed with persistence $u_t^S = \rho_s u_{t-1}^S + \varepsilon_t^S$. For illustration, Fed and markets may hold different beliefs about the policy rule coefficients, so that policy is the primitive of disagreement (the argument is generalized below and in Section 5).

Evaluate a simple 3-equation New Keynesian structure, composed by the newly obtained D–DIS, a standard New Keynesian Phillips Curve (NKPC) and a simplified policy rule considering only inflation.[32] Markets and Fed disagree on the inflation coefficient, representing the strength of response of Fed's policy: I denote each forecaster's prior by ϕ_π^i, with $i = (F, M)$.[33]
Imposing equilibrium ($c_t = y_t$) obtains a disagreement-enhanced dynamic IS curve (D–DIS). By replacing the interest rate rule in the D-DIS, the system takes the following form:[34]

$$\begin{aligned} y_t &= -\frac{1}{\sigma}\Big[\phi_\pi \, \mathbb{E}_t^F(\pi_t) + u_t^S - \mathbb{E}_t^F(\pi_{t+1})\Big] + \mathbb{E}_t^F(y_{t+1}) \; - \; \mathbb{E}_t^F(W_{t+1}) + W_t, \\ \pi_t &= \beta \; \mathbb{E}_t^M(\pi_{t+1}) + \kappa \, y_t. \end{aligned} \tag{16}$$

Relative to the standard three-equation New Keynesian model, the only new term is W_t in the IS equation. A monetary surprise now affects output both through the usual movement in expected real rates and through the way markets re-interpret the whole future policy path, captured by the induced change in W_t. The propagation of the shock to u_t^S will depend both on the structural parameters and on the response of the endogenous variables (including W_t) which is function of the expectation gap between Fed and markets.
Let us assume linear solutions of the form $y_t = \Psi_y u_t^S$, $\pi_t = \Psi_\pi u_t^S$, $W_t = \Psi_w u_t^S$.

[32] A more general policy rule leaves results unvaried. Appendix E.3.2 shows the case for a rule including an output gap response.

[33] Note that the *source* of disagreement is immaterial at this point. Analogous conclusions follow with alternative divergences. For instance, in Appendix E.3.5, markets and Fed disagree on the persistence ρ_s^i of the shock. Section 5 addresses this systematically, with multiple simultaneous sources of disagreement.

[34] See Appendix E.3.1 for the full derivation. Notice that ϕ_π is the realized rule's coefficient, generally distinct from each forecaster's prior, used instead to project forward rates.

Solving the system yields:[35]

$$\begin{aligned}\Psi_\pi &= \frac{\kappa}{1-\beta\rho_s}\Psi_y, \\ \Psi_y &= \frac{(\sigma(1-\rho_s)\Psi_w - 1)(1-\beta\rho_s)}{(1-\beta\rho_s)(1-\rho_s)\sigma + \kappa(\phi_\pi - \rho_s)}\end{aligned} \tag{17}$$

where $\Psi_w \equiv \partial W_t / \partial u_t^S$ represents the (impulse) response of the wedge to the shock. Equation (17) showcases the role of disagreement in shaping the response of output to monetary shocks: in the absence of a response of disagreement ($\Psi_w = 0$), the expression collapses to the familiar solution (Galí, 2015); when disagreement responds endogenously, instead, the effect of the shock on output includes an additional term. Rewriting Ψ_y:

$$\Psi_y = \underbrace{-\frac{(1-\beta\rho_s)}{\sigma(1-\beta\rho_s)(1-\rho_s) + \kappa(\phi_\pi - \rho_s)}}_{\text{Baseline NK response}} + \frac{(1-\beta\rho_s)(1-\rho_s)\sigma\Psi_w}{\sigma(1-\beta\rho_s)(1-\rho_s) + \kappa(\phi_\pi - \rho_s)} \tag{18}$$

This decomposition provides a transparent interpretation of the evidence documented in the empirical section. The first term represents the conventional output response in New Keynesian models, and it is negative under all common parametrizations; the second term emerges from the response of disagreement to the policy shock, and its sign determines the attenuation (or the amplification) of the monetary effects on output.[36]

In particular, if $\Psi_w > 0$, the second term is positive and output responds *less* to policy innovations than predicted by conventional theory. Intuitively, after a positive surprise, an increase in disagreement acts as a counter-force, dampening the contractionary effect of the tightening: markets expect future partial reversal, partially offsetting the contraction and muting the response of output. By contrast, a decrease would reflect a closer alignment in beliefs, amplifying the response to the shock.
Next, I connect the state-dependent restrictions emerging from the empirical results to the theoretical impulse response.

Empirical Restrictions The effects of monetary policy surprises documented empirically imply a state-contingent interpretation of the second term in (18): in Figure 6, the right panel ($\boldsymbol{\beta}_h$) indicates the presence of monetary attenuation *proportional* to the level of (ex post) disagreement. Each horizons' estimate measures the horizon-specific slope in the linear relationship between $\mathbf{y}_{t+h}$ and $s_t|\bar{D}_t|$,

[35] The derivation of the closed-form solution is detailed in Appendix E.3.3. Relatedly, other solutions under alternative deviations from full information (e.g., $\rho_s^F \neq \rho_s^M$) are provided too. Finally, a discussion of determinacy of the equilibrium can be found in Appendix E.4.

[36] Drawing a direct comparison of the terms in (18) to the local projections, the first term relates to α_h, i.e., the response of the output in the absence of disagreement; the second term maps to β_h, as I argue next.

i.e., the offsetting force of disagreement on monetary shocks. In other words, these results can be interpreted as revealing that in *states* of high disagreement the attenuation effect is stronger, dampening and sometimes even reversing the conventional direction of the impulse response.[37]

Mathematically, attenuation imposes the following restriction:

$$\Psi_y^L < \Psi_y^H \iff \Psi_w^L < \Psi_w^H, \tag{19}$$

that is, in states of *H*igh disagreement, the effect of a tightening shock is more muted than when disagreement is *L*ow, resulting in a lesser (if any) contraction of output. Next, I expand on the consequences of this restriction within the model.

To bridge between the theoretical predictions and the empirical restrictions, I characterize the local dynamics around differential levels of disagreement. Intuitively, this can be thought of as log-linearizing around *two* different disagreement steady states. That is, I analyze the impulse response functions of output conditioning on high/low levels of disagreement ($\Psi_y^{H,L}$):

$$\Psi_y^{H,L} = \underbrace{-\frac{(1-\beta\rho_s)}{\sigma(1-\beta\rho_s)(1-\rho_s)+\kappa(\phi_\pi-\rho_s)}}_{\text{Baseline NK response}} + \frac{(1-\beta\rho_s)(1-\rho_s)\sigma\Psi_w^{H,L}}{\sigma(1-\beta\rho_s)(1-\rho_s)+\kappa(\phi_\pi-\rho_s)} \tag{20}$$

In the model, equation (19) says that the wedge responds more in H than in L states. Formally, define $A \equiv -\frac{(1-\beta\rho_s)}{\sigma(1-\beta\rho_s)(1-\rho_s)+\kappa(\phi_\pi-\rho_s)}$, $B \equiv \frac{(1-\beta\rho_s)(1-\rho_s)\sigma}{\sigma(1-\beta\rho_s)(1-\rho_s)+\kappa(\phi_\pi-\rho_s)}$.

I rewrite (20) compactly:

$$\Psi_y^{H,L} = A + B\Psi_w^{H,L} \tag{21}$$

Under standard calibrations, $A < 0 < B$, allowing to show explicitly the implications of the positive $\boldsymbol{\beta}_h$ as in (19):

$$\boldsymbol{\beta}_h > 0 \iff \Psi_y^H - \Psi_y^L > 0 \iff B(\Psi_w^H - \Psi_w^L) > 0 \iff \Psi_w^H > \Psi_w^L \tag{22}$$

The economic intuition for this result can be found in the intertemporal substitution at the core of New Keynesian dynamics: a larger change in the wedge translates in a larger revision by the market, resulting in a more muted reallocation of intertemporal demand due to the persistent (heightened) disagreement when in high disagreement states. For example, in response to a positive monetary surprise ($u_t^S > 0$), output will react *less* ($\Psi_y^L < \Psi_y^H$) when market's expectations about future rates are revised more ($\Psi_w^L < \Psi_w^H$). This can be interpreted as the market regarding the surprise hike as a "mistake", and therefore expecting the Fed to "adjust" the policy trajectory in the future. This interpretation is close in spirit

[37] See the reported specific episodes associated with different levels of disagreement. Figure D7 in the Appendix shows examples on the occurrence of extreme minimum (1997m10) and maximum (2009m8) levels of disagreement, emphasizing highly state-dependent effects.

to Caballero and Simsek (2022)'s notion of *opinionated* markets, which, although remarkably different in mechanism, also features a framework where financial markets and Fed adjust expectations of future policy based on an understanding of the other's beliefs being biased.

4.2.1 Testing the model's predictions

I test the empirical validity of the model's prediction (19) by investigating the behavior of the wedge W_t in different disagreement states. I proxy for the infinite sum of disagreements using the available observations in my sample. My baseline approach constructs a wedge series that most closely spans the duration of the full sample period to align the theoretical predictions with the empirical findings. Data availability in the early years of the sample is constrained, often limiting the disagreement horizon to only two or three quarters into the future. This presents a tradeoff between model consistency – which would call for the longest possible sequence – and sample representativeness – from which the key empirical results are drawn. As a baseline, I exclude observations featuring less than three consecutive quarters of disagreement. Intuitively, this finite-horizon wedge is the empirical counterpart of W_t in the model. Appendix F presents alternative parametrizations along the trade-off between disagreement length and sample duration.

To facilitate the comparison with the empirics, I test wedge dynamics by estimating the same local projections as in the baseline results, the only difference being in the left-hand side variable:[38]

$$W_{t+h} = \mathbf{c}_h + \boldsymbol{\alpha}_h\, s_t + \boldsymbol{\beta}_h\, s_t|\bar{D}_t| + \textsc{Lags}_{t-l}^{L} + \boldsymbol{\varepsilon}_{t+h} \tag{23}$$

Notice, in fact, that differentiating (23) with respect to s_t we obtain:

$$\frac{\partial W_{t+h}}{\partial s_t} = \boldsymbol{\alpha}_h + \boldsymbol{\beta}_h|\bar{D}_t| \tag{24}$$

Where the left-hand side represents precisely the empirical analog of our object of interest: $\frac{\partial W_{t+h}}{\partial s_t} \equiv \Psi_w$. Then, the coefficients $\boldsymbol{\beta}_h$ capture how Ψ_w varies with the level of disagreement at each horizon h, namely its state dependence. In terms of the model, $\boldsymbol{\beta}_h$ tells us whether, at a given horizon, moving from low to high disagreement makes markets lean more strongly *against* the surprise or move more in the same direction.
The results of the estimation are displayed in Figure 15. They show that the expectation wedge decreases in times of agreement and increases as a function of the prevailing level of disagreement, with a delay of three quarters.

[38] Plus the exclusion of $h = 0$ – a nowcast – from the left-hand side to avoid simultaneity.

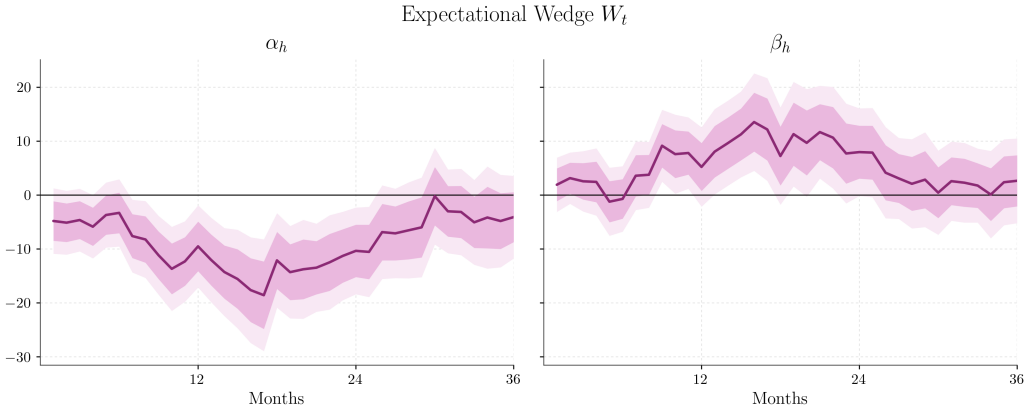


Figure 15: Impulse Response Functions of Wedge $W_t = \sum_{j=1}^{8} D_{t+j}$ | `min(3 quarters)` to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (23). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1992-2019, $N = 331$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

This finding has profound consequences on at least two dimensions. On one hand, it confirms the theory's prediction that the second term in the decomposition of output's impulse response – equation (20) – is consistent with a story of state-dependent attenuation. On the other hand, it implies a differential updating regime in the signal extraction by the Market; that is, it corroborates the notion that financial markets process monetary shocks in their expectations in a way that depends on the significance of the existing gap between their expectations and the Fed's.

More precisely, the local projections in equation (23) provide a theory consistent mechanism based on Fed-markets disagreement to rationalize the dampened monetary effects. That is, they confirm that attenuation obtains through the dynamics of the wedge, consistently with the fundamental intertemporal substitution mechanism captured by the Euler Equation.

Moreover, these findings inform the modeling of the learning that takes place around each FOMC announcement, implying that the updating of expectations by the market is a function of the prevailing level of disagreement at the time of the shock. In fact, this will be the core justification of the state-dependent belief revisions I describe next in the microfoundations of market's expectation formation. Informally, it suggests that, in states of low disagreement, market's revisions and monetary surprise have the same sign, while in high disagreement regimes the update occurs in the opposite direction of the realized surprise.[39]

Wedge Dynamics Note again that the results in Figure 15 are moot on the overall dynamics of the wedge in response to monetary shocks. Studying the *total response* of the wedge – that is, its *unconditional* response – can instead help

[39] Recall that $\mathbb{E}_t^M(i_{t+h}) \propto W_t^{-1}$, i.e., markets' expectations and wedge are inversely related.

pinning down the average revision behavior of financial markets. Equation (23) also estimates this object: the average estimate of $\hat{\boldsymbol{\alpha}}_h + \hat{\boldsymbol{\beta}}_h|\bar{D}_t|$ coincides with the results of a hypothetical regression of the wedge on the shock alone.[40] Figure 16 depicts the total marginal effect of a high-frequency monetary surprise evaluated at the average disagreement level.

Throughout the sample, the effect on the expectational wedge appears negative, although barely significant. This represents another useful moment for a model of learning to match, as it indicates that, on *average*, we observe a slight reduction in the disagreement gap about future rates, consistent with financial markets learning from the monetary surprises realizing at the time of each FOMC announcement.

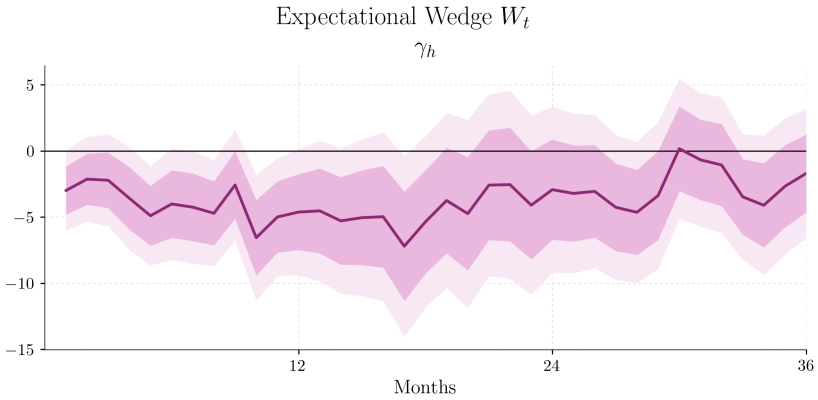


Figure 16: Impulse Response Functions of Wedge $W_t = \sum_{j=1}^{8} D_{t+j}$ | `min(3 quarters)` to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (23). Total average effect. Confidence bands at 68 and 90% levels. Sample: 1992-2019, $N = 332$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Although insightful, for the study of monetary attenuation it is even more revealing to observe the time-varying behavior of the total effect as a function of the degree of disagreement around every monetary policy decision. In fact, Figure 15 suggests a potentially differential updating behavior based on the prevailing level of Fed-markets expectations alignment. Then, I plot next the time-series variation of the wedge's total response at $h = 16$.[41]

[40] Denoting by $\hat{\boldsymbol{\gamma}}_h$ the coefficient of s_t, $\hat{\boldsymbol{\gamma}}_h \equiv \hat{\boldsymbol{\alpha}}_h + \hat{\boldsymbol{\beta}}_h$, as I normalize $|\bar{D}_t|$ to have mean one.

[41] The choice of horizon is due to estimation precision, as the interval 5-8 quarters is the less noisy. I employ here $L = 6$ lags, less than the baseline 12, to allow for tighter confidence bands. Qualitative results are unchanged using $L = 12$ and different horizons' responses (Appendix F).

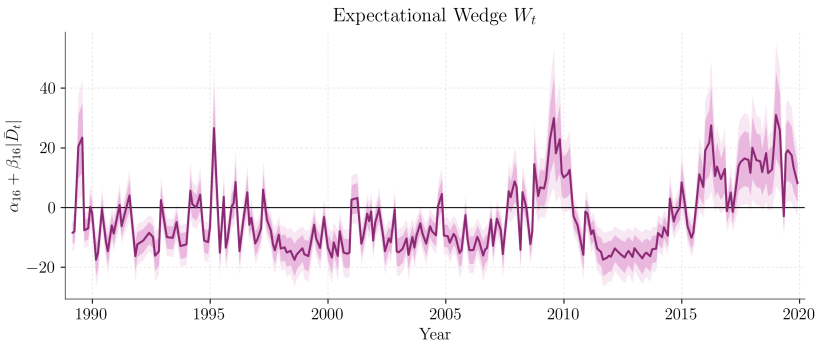


Figure 17: Impulse Response Functions of Wedge $W_t = \sum_{j=1}^{8} D_{t+j}$ | `min(3 quarters)` to an orthogonalized monetary policy surprise evaluated at each $|\bar{D}_t|$ at horizon at $h = 16$, estimated from lag-augmented LPs, as in equation (23). $L = 6$ lags. Confidence bands at 68 and 90% levels. Sample: 1991-2019, $N = 323$ observations (excludes 2001M9, for the 9/11 Terrorist Attacks).

Figure 17 confirms that the wedge behavior is also highly state dependent. Average and low levels of disagreement are associated with average decreases in the Fed-markets expectation gap, while FOMC announcements when disagreement was high saw significant *increases* in the distance between the expected paths of policy, further corroborating the pervasiveness of the channel I study. Economically, one interpretation is that when disagreement is low, markets treat surprises as information about a shifted path and move toward the Fed, whereas when disagreement is high, they treat surprises as transitory deviations and move away from the Fed, which is exactly the configuration that generates state-dependent attenuation in the output response.[42]

In the next section, I operationalize these insights in a microfounded model of expectation formation featuring monetary disagreement.

[42]This interpretation can be thought as a generalization of Caballero and Simsek (2022)'s notion of opinionated markets: *how* opinionated markets are *depends* on the current prevailing disagreement. In the limit (infinite disagreement), the market is "perfectly" opinionated and discounts completely the significance of the shock.

5 Microfoundations of Disagreement and Learning

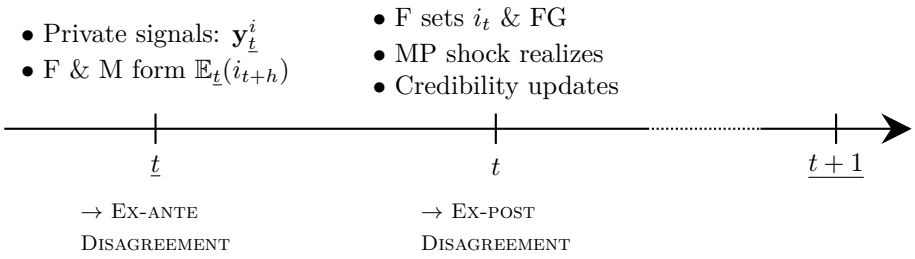


Figure 18: Model Timeline and information flow.

This section microfounds disagreement and the markets' learning dynamics. I embed heterogeneous information, time-varying perceived rule coefficients, and imperfect forward-guidance credibility in a model when the market updates around FOMC announcements using an Extended Kalman Filter. The framework links policy surprises and guidance to belief revisions, mapping into the expectational wedge and ultimately to state-dependent transmission. It nests the information channel and the response to news channel, yielding sharp sign restrictions under which they verify, and delivers estimable objects like updating gains and attribution shares.

Time is discrete and each period is divided in two sub-stages, $\underline{t}$ and t, depicted in Figure 18. Economic fundamentals are denoted by an unobservable $n \times 1$ vector $\mathbf{x}$, that represents the n variables relevant for the conduct of monetary policy (output gap, inflation, ...). At $\underline{t}$, the Market (M) and the Federal Reserve (F) observe private information and form expectations about fundamentals and interest rates. At time t, fundamentals realize, the interest rate decision takes place and the Federal Reserve issues a statement (guidance).

Let the n-dimensional latent state $\mathbf{x}_t$ evolve as:

$$\mathbf{x}_t = \rho\, \mathbf{x}_{t-1} + \boldsymbol{\eta}_t, \qquad \boldsymbol{\eta}_t \sim \text{i.i.d.}(0, \Sigma^{\eta}_t) \tag{25}$$

Throughout, $\rho \in (-1, 1)$ is scalar persistence. The corresponding unobserved vector of policy rule coefficients $\boldsymbol{\Phi}_t$:

$$\boldsymbol{\Phi}_t = \boldsymbol{\Phi}_{t-1} + \boldsymbol{\omega}_t, \qquad \boldsymbol{\omega}_t \sim \text{i.i.d.}(0, \Sigma^{\omega}_t) \tag{26}$$

At $\underline{t}$, each agent $i \in \{F, M\}$ observes a vector of noisy signals about current fundamentals:

$$\mathbf{y}^i_{\underline{t}} = \mathbf{x}_t + \boldsymbol{\varepsilon}^{y,i}_{\underline{t}}, \qquad \boldsymbol{\varepsilon}^{y,i}_{\underline{t}} \sim \text{i.i.d.}(0, \Sigma^{\varepsilon^i}_{\underline{t}}) \tag{27}$$

At t, the interest rate is set according to a monetary rule of the form

$$i_t = \mathbf{\Phi}_t' \, \mathbb{E}_t^F(\mathbf{x}_t) + u_t^G + u_t^S \tag{28}$$

where $\mathbf{\Phi}_t$ is an $n \times 1$ time-varying vector of coefficients, u_t^G represents an *intended* deviation from the rule (e.g., forward guidance) and u_t^S is a zero-mean monetary policy shock. Simultaneously, the Fed releases a (noisy) public signal about the deviation from the rule:

$$\begin{aligned} g_t &= u_t^G + \nu_t, & \nu_t &\sim \text{i.i.d.}(0, \sigma_{\nu,t}^2), \\ u_t^G &= \rho_u \, u_{t-1}^G + \zeta_t, & \zeta_t &\sim \text{i.i.d.}(0, \sigma_{\zeta,t}^2) \end{aligned}$$

In other words, the Fed signals how much it deviates from the rule.[43] The announced deviations are perceived as uncertain by the Market, who attaches noise ν_t to each horizon's announced deviation. The perceived variance of the noise, $\sigma_{\nu,t}^2$, can be thought of as representing the *credibility* of the Fed's guidance. Times of high disagreement correspond to periods where the Market attaches a lower weight to the commitments expressed by the Fed, therefore relying more on its own assessment and (private) information. Thus, for our purposes, one can think of $\sigma_{\nu,t}^2$ as a positive function of the *magnitude* of the expectation gap between markets and Fed (the guidance-side analogue of Assumption 1 below, which imposes the same on the rate-signal noise $R_{i,t}$), reflecting the decreasing reliance on guidance as the two paths drift apart.[44]

Learning The Market's learning deviates from canonical structures. New information is released in two sub-periods, $\underline{t}$ and t, representing respectively the (private) information that the private sector develops in between FOMC announcements (e.g., news analysis, research), and the information that is produced as a consequence of the monetary event (e.g., interest announcement, forward guidance). Because i_t is the product of two unknown vectors, the observation equation is bilinear and the classical (repeated) linear-Gaussian Kalman filter does not return the exact Bayesian posterior. Therefore, I model the Market's learning with an Extended Kalman Filter (EKF).[45] Although I defer to Appendix E.6 a detailed

[43] This type of forward guidance is conceptually close to the notion of *Odyssean* guidance.

[44] One way to characterize this dynamic (i.e., endogenize credibility) is to assume that after i_t is realized, the Market updates $\sigma_{\nu,t}^2$ according to the Fed's (last) forecast error:

$$\begin{aligned} \text{FE}_t^F &\equiv i_t - \mathbb{E}_{\underline{t}}^F(i_t) \\ \sigma_{\nu,t}^2 &= (1-r)\,\sigma_{\nu,t-1}^2 + r\,(\text{FE}_t^F)^2 \qquad r \in (0,1). \end{aligned} \tag{29}$$

Equation (29) captures the idea that credibility declines (increases) when the Fed proves to be wrong (right) in anticipating its monetary policy decisions. The gain coefficient r represents the Market's learning velocity. Credibility-as-updating-weight formulations of this kind go back to Bomfim and Rudebusch (2000).

[45] Excellent treatments of this and of related methodologies are in Simon (2006), or Särkkä and Svensson (2023).

derivation of each posterior and of its variance, I summarily establish the key equations and the logic of the learning.

Before the FOMC meeting takes place, private information is distributed to markets and Fed, representing the differential data and analysis preceding the announcement. As specified in equation (27), this is modeled linearly and featuring Gaussian disturbances.[46] Then, at $\underline{t}$, the Market and the Fed update their expectations on the state vector ($\mathbf{z}_t$) of unknowns ($\mathbf{x}_t, \mathbf{\Phi}_t, u_t^G$), hence forming expectations on future interest rates as follows. For $j \in \{M, F\}$:

$$\mathbf{z}_t \equiv \left(\mathbf{x}_t', \, \mathbf{\Phi}_t', \, u_t^G\right)', \tag{30}$$

$$\hat{\mathbf{z}}^j_{\underline{t}|\underline{t}} = \hat{\mathbf{z}}^j_{\underline{t}|t-1} + K^j_{\underline{t}} \, \tilde{\mu}^j_{\underline{t}}, \tag{31}$$

$$\mathbb{E}^j_{\underline{t}}\left(i_{t+h}\big(\hat{\mathbf{z}}^j_{\underline{t}|\underline{t}}\big)\right) = \hat{\mathbf{\Phi}}^{j\,\prime}_{\underline{t}} \, \rho^h \, \hat{\mathbf{x}}^j_{\underline{t}} + \rho^h_u \hat{u}^{G,j}_{\underline{t},h} + C^j_{\underline{t},h}. \tag{32}$$

where $K^j_{\underline{t}} \in \mathbb{R}^{(2n+1)\times n}$ is the Kalman gain mapping the innovation into the state update, $\tilde{\mu}^j_{\underline{t}}$ represents the unexpected component of the measurements (the private signals), and $C^j_{\underline{t},h}$ a covariance adjustment term capturing the product of random variables.

The second batch of observations is released at t, when the interest rate is announced and the Fed provides forward guidance on future deviations – u^G_{t+h} from the standard Taylor rule. As I discuss formally in Appendix E.6, this is the stage where the EKF applies. In fact, the interest rate realization constitutes a nonlinear signal, which needs a specific extraction procedure, while the guidance announcements are specified in the conventional form of "truth plus noise" signals, and therefore can be processed separately. Conceptually, the Market updates by linearizing the interest rate observation around its intermediate posteriors ($\bar{z}$).

By definition, the innovation with respect to the interest rate observation *is* the monetary policy surprise:

$$\begin{aligned} \text{MPS}_t &\equiv i_t - \mathbb{E}^M_{\underline{t}}(i_t) \\ &= \left[\hat{\mathbf{\Phi}}^{F\prime}_{\underline{t}} \, \hat{\mathbf{x}}^F_{\underline{t}} - \hat{\mathbf{\Phi}}^{M\prime}_{\underline{t}} \, \hat{\mathbf{x}}^M_{\underline{t}}\right] + (u^G_t - \hat{u}^{G,M}_{\underline{t},0}) - C^M_{\underline{t},0} + u^S_t \end{aligned} \tag{33}$$

Monetary policy surprises arise not only from true, exogenous policy shocks, but also from markets' imperfect knowledge of fundamentals, the Fed's rule and guidance. This feature captures the notion that (ex ante) disagreement about *any* element of the policy rule affects the identification of monetary shocks: high-frequency surprises embed belief revisions about the reaction function and forward-guidance

[46] Strictly speaking, Gaussian errors are not required to justify Kalman's optimality: in a linear state–space model with zero-mean noises of known covariances, the Kalman filter is the best linear minimum-variance estimator. Gaussianity is only needed for global MMSE/Bayes optimality.

terms.[47]

Together with the guidance, MPS_t is used to update the Market's posterior at t:[48]

$$\hat{\mathbf{z}}_{t|t} = \bar{z} + K_i \,\text{MPS}_t, \tag{34}$$

$$\mathbb{E}_t^M(i_{t+h}(\hat{\mathbf{z}}_{t|t})) = \hat{\mathbf{\Phi}}_t^{M\prime}\,\rho^h \hat{\mathbf{x}}_t^M + \rho_u^h \hat{u}_{t,0}^{G,M} + C_{t,h}^M \tag{35}$$

therefore revising its beliefs about the fundamentals, the policy coefficients and the deviations from the rule optimally (Bayesian). Notice that, in contrast to standard linear filtering, the sign of the shock and of the update can differ: for instance, in response to a positive surprise ($\text{MPS}_t > 0$), the Market might revise their expectations of the policy rate downward ($K_i < 0$) depending on the contemporaneous revisions of fundamentals and guidance. I elaborate on these and related issues next in Section 5.1.

The Fed only updates at $\underline{t}$, as it perfectly observes *its* intended deviation. This is consistent with the empirical analysis where I employ Tealbook forecasts, which do not vary between $\underline{t}$ and t (the FOMC announcement window): $\mathbb{E}_t^F(i_{t+h}) \equiv \mathbb{E}_{\underline{t}}^F(i_{t+h})$ and is equivalent to equation (32). Then, ex post disagreement (henceforth, simply *disagreement*) is:

$$\begin{aligned} D_{t+h|t} &= \mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^M(i_{t+h}) \\ &= \rho^h\Big(\hat{\mathbf{\Phi}}_t^{F\prime}\,\hat{\mathbf{x}}_t^F - \hat{\mathbf{\Phi}}_t^{M\prime}\,\hat{\mathbf{x}}_t^M\Big) + \rho_u^h\big(\hat{u}_{t,0}^{G,F} - \hat{u}_{t,0}^{G,M}\big) + \big(C_{t,h}^F - C_{t,h}^M\big) \end{aligned} \tag{36}$$

Thus, disagreement has four sources: (*i*) differing assessments of future fundamentals; (*ii*) differing views of the monetary policy rule; (*iii*) imperfect credibility of forward guidance; and (*iv*) a composite covariance term induced by the policy rule's multiplicative structure. Monetary policy surprises update markets' beliefs about the future policy path — through the posteriors $\hat{\mathbf{z}}_{t|t}$ — and thereby feed into disagreement.

Remarks. Building on (36), the framework nests familiar benchmarks by removing frictions in turn. If fundamentals $\mathbf{x}_t$ are observed, disagreement reflects only heterogeneity in perceived policy rules and guidance beliefs:

$$D_{t+h|t} = \rho^h\,\mathbf{x}_t\Big(\hat{\mathbf{\Phi}}_t^{F\prime} - \hat{\mathbf{\Phi}}_t^{M\prime}\Big) + \rho_u^h\big(\hat{u}_{t,0}^{G,F} - \hat{u}_{t,0}^{G,M}\big).$$

[47] It is worth pointing out the reason why $C_{\underline{t},h}^M$ shows in MPS_t but not the Fed's analog, $C_{\underline{t},h}^F$: as it is clear from (28), the current Taylor coefficients are modeled as a *choice* of the Fed, not as an unknown. On the contrary, *future* coefficients are indeed unknown to the Fed, representing possible changes in the FOMC composition, the periodic framework renewals, or political pressure, as documented in Drechsel (2024).

[48] Being it linear, the order is irrelevant and guidance is embedded first in the linearized intermediate posterior $\bar{z}$.

If, in addition, the coefficients are common knowledge, the rule term drops and disagreement is driven solely by differential beliefs about guidance shocks:

$$D_{t+h|t} = \rho_u^h \big(\hat{u}_{t,0}^{G,F} - \hat{u}_{t,0}^{G,M}\big).$$

With perfect credibility on top of observability and known coefficients, those belief gaps vanish and

$$D_{t+h|t} = 0.$$

Intuitively, disagreement is a wedge generated by three frictions: unobserved fundamentals, uncertainty about the policy rule, and imperfect credibility; removing them collapses the wedge and restores the common-expectations benchmark.

5.1 Signal Extraction and Macro Propagation

Next, I ask how the model generates state-dependent propagation to reconnect with the evidence of attenuation. As shown in Section 4.2.1, the model places wedge dynamics at the core of monetary transmission: a shock perturbs financial markets' beliefs in a way that can *undo* the familiar pass-through to the economy. More precisely, at t (around an FOMC announcement), the wedge dynamics implied by the Market's signal extraction are characterized by:[49]

$$\Delta W_t \equiv W_t - W_{\underline{t}} = -\frac{1}{\sigma}\Bigg\{\left[\frac{\bar{\mathbf{\Phi}}'}{1-\rho}\,K_{g,t}^{\mathbf{x}} + \frac{\bar{\mathbf{x}}'}{1-\rho}\,K_{g,t}^{\mathbf{\Phi}} + \frac{1}{1-\rho_u}\,K_{g,t}^{u}\right]\tilde{\mu}_t^g + \left[\frac{\bar{\mathbf{\Phi}}'}{1-\rho}\,K_{i,t}^{\mathbf{x}} + \frac{\bar{\mathbf{x}}'}{1-\rho}\,K_{i,t}^{\mathbf{\Phi}} + \frac{1}{1-\rho_u}\,K_{i,t}^{u}\right]\tilde{\mu}_t^i\Bigg\} \tag{37}$$

where $\tilde{\mu}_g$ and $\tilde{\mu}_i$ denote the innovations in forward guidance and the policy-rate announcement, respectively, and $K_g^{(\cdot)}$ and $K_i^{(\cdot)}$ are the state-specific Kalman gains that weight these innovations in the update step. Notice that (37) reflects the effects of disagreement on learning through the state-dependent weighting of the shocks.

For illustration, focus on the pure model-implied monetary innovation $\tilde{\mu}_t^i$. When disagreement is high, the gains associated with the interest rate announcement ($K_i^{(\cdot)}$) fall, rationally attributing more importance on the Market's prior.[50] In fact, higher disagreement leads the Market to attribute more of a policy move to transitory, idiosyncratic factors than to a genuine shift in the path it should track

[49] Equation (37) uses the strict definition of the wedge, featuring the *infinite* sum of Fed-Market disagreements; a finite-horizon version is qualitatively identical, and only affects the gains' parametric weights. See Appendix E.7 for details.

[50] Formally, this reflects in higher $K_x^{(\cdot)}$, which do not show explicitly in (37), but are embedded in the weights $[\bar{\mathbf{\Phi}}, \bar{\mathbf{x}}, 1]$ through the Jacobian underlying the linearization step of the EKF. See E.6 for details.

— it becomes, in the language of Caballero and Simsek (2022), more *opinionated*. A surprise read as noise is less informative, so the Market revises its beliefs by less, the wedge responds by less, and transmission is attenuated. This state-contingent filtering also addresses the natural rationality objection to the expected-rate channel behind the evidence of Section 3 — why would markets' expected paths persistently deviate from the policy that realizes? — since with disagreement-dependent gains such deviations are the rational response to announcements perceived as noisy, not a bias that experience should erase. Assumption 1 formalize this.

Assumption 1. The Market's perceived variance of the transitory, unanticipated component of the policy rate is increasing in the magnitude of the predetermined Fed–Market disagreement.[51]

The magnitude form is not imposed: credibility erodes with the Fed's forecast errors, and the Market expects those errors to be larger the further the Fed's projection sits from its own, so it is the *size* of the gap, not its sign, that governs attenuation. This restriction is consistent with the updating evidence: augmenting the wedge local projection (23) with a signed interaction (Appendix D.13), the wedge's response to surprises shows no robust dependence on the direction of the gap — the signed loading is indistinguishable from zero at every horizon, while the magnitude loading carries attenuation-consistent state dependence.[52] The same logic applies to the guidance block: its signals are discounted more under high disagreement.

Denoting by $d = (H, L)$ the states of high and low disagreements, suppose $\tilde{\mu}_t^i > 0$ (positive rate innovation, i.e. restrictive shock):

$$K_{i,t}^{(\cdot),H} < K_{i,t}^{(\cdot),L} \quad \Longrightarrow \quad \Delta W_t^H > \Delta W_t^L \quad \equiv \quad \Psi_w^H > \Psi_w^L \tag{38}$$

mirroring the New Keynesian attenuation derived in (20) and tested in (23). Hence, (38) establishes that the model's microfoundations naturally deliver state-dependent propagation of monetary shocks. Moreover, the closed-form expressions for the Market's signal extraction pin down the conditions under which each state variable is revised, and in which direction. The following lemma summarizes.

[51] A sufficient primitive is the same forecast-error mechanism that governs guidance credibility in (29), here applied to the rate-signal noise: $\sigma^2_{u^S,t} = (1-r)\,\sigma^2_{u^S,t-1} + r\,(\mathrm{FE}_t^F)^2$, with $\mathrm{FE}_t^F = i_t - \mathbb{E}_{\underline{t}}^F(i_t) = \mathrm{MPS}_t - D_{t|\underline{t}}$. Squaring and taking $\mathbb{E}_{\underline{t}}^M$, the cross term vanishes because $\mathbb{E}_{\underline{t}}^M(\mathrm{MPS}_t) = 0$ and $D_{t|\underline{t}}$ is predetermined, leaving $\mathbb{E}_{\underline{t}}^M\left[(\mathrm{FE}_t^F)^2\right] = \mathbb{E}_{\underline{t}}^M(\mathrm{MPS}_t^2) + D_{t|\underline{t}}^2$ — even in and increasing in $|D_{t|\underline{t}}|$, so credibility erodes with the *magnitude* of the gap.

[52] Within the announcement window, the signed loading is at best marginally significant and not robust to controls or horizon-composition choices, while the powered horizon projection shows no signed loading at any horizon (Figure D30, Appendix D.13). Section 4.2.1 estimates the magnitude-only version of (23).

Lemma: Signs of Updates. Under the learning model,[53]

$$\begin{aligned}
\operatorname{sign}\big(\Delta\,\mathbb{E}_t^M(\mathbf{x}_t)\big) &= \operatorname{sign}(\tilde{\mu}_t^i)\ \cdot\ \operatorname{sign}\big(P_t^{\mathbf{x}}\bar{\mathbf{\Phi}} + P_t^{\mathbf{x}\mathbf{\Phi}}\bar{\mathbf{x}} + P_t^{\mathbf{x}u}\big),\\
\operatorname{sign}\big(\Delta\,\mathbb{E}_t^M(\mathbf{\Phi}_t)\big) &= \operatorname{sign}(\tilde{\mu}_t^i)\ \cdot\ \operatorname{sign}\big(P_t^{\mathbf{\Phi}\mathbf{x}}\bar{\mathbf{\Phi}} + P_t^{\mathbf{\Phi}}\bar{\mathbf{x}} + P_t^{\mathbf{\Phi}u}\big),\\
\operatorname{sign}\big(\Delta\,\mathbb{E}_t^M(u_t^G)\big) &= \operatorname{sign}(\tilde{\mu}_t^i)\ \cdot\ \operatorname{sign}\big(P_t^{u\mathbf{x}}\bar{\mathbf{\Phi}} + P_t^{u\mathbf{\Phi}}\bar{\mathbf{x}} + P_t^{u}\big).
\end{aligned}$$

where $P_t^{\mathbf{ab}}$ are the $(\mathbf{a},\mathbf{b})$ block-entries of the prior covariance matrix $P_{t|\underline{t}} = \operatorname{Var}\big(\mathbf{z}_t' \mid \Omega_{\underline{t}}\big)$, $[\bar{\mathbf{x}}, \bar{\mathbf{\Phi}}, 1]$ are prior means.[54] The sign conditions are estimable (next section) and disciplined by matrix P. They nest the coefficients-only updates (Bauer and Swanson, 2023b) or the fundamentals-only updates (Melosi, 2017, Miranda-Agrippino and Ricco, 2021) when cross-covariances are zero, and they enable transparent attribution of revisions across $\mathbf{x_t}$, $\mathbf{\Phi}_t$, and u_t^G.

For instance, consider a single state's update, Market's beliefs of the policy rule coefficients in a simple univariate rule that only responds to output. The lemma implies that, in response to a positive (*hawkish*) monetary surprise, markets raise their assessment of Fed's reactivity to fundamentals (ϕ_t) if and only if

$$P_t^{x\phi} > -\frac{P_t^{\phi\phi}\bar{x} + P_t^{\phi u}}{\bar{\phi}}.$$

Economically, when markets believe the economy is expanding ($\bar{x} > 0$), they revise their expectations of ϕ upward after observing a higher-than-expected rate ($\mu^i > 0$), provided the cross-covariance $P^{x\phi}$ is not negative enough to flip the sign.[55] Research has explored the empirical plausibility of negative comovements between fundamentals and rule coefficients. Recessionary and ZLB periods are associated with more aggressive or less gradual policy (Cogley and Sargent, 2001, Aastveit et al., 2024, Bauer et al., 2024). These can stem from weaker policy power during contractions (Tenreyro and Thwaites, 2016), or from omitted time variation in the natural rate of interest, which makes Taylor rule regressions load less on the systematic component and more on the intercept of monetary response (Holston et al., 2017). Moreover, evidence of asymmetric central bank preferences that weight negative gaps more than positive ones (Surico, 2004) can also generate inverse comovements.

Beyond their empirical relevance, these results discipline influential theoretical narratives. Consider the popular information channel, i.e., the idea that markets interpret monetary shocks as revealing (private) information about future fundamentals, yielding revisions in counterintuitive directions. These narratives have been corroborated by survey data showing that hawkish surprises are often fol-

[53] Applying the sign operator componentwise.

[54] More precisely, each bracket equals the prior covariance between the state and the linearized rate $\bar{\mathbf{\Phi}}\,\mathbf{x}_t + \bar{\mathbf{x}}\,\mathbf{\Phi}_t + u_t^G$; the EKF update is a covariance–variance projection, so its sign is $\operatorname{sign}(\tilde{y}_i)$ times the sign of that covariance.

[55] Assuming (naturally) a strictly positive $\bar{\phi}$.

lowed by upward revisions in output forecasts.[56] The model nests this notion and characterizes the structural conditions under which it is verified. As in the previous example, the lemma implies that, given a $\tilde{\mu}_t^i > 0$, the information channel requires

$$P_t^{\mathbf{x}}\bar{\mathbf{\Phi}} + P_t^{\mathbf{x}\mathbf{\Phi}}\bar{\mathbf{x}} + P_t^{\mathbf{x}u} > 0 \tag{39}$$

that is, shocks and the state-specific Kalman gains must share the same sign (componentwise). Then, the conventional information channel story is a special case where the block cross-covariances are assumed equal to zero ($P_t^{\mathbf{x}}\bar{\mathbf{\Phi}} = P_t^{\mathbf{x}u} = 0$), so the sign of the update matches the sign of the shock, since $P_t^{\mathbf{x}}\bar{\mathbf{\Phi}}$ is strictly positive.

5.2 Illustration: Estimation and Wedge Primitives

Finally, for exposition, I begin by estimating a two–dimensional state-space for output gap and the corresponding policy rule coefficient (abstracting from forward guidance). This simplification is (nearly) without loss for the rate–surprise block because, given the timing assumed in Figure 18, the filter is invariant to the order between the processing of the interest rate announcement and the signal g_t.[57] Dropping u_t^G tightens posteriors and clarifies mechanisms, delivering cleaner attribution and propagation results. The full three–state EKF is estimated in Appendix E.8. The exercise quantifies how FOMC news reallocates markets' beliefs across the two primitives, i.e., the state of the economy x_t and the slope of the policy rule ϕ_t. The key measurement is the policy rate at announcement time, which, consistently with the general framework discussed above, is a nonlinear function in the states: $i_t = \phi_t x_t + u_t^S$. I treat (x_t, ϕ_t) as latent and let both evolve as stochastic processes. The two stages are unvaried: at $\underline{t}$, markets form expectations about current activity; as before, I consider the yield curve slope (10Y-2Y) as priced-in expectations of the future state of the economy, i.e., a proxy of $y_{\underline{t}}^M$. At t, the interest rate is announced, and inference is drawn through an EKF that updates beliefs using the meeting-specific intermediate priors, formally represented by the Jacobian $H_{i,t} = [\bar{\phi}_t, \bar{x}_t]$.[58]

The EKF delivers time-varying Kalman gains $K_{i,t} = [K_{i,t}^x, K_{i,t}^\phi]'$ and posterior means $(\hat{x}_t, \hat{\phi}_t)$. These are sufficient to map a given surprise into revisions of the wedge through the equivalent expression of equation (37) for the bi-dimensional system. In turn, wedge dynamics impact monetary transmission as per equation

[56] See Bauer and Swanson (2023b) for an exhaustive discussion (and debunking) of this idea.

[57] Technically, u_t^G's residual uncertainty only inflates S_i, leaving the sign conditions below unchanged. Equivalently, with the guidance signal omitted, the unresolved announced deviation is absorbed into the rate equation's measurement variance $R_{i,t}$, so the credibility channel operates there directly: $S_{i,t}$ is affine and strictly increasing in it, and the gain falls.

[58] This extra measurement pins down fundamentals. Without x_t^M, the rate equation $i_t = \phi_t x_t + u_t$ largely identifies only the product $\phi_t x_t$. Observing daily data on the yield slope disciplines the precision of x_t^M and the normal comovement of x_t^M and i_t; conditional on that anchor, meeting-day rate jumps load on ϕ_t.

(18).

Denote by $\tilde{\mu}_t^i$ the innovation in the policy rate. This is used by markets to update their expectations of current activity and policy responsiveness. Each update is in turn mapped into a revised expectation of *future* policy rates, therefore impacting the wedge. I summarize the impact of markets' revisions on the expectational wedge as

$$\Delta W_t = \frac{\sigma^{-1}}{1-\rho}\Big(\Delta W_t^x + \Delta W_t^\phi\Big)$$
$$\Delta W_t^x = \bar{\phi}_t K_{i,t}^x \tilde{\mu}_t^i, \qquad \Delta W_t^\phi = \bar{x}_t K_{i,t}^\phi \tilde{\mu}_t^i \tag{40}$$

These model-implied, time-varying objects encode the determinants of learning: (*i*) the leverage of each state in the policy rule ($\bar{\phi}_t$, $\bar{x}_t$); (*ii*) the informativeness of the rate surprise ($K_{i,t}^x$, $K_{i,t}^\phi$); and (*iii*) the sign and size of the shock $\tilde{\mu}_t^i$.

The decomposition in (40) allows to report time-varying contributions to the wedge dynamics:

$$\mathcal{C}_t^x = \frac{\Delta W_t^x}{\Delta W_t^x + \Delta W_t^\phi}, \qquad \mathcal{C}_t^\phi = 1 - \mathcal{C}_t^x \tag{41}$$

which attribute the wedge revision to fundamentals versus policy rule in each meeting. A value of $\mathcal{C}_t^x = 1$ means the wedge dynamics are entirely due to belief updates about x_t; notice that, due to potentially negative cross-covariances, $\mathcal{C}_t^x$ and $\mathcal{C}_t^\phi$ are not bounded to the unit interval, but can take negative or above-unity values. The sign of each revision is pinned down by $K_{i,t}$ and $H_{i,t}$, and the magnitude scales with the prevailing disagreement.

The EKF nests standard mechanisms of "coefficients-only'' and "fundamentals-only'' update, and makes their underlying conditions explicit. While these are respectively *related to* the response to news channel (Bauer and Swanson, 2023a,b) and the information channel of monetary shocks (Nakamura and Steinsson, 2018), the next section elaborates on a relevant distinction based on the difference between wedge revisions and belief updates.
By generalizing both mechanisms in a single system linking shocks, learning, and transmission, the model allows to attribute changes in the policy rate expectations (hence in the wedge) to their primitives. Figure 19 depicts the model-implied wedge movements at meeting, together with the relative contributions of output and coefficient revisions.

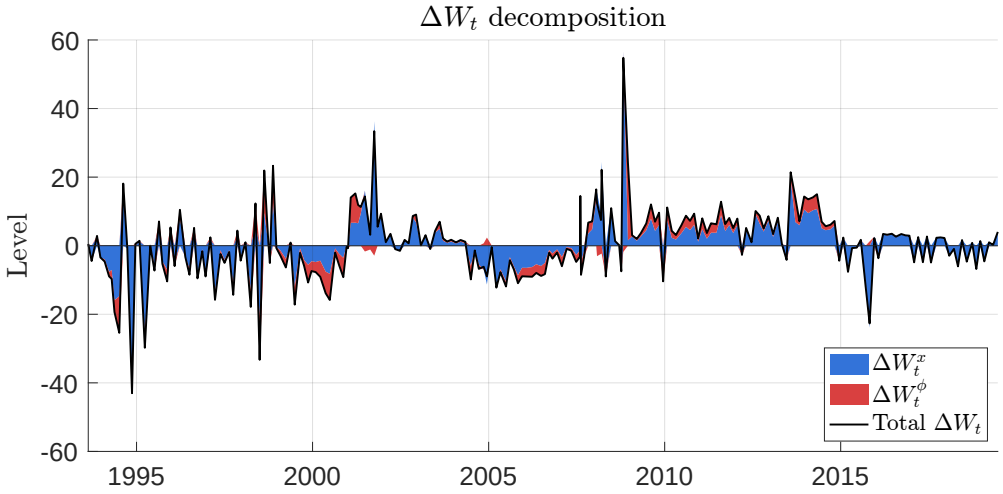


Figure 19: ΔW_t decomposition into fundamentals (blue) and policy-slope revisions (red).

The total wedge movement (black) closely tracks the x–channel (blue), with occasional but short-lived ϕ contributions (red). Signs often coincide—both channels push the wedge in the same direction—yet around stress/ZLB episodes the red bars emerge and sometimes offset the blue, producing smaller totals despite sizable surprises. The mapping from beliefs to the wedge is scaled by the linearization weights, with $\bar{\phi}_t$ being an order of magnitude higher than $\bar{x}_t$. Thus, comparable belief revisions load more strongly on ΔW_t through the x–term. This is exactly what we see in Figure 19.

A useful way to read it is as follows. Let us consider again the post-GFC era: negative realizations of ΔW_t in indicate that, after dovish surprises, markets systematically mark down the expected path of future rates. This, so the (signed) wedge becomes more negative. The decomposition clarifies why. The blue area shows that this is largely driven by revisions to fundamentals x_t: news about the depth and persistence of slack leads investors to infer weaker demand and lower real rates for longer, which they translate into a lower policy path. The red area, ΔW_t^ϕ, captures revisions to the perceived rule: conditional on that slack, markets infer a more aggressive reaction coefficient ϕ_t, i.e. stronger accommodation (QE, "lower for longer'' rates, etc.), thereby amplifying the wedge movement. Although a simplification – the filter is estimated without data on inflation and forward guidance – this matches precisely the dynamics of the wedge documented in the previous section and the empirically verified learning by the market, gradually lowering its expected policy path.

Quantitatively, the height of the blue and red bars is pinned down by the EKF's variance–covariance matrix: ΔW_t^x scales with the gain on x_t, $K_{i,t}^x$, while ΔW_t^ϕ scales with the gain on ϕ_t, $K_{i,t}^\phi$, both multiplied by the same rate surprise and the corresponding linearization weight. Periods in which the blue component dominates are those where the filter views fundamentals as relatively uncertain and

the rule as tightly pinned down, whereas visible red spikes occur when posterior uncertainty (or covariance) about the policy slope is high, so a given surprise is interpreted partly as news about ϕ_t rather than x_t. The next section clarifies on the difference between propagation (wedge) attribution and beliefs revision.

Belief Updates Note that the wedge–attribution shares answer *which channel moves the wedge*, not *which belief moved the most*. To study belief revisions directly, I therefore focus on the Kalman gains on the rate equation, $K^x_{i,t}$ and $K^\phi_{i,t}$, which map a given policy-rate innovation $\tilde{\mu}^i_t$ into state updates, $\Delta x_t = K^x_{i,t}\tilde{\mu}^i_t$ and $\Delta\phi_t = K^\phi_{i,t}\tilde{\mu}^i_t$. Figure 20 plots these gains with separate vertical axes (blue for x, red for ϕ), emphasizing that they are expressed in different units and their levels are not directly comparable across states.

The pattern is intuitive. Because x_t is relatively volatile and imperfectly observed, the filter assigns it a sizeable and time-varying gain (typically around 0.4–0.9), so rate surprises are interpreted mainly as news about current conditions. By contrast, ϕ_t is modeled as a smooth latent coefficient with low process variance, so $K^\phi_{i,t}$ is small and close to zero in normal times: the rule's slope is revised only when shocks are large and persistent, as around the GFC and early ZLB/QE period. Occasional spikes in $K^\phi_{i,t}$ mark episodes in which the surprise is read as a change in the reaction function itself (e.g. regime shifts in forward guidance), rather than as pure news about the state.
Empirically, this separation is disciplined by having two observables with distinct loadings—x^M_t loads primarily on fundamentals, while the policy rate loads on the product $\phi_t x_t$—so that rate surprises that cannot be reconciled with contemporaneous x^M_t (given its estimated precision) are interpreted as revisions to ϕ_t rather than to x_t.

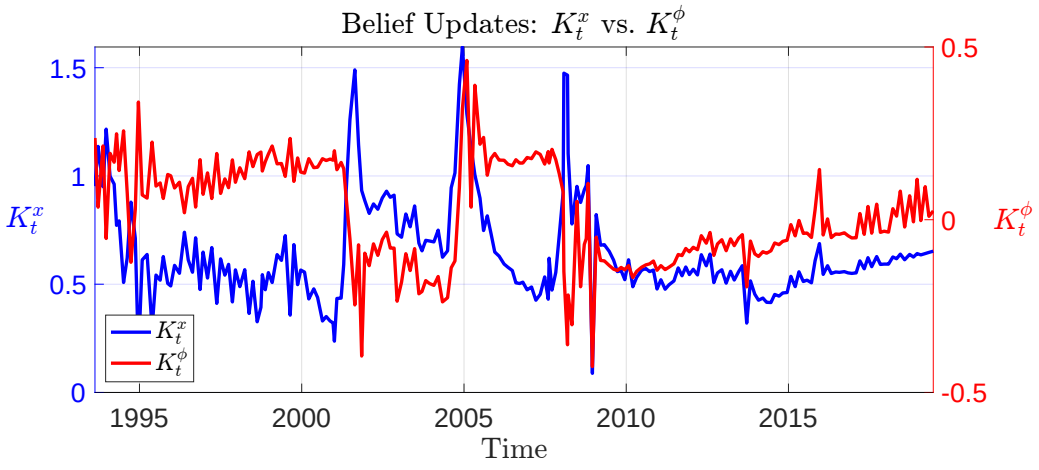


Figure 20: Belief Updates attribution: x (blue) vs. ϕ (red).

A useful way to see this is through the post–GFC episode. The filter generates policy surprises that are predominantly dovish, consistently with the market being persistently surprised by lower interest rates (Figure 1), with futures initially price a faster lift–off than the Fed. When $K^{\phi}_{i,t}$ is slightly negative, each dovish surprise raises the estimated slope ϕ_t ($\Delta\phi_t = K^{\phi}_{i,t}\tilde{\mu}^i_t > 0$): for a given amount of slack, markets infer that the Fed cuts more, or keeps rates lower for longer, than previously thought. At the same time, a large $K^{x}_{i,t}$ pushes $\hat{x}_t$ further into negative territory, reflecting worse-than-expected fundamentals. Over time, as ZLB/QE policy becomes anticipated and ϕ_t is re-estimated, $K^{\phi}_{i,t}$ shrinks toward zero and revisions taper off. The EKF thus provides a structural state-space counterpart to the empirical evidence that repeated dovish surprises at the ZLB lead markets to reprice both slack and a steeper, more aggressive policy rule.

6 Conclusions

This paper documents and explains how Fed–Market disagreement about the policy path shapes the transmission of U.S. monetary policy. I construct a real-time, meeting-level measure of disagreement that spans the period 1990-2019. Adding this measure to standard VARs improves short- and medium-horizon forecasts of activity, inflation, and the policy rate. Using lag-augmented local projections, I show that disagreement attenuates the effects of monetary shocks on real activity and can even flip their sign, providing a belief-based rationale for outstanding puzzles like the activity puzzle. Quantitatively, a 25 basis points increase in average disagreement reduces the magnitude of the industrial production response at a 12-month horizon by about 61%. The attenuation depends on the magnitude, not the sign, of the gap, and it is robust to alternative shock series, richer systems including financial conditions and unemployment, and controls for uncertainty.

I then assess the information-channel premise. Aligning futures-implied expectations to Tealbook cutoffs, markets forecast the federal funds rate more accurately than the Fed across 1–8 quarters. Rolling information-advantage tests show no systematic Fed edge, with only brief episodes mostly concentrated in the ZLB era. Together, these results weaken the view that puzzling announcement responses reflect the Fed's private news.

To organize the evidence, I derive a generalized New Keynesian framework in which disagreement appears as a wedge in the Euler equation. The wedge scales demand elasticities and delivers state-dependent impulse responses that mirror the empirical patterns. I then endogenize disagreement through meeting-by-meeting learning. A two-stage filter around each announcement uses the rate surprise and guidance to update beliefs about fundamentals, the policy rule, and credibility. This yields time-varying weights on news and a decomposition of the wedge into its sources. The framework nests "information" and "response-to-news" channels as special cases and tells us which channel moved expectations at each meeting.

Two lessons follow. First, communication and credibility are policy levers: clear, stable guidance narrows the beliefs gap and restores standard transmission. Second, monitoring disagreement should be routine: measuring the gap and its sources improves policy design and the evaluation of guidance. The core point is simple: beliefs about the policy path are a first-order state variable, and aligning them gives monetary policy the authority to shape expectations and the economy.

References

Aastveit, Knut Are, Jamie Cross, Francesco Furlanetto, and Herman K van Dijk (2024). *Asymmetric gradualism in US monetary policy*. Tech. rep. Tinbergen Institute Discussion Paper.

Acosta, Miguel, Connor M. Brennan, and Margaret M. Jacobson (May 2024). "Constructing high-frequency monetary policy surprises from SOFR futures". en. In: *Finance and Economics Discussion Series* 2024-034, pp. 1–9. ISSN: 1936-2854, 2767-3898. DOI: 10.17016/feds.2024.034. URL: https://www.federalreserve.gov/econres/feds/constructing-high-frequency-monetary-policy-surprises-from-sofr-futures.htm (visited on 06/23/2025).

Amodeo, Francesco (2024). "Strategic Forecasting". In: *Available at SSRN 5145169*.

— (2025). "A Corrected Measure of Orthogonalised Monetary Policy Surprises". In.

Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoît Mojon (2019). "Forward Guidance and Heterogeneous Beliefs". In: *American Economic Journal: Macroeconomics* 11.3, pp. 1–29. DOI: 10.1257/mac.20180141.

Andrade, Philippe and Hervé Le Bihan (2013). "Inattentive professional forecasters". In: *Journal of Monetary Economics* 60.8, pp. 967–982.

Angeletos, George-Marios and Chen Lian (2018). "Forward Guidance without Common Knowledge". In: *American Economic Review* 108.9, pp. 2477–2512. DOI: 10.1257/aer.20161996.

Angrist, Joshua D and Jörn-Steffen Pischke (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Baker, Scott R, Nicholas Bloom, and Steven J Davis (2016). "Measuring economic policy uncertainty". In: *The quarterly journal of economics* 131.4, pp. 1593–1636.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam (2024). "Perceptions about monetary policy". In: *The Quarterly Journal of Economics* 139.4, pp. 2227–2278.

Bauer, Michael D and Eric T Swanson (2023a). "A reassessment of monetary policy surprises and high-frequency identification". In: *NBER Macroeconomics Annual* 37.1, pp. 87–155.

— (2023b). "An alternative explanation for the "fed information effect"". In: *American Economic Review* 113.3, pp. 664–700.

Benchimol, Jonathan, Yossi Saadon, and Nimrod Segev (2023). "Stock market reactions to monetary policy surprises under uncertainty". In: *International Review of Financial Analysis* 89, p. 102783.

Bernanke, Ben S and Kenneth N Kuttner (2005). "What explains the stock market's reaction to Federal Reserve policy?" In: *The Journal of finance* 60.3, pp. 1221–1257.

Bhandari, Anmol, Jaroslav Borovička, and Paul Ho (2016). *Identifying ambiguity shocks in business cycle models using survey data*. Tech. rep. National Bureau of Economic Research.

Bianchi, Francesco, Sydney C Ludvigson, and Sai Ma (2022). "Belief distortions and macroeconomic fluctuations". In: *American Economic Review* 112.7, pp. 2269–2315.

Blanchard, Olivier Jean and Charles M Kahn (1980). "The solution of linear difference models under rational expectations". In: *Econometrica: Journal of the Econometric Society*, pp. 1305–1311.

Blinder, Alan S (2023). "Landings, soft and hard: The federal reserve, 1965–2022". In: *Journal of Economic Perspectives* 37.1, pp. 101–120.

Blinder, Alan S, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen (2008). "Central bank communication and monetary policy: A survey of theory and evidence". In: *Journal of economic literature* 46.4, pp. 910–945.

Bloom, Nicholas (2014). "Fluctuations in uncertainty". In: *Journal of economic Perspectives* 28.2, pp. 153–176.

Bomfim, Antulio N. and Glenn D. Rudebusch (2000). "Opportunistic and Deliberate Disinflation under Imperfect Credibility". In: *Journal of Money, Credit and Banking* 32.4, pp. 707–721.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020). "Overreaction in macroeconomic expectations". In: *American Economic Review* 110.9, pp. 2748–82.

Bordo, Michael D and Klodiana Istrefi (2023). "Perceived FOMC: The making of hawks, doves and swingers". In: *Journal of Monetary Economics* 136, pp. 125–143.

Caballero, Ricardo J. and Alp Simsek (July 2022). "Monetary Policy with Opinionated Markets". en. In: *American Economic Review* 112.7, pp. 2353–2392. ISSN: 0002-8282. DOI: 10.1257/aer.20210271. URL: https://pubs.aeaweb.org/doi/10.1257/aer.20210271 (visited on 07/02/2024).

Caldara, Dario and Edward Herbst (2019). "Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs". In: *American Economic Journal: Macroeconomics* 11.1, pp. 157–192.

Campbell, Jeffrey R, Charles L Evans, Jonas DM Fisher, Alejandro Justiniano, Charles W Calomiris, and Michael Woodford (2012). "Macroeconomic effects of federal reserve forward guidance [with comments and discussion]". In: *Brookings papers on economic activity*, pp. 1–80.

Campbell, Jeffrey R., Jonas D. M. Fisher, Alejandro Justiniano, and Leonardo Melosi (Apr. 2017). "The Event-Study Activity Puzzle". In: *Economic Perspec-*

tives 41.3. URL: https://www.chicagofed.org/publications/economic-perspectives/2017/3.

Campbell, John Y. and N. Gregory Mankiw (1989). "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence". In: *NBER Macroeconomics Annual 1989, Volume 4*. Ed. by Olivier J. Blanchard and Stanley Fischer. Cambridge, MA: MIT Press, pp. 185–246. DOI: 10.1086/654107.

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

Cieslak, Anna and Andreas Schrimpf (2019). "Non-monetary news in central bank communication". In: *Journal of International Economics* 118, pp. 293–315. DOI: 10.1016/j.jinteco.2019.01.012.

Clarida, Richard, Jordi Gali, and Mark Gertler (1999). "The science of monetary policy: a new Keynesian perspective". In: *Journal of economic literature* 37.4, pp. 1661–1707.

Cogley, Timothy and Thomas J Sargent (2001). "Evolving post-world war II US inflation dynamics". In: *NBER macroeconomics annual* 16, pp. 331–373.

Couture, Cody (2021). "Differences between central bank and private sector policy rate forecasts: Causes and implications for monetary policy". In: *Available at SSRN 4050133*.

Crawford, Vincent P and Joel Sobel (1982). "Strategic information transmission". In: *Econometrica: Journal of the Econometric Society*, pp. 1431–1451.

D'Agostino, Antonello and Karl Whelan (2008). "Federal Reserve information during the great moderation". In: *Journal of the European Economic Association* 6.2-3, pp. 609–620.

Diebold, Francis X and Robert S Mariano (2002). "Comparing predictive accuracy". In: *Journal of Business & economic statistics* 20.1, pp. 134–144.

Diercks, Anthony and Uri Carl (2019). "A simple macro-finance measure of risk premia in fed funds futures". In.

Dong, Ding, Zheng Liu, Pengfei Wang, and Min Wei (2025). *Inflation Disagreement Weakens the Power of Monetary Policy*. Working Paper 2024-27. Federal Reserve Bank of San Francisco. DOI: 10.24148/wp2024-27.

Drechsel, Thomas (2024). *Estimating the effects of political pressure on the fed: a narrative approach with new data*. Tech. rep. National Bureau of Economic Research.

Eggertsson, Gauti B et al. (2003). "Zero bound on interest rates and optimal monetary policy". In: *Brookings papers on economic activity* 2003.1, pp. 139–233.

Ellison, Martin and Thomas J Sargent (2012). "A Defence of the FOMC". In: *International Economic Review* 53.4, pp. 1047–1065.

Eusepi, Stefano and Bruce Preston (2010). "Central bank communication and expectations stabilization". In: *American Economic Journal: Macroeconomics* 2.3, pp. 235–271.

Flavin, Marjorie A. (Oct. 1981). "The Adjustment of Consumption to Changing Expectations about Future Income". In: *Journal of Political Economy* 89.5, pp. 974–1009. DOI: 10.1086/261016.

Gabaix, Xavier (2020). "A behavioral New Keynesian model". In: *American Economic Review* 110.8, pp. 2271–2327.
Galí, Jordi (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.
Gavin, William T and Rachel J Mandal (2003). "Evaluating FOMC forecasts". In: *International Journal of Forecasting* 19.4, pp. 655–667.
Gemmi, Luca and Rosen Valchev (2023). *Biased surveys*. Tech. rep. National Bureau of Economic Research, WP 31607.
Gerlach, Stefan and Rebecca Stuart (2019). "Plotting interest rates: The FOMC's projections and the economy". In: *Journal of Macroeconomics* 60, pp. 198–211.
Gertler, Mark and Peter Karadi (2015). "Monetary policy surprises, credit costs, and economic activity". In: *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.
Giacomini, Raffaella and Barbara Rossi (2010). "Forecast comparisons in unstable environments". In: *Journal of Applied Econometrics* 25.4, pp. 595–620.
Gilchrist, Simon and Egon Zakrajšek (2012). "Credit spreads and business cycle fluctuations". In: *American economic review* 102.4, pp. 1692–1720.
Gonçalves, Sílvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento (2024). "State-dependent local projections". In: *Journal of Econometrics* 244.2, p. 105702.
Gurkaynak, Refet S, Brian Sack, and Eric T Swanson (2005). "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements". In.
Hall, Robert E. (Dec. 1978). "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence". In: *Journal of Political Economy* 86.6, pp. 971–987. DOI: 10.1086/260724.
Hamilton, James D (1994). *Time series analysis*. Princeton University Press.
Hamilton, James D, Seth Pruitt, and Scott Borger (2011). "Estimating the market-perceived monetary policy rule". In: *American Economic Journal: Macroeconomics* 3.3, pp. 1–28.
Hoesch, Lukas, Barbara Rossi, and Tatevik Sekhposyan (2023). "Has the information channel of monetary policy disappeared? Revisiting the empirical evidence". In: *American Economic Journal: Macroeconomics* 15.3, pp. 355–387.
Holston, Kathryn, Thomas Laubach, and John C Williams (2017). "Measuring the natural rate of interest: International trends and determinants". In: *Journal of international economics* 108, S59–S75.
— (2023). *Measuring the natural rate of interest after COVID-19*. Tech. rep. Staff Reports.
Imbens, Guido W (2003). "Sensitivity to exogeneity assumptions in program evaluation". In: *American Economic Review* 93.2, pp. 126–132.
Jarociński, Marek and Peter Karadi (2020). "Deconstructing monetary policy surprises—the role of information shocks". In: *American Economic Journal: Macroeconomics* 12.2, pp. 1–43.
Jeong, Jaemin, Eunseong Ma, and Choongryul Yang (2026). *Attention-Dependent Monetary Transmission to Household Beliefs*. Working Paper 2026-9. Also Fi-

nance and Economics Discussion Series 2025-084, Board of Governors of the Federal Reserve System. Federal Reserve Bank of Atlanta. DOI: 10.29338/wp2026-09.

Jimenez-Gimpel, Vicente (Mar. 2026). "Forward Guidance for Skeptical Markets". Working paper, MIT Department of Economics.

Kohlhas, Alexandre N and Ansgar Walther (2021). "Asymmetric attention". In: *American Economic Review* 111.9, pp. 2879–2925.

Kuttner, Kenneth N (2001). "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market". In: *Journal of monetary economics* 47.3, pp. 523–544.

Laster, David, Paul Bennett, and In Sun Geoum (1999). "Rational bias in macroeconomic forecasts". In: *The Quarterly Journal of Economics* 114.1, pp. 293–318.

Lorenzoni, Guido (2009). "A theory of demand shocks". In: *American economic review* 99.5, pp. 2050–2084.

Lubik, Thomas A and Christian Matthes (2011). "Economic Brief". In: *Small* 11, p. 02.

Lucas, Robert E (1972). "Expectations and the Neutrality of Money". In: *Journal of economic theory* 4.2, pp. 103–124.

Malmendier, Ulrike and Stefan Nagel (2016). "Learning from inflation experiences". In: *The Quarterly Journal of Economics* 131.1, pp. 53–87.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2003). "Disagreement about inflation expectations". In: *NBER Macroeconomics Annual*. Vol. 18. MIT Press, pp. 209–248.

Melosi, Leonardo (2017). "Signalling effects of monetary policy". In: *The Review of Economic Studies* 84.2, pp. 853–884.

Miranda-Agrippino, Silvia and Giovanni Ricco (2021). "The transmission of monetary policy shocks". In: *American Economic Journal: Macroeconomics* 13.3, pp. 74–107.

Montiel Olea, José Luis and Mikkel Plagborg-Møller (2021). "Local projection inference is simpler and more robust than you think". In: *Econometrica* 89.4, pp. 1789–1823.

Moscarini, Giuseppe (2007). "Competence implies credibility". In: *American Economic Review* 97.1, pp. 37–63.

Nakamura, Emi and Jón Steinsson (2018). "High-frequency identification of monetary non-neutrality: the information effect". In: *The Quarterly Journal of Economics* 133.3, pp. 1283–1330.

Orphanides, Athanasios (2001). "Monetary Policy Rules Based on Real-Time Data". In: *American Economic Review* 91.4, pp. 964–985.

Ottaviani, Marco and Peter Norman Sørensen (2006). "Reputational cheap talk". In: *The Rand journal of economics* 37.1, pp. 155–175.

Piazzesi, Monika and Eric T Swanson (2008). "Futures prices as risk-adjusted forecasts of monetary policy". In: *Journal of Monetary Economics* 55.4, pp. 677–691.

Plagborg-Møller, Mikkel and Christian K Wolf (2021). "Local projections and VARs estimate the same impulse responses". In: *Econometrica* 89.2, pp. 955–980.

— (2022). "Instrumental variable identification of dynamic variance decompositions". In: *Journal of Political Economy* 130.8, pp. 2164–2202.

Ramey, Valerie A (2016). "Macroeconomic shocks and their propagation". In: *Handbook of macroeconomics* 2, pp. 71–162.

Romer, Christina D and David H Romer (2004). "A new measure of monetary shocks: Derivation and implications". In: *American economic review* 94.4, pp. 1055–1084.

— (2008). "The FOMC versus the Staff: Where Can Monetary Policymakers Add Value?" In: *American Economic Review* 98.2, pp. 230–235.

Romer, David H and Christina D Romer (2000). "Federal Reserve Information and the Behavior of Interest Rates". In: *American Economic Review* 90.3, pp. 429–457. DOI: 10.1257/aer.90.3.429.

Rondina, Giacomo and Todd B Walker (2013). "Informational fragility of dynamic rational expectations equilibria". In: *2013 Meeting Papers*. Vol. 83. Society for Economic Dynamics.

Rossi, Barbara and Tatevik Sekhposyan (2016). "Forecast rationality tests in the presence of instabilities, with applications to Federal Reserve and survey forecasts". In: *Journal of Applied Econometrics* 31.3, pp. 507–532.

Särkkä, Simo and Lennart Svensson (2023). *Bayesian filtering and smoothing*. Vol. 17. Cambridge university press.

Sastry, Karthik A. (2026). "Disagreement about Monetary Policy". In: *American Economic Journal: Macroeconomics* 18.2, pp. 486–522.

Schmeling, Maik, Andreas Schrimpf, and Sigurd AM Steffensen (2022). "Monetary policy expectation errors". In: *Journal of Financial Economics* 146.3, pp. 841–858.

Simon, Dan (2006). "Using nonlinear Kalman filtering to estimate signals". In: *Embedded Systems Design* 19.7, p. 38.

Sims, Christopher A (1980). "Macroeconomics and reality". In: *Econometrica: journal of the Econometric Society*, pp. 1–48.

— (2003). "Implications of rational inattention". In: *Journal of monetary Economics* 50.3, pp. 665–690.

Surico, Paolo (2004). *Inflation targeting and nonlinear policy rules: The case of asymmetric preferences*. Tech. rep. CESifo working paper.

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

Tenreyro, Silvana and Gregory Thwaites (2016). "Pushing on a string: US monetary policy is less powerful in recessions". In: *American Economic Journal: Macroeconomics* 8.4, pp. 43–74.

Vissing-Jorgensen, Annette (2020). *Informal central bank communication*. Tech. rep. National Bureau of Economic Research.

Woodford, Michael (2001). *Imperfect common knowledge and the effects of monetary policy.*

— (2005). *Central bank communication and policy effectiveness.*

Zohar, Osnat (2024). "Cyclicality of uncertainty and disagreement". In: *Journal of Monetary Economics* 143, p. 103544.

A Data

A.1 Macro Variable Construction

A.2 Expectations Extraction

A.2.1 Federal Reserve's Expectations

The construction of FFR expectations follows in part Caballero and Simsek (2022). I start from their replication materials (link) and apply the same principles, extending the sample to all available released data as of writing (September 2025).

Federal funds rate data (FFR) is public and downloaded from FRED (ticker `FEDFUNDS`). I use the monthly series. FOMC meeting dates come from Nakamura and Steinsson (2018), extended by Acosta et al. (2024).

Two main inputs are used:

- **Digitized Greenbook/Tealbook files.** The main Greenbook dataset contains quarterly projections for macro variables (see here to download). These files are organized by FOMC meeting with quarterly horizons. A supplemental dataset lists the staff's quarterly assumptions for financial variables, including the federal funds rate (FFR); I use that series up to its last meeting, which is the sixth meeting of 2008 (see here to download).[59]
- **Manual collection from PDFs.** For meetings not digitized, I use Caballero and Simsek (2022)'s transcriptions and expand them with all new FFR forecasts from the original Tealbooks PDFs. In total, these cover from the seventh meeting of 2008 through the latest release as of writing, dated November 26, 2019. See here.[60]

Merging the hand-collected observations with the digital files yields a single series of FFR assumptions at the Tealbook/meeting frequency with quarterly horizons. For most of the sample the timestamp of the Greenbook/Tealbook files is a few days (6.6 calendar days is the average over the sample) before the corresponding FOMC meeting.

Illustration Consider the FOMC meeting on December 16 2015. The source is the Tealbook dated December 9 2015, whose table of the "Staff Economic Projections Compared with the September Tealbook" reports the FFR for 2015, 2016,

[59] Some forecasts include a "+" or "–" marker indicating that the number might be slightly higher or lower than shown; I strip these markers and keep the numeric entries.

[60] The PDFs list annual numbers for the FFR for the current year and for several subsequent years. To obtain quarterly horizons comparable to the digital files, I convert those yearly averages to quarterly paths using a simple linear interpolation rule tied to the meeting date.

2017 and 2018. I linearly interpolate quarter by quarter between the level implied for the last quarter of the preceding year and the level consistent with the stated annual average of the next year.

This approach is applied to each meeting in the hand-collected span. Each constructed observation is tagged with its meeting date and the projection quarter. When a meeting is covered by both the digital financial file and a PDF, I use the digital value in the main series and keep the alternative for audit.

The SEP Beginning with the October 2007 FOMC, participants submit individual projections that are summarized in the SEP. Since 2012 the SEP also reports each participant's projection for the federal funds rate, the "dot plot." Following Caballero and Simsek (2022), I extend their coverage to include every available meeting from 2012 through 18 September 2024 from the SEPs posted on the Fed's historical archive (link) and, for later meetings, use the medians released with the advance SEP materials on the FOMC calendar pages (link). The SEP has a yearly forecasting horizon, for which I apply the same interpolation described above for the manually imputed Tealbooks.

A.2.2 Financial Markets' Expectations

Sources. I use federal funds futures prices from Bloomberg. The series are daily, run up to 35 months ahead, and are identified by tickers `FF1 Comdty`–`FF36 Comdty`. My sample spans September 1988 to March 2025. Prices are exported from the Bloomberg Terminal.

Construction. Futures settle at month-end at 100 minus the average effective FFR for that month. I convert prices to implied monthly forward rates using this convention; see the main text for the formula. For the baseline series, for each meeting we take end-of-day prices *after* the announcement so the forwards reflect post-meeting expectations. I then match the horizons to the Fed projections.

A.3 Monetary Policy Surprises

This appendix describes the construction of the monthly orthogonalized monetary policy surprise used in the paper. I follow Bauer and Swanson (2023b) to measure event-level news and adopt the Amodeo (2025) "purge-then-sum" timing. The event set comprises all scheduled FOMC statement releases from 1989–2023. For each announcement at time t, I compute 30-minute changes in money-market futures *rates* for the front four quarterly contracts covering the next four quarters. I then extract the first principal component of the four intraday rate changes, denoted $\texttt{MPS}_t$, in basis points. The principal component summarizes revisions to

both the current target and the expected path over the following year, consistent with evidence that longer-horizon expectations move on FOMC news (Gurkaynak et al., 2005).

To address the documented predictability of high-frequency surprises, I orthogonalize each $\texttt{MPS}_t$ with respect to information known before the start of the event window on day t. Let X_t collect six predictors defined as in Bauer and Swanson (2023b): the surprise component of the most recent nonfarm payrolls release (actual minus survey median), the 12-month log change in nonfarm payroll employment as of that release, the three-month log change in the S&P 500 up to $t!-!1$, the three-month change in the yield-curve slope up to $t!-!1$, the three-month log change in the Bloomberg Commodity Spot Price index up to $t!-!1$, and the option-implied skewness of the 10-year Treasury yield averaged over the preceding month and evaluated at $t!-!1$. I estimate

$$\texttt{MPS}_t \;=\; \alpha + \beta' X_t + u_t, \qquad \texttt{MPS_ORTH}_t \equiv \widehat{u}_t, \tag{A.1}$$

by OLS across all announcements in the sample and retain the residual $\texttt{MPS_ORTH}_t$ as the orthogonalized event-level surprise. Timing is strict: all predictors are constructed from data with timestamps that precede the announcement window. As emphasized in Amodeo (2025), orthogonalization must be performed at the announcement level before aggregation; aggregating first and purging later can leave predictable components in the monthly measure. Finally, I aggregate to the monthly instrument by summation over the set $\mathcal{T}(m)$ of announcements in month m,

$$\texttt{MPS_ORTH}_m \;=\; \sum_{t \in \mathcal{T}(m)} \texttt{MPS_ORTH}_t, \tag{A.2}$$

so that ex-ante orthogonality holds at the event level and is preserved at the monthly frequency. When a month contains multiple meetings, each event is purged using its own X_t and then included in (A.2). Months without an announcement are set to missing. Units are basis points throughout. In implementation I drop events lacking any predictor, check intraday series for obvious glitches, and keep all remaining observations. The resulting $\texttt{MPS_ORTH}_m$ series is the baseline monetary shock used in the paper; robustness checks with alternative high-frequency measures are reported in the main text.

A.4 Risk Premia

When comparing the Fed's projections with market expectations, the object of interest is the *physical* (risk-averse) expectation of the policy rate, as asset prices may embed a risk premium. Although it has been discussed that risk premia are very small — especially for short-term rate futures — relative to the documented disagreements, I adjust extracted expectations using the recommended corrections in Piazzesi and Swanson (2008) and Diercks and Carl (2019). Let $f_{t,h}$ denote the

futures-implied rate h months ahead. Then

$$f_{t,h} \;=\; \mathbb{E}_t(i_{t+h}) \;+\; \mathrm{TP}_{t,h} \;=\; \mathbb{E}_t^{\mathrm{RN}}(i_{t+h}), \tag{A.3}$$

where $\mathrm{TP}_{t,h}$ is the (time-varying) term/risk premium and $\mathbb{E}_t^{\mathrm{RN}}$ the risk-neutral expectation used for pricing. For accuracy comparisons, one should adjust $f_{t,h}$ to recover $\mathbb{E}_t(i_{t+h})$.

Empirically, front-end premia are small but not literally zero. Using Fed funds futures, Piazzesi and Swanson (2008) document positive, countercyclical excess returns, implying that raw futures *overstate* subsequent realized short rates on average. Recent macro-finance estimates for the front end suggest modest, fairly persistent premia (order of 1–3 basis points per month at 6–12 month horizons), with episodes of near-zero premia. In addition, BIS evidence indicates that much of the apparent "excess return" at very short maturities reflects expectation errors in rare easing episodes rather than a large systematic premium, reinforcing the view that premia are economically small in normal times.

Implementation. I adopt two complementary adjustments: (i) a *rule-of-thumb* correction

$$\widehat{\mathbb{E}}_t(i_{t+h}) \;=\; f_{t,h} \;-\; \kappa\, h, \qquad \kappa \in [1,3] \text{ bp per month},$$

benchmarked to macro-finance estimates; and (ii) a *horizon truncation* that restricts disagreement measures to short maturities where $\mathrm{TP}_{t,h}$ is negligible. Results are robust to either approach and virtually identical when using $\kappa = 1$.

Discussion. These choices ensure that the Fed–market disagreement series reflect differences in expectations rather than compensation for bearing short-rate risk. They also align the comparison target with the object relevant for forecast evaluation, while leaving asset-pricing interpretations (which are naturally risk-neutral) unaffected.

A.5 Descriptive Statistics for Disagreement

This appendix documents unconditional distributional features of average disagreement $\bar{D}_t$ and its absolute value $|\bar{D}_t|$ over 1988:10–2019:12. Panel A reports location, dispersion, and shape. Panel B shows the same by NBER regime and repeats the mean-difference p-values from the main table. Figures referenced plot boxplots, kernel densities by regime, Q–Q plots, and autocorrelograms (see Fig. A1–A6).

Table A1: Additional Descriptive Statistics for Disagreement

Panel A. Overall (N=370)						
	Median	SD	IQR	Skew	Kurt	JB p
$\bar{D}_t$	-0.033	0.490	0.604	0.216	2.893	0.218
$\lvert\bar{D}_t\rvert$	0.288	0.310	0.440	0.939	3.018	0.000
Panel B. By business-cycle regime						
	Exp. Median ($N = 333$)	Rec. Median ($N = 37$)	Rec−Exp (mean)	p (NW)		
$\bar{D}_t$	-0.040	0.214	-0.069	0.734		
$\lvert\bar{D}_t\rvert$	0.266	0.435	0.079	0.275		

Notes: Monthly data, 1988:10–2019:12. Units: pp. Panel A reports medians, dispersion, and shape; JB p is the Jarque–Bera normality test. Panel B reports medians by NBER regime and repeats the mean-difference test (Newey–West, 12 lags) from the main table.

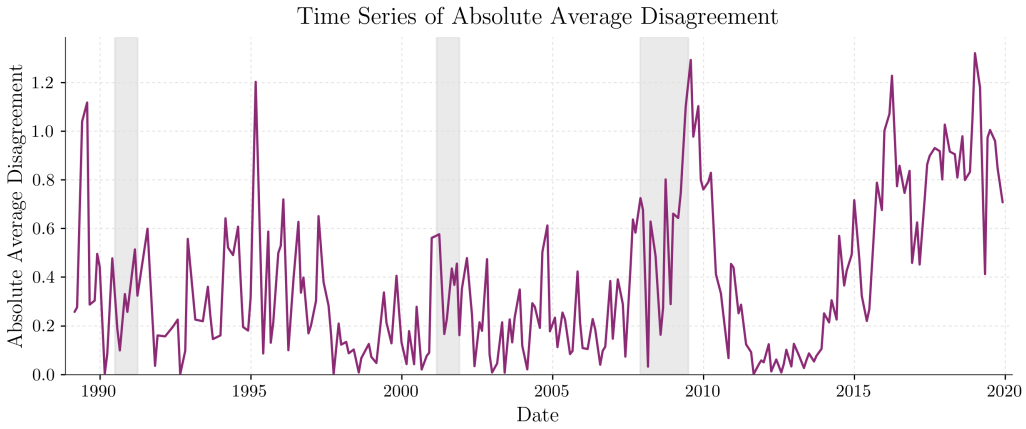


Figure A1: Fed-Market Absolute Average Disagreement Series. Monthly frequency. Sample: 1988:10-2019:12. Shaded grey areas correspond to NBER recession dates.

B Comparing Accuracy

Figures B1 narrows the non-overlapping bins of Figure 12 to one-year. Similar suggestive patterns of rank-fluctuation across forecaster emerge. Figure B2 adds to the comparison the consensus of the Survey of Professional Forecasters, again confirming that there is not a consistently better forecaster.

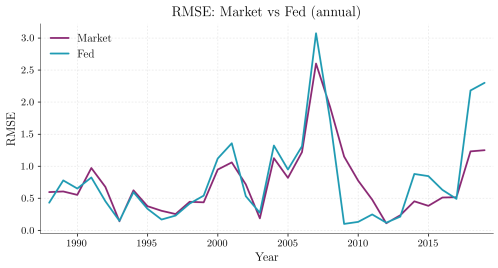


Figure B1: One-year root mean squared forecast errors for the federal funds rate by forecaster — Market and Fed — averaged over non-overlapping bins.

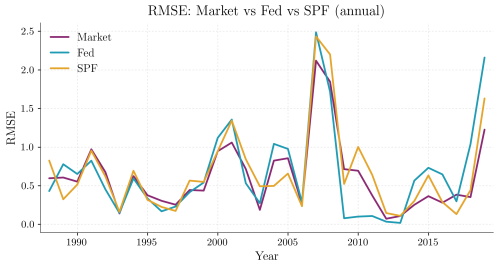


Figure B2: One-year root mean squared forecast errors for the federal funds rate by forecaster — Market, Fed and SPF's consensus — averaged over non-overlapping bins.

Tables B1-B2 below report alternative timing assumption in testing comparative accuracy through Diebold-Mariano tests. Respectively, they use as cutoff dates the trading day before the FOMC meeting, and the actual FOMC meeting day. Essentially, this amounts to providing the market with *more* information, as, on average, the Tealbooks are released internally 6.6 days. Thus, it is unsurprising to see that markets' forecasts confirm their predictive superiority, but I display the results anyway as a sanity check that the procedure works as expected.

Null: H_0	N	Statistic	PValue	Winner
Equal MSE at $h = 1$	234	-4.842	0.000	M
Equal MSE at $h = 2$	226	-4.155	0.000	M
Equal MSE at $h = 3$	179	-3.488	0.000	M
Equal MSE at $h = 4$	133	-3.857	0.000	M
Equal MSE at $h = 5$	119	-2.896	0.004	M
Equal MSE at $h = 6$	118	-2.697	0.007	M
Equal MSE at $h = 7$	107	-2.638	0.008	M
Equal MSE at $h = 8$	88	-2.842	0.004	M

Table B1: Diebold–Mariano tests of equal predictive accuracy (Market vs Fed) across horizons with Market's forecast on the day before the FOMC meeting. Reported are sample size N, test statistic, p-value, and the lower-MSE winner.

Null: H_0	N	Statistic	PValue	Winner
Equal MSE at $h = 1$	240	-4.923	0.000	M
Equal MSE at $h = 2$	231	-3.883	0.000	M
Equal MSE at $h = 3$	186	-3.711	0.000	M
Equal MSE at $h = 4$	134	-4.058	0.000	M
Equal MSE at $h = 5$	119	-3.075	0.002	M
Equal MSE at $h = 6$	118	-2.843	0.004	M
Equal MSE at $h = 7$	108	-2.754	0.006	M
Equal MSE at $h = 8$	90	-2.933	0.003	M

Table B2: Diebold–Mariano tests of equal predictive accuracy (Market vs Fed) across horizons with Market's forecast on the day of the FOMC meeting. Reported are sample size N, test statistic, p-value, and the lower-MSE winner.

B.1 Mechanics of Diebold-Mariano Tests

We have actual data $\{y_t\}$ and two forecasts $\{f_{1,t}\}, \{f_{2,t}\}$. For MSE loss, define

$$e_{i,t} = (y_t - f_{i,t})^2, \quad d_t = e_{1,t} - e_{2,t}, \quad \overline{d} = \frac{1}{T}\sum_{t=1}^{T} d_t$$

Long-run variance:

$$\hat{\gamma}(l) = \frac{1}{T}\sum_{t=l+1}^{T}(d_t - \overline{d})(d_{t-l} - \overline{d}), \quad \hat{J} = \hat{\gamma}(0) + 2\sum_{l=1}^{L} w_l\,\hat{\gamma}(l),$$

L is `maxlag` (SIC); w_l depends on `kernel` (uniform $w_l = 1$, Bartlett $w_l = 1 - \frac{l}{L+1}$)

Test statistic:

$$DM \;=\; \frac{\overline{d}}{\sqrt{\hat{J}/T}}, \quad p\text{-value} \;=\; 2\left[1-\Phi\big(|DM|\big)\right]$$

where $\Phi(\cdot)$ is the standard normal CDF

Null Hypothesis: Equal forecast accuracy; reject if p is small.

C Granger Causality in VAR Forecasting

Idea: determines whether a time series affects the predictions another. If adding it reduces forecast error, we say it *Granger-causes* the variable

Block-Wise Test: Let $y_{1,t}$ denote the m_1 "cause" variables and $y_{2,t}$ denote the m_2 "effect" variables. Consider a stationary VAR(p) model for $[\, y_{1,t},\, y_{2,t}\,]$:

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} = c + \delta t + \beta x_t + \begin{bmatrix} \Phi_{11,1} & \Phi_{12,1} \\ \Phi_{21,1} & \Phi_{22,1} \end{bmatrix} \begin{bmatrix} y_{1,t-1} \\ y_{2,t-1} \end{bmatrix} + \cdots + \begin{bmatrix} \Phi_{11,p} & \Phi_{12,p} \\ \Phi_{21,p} & \Phi_{22,p} \end{bmatrix} \begin{bmatrix} y_{1,t-p} \\ y_{2,t-p} \end{bmatrix} + \begin{bmatrix} e_{1,t} \\ e_{2,t} \end{bmatrix}$$

The block-wise Granger causality test hypotheses are:

$$H_0 : \Phi_{21,1} = \cdots = \Phi_{21,p} = 0_{m_2 \times m_1} \quad H_1 : \exists\, j \in \{1, \ldots, p\} \text{ s.t. } \Phi_{21,j} \neq 0_{m_2, m_1}$$

C.1 Robustness of Table 2

Gertler and Karadi (2015)'s VAR Estimating the VAR on the augmented system $[FFR, D, CPI, IP, EBP]$, qualitatively corroborates Table 2's findings.

Null: H_0	Distribution	Statistic	PValue	Decision
Exclude lagged D_t^1	F(24,322)	1.680	0.026	Reject
Exclude lagged D_t^2	F(24,322)	1.744	0.018	Reject
Exclude lagged D_t^3	F(28,291)	1.919	0.004	Reject
Exclude lagged D_t^4	F(44,207)	2.214	0.000	Reject
Exclude lagged D_t^5	F(44,117)	2.110	0.001	Reject
Exclude lagged D_t^6	F(48,107)	1.723	0.011	Reject
Exclude lagged D_t^7	F(48,107)	1.273	0.153	Fail to Reject
Exclude lagged D_t^8	F(48,100)	1.405	0.078	Fail to Reject
Exclude lagged $\bar{D}_t$	F(24,322)	1.910	0.007	Reject

Table C1: Block Granger Causality: 1–8 quarters and average disagreement. Null H_0: $\Phi_{n1,p} = 0\ \forall p$ (no predictive power) at $\alpha = 0.05$. Based on a VAR with constant including $[FFR, D_t, CPI, IP, EBP]$, where D_t is either vintage D_t^h or $\bar{D}_t$. Lags selected by $\max(AIC, SBC)$.

C.2 Univariate Granger Causality

Standard 3-variable VAR Tables C2-C4 estimate the same 3-variable system as in Table 2, but change the null to test for variable-specific results.

Null: H_0 (FFR_t)	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude lagged D_t^1	F(7)	5.772	0.000	Reject
Exclude lagged D_t^2	F(6)	6.978	0.000	Reject
Exclude lagged D_t^3	F(7)	5.306	0.000	Reject
Exclude lagged D_t^4	F(11)	1.955	0.034	Reject
Exclude lagged D_t^5	F(11)	3.211	0.001	Reject
Exclude lagged D_t^6	F(11)	3.228	0.001	Reject
Exclude lagged D_t^7	F(9)	2.119	0.031	Reject
Exclude lagged D_t^8	F(9)	1.989	0.044	Reject
Exclude lagged $\bar{D}_t$	F(6)	5.078	0.000	Reject

Table C2: Granger Causality tests of disagreement measures in the i equation.

Null: H_0 (CPI_t)	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude lagged D_t^1	F(7)	1.636	0.124	Fail to Reject
Exclude lagged D_t^2	F(6)	0.811	0.562	Fail to Reject
Exclude lagged D_t^3	F(7)	0.812	0.578	Fail to Reject
Exclude lagged D_t^4	F(11)	5.615	0.000	Reject
Exclude lagged D_t^5	F(11)	3.936	0.000	Reject
Exclude lagged D_t^6	F(11)	3.855	0.000	Reject
Exclude lagged D_t^7	F(9)	1.038	0.413	Fail to Reject
Exclude lagged D_t^8	F(9)	0.601	0.794	Fail to Reject
Exclude lagged $\bar{D}_t$	F(6)	1.006	0.421	Fail to Reject

Table C3: Granger Causality tests of disagreement measures in the CPI equation.

Null: H_0 (IP_t)	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude lagged D_t^1	F(7)	2.167	0.037	Reject
Exclude lagged D_t^2	F(6)	1.148	0.334	Fail to Reject
Exclude lagged D_t^3	F(7)	1.636	0.125	Fail to Reject
Exclude lagged D_t^4	F(11)	1.916	0.038	Reject
Exclude lagged D_t^5	F(11)	1.863	0.049	Reject
Exclude lagged D_t^6	F(11)	1.760	0.066	Fail to Reject
Exclude lagged D_t^7	F(9)	1.507	0.150	Fail to Reject
Exclude lagged D_t^8	F(9)	0.968	0.469	Fail to Reject
Exclude lagged $\bar{D}_t$	F(6)	1.658	0.131	Fail to Reject

Table C4: Granger Causality tests of disagreement measures in the IP equation.

Gertler and Karadi (2015)'s VAR Univariate testing can be implemented also on the augmented system $[FFR, D, CPI, IP, EBP]$. Tables C5-C8 display the results for each variable's results.

Null: H_0 (FFR_t)	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude lagged D_t^1	F(6)	5.787	0.000	Reject
Exclude lagged D_t^2	F(6)	6.978	0.000	Reject
Exclude lagged D_t^3	F(7)	5.306	0.000	Reject
Exclude lagged D_t^4	F(11)	1.955	0.034	Reject
Exclude lagged D_t^5	F(11)	3.211	0.001	Reject
Exclude lagged D_t^6	F(12)	2.765	0.002	Reject
Exclude lagged D_t^7	F(12)	1.853	0.045	Reject
Exclude lagged D_t^8	F(12)	1.612	0.095	Fail to Reject
Exclude lagged $\bar{D}_t$	F(6)	5.078	0.000	Reject

Table C5: Granger causality tests of disagreement measures in the i equation.

Null: H_0 (CPI_t)	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude lagged D_t^1	F(6)	1.691	0.122	Fail to Reject
Exclude lagged D_t^2	F(6)	0.811	0.562	Fail to Reject
Exclude lagged D_t^3	F(7)	0.812	0.578	Fail to Reject
Exclude lagged D_t^4	F(11)	5.615	0.000	Reject
Exclude lagged D_t^5	F(11)	3.936	0.000	Reject
Exclude lagged D_t^6	F(12)	3.685	0.000	Reject
Exclude lagged D_t^7	F(12)	1.961	0.032	Reject
Exclude lagged D_t^8	F(12)	0.740	0.711	Fail to Reject
Exclude lagged $\bar{D}_t$	F(6)	1.006	0.421	Fail to Reject

Table C6: Granger causality tests of disagreement measures in the CPI equation.

Null: H_0 (IP_t)	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude lagged D_t^1	F(6)	1.360	0.230	Fail to Reject
Exclude lagged D_t^2	F(6)	1.148	0.334	Fail to Reject
Exclude lagged D_t^3	F(7)	1.636	0.125	Fail to Reject
Exclude lagged D_t^4	F(11)	1.916	0.038	Reject
Exclude lagged D_t^5	F(11)	1.863	0.049	Reject
Exclude lagged D_t^6	F(12)	1.558	0.111	Fail to Reject
Exclude lagged D_t^7	F(12)	1.152	0.324	Fail to Reject
Exclude lagged D_t^8	F(12)	0.837	0.612	Fail to Reject
Exclude lagged $\bar{D}_t$	F(6)	1.658	0.131	Fail to Reject

Table C7: Granger causality tests of disagreement measures in the IP equation.

Null: H_0 (EBP_t)	**Distribution**	**Statistic**	**PValue**	**Decision**
Exclude lagged D_t^1	F(6)	0.145	0.990	Fail to Reject
Exclude lagged D_t^2	F(6)	0.506	0.804	Fail to Reject
Exclude lagged D_t^3	F(7)	2.219	0.033	Reject
Exclude lagged D_t^4	F(11)	2.331	0.010	Reject
Exclude lagged D_t^5	F(11)	3.188	0.001	Reject
Exclude lagged D_t^6	F(12)	2.334	0.009	Reject
Exclude lagged D_t^7	F(12)	0.822	0.627	Fail to Reject
Exclude lagged D_t^8	F(12)	0.870	0.579	Fail to Reject
Exclude lagged $\bar{D}_t$	F(6)	0.501	0.808	Fail to Reject

Table C8: Granger causality tests of disagreement measures in the EBP equation.

D Empirics

D.1 Total Effects: $\alpha_h + \beta_h$

Baseline total effect from regression (2).

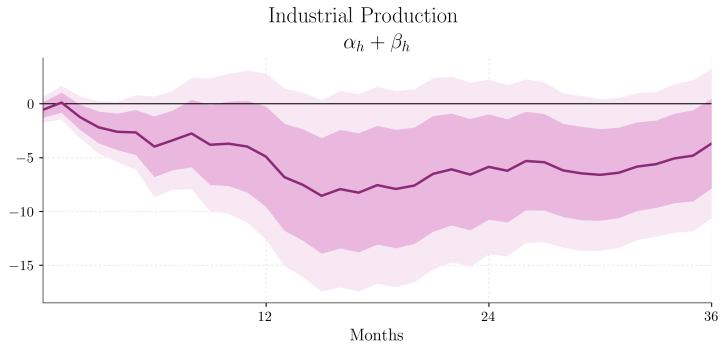


Figure D1: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Total effect. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

D.2 Evidence of Attenuation: Robustness

LPs with ex-post disagreement I re-estimate the baseline LPs using *ex-post* disagreement, measured after FOMC announcements, as the interaction state.

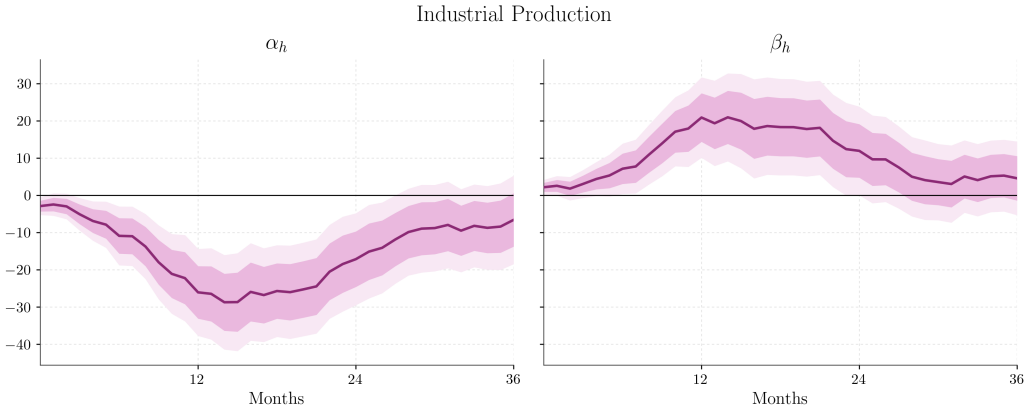


Figure D2: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

The resulting IRFs in Figure D2 are very similar to the baseline: the output response α_h is contractionary, while $\beta_h > 0$ implies that larger (absolute) ex-post disagreement attenuates the real effects of a tightening. Econometrically, this specification is valid under the same identifying assumption, although the interpretation of its results vary.

Gertler and Karadi (2015)'s SVAR Estimating the LPs on the augmented system $[FFR, D, CPI, IP, EBP]$

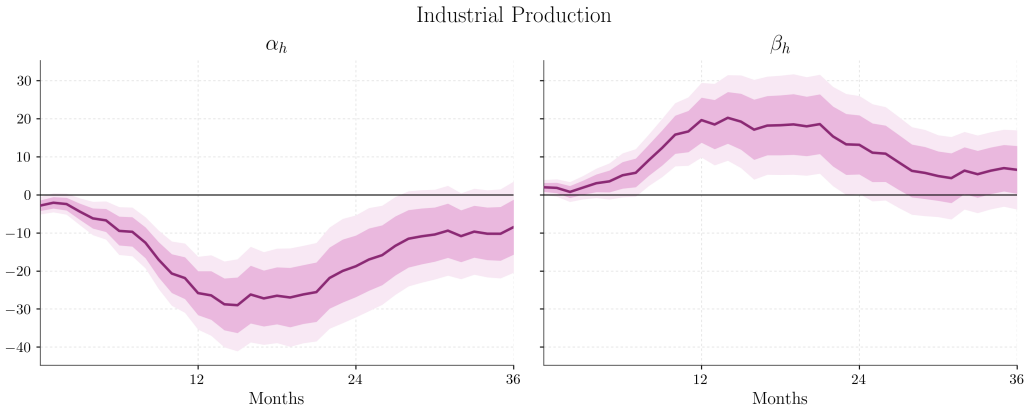


Figure D3: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Estimating the LPs on the original GK system $[1Y, D, CPI, IP, EBP]$

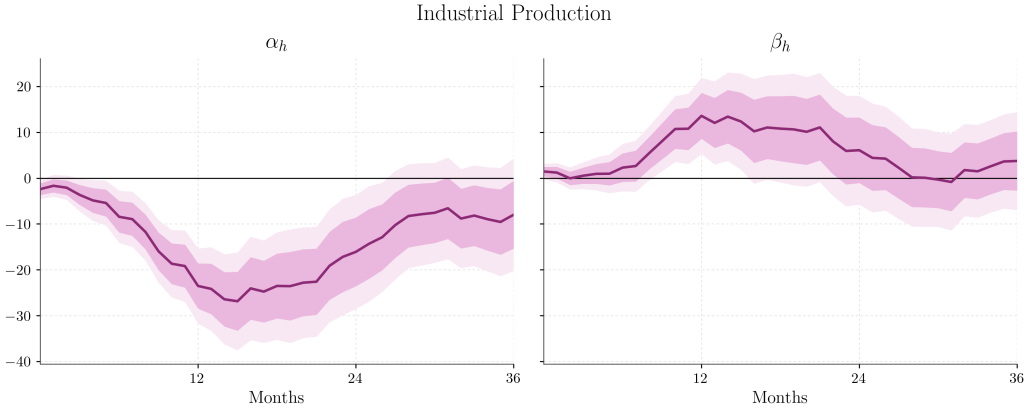


Figure D4: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Bauer and Swanson (2023b)'s SVAR Estimating the LPs on the augmented system $[1Y, D, CPI, IP, EBP, UNEMP, Comm]$ we get:

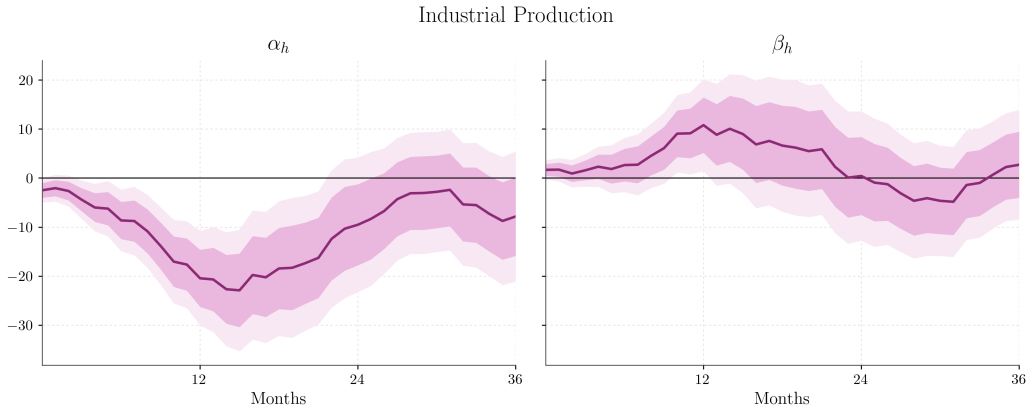


Figure D5: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Estimating the LPs on the augmented system $[FFR, D, CPI, IP, EBP, UNEMP, Comm]$ we get:

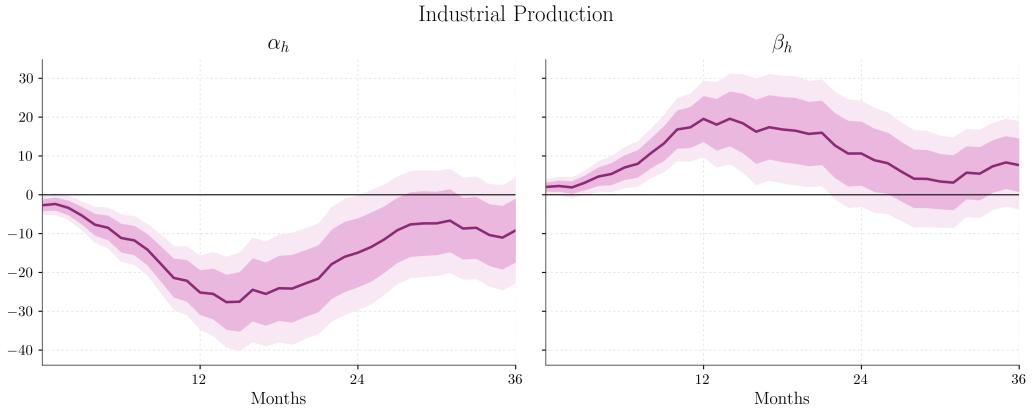


Figure D6: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (2). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

D.3 Elasticities of Monetary Effects with respect to Disagreement

Let absolute average disagreement in percentage points be $|\bar{d}_t|$ (pp). We rescale it by its sample mean so the regressor is unit-free and equals one at the mean:

$$|\bar{D}_t| = \frac{|\bar{d}_t|}{0.383}.$$

The total effect at $h = 12$ of a one-unit monetary shock is linear in normalized disagreement,

$$\mathrm{TE}_{12}(|\bar{D}_t|) = \alpha_{12} + \beta_{12}|\bar{D}_t| = \alpha_{12} + \beta_{12}\frac{|\bar{d}_t|}{0.383},$$

with estimates $\alpha_{12} = -20.659$, $\beta_{12} = 16.074$.

Baseline and magnitude At zero disagreement, $\mathrm{TE}_{12}(0) = \alpha_{12} < 0$, so the level response is a contraction. By "transmission strength'' we mean the *size* of the response, hence we study the absolute value $|\mathrm{TE}_{12}(|\bar{D}_t|)|$. Because $\alpha_{12} < 0$ and $\beta_{12} > 0$, the level becomes less negative as disagreement rises. While the level remains negative, taking the absolute value simply flips the sign:

$$|\mathrm{TE}_{12}(|\bar{D}_t|)| = -\mathrm{TE}_{12}(|\bar{D}_t|) = |\alpha_{12}| - \beta_{12}|\bar{D}_t|.$$

Thus each unit increase in $|\bar{D}_t|$ subtracts β_{12} from the magnitude; disagreement weakens transmission.

Percent reduction per unit of disagreement To express the attenuation in percentage terms, compare the change in magnitude to the baseline magnitude at zero disagreement. Linearizing around $|\bar{D}_t| = 0$ gives the first-order change $d|\mathrm{TE}_{12}| = -\beta_{12}\, d|\bar{D}_t|$. Dividing by the baseline $|\mathrm{TE}_{12}(0)| = |\alpha_{12}|$ yields

$$\frac{\Delta|\mathrm{TE}_{12}|}{|\mathrm{TE}_{12}(0)|} \approx \frac{\beta_{12}}{|\alpha_{12}|}\,\Delta|\bar{D}_t| = \frac{\beta_{12}}{|\alpha_{12}|}\frac{\Delta|\bar{d}_t|}{0.383}.$$

Because 1 basis point = 0.01 pp,

$$\%\text{ reduction in magnitude } \approx\ 100 \cdot \frac{\beta_{12}}{|\alpha_{12}|} \cdot \frac{1}{0.383} \cdot \frac{X}{100} = 2.0314\% \times X \quad \text{for } X \text{ bps.}$$

where the formula holds within the domain of the Taylor approximation.[61]

[61] Domain of validity: the linear formula uses the local derivative at $|\bar{D}_t| = 0$ and is valid while the response remains negative, i.e.

$$|\bar{d}_t| < |\bar{d}_t^*| := 0.383 \cdot \frac{|\alpha_{12}|}{\beta_{12}} = 0.4923 \text{ pp } (49.23 \text{ bps}),$$

D.4 Monetary Episodes: Extra

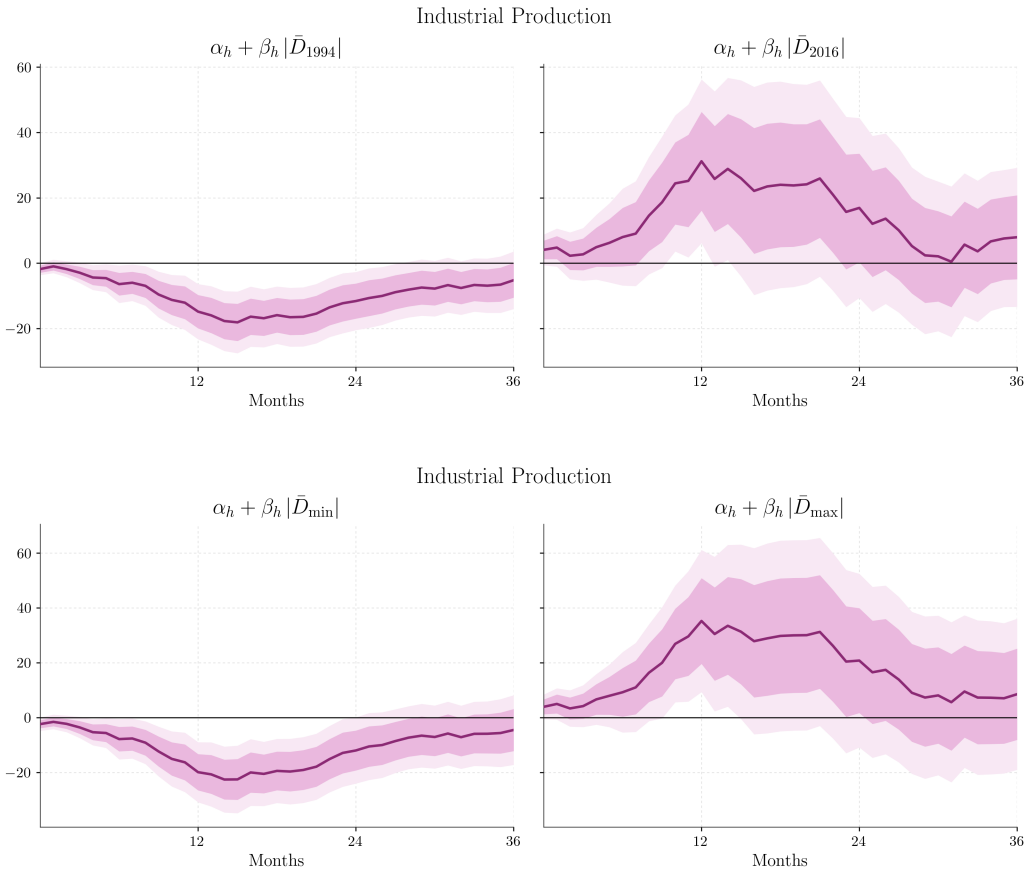


Figure D7: Impulse responses of industrial production to an orthogonalized monetary policy surprise, highlighting heterogeneity in the total effect $\alpha_h + \beta_h|D_t|$. Top: 1994m1 vs. 2016m8. Bottom: minimum- vs. maximum-disagreement months (1997m10 vs. 2009m8).

D.5 Disagreement Episodes: Full List

Table D1 lists all contiguous monthly runs of at least six months in which the average Fed–market disagreement $\bar{D}_t$ lies in the top or bottom signed quintile of its unconditional sample distribution (1988:10–2019:12), allowing four-month gaps to bridge brief reversions. Quintile thresholds: $\bar{D}_t > +0.47$ (Fed-hawk) and $\bar{D}_t < -0.29$ (Market-hawk). Figure callouts in the body Figure 5 follow an *actor-first* principle: every label begins with the higher-path actor ("Fed" for top quintile, "Market" for bottom quintile), so direction is encoded in the first word. Episodes 1 (1989:06–1989:12) and 8 (1999:06–1999:11) are shaded in the figure but unlabeled to reduce visual density; both are listed below.

since at $|\bar{d}_t^*|$ we have $\mathrm{TE}_{12} = 0$ and $|\mathrm{TE}_{12}|$ has a kink. Beyond this point the sign flips, so "percent reduction of a contraction'' is not well-defined.

Table D1: Top- and bottom-signed-quintile episodes of $\bar{D}_t$.

Type	Period	Mo.	Mean $\bar{D}_t$	Peak date	Narrative
Top quintile (Fed > Market)					
Fed-hawk	1989:06–1989:12	7	+0.66	1989:08	The FOMC eased steadily through the second half of 1989 (roughly 150 bp of cuts from the February peak); markets, reading the inverted yield curve as a recession signal, priced an even sharper easing.
Fed-hawk	1995:08–1996:02	7	+0.43	1996:02	After the February 1995 terminal hike to 6%, the FOMC delivered three 25 bp insurance cuts; markets priced a deeper and faster easing as growth softened.
Fed-hawk	2014:07–2015:03	9	+0.51	2015:01	FOMC participants signaled a 2015 liftoff with a relatively steep path; markets priced a flatter trajectory amid the late-2014 oil collapse, dollar strength, and persistently below-target inflation.
Fed-hawk	2015:09–2019:12	52	+0.86	2019:01	Through the post-liftoff normalization, FOMC projections stayed above market-implied paths; the gap peaked after Powell's October 2018 "long way from neutral" remark and closed only with the January 2019 pivot.
Bottom quintile (Market > Fed)					
Mkt-hawk	1991:04–1991:09	6	−0.45	1991:08	Markets priced a stronger 1990–91 recovery and expected the Fed to halt easing; the FOMC instead extended cuts to offset credit-crunch drag from the S&L resolution and bank balance-sheet repair.
Mkt-hawk	1994:02–1995:05	16	−0.50	1995:03	After the 1994 tightening cycle (including a 75 bp hike in November), bond markets extrapolated the pace and priced terminal rates well above 6%; the FOMC paused at 6% in February 1995.
Mkt-hawk	1996:06–1997:07	14	−0.38	1997:04	With unemployment below 5.5%, markets anticipated an inflation-driven tightening; the FOMC, persuaded by accelerating productivity, delivered only a single 25 bp hike in March 1997.
Mkt-hawk	1999:06–1999:11	6	−0.25	1999:11	Markets priced a steeper continuation of the mid-1999 tightening cycle on rising oil prices and dot-com demand; the FOMC delivered only three 25 bp hikes totaling 75 bp.
Mkt-hawk	2003:12–2004:11	12	−0.28	2004:11	Markets repeatedly anticipated faster hikes from the 1% floor as the recovery firmed; the FOMC held to its "measured pace" guidance, delivering only 25 bp increments from June 2004.
Mkt-hawk	2008:04–2011:02	35	−0.61	2009:08	Markets kept pricing premature normalization across this multi-year, largely ZLB-era window—re-tightening as the crisis eased, then earlier liftoff—against the FOMC's lower-for-longer guidance.

D.6 ZLB Exclusion

This appendix assesses whether the paper's main local-projection results are driven by the zero lower bound (ZLB) period. Using the same monthly specification and horizon-by-horizon LP framework, I (*i*) define the ZLB window as 2008m12–2015m12 and construct a ZLB indicator, (*ii*) re-estimate impulse responses allowing both the level effect of the shock and the marginal effect of disagreement to differ at the ZLB via interactions with the ZLB dummy, and (*iii*) report the implied IRFs for Non-ZLB. I further run diagnostics in the form of a horizon-by-horizon test of equality of the disagreement-slope across regimes.

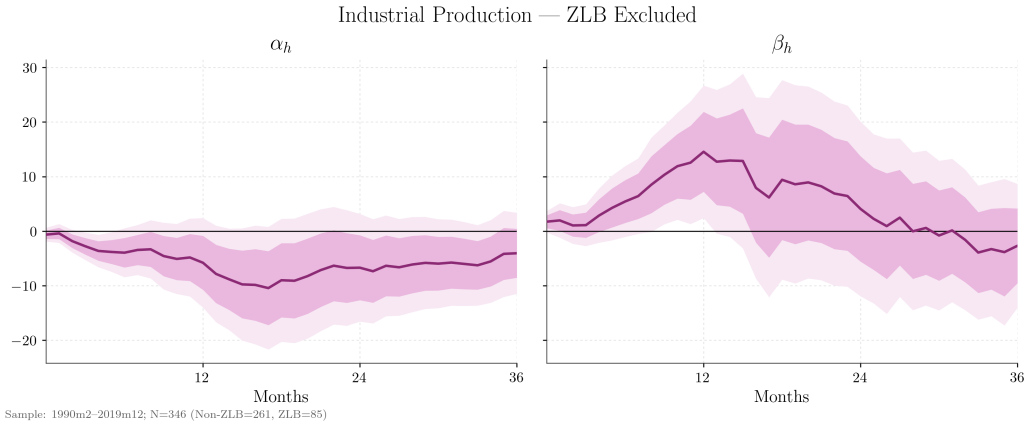


Figure D8: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock excluding 2008:12–2015:12, equation (3). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 261$ observations (excludes ZLB era and 2001M9 for the 9/11 Terrorist Attacks).

Excluding the ZLB period leaves the core mechanism intact. Unsurprisingly, dropping roughly a quarter of the sample makes the LP estimates noisier and the confidence bands wider, especially at medium horizons. The right panel still shows a positive coefficient: higher Fed–market disagreement continues to attenuate the real-side response to an orthogonalized monetary policy shock. Interestingly, the left panel indicates that the average level effect is materially weaker off the ZLB, while the disagreement-dependent slope remains the more stable feature of the data. Tests of regime-stability estimation confirm the latter.

Separately, Appendix D.12.4 tests whether the monetary shock s_t is drawn from the same distribution in the Zero Lower Bound period versus outside it, supporting the state-dependent claims of state-dependent transmission which hinged on treatment homogeneity.

D.7 SEP-Implied Disagreement

A natural and frequently anticipated concern is whether the attenuation result documented in Section 3 survives when the Fed's expectations are measured from the publicly observable Summary of Economic Projections (SEP) rather than from the staff Tealbook forecasts. The SEP is, by construction, a coarser instrument: it is published only at the four projection FOMC meetings per year, started reporting federal funds rate paths only in January 2012, and aggregates the individual policy outlooks of *all* FOMC participants — a collection of *intentions* more than a unified *forecast*, as argued in Gerlach and Stuart (2019) and Bordo and Istrefi (2023) and discussed in Section 2. For those and other reasons, the Tealbook is the most appropriate baseline. Nevertheless, confronting the two on the overlapping sample is the most direct way to gauge whether the documented attenuation channel is an artifact of the Tealbook measurement choice.

I construct $|\bar{D}_t^{\text{SEP}}|$ following exactly the same procedure described in Section 2, replacing the Tealbook FFR projections with the SEP median released at each scheduled SEP meeting and applying the same monthly-frequency interpolation (Appendix A.2). The resulting series covers 2012:01–2023:12, with $N = 110$ usable observations at horizon zero (and $N = 90$ at horizon zero in the pooled regression of Table D2, declining to $N = 80$ at the longer horizons after the lag construction on the shorter overlap window).[62] I then re-estimate equation (2) with $|\bar{D}_t^{\text{SEP}}|$ in place of $|\bar{D}_t|$, leaving the rest of the specification untouched. Figure D9 reports, in two rows, the $\boldsymbol{\alpha}_h$ and $\boldsymbol{\beta}_h$ estimates for (A) the headline Tealbook specification restricted to the 2012:01–2019:12 overlap window, and (B) the SEP-based specification on its full available sample.

The substantive content of the figure is in the comparison of its two rows against the headline Figure 6: changing the measurement choice (Tealbook to SEP) while *also* changing the sample window (1990–2019 to 2012–2023) confounds two distinct effects, and the figure isolates each in turn. Panel A of Figure D9 holds the disagreement measure fixed (Tealbook) and varies only the sample relative to Figure 6; panel B holds the sample fixed (post-2012) and varies only the measure relative to panel A. The medium-horizon attenuation visible in Figure 6's $\boldsymbol{\beta}_h$ — the headline mechanism — slightly softens in panel A on the shorter window, then survives in panel B with the SEP measure on essentially the same horizons. The point estimates are less precise at the longest horizons, where the post-2012 identifying variation is mechanically thinner, but the medium-horizon interaction keeps its sign and shape under the SEP measure, and the divergence from the full-sample headline is sample-driven rather than measurement-driven.

[62] The monthly panel excludes the pandemic months 2020:03–2020:12, which drops the June, September, and December 2020 SEP rounds (the March 2020 SEP was cancelled); the local projections are estimated on a calendar-preserving time axis so that no lead or lag operator bridges the resulting gap. Re-interpolating the SEP series on the true calendar leaves the estimates within the reported confidence bands at every horizon.

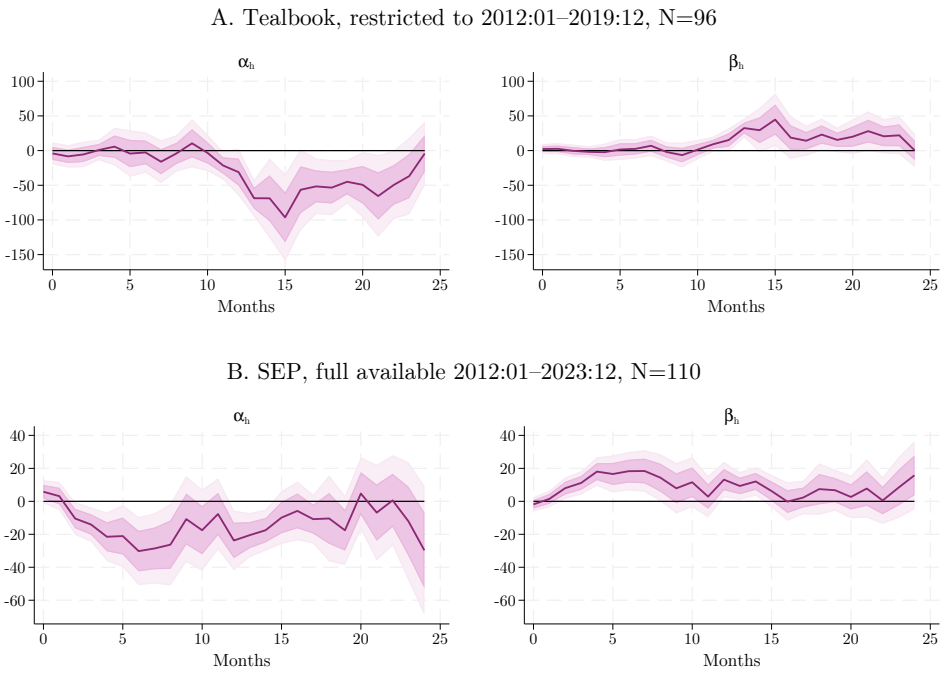


Figure D9: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, with Tealbook-based and SEP-based disagreement on the SEP-available window. Panel *A*: paper's main specification with Tealbook-implied $|\bar{D}_t|$ restricted to the SEP overlap window. Panel *B*: same specification with SEP-implied $|\bar{D}_t^{\text{SEP}}|$ on the full SEP-available sample. Left column: $\boldsymbol{\alpha}_h$ (response in the absence of disagreement). Right column: $\boldsymbol{\beta}_h$ (marginal interaction with the disagreement measure). Confidence bands at 68% and 90%.

For a more quantitative outlook, Table D2 reports a horizon-by-horizon Wald test of $H_0 : \boldsymbol{\beta}_h^T = \boldsymbol{\beta}_h^S$ obtained from a single *pooled* local projection on the 2012–2019 window that includes both $s_t \times |\bar{D}_t^T|$ and $s_t \times |\bar{D}_t^S|$. At no horizon $h \in \{0, \ldots, 24\}$ is the equality rejected at conventional levels (p-values range from 0.26 to 0.99, with $p = 0.26$ at the focal horizon $h = 12$). The pooled correlation of the two measures on the overlap window is $\rho(|\bar{D}_t^T|, |\bar{D}_t^S|) = 0.65$, indicating that they share the bulk of their joint variation. The test is necessarily underpowered on this short overlap window ($N = 80$) with two correlated regressors, so the failure to reject must be read as a non-rejection rather than as evidence of strict equivalence; nevertheless, the failure is uniform across horizons.

I read the combined evidence as a qualitative corroboration of the attenuation channel under the SEP measure. The substantive conclusion is that the SEP and Tealbook measures of policy-rate disagreement convey the same message about monetary transmission, with magnitude precision that scales with the sample size rather than with the measurement choice.

Table D2: SEP–vs–Tealbook disagreement: pooled overlap test of equal interaction coefficients

h	$\hat{\beta}_h^T$	$(s.e.)$	$\hat{\beta}_h^S$	$(s.e.)$	F	p-value	N
0	−0.60	(3.17)	−2.83	(4.63)	0.14	0.71	90
3	1.29	(4.25)	−4.53	(7.97)	0.30	0.59	89
6	−2.89	(6.35)	−7.49	(14.19)	0.07	0.80	86
9	−8.22	(11.78)	6.11	(18.35)	0.35	0.56	83
12	−4.58	(11.88)	23.26	(17.61)	1.34	0.26	80
15	20.71	(11.63)	11.94	(18.99)	0.12	0.73	80
18	21.98	(9.31)	6.35	(15.35)	0.60	0.45	80
21	5.46	(8.48)	5.67	(11.66)	0.00	0.99	80
24	3.11	(5.97)	2.93	(8.27)	0.00	0.99	80

Notes: Horizon-by-horizon Wald F-tests of $H_0 : \beta_h^T = \beta_h^S$ from a single pooled local projection that includes both interaction terms $s_t \times |\bar{D}_t^T|$ and $s_t \times |\bar{D}_t^S|$ on the 2012:01–2019:12 overlap window, where T denotes Tealbook-based and S denotes SEP-based disagreement. The specification follows the paper's main block with the lag length reduced to six for the pooled regression to preserve degrees of freedom; the LPs underlying the panels of Figure D9 retain the headline twelve lags. Heteroskedasticity-robust standard errors in parentheses.

Lastly, independently of power, the Tealbook is confidential and therefore independent of the market leg by construction, whereas the SEP is public and partly embedded in the futures prices from which the market expectation is extracted.

D.8 Additional Uncertainty Proxies and Diagnostics

Figure D10 compares the three monthly uncertainty controls used in the paper: the VIX, the Baker–Bloom–Davis Economic Policy Uncertainty index, and SPF dispersion. Because these series are measured in different units, each line is divided by its own mean over the displayed window. Table D3 reports descriptive statistics and pairwise correlations. The three proxies comove positively, especially around recessions, but far from perfectly, which makes them useful as complementary controls rather than interchangeable ones.

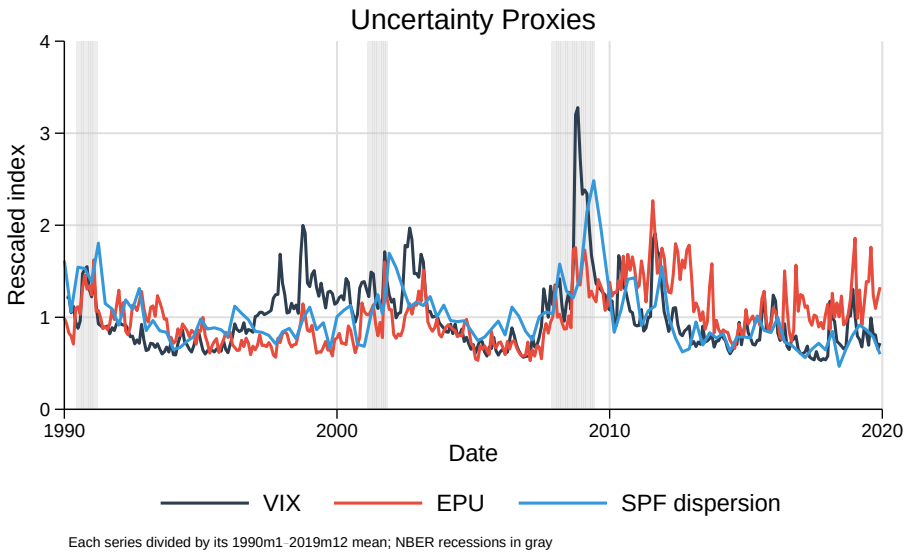


Figure D10: Monthly uncertainty proxies. Each series is divided by its own mean over 1990m1–2019m12 so to compare movements rather than raw units. NBER recessions in gray.

Table D3: Monthly Uncertainty Proxies

	VIX	EPU	SPF
VIX	1.000	0.389	0.528
EPU	0.389	1.000	0.289
SPF	0.528	0.289	1.000

Notes: The upper panel reports descriptive statistics over the paper window, 1990m1–2019m12. The lower panel reports pairwise correlations on the common monthly sample ($N = 359$), determined by the availability of the monthly VIX.

Economic Policy Uncertainty. I next replace VIX_t in equation (3) with the monthly Baker–Bloom–Davis Economic Policy Uncertainty index, EPU_t. Figure D11 shows that the disagreement slope $\boldsymbol{\beta}_h$ remains positive, so the main attenuation result survives this alternative uncertainty control. Relative to the VIX specification, the left-panel coefficient $\boldsymbol{\alpha}_h$ again becomes more negative, indicating that the average real-side effect is more sensitive to the chosen uncertainty proxy than the disagreement-dependent slope itself.

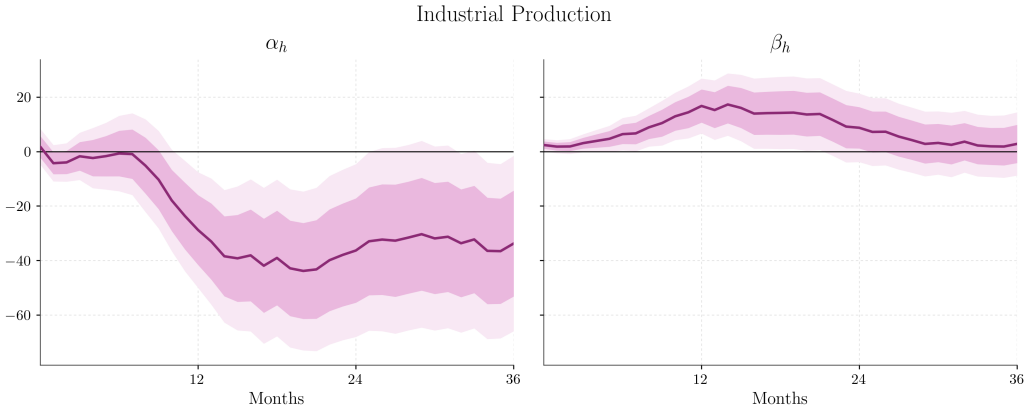


Figure D11: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections controlling for the monthly Economic Policy Uncertainty index, EPU_t, and its lags, equation (3). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1989–2019, $N = 370$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Survey of Professional Forecasters dispersion. I use the dispersion in the Survey of Professional Forecasters as a measure of macro uncertainty (Bloom, 2014, Bhandari et al., 2016). Here, SPF_t is a monthly proxy: at each survey I compute the interquartile range (P75–P25) of nominal GDP level forecasts at quarterly horizons $h = 0$ to $h = 4$, average across horizons to form a quarterly index, then linearly interpolate to monthly. Figure D12 shows that the marginal effect of disagreement (right panel) on the transmission of monetary policy is left mostly unchanged, while $\boldsymbol{\alpha}_h$ – the response of industrial production to a monetary shock in the absence of disagreement – is much noisier. Although the point estimates remain within the ballpark of the baseline specifications, the increase in noise is intuitive: SPF dispersion is highly persistent (first-order autocorrelation ~ 0.97), so that $s_t \upsilon_t$ is nearly collinear with s_t (correlation ~ 0.95), inflating standard errors. Economically, SPF captures slow-moving uncertainty; including SPF and its lags shifts identification of $\boldsymbol{\alpha}_h$ (the effect at average uncertainty) toward higher-frequency variation.

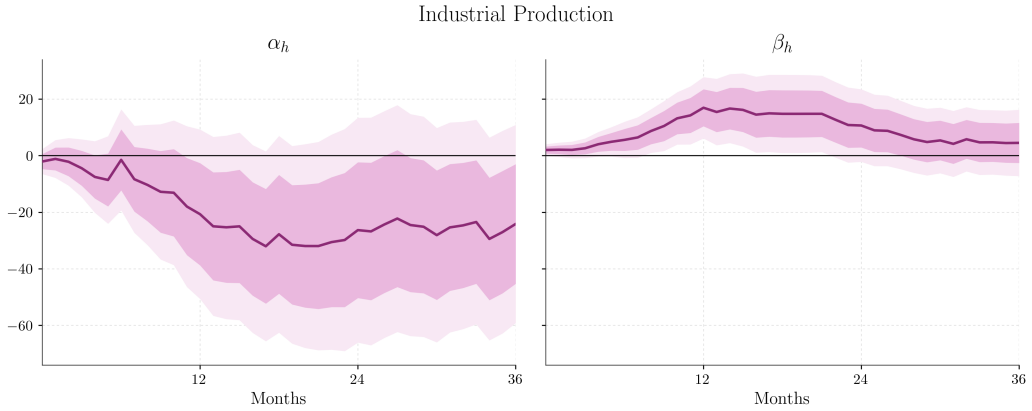


Figure D12: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections controlling for the dispersion of the Survey of Professional Forecasters, SPF_t, and its lags, equation (3). Left panel: $\boldsymbol{\alpha}_h$; right panel: $\boldsymbol{\beta}_h$. Confidence bands at 68 and 90% levels. Sample: 1989–2019, $N = 370$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

Together with the VIX specification in the main text, these results show that the attenuation associated with the level of disagreement in the economy is robust to widely used financial, policy, and survey-based measures of uncertainty.

D.9 Sensitivity to Omitted Variables: Imbens (2003)

A potential concern with the local projection estimates is omitted variable bias: an unobserved confounder U_t correlated with both the interaction term $s_t|\bar{D}_t|$ and the outcome $\mathbf{y}_{t+h}$ would bias the coefficient $\boldsymbol{\beta}_h$. I address this concern using the sensitivity analysis developed by Imbens (2003), which quantifies how strongly any hypothetical confounder must correlate with both regressor and outcome to overturn the baseline result.

Let $\tilde{Z}_t$ denote the residual from projecting the interaction term $s_t|\bar{D}_t|$ onto controls, and $\tilde{Y}_{t+h}$ the residual from projecting the outcome onto controls and the shock. The sensitivity analysis constructs a hypothetical unobserved covariate U_t and characterizes it by two partial R^2 values:

$$R^2_{\tilde{Z},U} \equiv \frac{\mathrm{Var}\big(\mathbb{E}[\tilde{Z}_t \mid U_t]\big)}{\mathrm{Var}(\tilde{Z}_t)} = 1 - \frac{\mathbb{E}\big[\mathrm{Var}(\tilde{Z}_t \mid U_t)\big]}{\mathrm{Var}(\tilde{Z}_t)},$$
$$R^2_{\tilde{Y},U} \equiv \frac{\mathrm{Var}\big(\mathbb{E}[\tilde{Y}_{t+h} \mid U_t]\big)}{\mathrm{Var}(\tilde{Y}_{t+h})} = 1 - \frac{\mathbb{E}\big[\mathrm{Var}(\tilde{Y}_{t+h} \mid U_t)\big]}{\mathrm{Var}(\tilde{Y}_{t+h})}.$$

These measure the share of *residual* variance in Z and Y—unexplained by observed controls—that U would explain. Following Imbens (2003), I search over the space of $(R^2_{\tilde{Z},U}, R^2_{\tilde{Y},U})$ pairs to identify which combinations would flip the sign of $\hat{\boldsymbol{\beta}}_h$.

I report the *symmetric* robustness value, imposing $R^2_{\tilde{Z},U} = R^2_{\tilde{Y},U} \equiv r^2$. This restriction identifies the point along the diagonal of the $(R^2_{\tilde{Z}}, R^2_{\tilde{Y}})$ space and corresponds to the worst-case confounder: one equally correlated with both regressor and outcome. The robustness value r^2_{min} is the minimal such r^2 for which the adjusted $\hat{\boldsymbol{\beta}}_h$ changes sign.[63]

Table D4: Imbens (2003) Sensitivity Analysis for $\boldsymbol{\beta}_{12}$

	$\hat{\beta}_{12}$	t-stat	r^2_{min}	VIX: $R^2_{\tilde{Z},U}$
INDUSTRIAL PRODUCTION	16.56	2.59	**12.81%**	2.62%

Notes: r^2_{min} is the minimal partial R^2 an omitted variable U must have with *both* the interaction term and the outcome (symmetric case) to flip the sign of $\hat{\beta}_{12}$. The VIX column reports the partial R^2 of VIX with the interaction term as a benchmark.

Table D4 reports $r^2_{\text{min}} = 12.81\%$ for industrial production at $h = 12$. Any omitted variable must explain at least 12.81% of the residual variance in *both* the interaction term $s_t|\bar{D}_t|$ and industrial production – after partialling out all controls – to reverse the sign of $\hat{\boldsymbol{\beta}}_{12}$. As a benchmark, VIX, a leading measure of uncertainty often invoked as a potential confounder, explains only 2.62% of the residual variance in the interaction term. An omitted variable would need to be nearly five times as influential as VIX to invalidate the result.

D.10 Asymmetric Disagreement with Asymmetric Shocks

Define

$$s_t^+ = \max(s_t, 0) \quad s_t^- = \min(s_t, 0) \quad \bar{D}_t^+ = \max(\bar{D}_t, 0) \quad \bar{D}_t^- = \min(\bar{D}_t, 0)$$

For horizon h, I run lag-augmented LPs (D.1) on Industrial Production separately for $h = 0, \ldots, H$ with standard macro controls (IP, CPI, EBP, FFR):

$$\begin{aligned} \mathbf{y}_{t+h} = c_h + \boldsymbol{\theta}_h \text{LAGS}_{t-1}^L \\ + \alpha_h^+ s_t^+ + \alpha_h^- s_t^- \\ + \phi_h^+ (s_t^+ \bar{D}_t^+) + \phi_h^- (s_t^+ \bar{D}_t^-) \\ + \lambda_h^+ (s_t^- \bar{D}_t^+) + \lambda_h^- (s_t^- \bar{D}_t^-) + \varepsilon_{t+h} \end{aligned} \tag{D.1}$$

[63] Technical note on the symmetric case: Imbens (2003) presents contour plots over the full (R^2_Z, R^2_Y) space. The asymmetric case is more complete but yields a 2D region rather than a single statistic. The symmetric restriction $R^2_Z = R^2_Y$ is standard in applied work because (*i*) it produces a single, interpretable robustness measure; (*ii*) it is conservative: it corresponds to the "worst-case" confounder configuration.

with $\text{LAGS}_{t-1}^{L} \equiv \left[\mathbf{y}_{t-1:t-L},\ s_{t-1:t-L}\right]'$ and $\boldsymbol{\theta}_h \equiv [\,\phi_h^{1:L},\ \gamma_h^{1:L}]$.

Interpretation of Coefficients

- α_h^+: marginal effect at $\bar{D}_t = 0$ of a unit increase in the policy rate when the Market is *positively* surprised.
- α_h^-: marginal effect at $\bar{D}_t = 0$ of a unit increase in the policy rate when the Market is *negatively* surprised.
- ϕ_h^+: marginal change (attenuation/amplification) in the per-unit effect of a policy rate increase when the Market is positively surprised (α_h^+) per-unit of positive $\bar{D}_t^+$ (F > M).
- ϕ_h^-: marginal change (attenuation/amplification) in the per-unit effect of a policy rate increase when the Market is positively surprised (α_h^+) per-unit of negative $\bar{D}_t^-$ (M > F).
- λ_h^+: marginal change (attenuation/amplification) in the per-unit effect of a policy rate increase when the Market is negatively surprised (α_h^-) per-unit of positive $\bar{D}_t^+$ (F > M).
- λ_h^-: marginal change (attenuation/amplification) in the per-unit effect of a policy rate increase when the Market is negatively surprised (α_h^-) per-unit of negative $\bar{D}_t^-$ (M > F).

Results and Interpretation First, I display the baseline sign-dependent effects, that is, the effects of a marginal increase in the policy rate when the Market's surprise is respectively positive and negative:

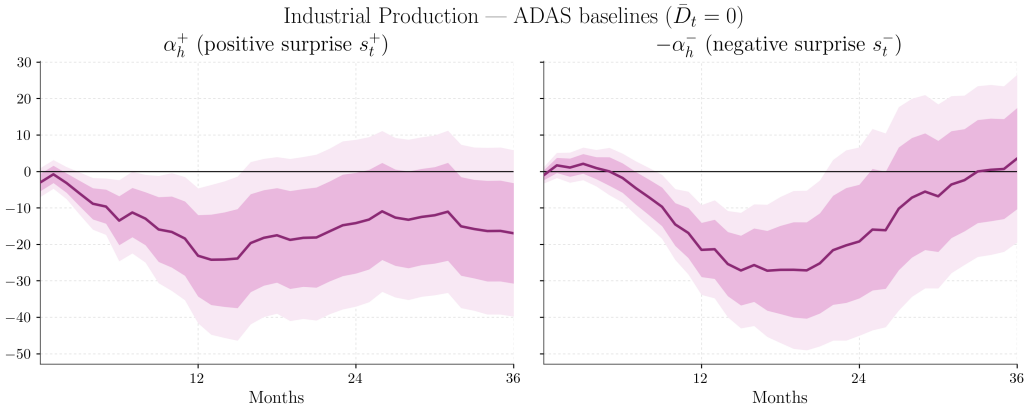


Figure D13: *Baselines*: sign-dependent IRFs when the Market's surprise is positive (left) and negative (right).

Second, the marginal changes in the baseline effects (attenuation/amplification) follow, both shock sign-dependent and disagreement sign-dependent. In particular, the first column displays the positive-disagreement cases and the second column displays the negative-disagreement cases, with first row collecting positive surprises and the second row the negative surprises.

Let us start from states when the Market is negatively surprised (s_t^-: right panel Figure D13, second row Figure D14). That is, we are testing the marginal effects of a *contractionary* increase in the policy rate *when the shock is expansionary*, i.e., negative. We observe an expected, pronounced decrease in industrial production as a baseline effect (evaluated at $\bar{D}_t = 0$), while disagreement attenuates the contraction the more the Fed's projections are above the Market's ($\bar{D}_t^+$); on the other hand, the marginal effects of disagreement when the Market's expectations are above the Fed's follow a non-monotonic trajectory: we observe a short-lived (0-2Q) amplification of the baseline contraction, followed by a delayed attenuation starting at a year after the shock (4-8Q).
Stated plainly, when facing unexpectedly dovish monetary decisions, increasing rates is conventionally contractionary in states of agreement, while it is less contractionary when Market's future expected stance on monetary policy is more accommodative, and it initially deepens but later mitigates when the Market's outlook is tighter than the Fed's.

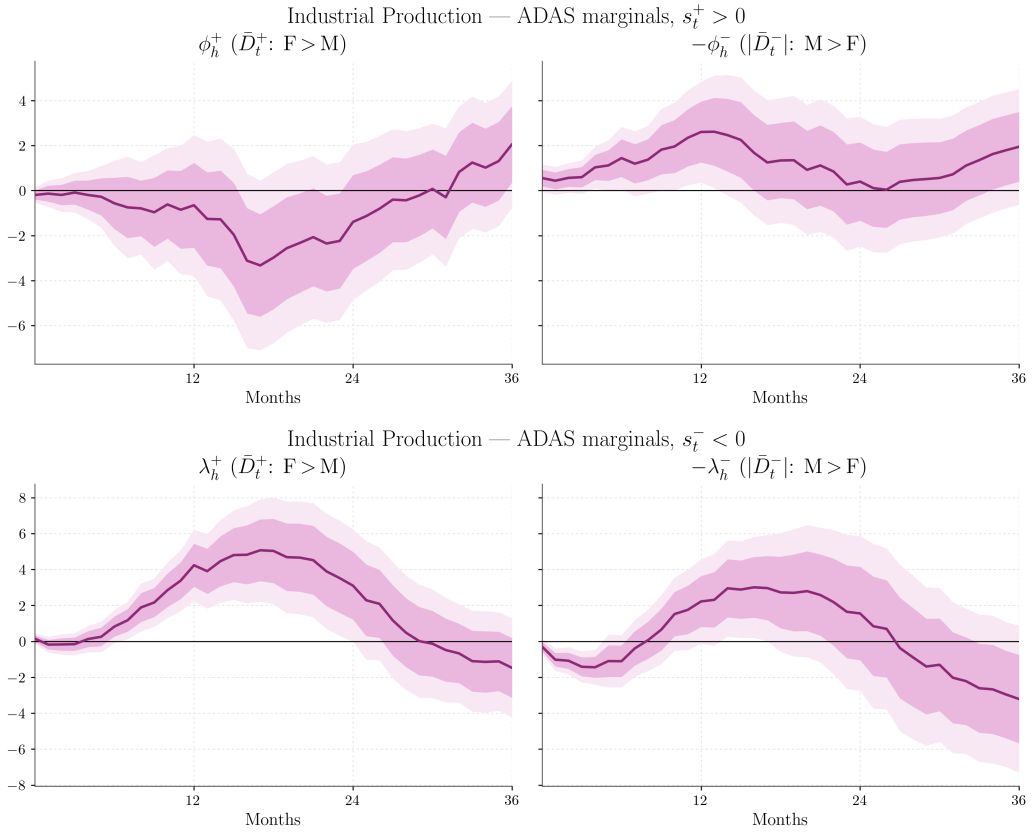


Figure D14: *Marginals*: shock & disagreement sign-dependent IRFs.
Top-left: $(s_t^+, \bar{D}_t^+)$; top-right: $(s_t^+, \bar{D}_t^-)$; bottom-left: $(s_t^-, \bar{D}_t^+)$; bottom-right: $(s_t^-, \bar{D}_t^-)$.

What are the implications for the model's structure? In times of expansionary surprises, on top of the conventional negative baseline New Keynesian effect, the behavior of the expectation wedge in times when the Fed's outlook is more hawkish than the Market's ($W_t > 0$) monotonically attenuates the contraction, meaning that $\Psi_w^H > \Psi_w^L$. Conversely, when the Market's outlook is tighter ($W_t < 0$), higher disagreement produces an initial (0–2Q) amplification of the contraction – representing $\Psi_w^H < \Psi_w^L$ – followed by a delayed (4–8Q) attenuation, implying $\Psi_w^H > \Psi_w^L$.

On the other hand, when the Market is positively surprised (s_t^+: left panel Figure D13, first row Figure D14), we observe a weaker, barely significant contraction in response to a policy rate increase (at $\bar{D}_t = 0$), while disagreement is much less influential in both directions: for positive disagreement scenarios, we obtain a short-lived, delayed (5-7Q) amplification of the contraction which is increasing in the difference between the Fed's projected path and the Market's; when the Market is more hawkish, on the contrary, we cannot detect monotonic effects of larger disagreements on the magnitude of the (already mild) contraction.

From a theory point of view, we first conclude that the baseline New Keynesian

effect is less pronounced when conditioning on already positive shocks. Then, evidence from Figure D14's top panels demonstrates similarly less striking consequences of differential disagreement gaps, with the contraction being weakly (and belatedly) amplified in cases of *more positive* disagreement, signifying $\Psi_w^H < \Psi_w^L$; and appearing orthogonal to the disagreement magnitude in times of negative disagreement, implying $\Psi_w^H \simeq \Psi_w^L$.[64]

D.11 Lead-Lag Correlation Maps: Details & Extras

This appendix details the procedure used to map the lead–lag relation between the monetary surprise s_t and disagreement D_t over horizons $h \in [-H, H]$.

Step 1: Shift & correlate. For each $h \in [-H, H]$ compute

$$\hat{\rho}_h = \operatorname{corr}\big(s_t,\ D_{t+h}\big)$$

where $D_{t+h} = L^h D_t$ if $h > 0$ and $F^{|h|} D_t$ if $h < 0$. (Here L and F denote the one-period lag and lead operators, respectively.)

Step 2: Exact 95% band (Fisher-z).

$$z_h = \frac{1}{2} \ln\Big(\frac{1+\hat{\rho}_h}{1-\hat{\rho}_h}\Big), \qquad \mathrm{SE} = 1/\sqrt{N-3}$$

$$z_h^{\pm} = z_h \pm 1.96\,\mathrm{SE}, \qquad \rho_h^{\pm} = \tanh\big(z_h^{\pm}\big)$$

N is the number of overlapping (s_t, D_{t+h}) pairs at horizon h.

Step 3: Plot. Draw $\hat{\rho}_h$ against h with the shaded confidence band $\big[\rho_h^-,\ \rho_h^+\big]$ and dotted axes at $h = 0$ and $\rho = 0$.

[64] Again notice that this does not represent a statement about the *absolute* effects of (negative) disagreements on contractionary shocks, but only about its dependence on the size of disagreement. That is, the restrictions imposed by the top-right panel of Figure D14 do not exclude *absolute* attenuation or amplification.

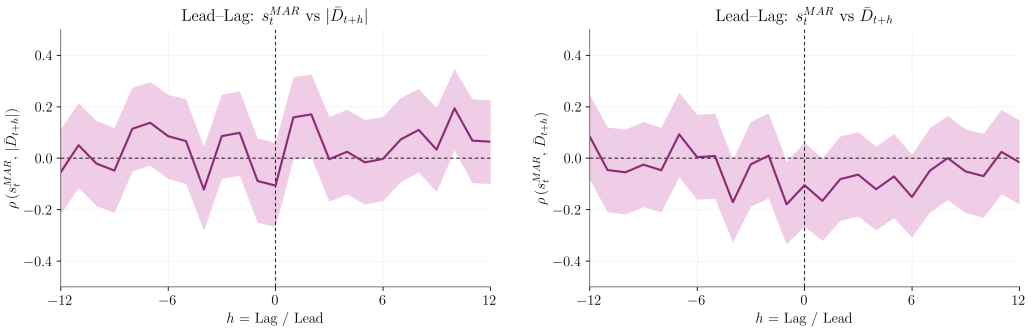


Figure D15: Left panel: Lead-Lag correlation map between monetary policy shocks from Miranda-Agrippino and Ricco (2021), and absolute average disagreement. Right panel: same, with average disagreement. 95% confidence bands computed with Fisher's z-transform (exact method). Estimation sample: 1990-2019, $N = 358$ observations (excludes 2001M9).

Interpretation. Values $\hat{\rho}_h > 0$ for $h > 0$ indicate that disagreement is Granger-predicted by the surprise (information in s_t precedes D_t). Values $\hat{\rho}_h > 0$ for $h < 0$ indicate that shocks are moved by precedent disagreement.

D.12 Treatment Homogeneity

This appendix tests whether the monetary policy shock differs across disagreement states. I re-estimate the local projections using interest-rate outcomes—the federal funds rate (FFR), 1-year, and 2-year Treasury yields—in a three-variable system and in standard specifications (Gertler and Karadi, 2015, Bauer and Swanson, 2023b). I also replace the orthogonalized surprises with leading alternatives (Gertler and Karadi, 2015, Jarociński and Karadi, 2020, Miranda-Agrippino and Ricco, 2021). Let β_h denote the incremental response per unit of disagreement $|D_t|$ and α_h the response extrapolated to $|D_t| = 0$, so that $\alpha_h + \beta_h$ is the response at the average disagreement state. Because interest rates are equilibrium objects—embedding the Fed's endogenous reaction as well as the shock's propagation—state dependence in these responses ($\beta_h \neq 0$) speaks to *transmission*, not to the treatment; state dependence of the identified *surprise itself*, which would instead threaten identification, is tested directly in Appendix D.12.3.

D.12.1 Rate-Outcome Evidence

3-variable SVAR In the differential response depicted in Figure D16, β_h is positive and sizable at intermediate horizons—comparable in magnitude to α_h, with the 90% band excluding zero around the peak ($\beta_{17} \approx 6.4$ against $\alpha_{17} \approx -5.8$). The reading is of two offsetting state components: at $|D_t| = 0$ the policy-rate path reverts below baseline within the year, while in high-disagreement states it stays persistently higher, so the average-state response $\alpha_h + \beta_h$ is close to zero throughout—which is also why the average rate response to the orthogonalized

surprise appears weak. This pattern goes in the *opposite* direction of what would be needed to mechanically generate attenuation in macro outcomes, and therefore strengthens the heterogeneous-*transmission* interpretation.

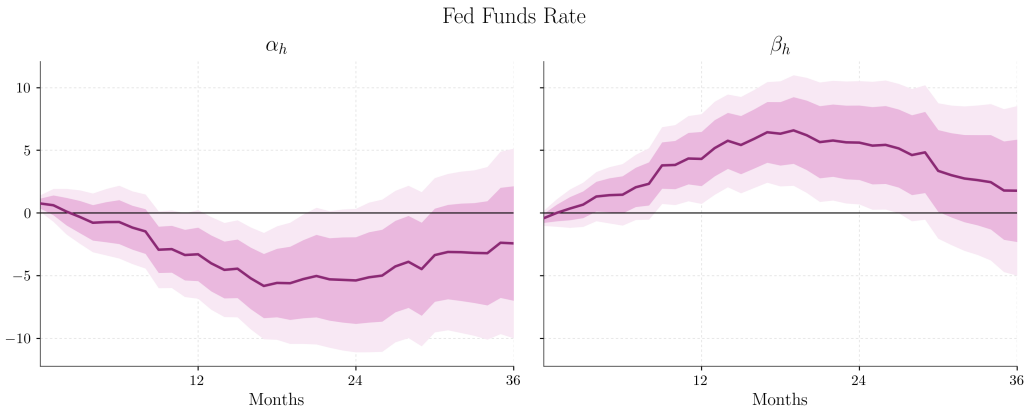


Figure D16: Fed Funds Rate. Local–projection IRFs from the 3-variable system shocked by orthogonalized monetary surprises (Bauer and Swanson (2023a), Amodeo (2025)).Left panel: response at $|D_t| = 0$, α_h. Right panel: differential response β_h. Sample: 1990–2019, $N = 358$.

Gertler and Karadi (2015)'s SVAR Figure D17. Across the policy-relevant window, β_h is close to zero with overlapping bands. Shocks are essentially homogeneous across states at the 1-year maturity.

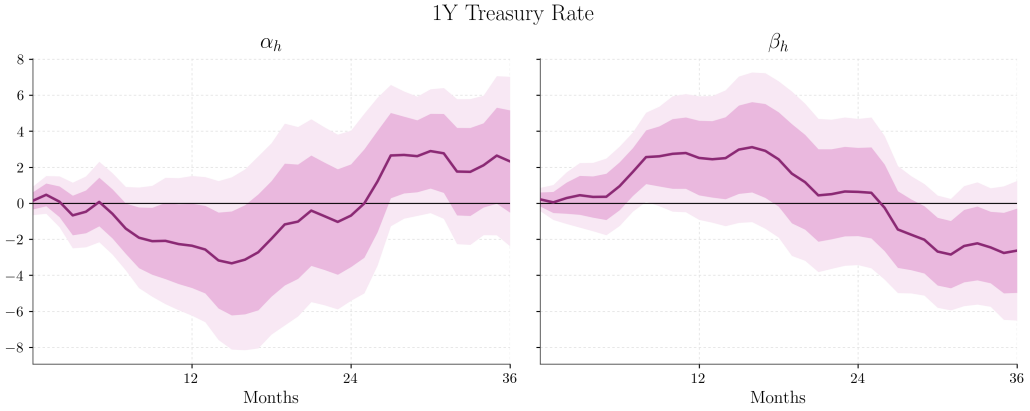


Figure D17: 1-Year Treasury Rate. Local–projection IRFs from the 4-variable system shocked by orthogonalized monetary surprises (Bauer and Swanson (2023a), Amodeo (2025)). Left panel: baseline response α_h. Right panel: differential response β_h. Sample: 1990–2019, $N = 358$.

Bauer and Swanson (2023b)'s SVAR Figure D18. β_h is near zero throughout, and its estimates are less statistically significant than in smaller systems. Shocks appear essentially indistinguishable across states at the 2-year maturity.



Figure D18: 2-Year Treasury Yield. Local–projection IRFs from the 6-variable system shocked by orthogonalized monetary surprises (Bauer and Swanson (2023a), Amodeo (2025)). Left panel: baseline response α_h. Right panel: differential response β_h. Sample: 1990–2019, $N = 358$.

In summary, any FFR differences work against a spurious-attenuation story. At the 1-year and 2-year horizons—those most relevant and widely used in the literature (Gertler and Karadi (2015), Bauer and Swanson (2023a))—the shocks are comparable across disagreement states, while attenuation in real activity and prices persists. The data favor heterogeneous *transmission*, not heterogeneous *shocks*.

D.12.2 Alternative Shocks

This section verifies that results are not driven by the shock measure. For each leading identification, the 3-variable LP with industrial production delivers the same qualitative pattern: α_h shows a contraction followed by recovery, while the $|D_t|$-dependent effects β_h have a similar shape.

Bauer and Swanson (2023b)'s unadjusted orthogonalized shocks

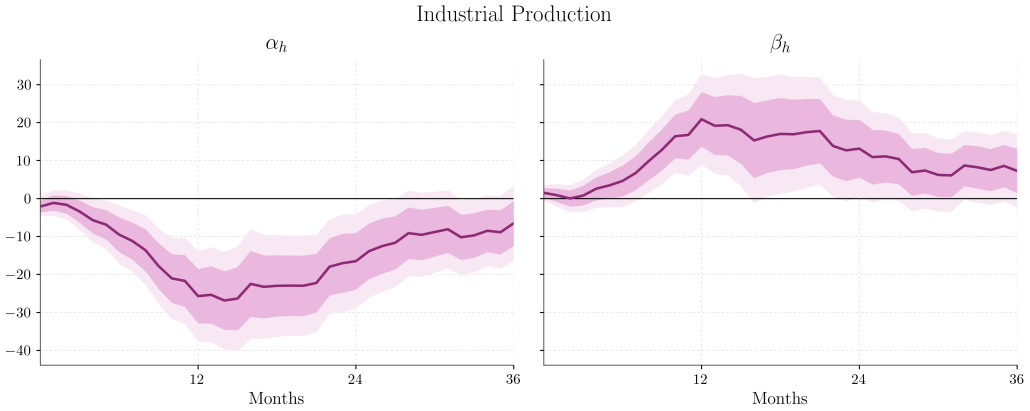


Figure D19: Industrial Production, 3-variable system shocked by orthogonalized monetary surprises as in Bauer and Swanson (2023b). Left: α_h; Right: β_h. Sample: 1990–2019, $N = 358$.

Gertler and Karadi (2015)'s shock

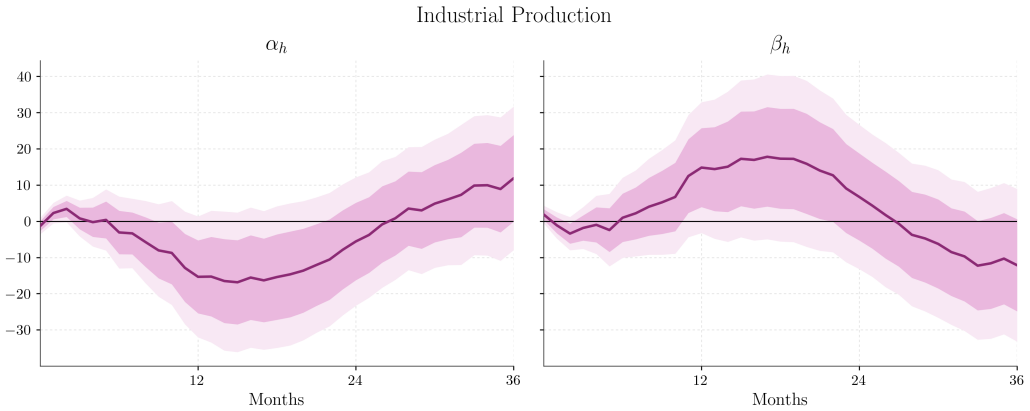


Figure D20: Industrial Production. 3-variable system shocked by high-frequency monetary surprises as in Gertler and Karadi (2015). Left: α_h; Right: β_h. Sample: 1991–2012, $N = 257$.

Miranda-Agrippino and Ricco (2021)'s MPI shock

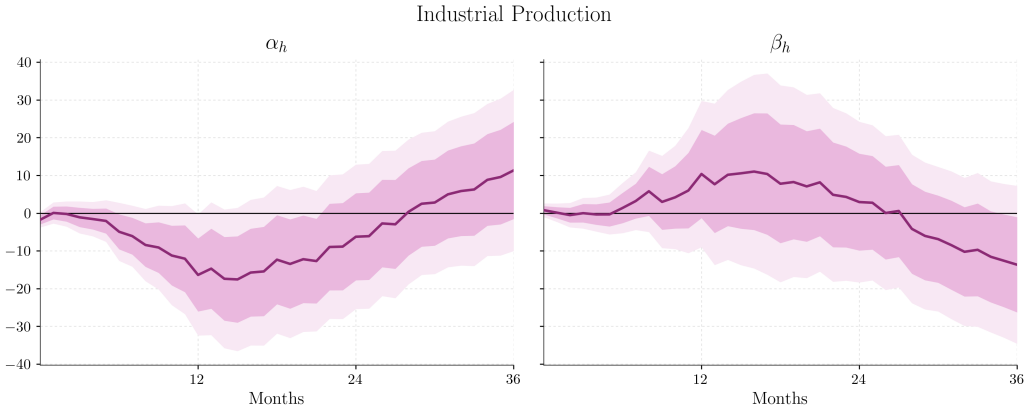


Figure D21: Industrial Production, 3-variable system shocked by high-frequency Monetary Policy Instrument (MPI) as in Miranda-Agrippino and Ricco (2021). Left: α_h; Right: β_h. Sample: 1992–2009, $N = 215$.

Jarociński and Karadi (2020)'s First Principal Component shock

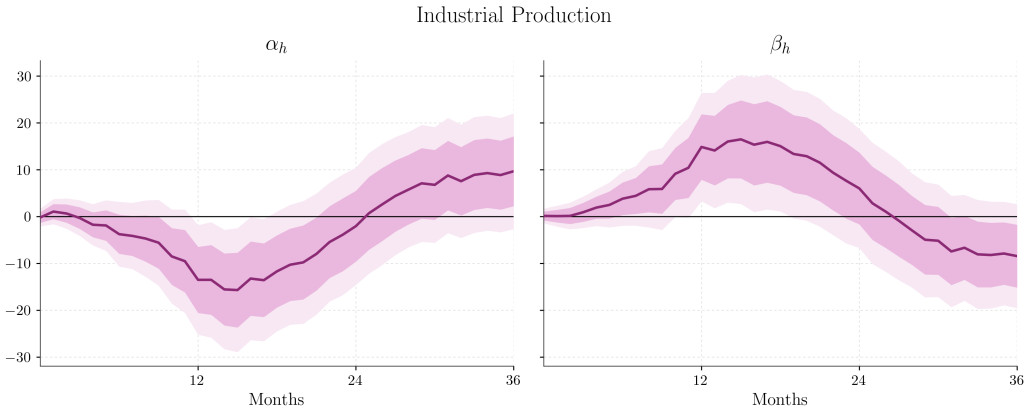


Figure D22: Industrial Production, 3-variable system shocked by high-frequency monetary surprises as in Jarociński and Karadi (2020) (first principal component). Left: α_h; Right: β_h. Sample: 1991–2018, $N = 331$.

Jarociński and Karadi (2020)'s Identified Monetary Policy shock

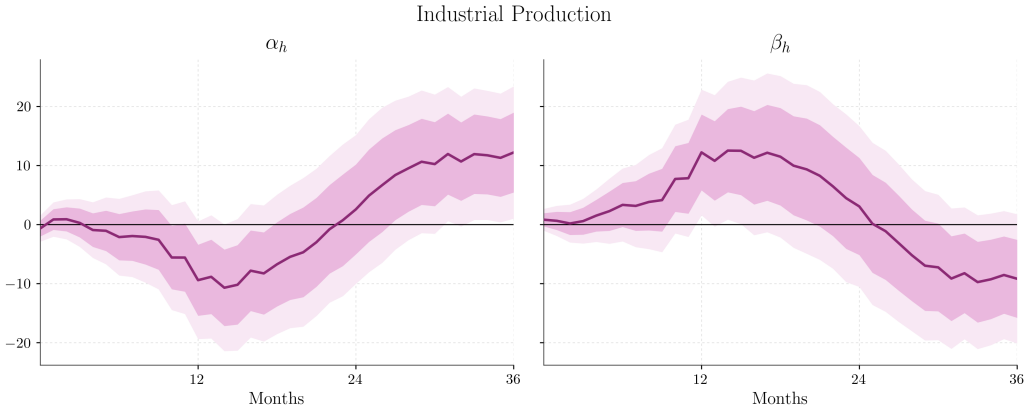


Figure D23: Industrial Production, 3-variable system shocked by high-frequency monetary surprises as in Jarociński and Karadi (2020) (monetary shock "poor man" version). Left: α_h; Right: β_h. Sample: 1991–2018, $N = 331$.

D.12.3 Shocks Across Disagreement States

This appendix tests whether the monetary shock is drawn from the same distribution across quartiles of absolute average Fed–Market disagreement $|\bar{D}_t|$. Claims about state-dependent transmission require treatment homogeneity: if the shock's distribution changes with the state, estimated heterogeneity could reflect different treatments rather than different propagation.

Results reveal that disagreement states do not conceal distinct shock regimes. Then, it is unlikely for the results to be driven by the shocks themselves differing across states, supporting the paper's main point of differential policy transmission.

Figure D24 shows the histograms of the baseline shocks across quartiles of the interacting variable (absolute average disagreement). The four panels look the similar: centered at zero with the same sharp spike at 0, thin and broadly symmetric tails, and no visible shift in mass across quartiles. Because bins and axes are kept identical, visual similarity is meaningful. Moreover, I next implement statistical formal tests of distributional similarity. Figure D25 displays the same notions in a more direct way using kernel densities.

Table D5 displays analytical results from the comparison of shocks across disagreement quartiles. Means are near zero in every quartile; standard deviations rises modestly with disagreement, consistent with mild heteroskedasticity. The mass at zero is statistically identical across groups.

Panel A reports sample size, means, dispersion (SD), and the share of exact zeros by quartile. Panel B reports p-values from nonparametric tests: equality of typical levels (Kruskal–Wallis), equality of zero shares, and pairwise KS tests of whole-distribution equality. When $p > 0.05$ we treat groups as statistically similar for that feature.

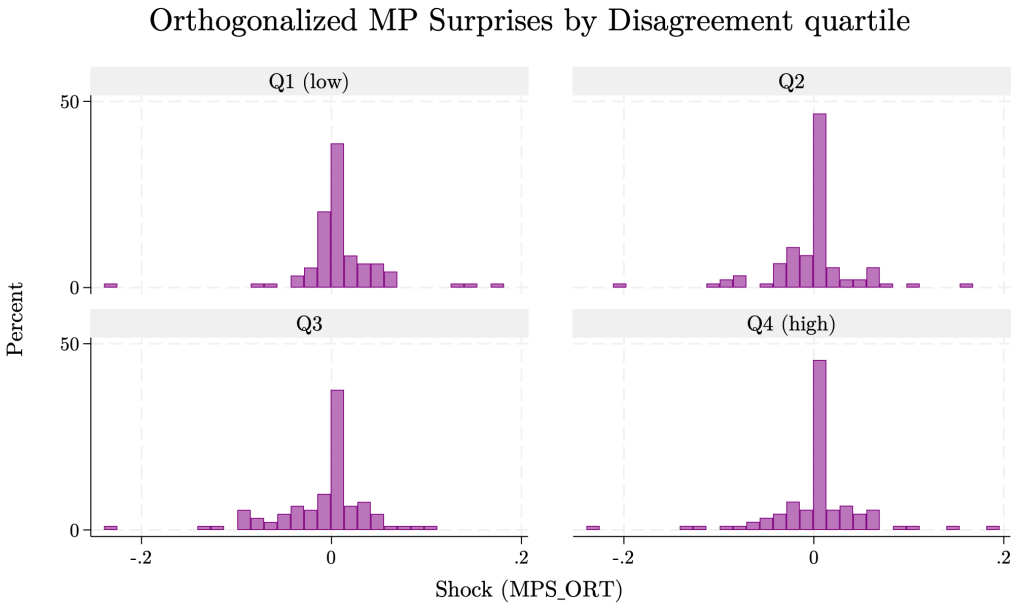


Figure D24: Histograms of s_t by disagreement quartile; percent scale.

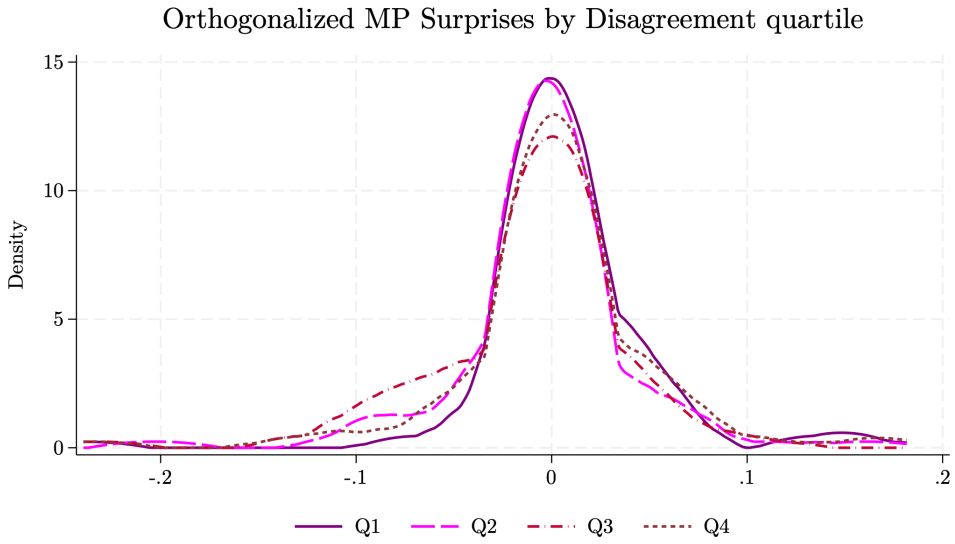


Figure D25: Kernel densities of s_t by disagreement quartile. A single bandwidth is used for all groups and the support is held fixed.

Kruskal–Wallis asks whether the typical level of the shock is the same across groups without assuming normality. The "zero-mass" test simply compares the share of exact zeros across groups. The two-sample Kolmogorov–Smirnov (KS) test compares the entire shape of two distributions—center, spread, and tails—

at once. Large p-values in these tests mean the data do not provide evidence of differences across disagreement states.
Kruskal–Wallis tests fails to reject equal locations. The proportion of zeros is constant as well. Finally, only one unadjusted KS comparison is low (Q1–Q3), but it is not significant after Holm–Bonferroni across six pairs.
In summary, graphical inspection and formal tests allow us to treat shocks as distributionally equivalent across disagreement states.

Panel A. Moments by quartile

	Q1 (low)	Q2	Q3	Q4 (high)
N	93	92	93	92
Mean	0.0069	-0.0030	-0.0105	0.0018
SD	0.0440	0.0435	0.0478	0.0515
$\Pr(s_t = 0)$	0.269	0.326	0.333	0.304

Panel B. Distributional tests (p-values unless noted)

Kruskal–Wallis $\chi^2(3)$	6.355	$p = 0.0956$
Zero-mass test $\chi^2(3)$	1.097	$p = 0.778$

Two-sample Kolmogorov-Smirnov p-values

	Q1	Q2	Q3	Q4
Q1	—	0.304	**0.027**	0.519
Q2	0.304	—	0.551	0.771
Q3	**0.027**	0.551	—	0.291
Q4	0.519	0.771	0.291	—

Table D5: Moments and distributional tests of s_t across disagreement quartiles. Quartiles by $|\bar{D}_t|$. KS entries are combined p-values. After family-wise error control (Holm–Bonferroni), no KS comparison is significant at 5%.

D.12.4 Shocks Across the ZLB

This appendix tests whether the monetary shock s_t is drawn from the same distribution in the Zero Lower Bound period versus outside it. Claims about state-dependent transmission across monetary regimes require treatment homogeneity: if the shock distribution changes mechanically in the ZLB, estimated heterogeneity could reflect different treatments rather than different propagation.

I define the ZLB as 2008m12–2015m12 and Non-ZLB otherwise. After dropping missing shocks, the sample is $N = 420$ (Non-ZLB: 335; ZLB: 85). Figure D26 plots comparable histograms by regime, and Figure D27 overlays kernel densities using

a common bandwidth and common support.

Overall, the shock looks centered at zero in both regimes, with a similar spike at exactly zero. Formal tests confirm equality of location and zero-mass, while rejecting equality of scale: the shock is statistically more dispersed outside the ZLB. Importantly, once standardized by the regime-specific standard deviation, a KS test fails to reject equality of distributions ($p = 0.822$), suggesting the two regimes are the same up to scale.

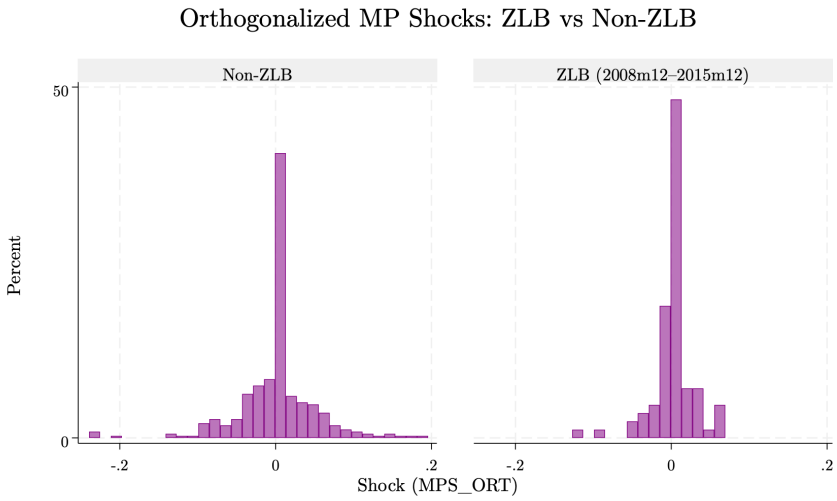


Figure D26: Histograms of s_t by regime (ZLB vs Non-ZLB); percent scale; common binning.

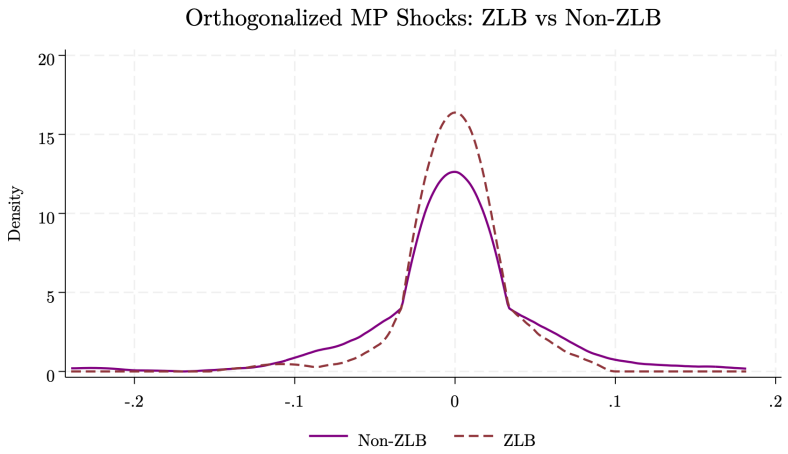


Figure D27: Kernel densities of s_t by regime. A single bandwidth is used for both groups.

Table D6 reports moments and distributional tests. Panel A shows regime-specific moments and the share of exact zeros. Panel B reports p-values for (i) equality of the zero share (chi-square test), (ii) equality of central tendency (two-sample t

test and Wilcoxon rank-sum), (iii) equality of the whole distribution (two-sample KS), and (iv) equality of dispersion (variance-ratio and Levene/Brown–Forsythe).

The evidence supports treatment homogeneity in level and zero-mass: we fail to reject equal means ($p = 0.982$), equal typical levels ($p = 0.985$), equal whole-distribution shapes at conventional levels ($p = 0.142$), and equal mass at zero ($p = 0.580$). The only robust difference is dispersion: the shock variance is significantly larger outside the ZLB (variance-ratio $p < 0.001$; Levene/BF $p = 0.0011$). Standardizing within regime removes this difference and yields no evidence against distributional equality (standardized KS $p = 0.822$). In short, ZLB vs Non-ZLB does not conceal distinct *shock types*; differences are primarily scale.

Panel A. Moments by regime

	Non-ZLB	ZLB (2008m12–2015m12)
N	335	85
Mean	-0.000	0.000
SD	0.051	0.027
$\Pr(s_t = 0)$	0.299	0.329

Panel B. Distributional tests (p-values unless noted)

Zero-mass test $\chi^2(1)$ (share $s_t = 0$)	0.306 $(p = 0.580)$
Two-sample t test (mean, conditional on $s_t \neq 0$)	$p = 0.982$
Wilcoxon rank-sum (conditional on $s_t \neq 0$)	$p = 0.985$
KS test (conditional on $s_t \neq 0$)	$p = 0.142$
Variance-ratio test (conditional on $s_t \neq 0$)	$F = 3.2$ $(p \sim 0)$
Levene/Brown–Forsythe (conditional on $s_t \neq 0$)	$p = 0.001$
KS test after within-regime standardization ($s_t/\hat{\sigma}_{\text{regime}}$)	$p = 0.822$

Table D6: Moments and distributional tests of s_t across ZLB vs Non-ZLB regimes. ZLB defined as 2008m12–2015m12. Most tests condition on $s_t \neq 0$ to separate the spike at zero from the continuous component.

D.13 Sign versus Magnitude of Disagreement

The baseline local projections conditions monetary transmission on the absolute disagreement $|\bar{D}_t|$. It is natural to ask whether the *sign* of the gap — the Fed projecting above versus below the market — matters for transmission, and not only its size. I address this in two steps: a sign-partitioned specification, and a direct joint test of sign versus magnitude.

Sign-partitioned specification. A first, intuitive cut splits disagreement into its positive and negative parts and lets each load the shock separately:

$$\begin{aligned} \bar{D}_t^+ &= \max(\bar{D}_t, 0), \quad \bar{D}_t^- = \min(\bar{D}_t, 0) \\ \mathbf{y}_{t+h} &= \mathbf{c}_h + \boldsymbol{\alpha}_h s_t + \boldsymbol{\beta}_h^+ (s_t \bar{D}_t^+) + \boldsymbol{\beta}_h^- (s_t \bar{D}_t^-) + \textsc{Lags}_{t-1}^L + \boldsymbol{\varepsilon}_{t+h} \end{aligned} \tag{D.2}$$

To clarify, I treat $\bar{D}_t^+$and $\bar{D}_t^-$ as the positive and negative parts of $\bar{D}_t$, so $|\bar{D}_t| = \bar{D}_t^+ - \bar{D}_t^-$. Increasing $\bar{D}_t^+$ or making $\bar{D}_t^-$ more negative *raises* the distance from zero, so for $s_t > 0$ attenuation corresponds to $\boldsymbol{\beta}_h^+ > 0$ and $\boldsymbol{\beta}_h^- < 0$[65].

As in previous specifications, $\boldsymbol{\alpha}_h$ is the baseline impulse response at zero disagreement, while $\boldsymbol{\beta}_h^+$ and $\boldsymbol{\beta}_h^-$ load the shock by the positive and negative components of disagreement. Figure D28 shows that attenuation of the response of industrial production (i.e., a less negative IRF for IP) obtains when $\boldsymbol{\beta}_h^+ > 0$ for $\bar{D}_t^+ > 0$ and when $\boldsymbol{\beta}_h^- < 0$ for $\bar{D}_t^- < 0$. Empirically, the point estimates align with this prediction on both sides of the distribution across horizons, implying that larger disagreement — whether the Fed projects above or below markets — systematically dampens monetary transmission to real activity.

To aid interpretation, D_t^+ and D_t^- are rescaled so that their coefficients can be read as attenuation at "typical" (average) levels of the partitioned disagreement. While confidence bands widen once rich lag structures are included — unsurprising given the proliferation of regressors — the shape and relative magnitude of $\widehat{\boldsymbol{\beta}}_h^+$ and $\widehat{\boldsymbol{\beta}}_h^-$ are remarkably similar, suggesting that the level of disagreement, rather than its sign, is the key state variable for transmission. In richer systems of equations (including EBP_t and/or $UNEMP_t, PCOMM_t$), this finding is even more predominant, and this symmetry in the effects emerges clearly, with the only difference being that negative disagreement appears to have more persistent attenuating effects.

[65] I am interested in increases in $|\bar{D}_t|$. On the negative side this means $\bar{D}_t^-$ becoming more negative. Since $|\bar{D}_t^-| = -\bar{D}_t^-, \partial \text{IRF}_{t+h}^{IP} / \partial |\bar{D}_t^-| = -\beta_h^- s_t$; hence $\beta_h^- < 0$ implies attenuation.

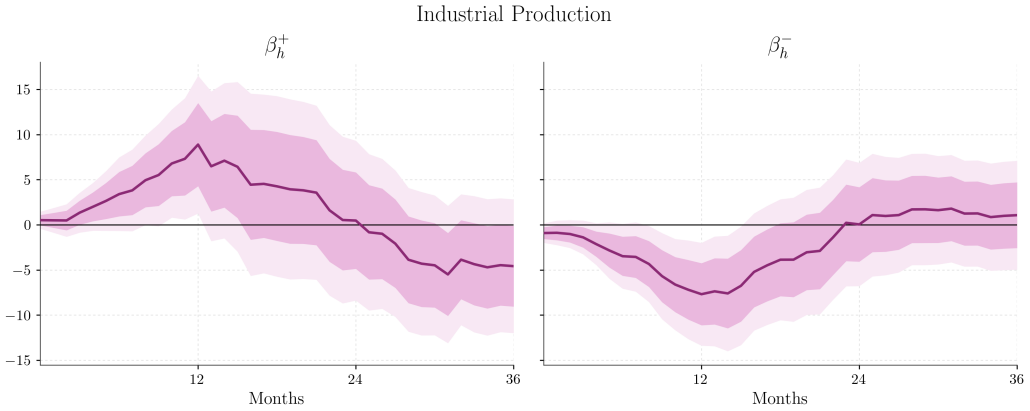


Figure D28: Impulse Response Functions of Industrial Production to an orthogonalized monetary policy shock, estimated from lag-augmented local projections, equation (D.2). Left panel: $\boldsymbol{\beta}_h^+$; right panel: $\boldsymbol{\beta}_h^-$. Confidence bands at 68 and 90% levels. Sample: 1990-2019, $N = 358$ observations (excludes 2001M9 for the 9/11 Terrorist Attacks).

This split is intuitive but low-powered: partitioning the sample discards within-state variation, so it cannot by itself rule out a sign effect.

A direct sign-versus-magnitude test. I therefore test the question in a single specification that includes both the signed and the absolute interaction:

$$\mathbf{y}_{t+h} = \mathbf{c}_h + \boldsymbol{\alpha}_h s_t + \boldsymbol{\beta}_h^O \left(s_t \bar{D}_t\right) + \boldsymbol{\beta}_h^E \left(s_t |\bar{D}_t|\right) + \textsc{Lags}_{t-1}^L + \boldsymbol{\varepsilon}_{t+h}. \tag{D.3}$$

Here $\boldsymbol{\beta}_h^O$ picks up the part of the state-dependent response that is an *odd* function (sign-dependent) in $\bar{D}_t$ — the component a sign mechanism would predict — and $\boldsymbol{\beta}_h^E$ the part that is *even* (magnitude-dependent). The even loading carries the attenuation, peaking near one year, while the odd loading is statistically indistinguishable from zero at every horizon (Figure D29).

Under the orthogonalization of Section 3.3.2 the shock is independent of $\text{sign}(\bar{D}_t)$, which is what identifies $\boldsymbol{\beta}^O$; the test does *not* require $\bar{D}_t$ to be symmetrically distributed (its mild skew is immaterial). The result is robust across the disagreement timings and the sample window.[66]

[66]The odd loading is insignificant at all horizons under day-before, book-date and FOMC-month timings, and on both the full and the ZLB-excluded samples. Its point estimate is essentially zero, with sign *opposite* to attenuation; the test is adequately powered against the relevant alternative: a sign-driven mechanism would predict an odd loading of the same order as the even one ($\widehat{\boldsymbol{\beta}}_{12}^E = 12.7$), while the 90% confidence interval for $\widehat{\boldsymbol{\beta}}_{12}^O$ is $[-8.5,\ 6.1]$ — a same-order component is excluded. The two interactions are nearly orthogonal in sample (correlation 0.005, variance-inflation factor 1.5), so the null odd loading is not a collinearity artifact.

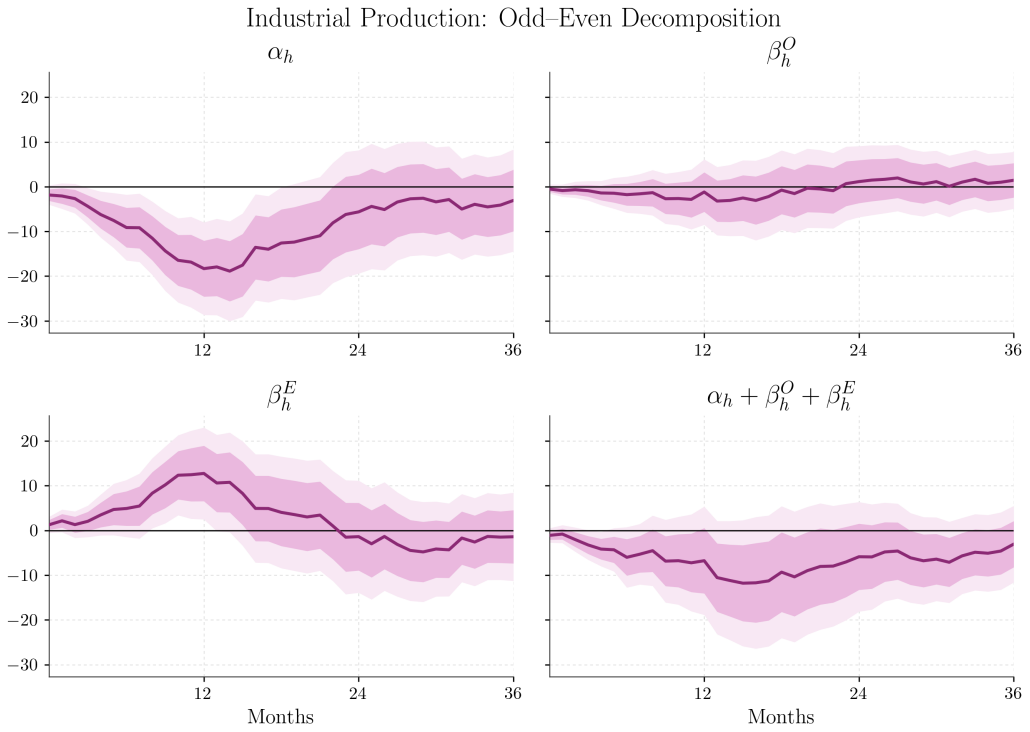


Figure D29: Odd–even decomposition of the state-dependent response of industrial production, from the joint specification (D.3). Panels: $\boldsymbol{\alpha}_h$ (response at zero disagreement), $\boldsymbol{\beta}_h^O$ (signed/odd loading), $\boldsymbol{\beta}_h^E$ (absolute/even loading), and the total $\boldsymbol{\alpha}_h + \boldsymbol{\beta}_h^O + \boldsymbol{\beta}_h^E$. The signed and absolute interactors are placed on a common scale (the mean absolute disagreement), so the loadings are directly comparable across panels; the $\boldsymbol{\alpha}_h$ panel shows that the baseline response is unchanged by the addition of the odd term. Confidence bands at 68 and 90% levels; axes shared. Sample: 1990–2019, $N = 358$.

Reconciliation with the theory wedge. This is consistent with the signed wedge of Section 4.1. The interaction coefficient measures how the *transmission* of a shock varies with the state — the cross-partial $\partial^2 \mathbf{y}_{t+h}/\partial s_t \, \partial \bar{D}_t$ — a distinct object from the effect of disagreement on the *level* of demand. A signed level effect is fully compatible with a transmission that is even in $\bar{D}_t$: the former concerns where demand sits, the latter how policy passes through, and the test isolates the latter. The specification conditions on the contemporaneous signed disagreement $\bar{D}_t$ among its controls, so the signed level channel is held fixed; the finding is that the shock's *transmission* carries no signed component ($\boldsymbol{\beta}_h^O \approx 0$).[67] In the learning model of Section 5.1, this even dependence arises under Assumption 1 — credibility declines with the *magnitude* of the gap — so that identical surprises move beliefs less when disagreement is large, irrespective of its sign; a mechanism-level check on the wedge's own updating is reported below.

[67] The result is insensitive to the level controls: $\widehat{\boldsymbol{\beta}}^O$ is statistically zero at every horizon whether or not the signed level $\bar{D}_t$ and its lags are included, because under $s_t \perp \bar{D}_t$ the shock interaction is near-orthogonal to the level term.

A mechanism-level check. Assumption 1 restricts how the *updating step* responds to surprises: the Market's revision of its expected path should scale with the surprise, damped by the magnitude — not the direction — of the pre-meeting gap. Both margins are testable directly on the wedge. Figure D30 augments the wedge local projection (23) with a *signed* surprise interaction alongside the *absolute* one and plots the two loadings across horizons. The signed loading (left panel) is statistically indistinguishable from zero at every horizon ($h = 1, \ldots, 36$; minimum $p = 0.13$): the wedge's response to a surprise does not depend on *which* side of the Market the Fed sits on. Because the baseline wedge response to a surprise is negative, damping implies a *positive* absolute loading — a larger gap pulls that response toward zero — and this is what the right panel shows: the absolute loading is positive and significant at the short horizons where the meeting-window revision of the wedge concentrates (notably $h = 1$–3). This is exactly the sign-versus-magnitude pattern Assumption 1 predicts.

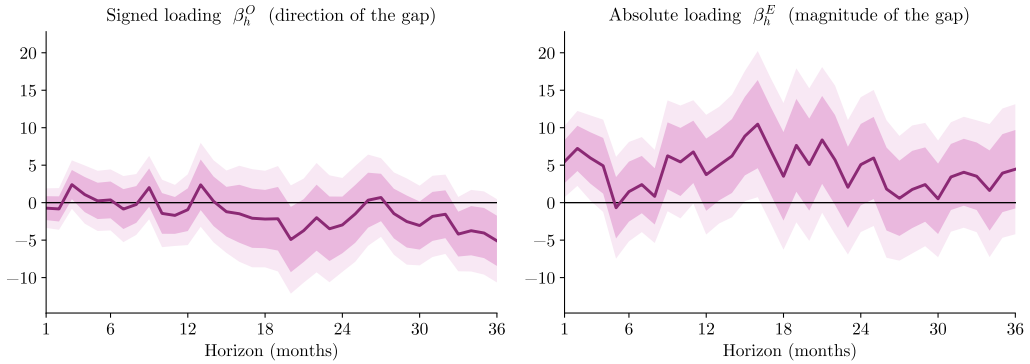


Figure D30: Wedge response to a monetary surprise, interacted with the *signed* (left) and *absolute* (right) pre-meeting Fed–Market gap, across horizons $h = 1, \ldots, 36$ (family local projection (23); $N = 332$; shaded 68% and 90% bands). The signed loading is indistinguishable from zero at every horizon; the absolute loading is positive and significant at short horizons, the pattern implied by magnitude-based damping.

D.14 Is Rate-Path Disagreement Just Inflation Disagreement?

A natural objection to the baseline design is that Fed–Market disagreement about the *rate path* may merely repackage disagreement about *inflation*: if the two sides read price pressures differently, any Taylor-type rule maps that difference into different projected paths, and the attenuation of Section 3 would concern disagreement over fundamentals rather than over the conduct of policy. This appendix measures Fed–Market inflation disagreement directly and reaches two conclusions. Inflation disagreement is a genuine component of rate-path disagreement at the one-year horizon — the correlation documented by Caballero and Simsek (2022). It is nonetheless not what drives the attenuation: conditioning the baseline local projection on inflation disagreement leaves the magnitude loading essentially un-

changed, while the inflation loading is indistinguishable from zero at every horizon.

Measurement. I pair the Greenbook CPI forecast with the mean CPI forecast from the Survey of Professional Forecasters (SPF) for the same target quarter, matched in real time — each meeting uses the latest survey whose submission deadline precedes the Greenbook cutoff — and average the absolute per-horizon gaps over the survey's forecast window,

$$D^{\pi}_{h,t} = \mathbb{E}^{F}_{t}(\pi_{t+h}) - \mathbb{E}^{S}_{t}(\pi_{t+h}), \quad |\bar{D}^{\pi}_{t}| = \frac{1}{|H_t|}\sum_{h\in H_t}\left|D^{\pi}_{h,t}\right|, \quad H_t \subseteq \{1,\dots,4\}, \tag{D.4}$$

with $\mathbb{E}^{S}_{t}$ the survey consensus and h in quarters; the series is interpolated to monthly frequency and mean-rescaled exactly as $|\bar{D}_t|$. A market-priced variant (Greenbook minus TIPS-breakeven forwards, available from 2004) delivers the same qualitative reading.

Inflation disagreement is part of the one-year rate gap. Regressing, meeting by meeting, the signed rate gap on the signed inflation gap at the same target quarter,

$$D_{h,t} = a_h + b_h\, D^{\pi}_{h,t} + u_{h,t}, \tag{D.5}$$

gives a loading that rises sharply with horizon (Table D7): indistinguishable from zero at one and two quarters, where the near-term path is effectively telegraphed and beliefs about inflation have little room to move it, and 0.80 ($R^2 = 0.38$) at four quarters, where policy is genuinely open. This is the pattern in Caballero and Simsek (2022) (their Table D.1); its magnitude is sample- and gauge-dependent — over 1990–2013 with the CPI gauge the four-quarter loading is 0.30 ($R^2 = 0.08$) — but the horizon gradient itself is robust. Rate disagreement therefore *contains* inflation disagreement at business-cycle horizons, which is precisely why whether transmission loads on it must be settled conditionally rather than by an unconditional correlation.[68]

[68] The two states are in fact nearly orthogonal in sample: the late-sample divergence between the FOMC's projections and market pricing pairs the largest rate disagreement with the smallest inflation disagreement. Orthogonal states make the conditional projection well-identified rather than collinear.

Horizon	$\hat{b}_h$	t	R^2	N
1	−0.04	−2.2	0.02	229
2	+0.02	0.5	0.00	224
3	+0.39	4.8	0.11	190
4	+0.80	7.4	0.38	74

Table D7: Inflation disagreement and the one-year rate gap. FOMC-meeting-level regressions (D.5) of the signed rate gap on the signed inflation gap at the same target quarter, 1990–2019; heteroskedasticity-robust t. The four-quarter loading matches Caballero and Simsek (2022), Table D.1.

Conditioning transmission on inflation disagreement. I augment the baseline (2) with the inflation-disagreement interaction, entered exactly as the rate-disagreement state — contemporaneously, with its own level and interaction lags absorbed into $\textsc{Lags}_{t-1}^{L}$:

$$\mathbf{y}_{t+h} = \mathbf{c}_h + \boldsymbol{\alpha}_h s_t + \boldsymbol{\beta}_h \left(s_t|\bar{D}_t|\right) + \boldsymbol{\gamma}_h \left(s_t|\bar{D}_t^{\pi}|\right) + \textsc{Lags}_{t-1}^{L} + \boldsymbol{\varepsilon}_{t+h}. \tag{D.6}$$

Because the survey series begins in 1990M3 the conditional sample is $N = 346$. Figure D31 therefore plots both loadings across horizons, with the baseline re-estimated on that common sample. The attenuation does not move — at the one-year peak $\widehat{\boldsymbol{\beta}}_{12}$ goes from 14.27 to 14.47 — and $\widehat{\boldsymbol{\gamma}}_h$ is insignificant at every one of the 37 horizons, with a point estimate that hovers around zero. Adding a Greenbook-minus-survey real-GDP-growth gap as a third interaction — controlling for disagreement about the full macroeconomic outlook — leaves $\widehat{\boldsymbol{\beta}}_{12} = 13.35$, with both outlook loadings insignificant.



Figure D31: The rate-path attenuation survives conditioning on inflation disagreement. Coefficient paths from the conditional local projection (D.6), industrial production, on the common sample 1990M3–2019M12 ($N = 346$); shaded 68% and 90% bands. Left: the rate-disagreement loading $\boldsymbol{\beta}_h$ retains the baseline attenuation hump, peaking near one year (at $h = 12$, 14.47 against a same-sample baseline of 14.27). Right: the inflation-disagreement loading $\boldsymbol{\gamma}_h$ is indistinguishable from zero at every horizon. Axes shared.

D.15 Disagreement in LPs: Formal Results

Let s_t be the MP shock and $D_t = \mathbb{E}_t^F - \mathbb{E}_t^M$ the disagreement. Define $D_t^+ = \max(D_t, 0)$, $D_t^- = \min(D_t, 0) \leq 0$. Then

$$D_t = D_t^+ + D_t^-, \quad |D_t| = D_t^+ - D_t^-, \quad D_t^+ D_t^- = 0 \text{ a.s.}$$

Let the *true* LP (after partialling out lags and s_t itself) be

$$y'_{t+h} = \beta_h \, s_t |D_t| + u_{t+h}, \qquad \mathbb{E}[u_{t+h} \mid s_t, D_t] = 0, \quad \mathbb{E}[s_t] = 0, \quad s_t \perp D_t.$$

Case (1): absolute-value interaction. Regress y'_{t+h} on $X = s_t|D_t|$.

$$\hat{\beta}_h = \frac{\mathrm{Cov}(X, y')}{\mathrm{Var}(X)} = \beta_h \frac{\mathrm{Var}(X)}{\mathrm{Var}(X)} \xrightarrow{p} \beta_h.$$

Intuition: identification is direct because the regressor equals the true interaction.

Case (2): signed interaction only. Estimate $y'_{t+h} = \beta'_h \, s_t D_t + \text{noise}$.

$$\beta'_h = \frac{\mathrm{Cov}(s_t D_t, \, y')}{\mathrm{Var}(s_t D_t)} = \beta_h \frac{\mathbb{E}[s_t^2 D_t \cdot |D_t|]}{\mathbb{E}[s_t^2 D_t^2]} = \beta_h \frac{\mathbb{E}[D_t \cdot |D_t|]}{\mathbb{E}[D_t^2]}.$$

If D_t is symmetric about 0 (so $\mathbb{E}[D_t \cdot |D_t|] = \mathbb{E}[\mathrm{sign}(D_t) D_t^2] = 0$), then $\beta'_h = 0$. *Intuition:* the effect is *even* in D_t (depends on magnitude, not sign); the signed regressor averages positive and negative sides to zero.

Case (3): split by sign. Estimate $y'_{t+h} = \beta_h^+ (s_t D_t^+) + \beta_h^- (s_t D_t^-) + \text{noise}$. Because $D_t^+ D_t^- = 0$ a.s.,

$$\mathrm{Cov}(s_t D_t^+, \, s_t D_t^-) = \mathbb{E}[s_t^2 D_t^+ D_t^-] = 0,$$

so $(s_t D_t^+, s_t D_t^-)$ are orthogonal regressors. Using $y'_{t+h} = \beta_h s_t (D_t^+ - D_t^-) + u$,

$$\beta_h^+ = \frac{\mathrm{Cov}(s_t D_t^+, y')}{\mathrm{Var}(s_t D_t^+)} = \beta_h \frac{\mathrm{Var}(s_t D_t^+)}{\mathrm{Var}(s_t D_t^+)} = \beta_h,$$

$$\beta_h^- = \frac{\mathrm{Cov}(s_t D_t^-, y')}{\mathrm{Var}(s_t D_t^-)} = \beta_h \frac{-\mathrm{Var}(s_t D_t^-)}{\mathrm{Var}(s_t D_t^-)} = -\beta_h.$$

Thus plotting β_h^+ and $-\beta_h^-$ reproduces the $|D|$-specification.

Intuition: $D_t^- \leq 0$; increasing D_t^- (less negative) *reduces* $|D_t|$, so its coefficient must be the negative of the positive-side coefficient. Splitting prevents the cancellation that kills Case (2).

Remarks. If symmetry or $s_t \perp \mathrm{sign}(D_t)$ fails, Case (2) can load nonzero: $\beta'_h = \beta_h \dfrac{E[\mathrm{sign}(D_t) D_t^2]}{\mathbb{E}[D_t^2]}$. Testing symmetry: Wald $H_0 : \ \beta_h^+ + \beta_h^- = 0$.

Proof $\mathbb{E}[D \cdot |D|] = \mathbb{E}[\text{sign}(D)D^2] = 0$

$|D|$ is *even*: $|-d| = |d|$. D and $\text{sign}(D)$ are *odd*. Thus $D \cdot |D| = \text{sign}(D)D^2$ is *odd*. For this illustration I assume a symmetric distribution: D has pdf f with $f(d) = f(-d)$ and $\mathbb{E}[D^2] < \infty$. (The joint specification (D.3) does not require symmetry: without it the signed-only projection recovers $\beta_h \, \mathbb{E}[D\,|D|]/\,\mathbb{E}[D^2]$ — an in-sample skew loading of ≈ 0.27, so a signed-only regression would recover only about a quarter of a true even effect — while the joint design identifies $\boldsymbol{\beta}^O$ directly, as noted in the text.)

Lemma (odd ⇒ zero mean). If g is odd ($g(-d) = -g(d)$),

$$\mathbb{E}[g(D)] = \int_{-\infty}^{\infty} g(d)f(d)\,dd = \int_{0}^{\infty} g(d)f(d)\,dd + \int_{-\infty}^{0} g(d)f(d)\,dd = 0.$$

Proof. In the second integral let $u = -d$: $\int_{-\infty}^{0} g(d)f(d)\,dd = \int_{0}^{\infty} g(-u)f(-u)\,(-du) = -\int_{0}^{\infty} g(-u)f(u)\,du = \int_{0}^{\infty}[-g(u)]f(u)\,du$ (oddness + symmetry) $= -\int_{0}^{\infty} g(u)f(u)\,du$. Sum with the first integral to get 0.

Example. Take $g(d) = d \cdot |d| = \text{sign}(d)d^2$, which is odd. By the lemma,

$$\mathbb{E}[D \cdot |D|] = \mathbb{E}[\text{sign}(D)D^2] = 0.$$

Implication for Case (2). With true model $y' = \beta\, s|D| + u$, $\mathbb{E}[u \mid s, D] = 0$, $E[s] = 0$, $s \perp D$,

$$\text{Cov}(sD,\ y') = \beta\ \text{Cov}(sD,\ s|D|) = \beta\{\mathbb{E}[s^2 D{\cdot}|D|] - \mathbb{E}[sD]\,\mathbb{E}[s|D|]\} = \beta\ \mathbb{E}[s^2]\ \underbrace{\mathbb{E}[D \cdot |D|]}_{0} = 0$$

Hence the OLS coefficient on sD is 0 in population.

Implication for Case (3). Let $D^+ = \max(D, 0)$, $D^- = \min(D, 0) \le 0$. Then $|D| = D^+ - D^-$ and $D^+D^- = 0$ a.s. Regress y' on $X_1 = sD^+$ and $X_2 = sD^-$. Orthogonality: $\text{Cov}(X_1, X_2) = E[s^2 D^+ D^-] = 0 \Rightarrow (X'X)$ is diagonal. Since $y' = \beta s(D^+ - D^-) + u$,

$$\beta^+ = \frac{\text{Cov}(X_1, y')}{\text{Var}(X_1)} = \beta, \qquad \beta^- = \frac{\text{Cov}(X_2, y')}{\text{Var}(X_2)} = -\beta.$$

Thus plotting β^+ and $-\beta^-$ reproduces the $|D|$ specification.

Intuition. The economic effect depends on the *magnitude* of disagreement (even in D). Using the signed interaction (sD) mixes equal-and-opposite contributions and cancels under symmetry; splitting by sign (or using $|D|$) prevents this cancellation.

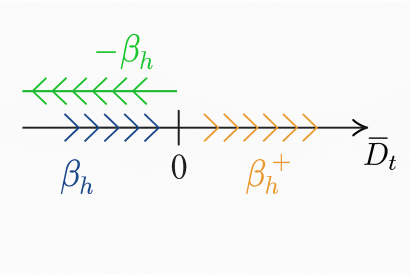


Figure D32: Graphical illustration of the interpretation logic of the estimated coefficients β_g

Economic Intuitions for the different specification

- **Magnitude channel.** Transmission depends on the *size* of Fed–market disagreement, not its sign. Empirically $\alpha_h < 0$ and $\beta_h > 0$: larger $|D_t|$ attenuates the impact of the same monetary shock s_t (interpretation: weaker coordination/credibility when forecasts diverge).
- **Why the signed interaction fails.** Replacing $|D_t|$ with D_t estimates the coefficient on $s_t D_t$. The attenuation effect is *even* in D_t, so positive and negative disagreements contribute with opposite signs and cancel. Under symmetry of D_t around zero and $s_t \perp D_t$, $\text{Cov}(s_t D_t,\, s_t|D_t|) = E[s_t^2 D_t \cdot |D_t|] = 0$, yielding a near-zero OLS coefficient.
- **Why splitting by sign works.** Writing $|D_t| = D_t^+ - D_t^-$ and regressing on $s_t D_t^+$ and $s_t D_t^-$ recovers the magnitude channel. Because $D_t^+ D_t^- = 0$ a.s., the regressors are orthogonal and OLS returns $\beta_h^+ = \beta_h$ and $\beta_h^- = -\beta_h$; plotting β_h^+ and $-\beta_h^-$ matches the $|D_t|$ specification.
- **When the signed interaction would matter.** A nonzero coefficient on $s_t D_t$ signals asymmetry (e.g., disagreement hurts more on one side).

D.16 Decomposing Attenuation: a GDP Accounting Perspective

This subsection asks whether the attenuation of monetary transmission operates through specific GDP components or proportionally across them. I map local-projection responses for consumption, investment, government purchases, and net exports into output using the log-linearized accounting identity:

$$Y_t = C_t + I_t + G_t + NX_t, \qquad \hat{y}_t \approx a_c \hat{c}_t + a_i \hat{i}_t + a_g \hat{g}_t + a_{nx} \widehat{nx}_t, \quad a_x = \bar{x}/\bar{y}.$$

Shares are set to sample averages $a_c = 0.66$, $a_i = 0.16$, $a_g = 0.20$, $a_{nx} = -0.02$. I estimate LPs for each component in the same system used for output, recover the baseline response α_h (shock effect at $|\bar{D}_t| = 0$) and the state-dependence term β_h (marginal effect per unit of disagreement), and compute contributions $a_x \hat{x}_h$ at $h = 12$ for illustration. The aggregate GDP response is:

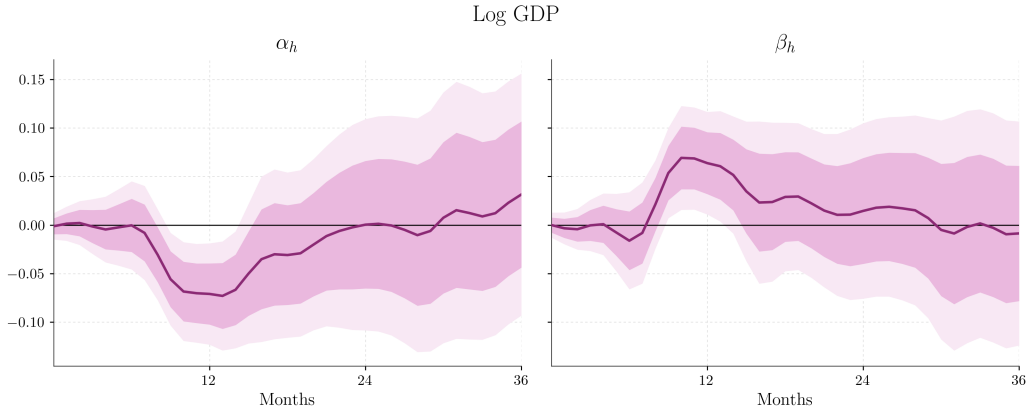


Figure D33: Aggregate GDP Response to the baseline Local Projection featuring disagreement.

Expectedly matching the evidence for industrial production. Thus, I proceed in the national accounting identity decomposition.

At the one year horizon, the baseline contraction is well accounted for by C and I; government spending is negligible and $-NX$ offsets part of the decline:

$$\alpha_{12}: \quad \hat{y}_{12} \simeq -0.06 \;\approx\; 0.66(-0.07)+0.16(-0.25)+0.20(0)+(-0.02)(-1.5) = -0.056,$$

implying approximate shares $C \sim 83\%$, $I \sim 71\%$, $G \sim 0\%$, and $-NX \sim -54\%$.

Moving to positive absolute disagreement, the attenuation term can be similarly decomposed, and it appears originating again primarily from C and I; G remains small and $-NX$ offsets:

$$\beta_{12}: \quad \hat{y}_{12} \simeq 0.05 \;\approx\; 0.66(0.05) + 0.16(0.20) + 0.20(0.01) + (-0.02)(1) = 0.047,$$

with relative shares $C \sim 69\%$, $I \sim 66\%$, $G \sim 5\%$, and $-NX \sim -40\%$.

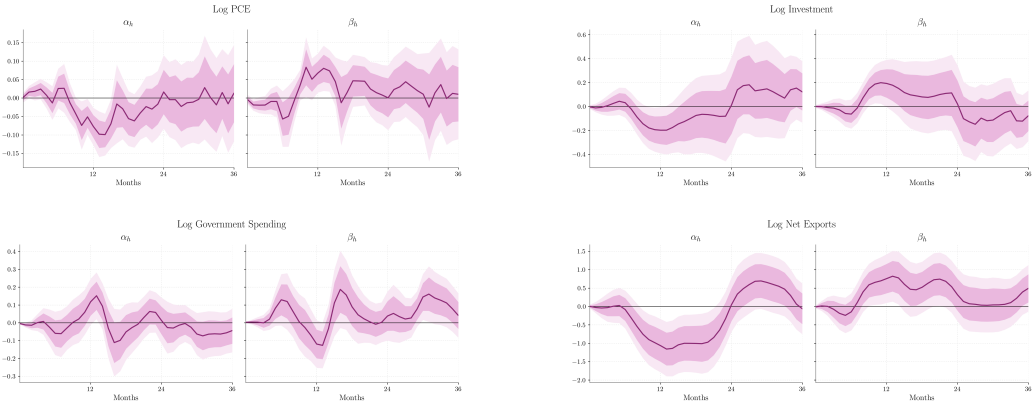


Figure D34: Top left: Consumption IRFs; Top right: Investment IRFs; Bottom left: Government Expenditure; Bottom right: Net Imports.

Two conclusions follow. First, most of the output response—both baseline and attenuation—runs through household demand and investment. Second, the pattern of contributions is broadly proportional across states, indicating that disagreement dampens transmission fairly homogeneously across GDP components rather than reshuffling composition.

D.17 Fed-Market Horse Races

This appendix details variations in the tests proposed by Hoesch et al. (2023) and Giacomini and Rossi (2010).

D.17.1 Information Advantage Tests: Variations

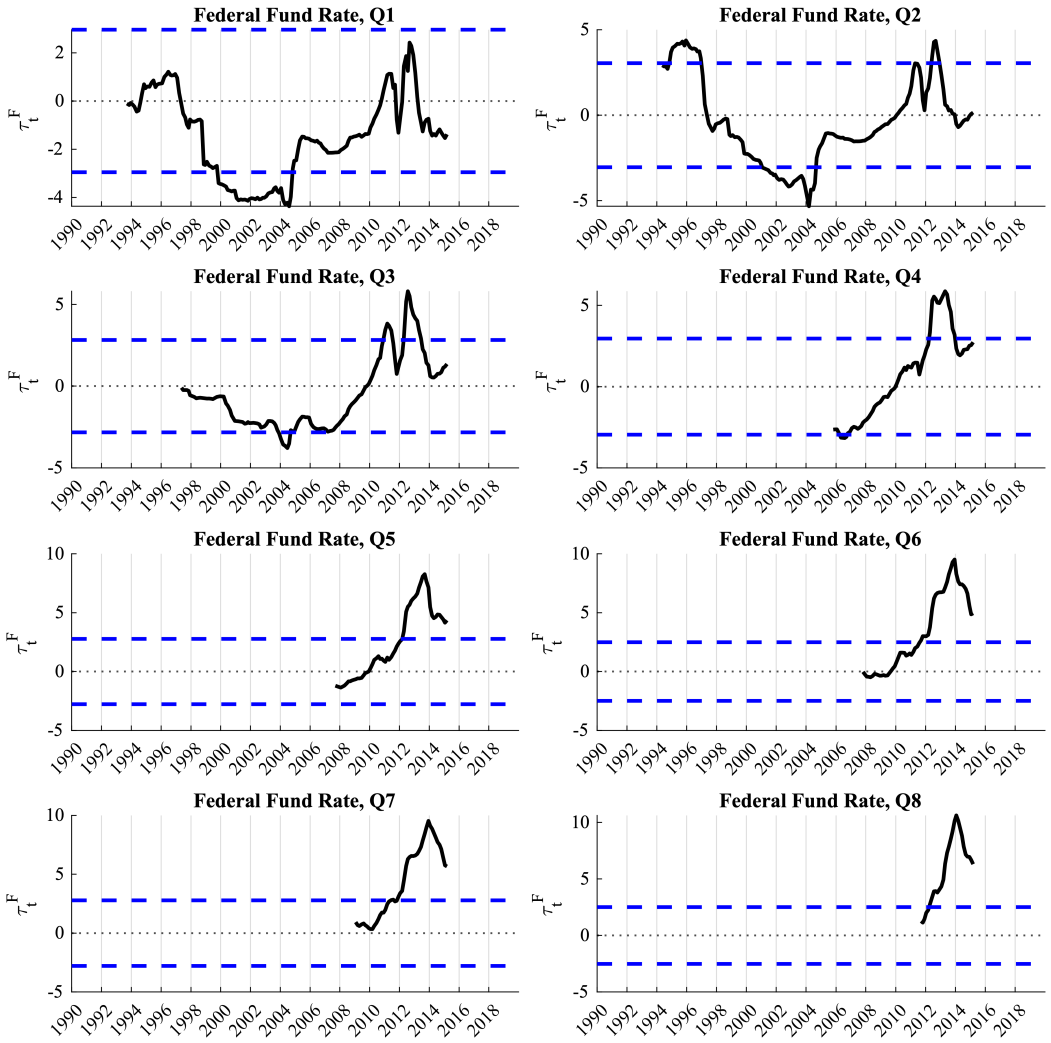


Figure D35: Information advantage—FFR. Rolling t-statistics τ_t^F from (4), estimated in $m = 60$–meeting windows (dated at the midpoint). Dashed red lines are the two-sided 5% Rossi and Sekhposyan (2016) critical values. HAC s.e. are Newey–West with bandwidth $m^{1/4}$. Market forecasts are extracted from Fed Funds Futures aligned to the Tealbook cutoff.

D.17.2 Accuracy Tests: Variations

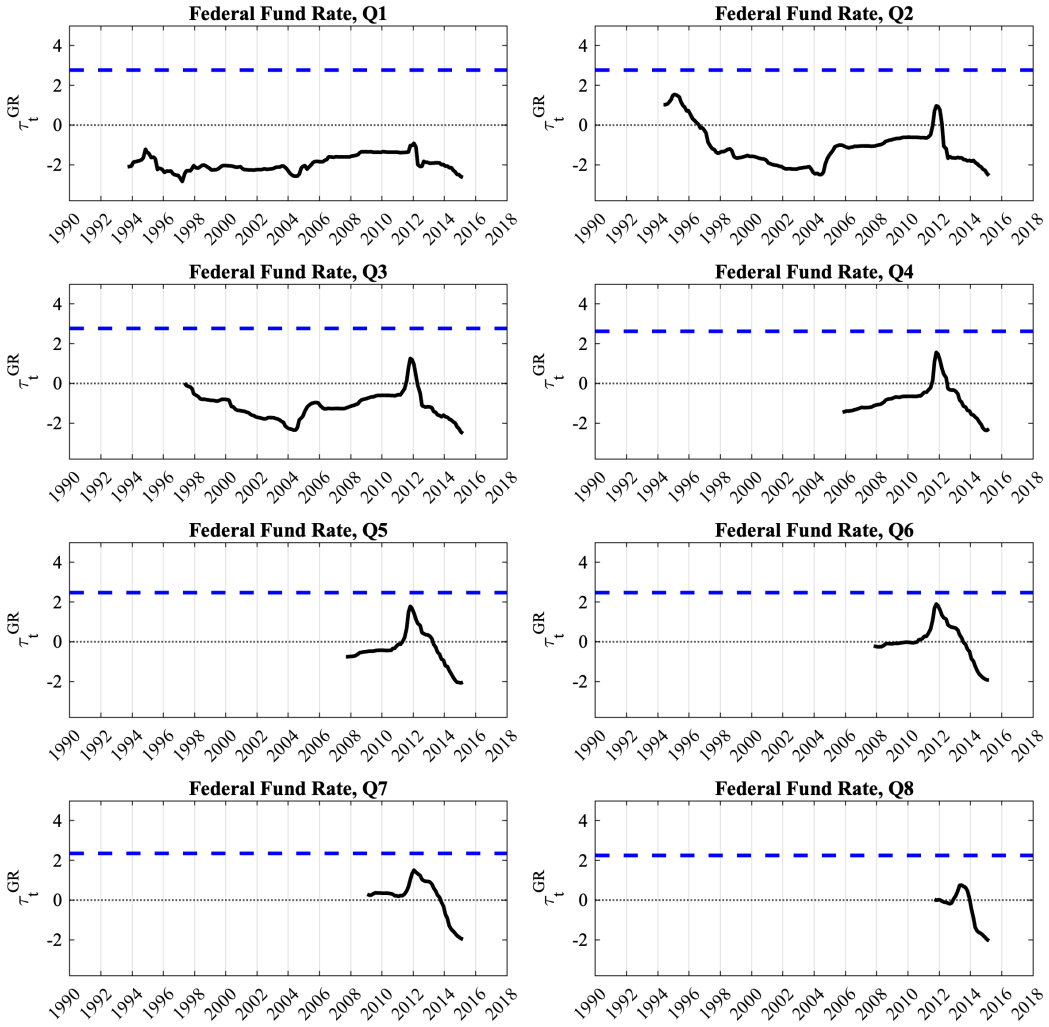


Figure D36: Forecast accuracy—FFR ($h = 1{:}8$). Rolling $\tau_{t,h}^{\mathrm{GR}}$ from regression 8 (positive $\approx$ Fed more accurate), estimated in $m = 60$–meeting windows and dated at the midpoint. Dashed red lines mark the one-sided 5% Giacomini and Rossi (2010) critical values. HAC s.e.: Newey–West with bandwidth $m^{1/4}$. Market forecasts are extracted from Fed Funds Futures aligned to the Tealbook cutoff.

E Model and Auxiliary Results

E.1 $P = \mathbb{E}$: From Market Prices to Expectations

Fix date t and a one–period futures contract on i_{t+1} with price P_t. There is a unit mass of traders $i \in [0,1]$, each with subjective belief $\mathbb{E}_t^i[i_{t+1}]$ about the payoff. Let $\xi_{i,t}$ denote trader i's position in the contract (positive = long). The contract is in zero net supply:

$$\int_0^1 \xi_{i,t}\, di = 0. \tag{E.1}$$

Assumption (linear demand). For any price P_t, trader i submits a demand schedule

$$\xi_{i,t}(P_t) = \beta_{i,t}\big(\mathbb{E}_t^i[i_{t+1}] - P_t\big), \qquad \beta_{i,t} > 0, \tag{E.2}$$

where $\beta_{i,t}$ summarizes i's risk tolerance, wealth, and trading intensity. This reduced form is implied, for example, by CARA utility and Gaussian beliefs, or more generally by mean–variance preferences.

Aggregation. Market clearing (E.1) and (E.2) imply

$$0 = \int_0^1 \beta_{i,t}\big(\mathbb{E}_t^i[i_{t+1}] - P_t\big)\, di = \int_0^1 \beta_{i,t}\, \mathbb{E}_t^i[i_{t+1}]\, di - P_t \int_0^1 \beta_{i,t}\, di.$$

Solving for the equilibrium futures price,

$$P_t = \frac{\displaystyle\int_0^1 \beta_{i,t}\, \mathbb{E}_t^i[i_{t+1}]\, di}{\displaystyle\int_0^1 \beta_{i,t}\, di}. \tag{E.3}$$

Hence P_t equals a risk–tolerance–weighted average of subjective expectations. If $\beta_{i,t}$ is identical across traders, then (E.3) simplifies to

$$P_t = \int_0^1 \mathbb{E}_t^i[i_{t+1}]\, di,$$

so the futures price coincides with the cross–sectional average belief, and, when all condition on the same information set, $P_t = \mathbb{E}_t^M[i_{t+1}]$.

E.2 Augmented Euler Equation under Inflation Disagreement

Maintain the assumptions in the text that $r_t^n = \rho$ is known and constant and that the Euler equation is formed under Market expectations. Iterating forward gives

$$c_t \;=\; -\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big[\,\mathbb{E}_t^M(i_{t+h}) - \mathbb{E}_t^M(\pi_{t+h+1}) - \rho\Big]. \tag{E.4}$$

Allow for disagreement on both the nominal rate and inflation. Define

$$D^i_{t+h|t} \equiv \mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^M(i_{t+h}), \qquad D^\pi_{t+h+1|t} \equiv \mathbb{E}_t^F(\pi_{t+h+1}) - \mathbb{E}_t^M(\pi_{t+h+1}).$$

Add and subtract $\mathbb{E}_t^F(i_{t+h})$ and $\mathbb{E}_t^F(\pi_{t+h+1})$ term by term:

$$c_t = -\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big[\,\mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^F(\pi_{t+h+1}) - \rho\Big] + \frac{1}{\sigma}\sum_{h=0}^{\infty}D^i_{t+h|t} - \frac{1}{\sigma}\sum_{h=0}^{\infty}D^\pi_{t+h+1|t}. \tag{E.5}$$

Let

$$c_t^{\textsc{Agreement}} \;\equiv\; -\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big[\,\mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^F(\pi_{t+h+1}) - \rho\Big],$$

and define the wedges

$$W_t^i \;\equiv\; \frac{1}{\sigma}\sum_{h=0}^{\infty}D^i_{t+h|t}, \qquad W_t^\pi \;\equiv\; -\frac{1}{\sigma}\sum_{h=0}^{\infty}D^\pi_{t+h+1|t}.$$

Then (E.5) is

$$c_t \;=\; c_t^{\textsc{Agreement}} \;+\; W_t^i \;+\; W_t^\pi. \tag{E.6}$$

Interpretation W_t^i is the rate-disagreement wedge in the text. W_t^π is an inflation-disagreement wedge with opposite sign to D^π because inflation lowers the real rate in the Euler equation. If $\mathbb{E}_t^F(\pi_{t+h+1}) > \mathbb{E}_t^M(\pi_{t+h+1})$ ($D^\pi > 0$), then $W_t^\pi < 0$ and consumption is lower than under agreement; the reverse holds if the Fed expects less inflation than the Market.

E.3 Model Results

E.3.1 Derivation of the disagreement-augmented IS

This appendix derives equation (16) from primitives. The household Euler in market beliefs is the starting point; the wedge W_t enters as a discounted sum of future disagreements. Define the agreement-equivalent allocation $\widetilde{c}_t \equiv c_t - W_t$.

$$c_t = -\frac{1}{\sigma}\Big[\,i_t - \mathbb{E}_t^M(\pi_{t+1}) - \rho\Big] + \mathbb{E}_t^M(c_{t+1}) \qquad \text{Market's Euler}$$

$$= -\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big\{\ \mathbb{E}_t^M(i_{t+h}) - \mathbb{E}_t^M(\pi_{t+h+1}) - \rho\ \Big\} + \underbrace{\lim_{j\to\infty}\mathbb{E}_t^M(c_{t+j})}_{=0} \qquad \text{iterating forward}$$

$$= -\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big\{\ \mathbb{E}_t^F(i_{t+h}) - D_{t+h|t} - \mathbb{E}_t^M(\pi_{t+h+1}) - \rho\ \Big\} \qquad \text{using } D_{t+h|t} \text{ def.}$$

$$= \underbrace{-\frac{1}{\sigma}\sum_{h=0}^{\infty}\Big\{\ \mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^F(\pi_{t+h+1}) - \rho\ \Big\}}_{\equiv\,\widetilde{c}_t} + \underbrace{\frac{1}{\sigma}\sum_{h=1}^{\infty} D_{t+h|t}}_{=W_t} \qquad \text{impos. } \mathbb{E}_t^M(\pi) = \mathbb{E}_t^F(\pi)$$

$$= -\frac{1}{\sigma}\Big[\ i_t - \mathbb{E}_t^F(\pi_{t+1}) - \rho\ \Big] + \mathbb{E}_t^F\left(\widetilde{c}_{t+1}\right) \ +\ W_t \qquad \text{LIE under } \mathbb{E}^F$$

$$= -\frac{1}{\sigma}\Big[\ i_t - \mathbb{E}_t^F(\pi_{t+1}) - \rho\ \Big] + \mathbb{E}_t^F(c_{t+1}) - \mathbb{E}_t^F(W_{t+1}) + W_t \qquad \text{using } \widetilde{c}_{t+1} = c_{t+1} - W_{t+1}$$

The last line is the disagreement-augmented Euler in one-period form.
Imposing equilibrium $c_t = y_t$ and substituting the realized rule $i_t = \phi_\pi\ \mathbb{E}_t^F(\pi_t) + u_t^S$ delivers (16). Here ϕ_π is the realized coefficient, which is generally defined and can or cannot coincide with the Fed's belief. This convention is the static counterpart of the EKF state $\mathbf{\Phi}_t$ in Section 5, where $\mathbf{\Phi}_t$ is the (unobserved) objective rule coefficient and each agent's posterior $\mathbf{\Phi}_t^j$ is an estimate of it.

E.3.2 Impulse Responses with General Rule

Below, I prove the general form of the impulse responses in the 3-equation New Keynesian system featuring disagreement through the expectational wedge. The proof holds identically for the less general case illustrated in the text, leading to equations (17). Similarly, I report below the case for a persistent monetary shock $u_t^S = \rho_s^i u_{t-1}^S + \varepsilon_t^S$, but analogous conclusions follow for an MIT shock. The same assumption of homogeneous inflation expectations between Fed and markets hold, but it is immediate to generalize to inflation disagreement (see E.2).

$$y_t = -\frac{1}{\sigma}\Big[\ i_t - \mathbb{E}_t^F(\pi_{t+1}) - r_t^n\ \Big] + \mathbb{E}_t^F(y_{t+1}) + \frac{1}{\sigma}\sum_{k=0}^{\infty} D_{t+k+1} \qquad r_{t+1}^n = 0$$
$$\pi_t = \beta\ \mathbb{E}_t^M(\pi_{t+1}) + \kappa\, y_t$$

Replacing with the interest rate rule $i_t = \phi_\pi\ \mathbb{E}_t^F(\pi_t) + \phi_y\ \mathbb{E}_t^F(y_t) + u_t^S$

$$y_t = -\frac{1}{\sigma}\big[\phi_\pi\ \mathbb{E}_t^F(\pi_t) + \phi_y\ \mathbb{E}_t^F(y_t) + u_t^S - \mathbb{E}_t^M(\pi_{t+1})\big] + \mathbb{E}_t^F(y_{t+1}) + W_t$$
$$\pi_t = \beta\ \mathbb{E}_t^M(\pi_{t+1}) + \kappa\, y_t$$

Conjecture $y_t = \Psi_y u_t^S, \pi_t = \Psi_\pi u_t^S, W_t = \Psi_w u_t^S$:

$$\begin{aligned}\Psi_y u_t^S &= -\sigma^{-1}\left[\phi_\pi \Psi_\pi u_t^S + \phi_y \Psi_y u_t^S + u_t^S - \Psi_\pi \rho_s^M u_t^S\right] + \Psi_y \rho_s^F u_t^S \; - \; \rho_s^F \Psi_w u_t^S + \Psi_w u_t^S \\ \Psi_\pi u_t^S &= \beta \rho_s^M \, \Psi_\pi u_t^S + \kappa \, \Psi_y u_t^S.\end{aligned} \tag{E.7}$$

For output, we conclude

$$\Psi_y = \frac{(1-\beta\rho_s^M)(\sigma(1-\rho_s^F)\Psi_w - 1)}{\sigma(1-\beta\rho_s^M)(1-\rho_s^F - \sigma^{-1}\phi_y) + \kappa(\phi_\pi - \rho_s^F)} \tag{E.8}$$

Equation (E.8) resembles (17), with the denominator featuring an additional (expected) response to output term. The closed-form solutions are detailed in E.3.3.

E.3.3 Equilibrium under Incomplete Information

Section 4 introduces a structural interpretation of the empirical findings of attenuation starting from the smallest deviation from commonly used full information New Keynesian frameworks. In this exercise, the objective is not to pinpoint or discuss the specific source of disagreement, – i.e., the type of deviation from full information –, but to demonstrate the more general notion that *any* source of disagreement can rationalize the observed attenuation via a testable model-implied relationship between the (impulse) response of output and that of beliefs (disagreement) – as argued in equation (18). Notice that this does *not* imply that the solution under different disagreement sources is the same – it is not.

To illustrate this point, Section 4.2 postulated a simple deviation from full information: Fed and Markets hold different beliefs on the policy coefficient governing the reaction to inflation (ϕ_π). Below, E.3.4 solves for the equilibrium (fixed point) under this conjecture, while E.3.5 displays the same exercise under the different hypothesis of different shock's perceptions. Section 5 will tie everything together in a generalized state-space where disagreement arises from multiple sources.

E.3.4 Disagreement on the Rule Coefficients

The shock is known and perceived homogeneously: $u_t^S = \rho_s u_{t-1}^S + \varepsilon_t^S$, but the markets has its own expectation of the policy rule coefficient, denoted by ϕ_π^M (I abstract from ϕ_y^M for illustration).

$$\begin{aligned}y_t &= -\frac{1}{\sigma}\left[\phi_\pi \, \mathbb{E}_t^F(\pi_t) + u_t^S - \mathbb{E}_t^F(\pi_{t+1})\right] + \mathbb{E}_t^F(y_{t+1}) \; - \; \mathbb{E}_t^F(W_{t+1}) + W_t, \\ \pi_t &= \beta \, \mathbb{E}_t^M(\pi_{t+1}) + \kappa \, y_t\end{aligned}$$

and $W_t = \frac{1}{\sigma}\sum_{h=1}^{\infty} D_{t+h|t} = \frac{1}{\sigma}\sum_{h=1}^{\infty}\left[\mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^M(i_{t+h})\right]$.

Set conjectures to $y_t = \Psi_y u_t^S, \pi_t = \Psi_\pi u_t^S, W_t = \Psi_w u_t^S$. Again, adjusting the notation, the reduced-form results in the paper (equation 18) hold:

$$\Psi_\pi = \frac{\kappa}{1-\beta\rho_s}\,\Psi_y, \qquad \Psi_y = \frac{(\sigma(1-\rho_s)\Psi_w - 1)(1-\beta\rho_s)}{(1-\beta\rho_s)(1-\rho_s)\sigma + \kappa(\phi_\pi - \rho_s)} \tag{E.9}$$

I use the W_t definition to retrieve Ψ_w:

$$\begin{aligned}\Psi_w u_t^S &= \frac{1}{\sigma}\sum_{h=1}^{\infty}\Big[\phi_\pi^F\ \mathbb{E}_t^F(\pi_{t+h}) + \mathbb{E}_t^F(u_{t+h}^S) - \phi_\pi^M\ \mathbb{E}_t^M(\pi_{t+h}) - \mathbb{E}_t^M(u_{t+h}^S)\Big] \\ &= \frac{1}{\sigma}(\phi_\pi^F - \phi_\pi^M)\sum_{h=1}^{\infty}\mathbb{E}_t^{\cdot}(\pi_{t+h}) = \frac{(\phi_\pi^F - \phi_\pi^M)\rho_s\,\Psi_\pi}{\sigma(1-\rho_s)}\,u_t^S.\end{aligned}$$

Hence

$$\Psi_w = \frac{(\phi_\pi^F - \phi_\pi^M)\rho_s}{\sigma(1-\rho_s)}\,\Psi_\pi.$$

Now I use the 3 equations to solve in closed-form for the 3 coefficients. Define

$$\mathcal{A} \equiv \frac{\kappa}{1-\beta\rho_s}, \qquad \mathcal{C} \equiv \frac{(\phi_\pi^F - \phi_\pi^M)\rho_s}{\sigma(1-\rho_s)}.$$

Algebraic manipulation delivers:

$$\begin{aligned}\Psi_y &= \frac{-\frac{1}{\sigma}}{(1-\rho_s)(1-\mathcal{C}\,\mathcal{A}) + \frac{\mathcal{A}}{\sigma}(\phi_\pi - \rho_s)}, \\ \Psi_\pi &= \mathcal{A}\,\Psi_y = \frac{\mathcal{A}(-\frac{1}{\sigma})}{(1-\rho_s)(1-\mathcal{C}\,\mathcal{A}) + \frac{\mathcal{A}}{\sigma}(\phi_\pi - \rho_s)}, \\ \Psi_w &= \mathcal{C}\,\Psi_\pi = \mathcal{C}\,\frac{\mathcal{A}(-\frac{1}{\sigma})}{(1-\rho_s)(1-\mathcal{C}\,\mathcal{A}) + \frac{\mathcal{A}}{\sigma}(\phi_\pi - \rho_s)}.\end{aligned} \tag{E.10}$$

which are verified substituting them into the equilibrium equations.

E.3.5 Disagreement on the Shock's Persistence

The shock is perceived heterogeneously: $u_t^S = \rho_s^i u_{t-1}^S + \varepsilon_t^S$, with $i = F, M$. We start from the same (simplified) NK system as in the body:

$$\begin{aligned}y_t &= -\frac{1}{\sigma}\big[\phi_\pi\,\mathbb{E}_t^F(\pi_t) + u_t^S - \mathbb{E}_t^F(\pi_{t+1})\big] + \mathbb{E}_t^F(y_{t+1}) - \mathbb{E}_t^F(W_{t+1}) + W_t, \\ \pi_t &= \beta\,\mathbb{E}_t^M(\pi_{t+1}) + \kappa\,y_t\end{aligned}$$

and $W_t = \frac{1}{\sigma}\sum_{h=1}^{\infty} D_{t+h|t} = \frac{1}{\sigma}\sum_{h=1}^{\infty}\big[\mathbb{E}_t^F(i_{t+h}) - \mathbb{E}_t^M(i_{t+h})\big]$. Set conjectures to $y_t = \Psi_y u_t^S, \pi_t = \Psi_\pi u_t^S, W_t = \Psi_w u_t^S$ and denote $r_F \equiv \rho_s^F, r_M \equiv \rho_s^M$. Adjusting the

notation, the reduced-form results in the paper (equation 17) hold:

$$\Psi_\pi = \frac{\kappa}{1-\beta r_M}\,\Psi_y, \qquad \Psi_y = \frac{(\sigma(1-r_F)\Psi_w - 1)(1-\beta r_M)}{(1-\beta r_M)(1-r_F)\sigma + \kappa(\phi_\pi - r_F)} \tag{E.11}$$

and I use the W_t definition to retrieve Ψ_w:

$$\begin{aligned}
\Psi_w\, u_t^M &= \frac{1}{\sigma}\sum_{h=1}^{\infty}\Big[\phi_\pi\,\mathbb{E}_t^F(\pi_{t+h}) + \mathbb{E}_t^F(u_{t+h}^M) - \phi_\pi\,\mathbb{E}_t^M(\pi_{t+h}) - \mathbb{E}_t^M(u_{t+h}^M)\Big] \\
\sigma\,\Psi_w\, u_t^s &= u_t^s\,(\phi_\pi\Psi_\pi + 1)\left[\sum_{n=1}^{\infty}\left(r_F^n - r_M^n\right)\right] = u_t^s\,(\phi_\pi\Psi_\pi + 1)\,\frac{r_F - r_M}{(1-r_F)(1-r_M)} \\
\Psi_w &= \frac{(r_F - r_M)\,(\phi_\pi\Psi_\pi + 1)}{\sigma\,(1-r_F)\,(1-r_M)}
\end{aligned} \tag{E.12}$$

Now I use the 3 equations to solve in closed-form for the 3 coefficients. Define

$$\mathcal{A} \equiv \frac{\kappa}{1-\beta r_M}, \qquad \mathcal{B} \equiv \frac{r_F - r_M}{\sigma(1-r_F)(1-r_M)}.$$

Algebraic manipulation delivers:

$$\begin{aligned}
\Psi_y &= \frac{-\frac{1}{\sigma} + (1-r_F)\mathcal{B}}{(1-r_F)(1-\mathcal{B}\,\phi_\pi\mathcal{A}) + \frac{\mathcal{A}}{\sigma}(\phi_\pi - r_F)} \\
\Psi_\pi &= \mathcal{A}\,\Psi_y = \frac{\mathcal{A}\big(-\frac{1}{\sigma} + (1-r_F)\mathcal{B}\big)}{(1-r_F)(1-\mathcal{B}\,\phi_\pi\mathcal{A}) + \frac{\mathcal{A}}{\sigma}(\phi_\pi - r_F)} \\
\Psi_w &= \mathcal{B}\big(1 + \phi_\pi\Psi_\pi\big) = \mathcal{B}\left(1 + \phi_\pi\,\frac{\mathcal{A}\big(-\frac{1}{\sigma} + (1-r_F)\mathcal{B}\big)}{(1-r_F)(1-\mathcal{B}\,\phi_\pi\mathcal{A}) + \frac{\mathcal{A}}{\sigma}(\phi_\pi - r_F)}\right)
\end{aligned} \tag{E.13}$$

which are verified substituting them into the equilibrium equations.

E.4 Determinacy: Blanchard-Kahn & Taylor Principle

This appendix derives the determinacy restriction for the baseline (log-linear) NK block used in the main text. In the model environment, market expectations are formed under $\mathbb{E}_t^M(\cdot)$. Since the baseline NK block has two jump variables (x_t, π_t) and no predetermined state, a unique bounded rational-expectations equilibrium requires two unstable roots in the Blanchard and Kahn (1980) sense.

Baseline NK block.

$$x_t = \mathbb{E}_t^M(x_{t+1}) - \frac{1}{\sigma}\Big(i_t - \mathbb{E}_t^M(\pi_{t+1}) - r_t^n\Big), \tag{E.14}$$

$$\pi_t = \beta\mathbb{E}_t^M(\pi_{t+1}) + \kappa x_t, \tag{E.15}$$

$$i_t = \phi_\pi \pi_t + \phi_x x_t. \tag{E.16}$$

Assume $\beta \in (0,1)$, $\sigma > 0$, $\kappa > 0$, and set $r_t^n \equiv 0$.

Matrix representation. Substitute (E.16) into (E.14) and collect terms:

$$\Big(1 + \frac{\phi_x}{\sigma}\Big) x_t + \frac{\phi_\pi}{\sigma} \pi_t = \mathbb{E}_t^M(x_{t+1}) + \frac{1}{\sigma} \mathbb{E}_t^M(\pi_{t+1}). \tag{E.17}$$

Let $y_t \equiv \begin{bmatrix} x_t \\ \pi_t \end{bmatrix}$. Stack (E.17)–(E.15):

$$L\, y_t \;=\; R\, \mathbb{E}_t^M(y_{t+1}), \qquad L \equiv \begin{bmatrix} 1 + \phi_x/\sigma & \phi_\pi/\sigma \\ -\kappa & 1 \end{bmatrix}, \quad R \equiv \begin{bmatrix} 1 & 1/\sigma \\ 0 & \beta \end{bmatrix}. \tag{E.18}$$

Since R is invertible,

$$\mathbb{E}_t^M(y_{t+1}) = M\, y_t, \qquad M \equiv R^{-1} L. \tag{E.19}$$

$$M = \begin{bmatrix} 1 + \phi_x/\sigma + \kappa/(\beta\sigma) & (\beta\phi_\pi - 1)/(\beta\sigma) \\ -\kappa/\beta & 1/\beta \end{bmatrix}. \tag{E.20}$$

BK determinacy. Because $y_t = (x_t, \pi_t)'$ are both jump variables, determinacy requires both eigenvalues of M to satisfy $|\lambda| > 1$. Equivalently, both eigenvalues μ of M^{-1} must satisfy $|\mu| < 1$. For a 2×2 matrix,

$$\mu^2 - \operatorname{tr}(M^{-1})\, \mu + \det(M^{-1}) = 0,$$

and stability of M^{-1} is equivalent to the three inequalities

$$\det(M) - \operatorname{tr}(M) + 1 > 0, \qquad \det(M) + \operatorname{tr}(M) + 1 > 0, \qquad \det(M) > 1.$$

Under $\beta \in (0,1)$, $\sigma, \kappa > 0$, and $\phi_\pi, \phi_x \geq 0$, the last two always hold because $\operatorname{tr}(M) > 0$ and $\det(M) = \frac{\kappa\phi_\pi + \phi_x + \sigma}{\beta\sigma} \geq \frac{1}{\beta} > 1$. Thus the knife-edge restriction is $\det(M) - \operatorname{tr}(M) + 1 > 0$.

Generalized Taylor principle. From (E.20),

$$\operatorname{tr}(M) = \Big(1 + \frac{\phi_x}{\sigma} + \frac{\kappa}{\beta\sigma}\Big) + \frac{1}{\beta} = \frac{\beta\phi_x + \sigma(\beta + 1) + \kappa}{\beta\sigma}, \tag{E.21}$$

$$\det(M) = \frac{\kappa\phi_\pi + \phi_x + \sigma}{\beta\sigma}. \tag{E.22}$$

$$\det(M) - \operatorname{tr}(M) + 1 = \frac{\kappa(\phi_\pi - 1) + (1 - \beta)\phi_x}{\beta\sigma}, \tag{E.23}$$

so determinacy is equivalent to

$$\kappa(\phi_\pi - 1) + (1 - \beta)\phi_x > 0 \qquad \Longleftrightarrow \qquad \phi_\pi + \frac{1-\beta}{\kappa}\phi_x > 1. \tag{E.24}$$

Condition (E.24) is the generalized Taylor principle for the baseline NK block. Economically, it says that when inflation or the output gap rises, the Taylor rule increases the nominal rate strongly enough that the *expected real* rate also rises: higher ϕ_π and/or ϕ_x make current demand less attractive. The "aggressive" real-rate response is what places both eigenvalues of M outside the unit circle, so that BK delivers a unique bounded equilibrium. In this sense, the Taylor principle is a parametric restriction on policy-rule coefficients that makes the BK condition hold by pushing the system's eigenvalues into the right configuration.

E.5 Learning

This extension of the baseline model details the evolution of beliefs about the rule coefficients and the underlying state as a function of the flow of information at every period. Below, I cover some of the complications emerging from the nonlinear nature of the policy rule.

Why one Kalman filter isn't enough The FOMC sets the short-term rate as

$$i_t \;=\; \mathbf{\Phi}_t'\mathbf{x}_t \;+\; u_t^G \;+\; u_t^S, \qquad u_t^G \;\sim \mathcal{N}(0, \sigma^2_{u^G}),\; u_t^S \sim \mathcal{N}(0, \sigma^2_{u^S}),$$

where $\mathbf{x}_t$ (the fundamentals) and $\mathbf{\Phi}_t$ (the policy-rule coefficients) are *both* unobserved. Because i_t is the product of two unknown vectors, the observation equation is bilinear and the classical linear-Gaussian Kalman filter does not return the exact Bayesian posterior.

Routes: 2-step filter vs. EKF A possible way to tackle the nonlinearity is with a 2-step practical approximation. Assume

1. *Independence*: the shocks driving $\mathbf{x}_t$ and $\mathbf{\Phi}_t$ are mutually independent and independent of u_t^S;
2. *Slow drift*: $\mathbf{\Phi}_t = \mathbf{\Phi}_{t-1} + v_t,\; v_t \sim \mathcal{N}(0, W),\; W$ small.

Under (A1)–(A2) the cross-covariances created by the bilinear term are $O(W)$; neglecting them yields two almost-decoupled recursions. Below, we present scalar versions for expositional clarity, assuming the information structure from the main text; the extension to vector $\mathbf{x}_t$ with conformable matrix $P_t^{\mathbf{x}}$ is straightforward.

Step 1: Kalman filter for the fast state x_t

Assuming $x_t = \rho\, x_{t-1} + w_t$ with $w_t \sim \mathcal{N}(0, \sigma_w^2)$:

$$\hat{x}_{t|t-1} = \rho\, \hat{x}_{t-1}, \tag{E.25}$$

$$P^x_{t|t-1} = \rho^2 P^x_{t-1} + \sigma_w^2, \tag{E.26}$$

$$K^x_t = \frac{P^x_{t|t-1}}{P^x_{t|t-1} + \sigma^2_{\varepsilon^S}}, \tag{E.27}$$

$$\hat{x}_t = \hat{x}_{t|t-1} + K^x_t\big(y^M_t - \hat{x}_{t|t-1}\big), \tag{E.28}$$

$$P^x_t = (1 - K^x_t) P^x_{t|t-1}. \tag{E.29}$$

Step 2: Recursive-least-squares (RLS) update for slow coefficients $\mathbf{\Phi}_t$

Let $\sigma_u^2 = \sigma^2_{u^S}$ denote the variance of the shock unobserved by agents when updating beliefs about $\mathbf{\Phi}_t$.

$$\mathrm{MPS}_t = i_t - \hat{u}^G_t - \hat{\mathbf{\Phi}}'_{t-1}\hat{x}_t, \tag{E.30}$$

$$P^\Phi_{t|t-1} = P^\Phi_{t-1} + W, \tag{E.31}$$

$$K^\Phi_t = \frac{P^\Phi_{t|t-1}\hat{x}_t}{\sigma_u^2 + \hat{x}_t' P^\Phi_{t|t-1}\hat{x}_t}, \tag{E.32}$$

$$\hat{\mathbf{\Phi}}_t = \hat{\mathbf{\Phi}}_{t-1} + K^\Phi_t\, \mathrm{MPS}_t, \tag{E.33}$$

$$P^\Phi_t = (\mathbf{I} - K^\Phi_t \hat{x}_t') P^\Phi_{t|t-1}. \tag{E.34}$$

If W is chosen small enough (policy parameters drift slowly), the error made by ignoring the cross-covariance is $O(W)$. Note that using the filtered estimate $\hat{x}_t$ in place of the true x_t introduces an additional errors-in-variables approximation whose magnitude scales with the posterior variance P^x_t.

Alternatively, linearizing around the predicted mean each period gives the *Extended Kalman filter* (EKF), which is optimal among estimators using first-order (linear) approximations to the nonlinear observation function, but technically more involved. I develop this technique in detail below.

E.6 Extended Kalman Filter

Throughout this appendix I set $n = 1$ for clarity of exposition: this can be thought of as the policy rule being only function of inflation or the output gap. For $n > 1$, replace scalar entries in each matrix below by blocks and by transposes in each vector.

Timing. At $\underline{t}$ the Market observes a private signal $y^M_{\underline{t}}$ about x_t. At t, the FOMC simultaneously releases guidance g_t and sets the policy rate i_t. We process

$y^M_{\underline{t}}$ first; at time t we process g_t (linear) and then i_t (nonlinear) to stabilize the EKF linearization point.

State and Prior. The state is

$$z_t = \begin{bmatrix} x_t \\ \phi_t \\ u^G_t \end{bmatrix} \tag{E.35}$$

with

$$\begin{aligned} x_t &= \rho\, x_{t-1} + \eta_t, \\ \phi_t &= \phi_{t-1} + \omega_t, \\ u^G_t &= \rho_G\, u^G_{t-1} + \zeta_t \end{aligned} \qquad \varepsilon_t \equiv \begin{bmatrix} \eta_t \\ \omega_t \\ \zeta_t \end{bmatrix} \sim \mathcal{N}(0, Q) \tag{E.36}$$

Then $z_t = Az_{t-1} + \varepsilon_t, A = \mathrm{diag}(\rho,\, 1,\, \rho_G), \varepsilon_t \sim \mathcal{N}(0, Q)$, with $Q = \mathrm{diag}(Q_x, Q_\phi, Q_G)$ in the baseline; $Q \in \mathbb{S}^3_+$ is the general case.[69] Given information Ω_{t-1}, the prediction (prior) is $\hat{z}_{t|t-1},\ P_{t|t-1} := \mathrm{Var}(z_t \mid \Omega_{t-1})$.

$$\hat{z}_{t|t-1} = A\, \hat{z}_{t-1|t-1}, \tag{E.37}$$

$$P_{t|t-1} = A\, P_{t-1|t-1}\, A' + Q \tag{E.38}$$

$$= \begin{bmatrix} \rho^2 P^{xx}_{t-1|t-1} + Q_x & \rho\, P^{x\phi}_{t-1|t-1} & \rho\rho_G\, P^{xu}_{t-1|t-1} \\ \rho\, P^{\phi x}_{t-1|t-1} & P^{\phi\phi}_{t-1|t-1} + Q_\phi & \rho_G\, P^{\phi u}_{t-1|t-1} \\ \rho\rho_G\, P^{ux}_{t-1|t-1} & \rho_G\, P^{u\phi}_{t-1|t-1} & \rho^2_G P^{uu}_{t-1|t-1} + Q_G \end{bmatrix} \tag{E.39}$$

Measurement Equations. All noises are mean–zero, mutually independent, and independent of the state (conditionally Gaussian with stated variances):

$$\begin{aligned} y^M_{\underline{t}} &= H_y z_t + \varepsilon^y_{\underline{t}}, & H_y &= \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}, & \varepsilon^y_{\underline{t}} &\sim \mathcal{N}(0, R_x), \\ g_t &= H_g z_t + \nu_t, & H_g &= \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}, & \nu_t &\sim \mathcal{N}(0, R_{g,\underline{t}}), \\ i_t &= h(z_t) + u^S_t, & h(z_t) &= \phi_t x_t + u^G_t, & u^S_t &\sim \mathcal{N}(0, R_i). \end{aligned}$$

E.6.1 Update A ($\underline{t}$): $y^M_{\underline{t}}$ (Linear, Exact)

Innovation and its Variance.

$$\tilde{\mu}_x \equiv y^M_{\underline{t}} - H_y \hat{z}_{t|t-1} = y^M_{\underline{t}} - \hat{x}_{t|t-1}, \tag{E.40}$$

$$S_y = H_y P_{t|t-1} H'_y + R_x = P^{xx}_{t|t-1} + R_x \tag{E.41}$$

[69] I impose independent transition shocks; the transition itself is diagonal so it does not introduce cross-covariances by construction.

Gain, Posterior Mean and Variance.

$$K_y = P_{\underline{t}|t-1} H_y' S_y^{-1} = \frac{1}{S_y} \begin{bmatrix} P^{xx}_{\underline{t}|t-1} \\ P^{\phi x}_{\underline{t}|t-1} \\ P^{ux}_{\underline{t}|t-1} \end{bmatrix} \tag{E.42}$$

$$\hat{z}^y_{\underline{t}|\underline{t}} = \hat{z}_{\underline{t}|t-1} + K_y \tilde{\mu}_y, \tag{E.43}$$

$$P^y_{\underline{t}|\underline{t}} = P_{\underline{t}|t-1} - K_y S_y K_y' \tag{E.44}$$

E.6.2 Update B (t): g_t (Linear, Exact)

Innovation and its Variance.

$$\tilde{\mu}_g \equiv g_t - H_g \hat{z}^y_{\underline{t}|\underline{t}} = g_t - \hat{u}^{G,y}_{\underline{t}|\underline{t}} \tag{E.45}$$

$$S_g = H_g P^y_{\underline{t}|\underline{t}} H_g' + R_g = P^{y,uu}_{\underline{t}|\underline{t}} + R_{g,\underline{t}} \tag{E.46}$$

Notice that if current guidance supersedes past guidance, $\tilde{\mu}_g \equiv g_t$.

Gain, Posterior Mean and Variance.

$$K_g = P^y_{\underline{t}|\underline{t}} H_g' S_g^{-1} = \frac{1}{S_g} \begin{bmatrix} P^{xu}_{\underline{t}|\underline{t}} \\ P^{\phi u}_{\underline{t}|\underline{t}} \\ P^{uu}_{\underline{t}|\underline{t}} \end{bmatrix} \tag{E.47}$$

$$\hat{z}^{(y,g)}_{t|t} = \hat{z}^y_{t|\underline{t}} + K_g \tilde{\mu}_g \tag{E.48}$$

$$P^{(y,g)}_{t|t} = P^y_{t|\underline{t}} - K_g S_g K_g' \tag{E.49}$$

Because both are linear–Gaussian with independent noises, processing $y^M_{\underline{t}}$ then g_t or vice versa yields the same $(\hat{z}^{(y,g)}_{t|t}, P^{(y,g)}_{t|t})$. Notice that Update B addresses the general case where u^G_t has some persistence ρ_G, while in the body new signals completely supersede old announcements.

E.6.3 Update C (t): i_t (Nonlinear, Extended Kalman Filter)

The rate measurement is nonlinear: $i_t = h(z_t) + u^S_t$ with $h(z_t) = \phi_t x_t + u^G_t$. Kalman updates require a linear measurement. The EKF replaces h by its (linear) first–order Taylor expansion around the current "pre-posterior" mean $\hat{z}^{(y,g)}_{t|t}$; the result is an innovation linear in the state deviation $z_t - \hat{z}^{(y,g)}_{t|t}$ plus Gaussian noise.

Linearization. Let $(\bar{z}, \bar{P}) \equiv (\hat{z}_{t|t}^{(y,g)}, P_{t|t}^{(y,g)})$ with $\bar{z} = \begin{bmatrix} \bar{x} & \bar{\phi} & \bar{u}^G \end{bmatrix}'$, collecting all "pre-posteriors" before the interest rate update. Using $h(z_t) = \phi_t x_t + u_t^G$, the Jacobian (row) of h w.r.t. z_t is

$$\frac{\partial h(z_t)}{\partial z_t'} = \begin{bmatrix} \frac{\partial h(z_t)}{\partial x_t} & \frac{\partial h(z_t)}{\partial \phi_t} & \frac{\partial h(z_t)}{\partial u_t^G} \end{bmatrix} = \begin{bmatrix} \phi_t & x_t & 1 \end{bmatrix} \tag{E.50}$$

I linearize at the pre-posterior, $\bar{z} = \begin{bmatrix} \hat{x}_{t|t}^{(y,g)}, \hat{\phi}_{t|t}^{(y,g)}, \hat{u}_{t|t}^{G,(y,g)} \end{bmatrix}'$. Then

$$H_i(\bar{z}) = \left. \frac{\partial h(z_t)}{\partial z_t'} \right|_{z_t = \bar{z}} = \begin{bmatrix} \bar{\phi} & \bar{x} & 1 \end{bmatrix} \tag{E.51}$$

Applying a first–order Taylor approximation:

$$h(z_t) \approx h(\bar{z}) + H_i(\bar{z})(z_t - \bar{z}) \tag{E.52}$$

representing the linear approximation around the pre-posterior. Then, we can resort to the familiar "innovation plus noise" updating routine.

Innovation (Surprise) and its Variance. The innovation represented by i_t is, as usual, the difference between the realization and its prior expectation (E.52), i.e. the linear approximation of $h(z)$ around the pre-posterior:

$$\tilde{\mu}_i \equiv i_t - h(\bar{z}) = H_i(\bar{z})(z_t - \bar{z}) + u_t^S \tag{E.53}$$

$$S_i = H_i \, \bar{P} \, H_i' + R_i. \tag{E.54}$$

Notice that $\tilde{\mu}_i$ is the *monetary policy surprise* implied by the state-space, and S_i its variance.[70]

EKF Gain, Posterior Mean and Variance.

$$K_i = \bar{P} \, H_i' S_i^{-1} = \frac{1}{S_i} \begin{bmatrix} \bar{P}^{xx} \, \bar{\phi} + \bar{P}^{x\phi} \, \bar{x} + \bar{P}^{xu} \\ \bar{P}^{\phi x} \, \bar{\phi} + \bar{P}^{\phi\phi} \, \bar{x} + \bar{P}^{\phi u} \\ \bar{P}^{ux} \, \bar{\phi} + \bar{P}^{u\phi} \, \bar{x} + \bar{P}^{uu} \end{bmatrix} \tag{E.55}$$

$$\hat{z}_{t|t} = \bar{z} + K_i \, \tilde{\mu}_i, \tag{E.56}$$

$$P_{t|t} = \bar{P} - K_i S_i K_i' \tag{E.57}$$

Credibility update. After i_t is realized, the Market updates its perceived noise variances according to the Fed's credibility, proxied by its (last) forecast error.

[70] The explicit form of S_i is:

$$S_i = \bar{\phi}^2 \, \bar{P}^{xx} + \bar{x}^2 \, \bar{P}^{\phi\phi} + \bar{P}^{uu} + 2\bar{\phi} \, \bar{x} \, \bar{P}^{x\phi} + 2\bar{\phi} \, \bar{P}^{xu} + 2\bar{x} \, \bar{P}^{\phi u} + R_i.$$

The driver is a *rate* forecast error, so it is the natural updater of the rate-signal noise R_i — the operative credibility object in the estimated two-state model:

$$\begin{aligned}&\mathrm{FE}_t^F \equiv i_t - \mathbb{E}_{\underline{t}}^F(i_t),\\ &R_{i,t} = (1-r)\,R_{i,\underline{t}} + r\,(\mathrm{FE}_t^F)^2, \qquad r \in (0,1), \qquad R_{i,\underline{t+1}} := R_{i,t}.\end{aligned}$$

The guidance-signal credibility $R_{g,t}(= \sigma^2_{\nu,t})$ follows the same recursion, governing the separate guidance gain.

E.7 Wedge Dynamics

E.7.1 Belief Updates and Wedge Revisions

The EKF splits a rate surprise $\tilde{\mu}_t^i$ into belief updates $\Delta x_t = K^x_{i,t}\tilde{\mu}_t^i$ and $\Delta\phi_t = K^\phi_{i,t}\tilde{\mu}_t^i$ using the estimated prior covariance and the local Jacobian $H_{i,t} = [\bar{\phi}_t, \bar{x}_t]$. Mapping these belief changes into expected rates is mechanical: for $i_t = \phi_t x_t$,

$$\Delta W_t \propto \bar{\phi}_t\,\Delta x_t + \bar{x}_t\,\Delta\phi_t = \bar{\phi}_t K^x_{i,t}\tilde{\mu}_t^i + \bar{x}_t K^\phi_{i,t}\tilde{\mu}_t^i.$$

Hence, two distinct notions emerge:

(i). belief attribution, answering the question "which belief moved more?" and measured by

$$\kappa_x = K^x_{i,t}/(K^x_{i,t} + K^\phi_{i,t}), \qquad \kappa_\phi = 1 - \kappa_x$$

(ii). wedge attribution, answering the question "which belief update moved expected rates more?" and measured by:

$$\mathcal{C}_t^x = \frac{\bar{\phi}_t K^x_{i,t}}{\bar{\phi}_t K^x_{i,t} + \bar{x}_t K^\phi_{i,t}}, \qquad \mathcal{C}_t^\phi = 1 - \mathcal{C}_t^x,$$

The difference arises solely from the local leverages $(\bar{\phi}_t, \bar{x}_t)$, which represent a map of how each belief update impacts into the wedge dynamics. Two examples follow to illustrate the conceptual significance of this distinction.

Example A. Let $\tilde{\mu}_t^i = 1$, $(\bar{x}_t, \bar{\phi}_t) = (0, 1.5)$, and gains $(K^x_{i,t}, K^\phi_{i,t}) = (0.05, 0.20)$. Then

$$\Delta x_t = 0.05, \quad \Delta\phi_t = 0.20, \quad \kappa_\phi = 0.80,$$

but

$$\Delta W_t \propto 1.5(0.05) + 0(0.20) = 0.075, \quad \mathcal{C}_t^\phi = 0.$$

The economy is at a neutral point with $\bar{x}_t \approx 0$, so revisions to the rule have no leverage on expected rates. Only news about fundamentals shifts the wedge via

$\bar{\phi}_t \, \Delta x_t$. Even if beliefs move mostly in ϕ, propagation is entirely via x because $\bar{x}_t \simeq 0$.

Example B. Let $\tilde{\mu}_t^i = 1$, $(\bar{x}_t, \bar{\phi}_t) = (1, 1.5)$, and gains $(K_{i,t}^x, K_{i,t}^\phi) = (0.05, 0.20)$. Then

$$\Delta x_t = 0.05, \quad \Delta \phi_t = 0.20, \quad \kappa_\phi = 0.80,$$

and

$$\Delta W_t \propto 1.5(0.05) + 1(0.20) = 0.275, \quad \mathcal{C}_t^\phi = \frac{0.20}{0.275} \approx 0.73.$$

The state is nonzero, $\bar{x}_t = 1$, so a change in the rule immediately maps into expected rates. Both attribution and propagation are therefore ϕ-driven at this operating point. Now both beliefs and propagation are ϕ–dominated because $\bar{x}_t$ provides leverage to $\Delta \phi_t$.

In summary, large $\mathcal{C}_t^x$ or $\mathcal{C}_t^\phi$ need not reflect which belief moved most; they reflect how the belief changes translate into expected rates at the prevailing $(\bar{x}_t, \bar{\phi}_t)$.

E.8 3-dimensional State-Space: Estimation

E.9 Market Updating Regimes [UNDER RESTRUCTURING]

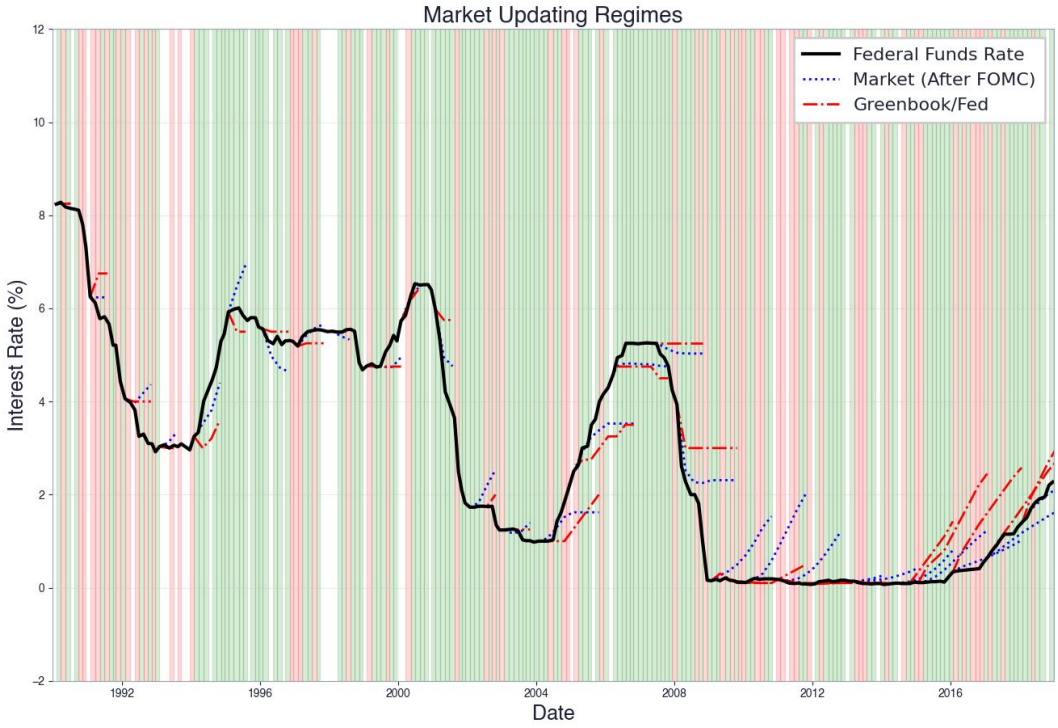


Figure E1: Market Updating Regime heuristic: shaded green areas correspond to meetings when market updates in the same direction as the shock (Bayesian); shaded red areas correspond to meetings when market updates in the opposite direction as the shock (Non-Bayesian); unshaded areas correspond to updates close to zero.

Figure E1 provides a heuristic overview of the market updating behavior over the full length of the sample. The vertical shaded areas correspond to the direction of the average shift of the available Fed Funds Future prices around each of the $N = 358$ monetary policy decisions in the sample. In green, are represented those FOMC announcements when the market revised its expectations of future policy in the same direction as the sign of the baseline shock series (`MPS_ORT`, Amodeo (2025)); in red, the complement case; in white, the statistically null shifts. Notice, for instance, two clusters of episodes: first, December 1993 - December 1996 represents the longest continuous streak of green ("Bayesian") updates. According to the model this would correspond to an unmitigated, effective conduct of monetary policy; unsurprisingly, this interval coincides with what Alan Blinder (Blinder (2023)), former vice chairman of the Federal Reserve, defined the "*perfect soft landing that helped make Alan Greenspan a central banking legend.*"
Second, December 2007 - July 2011 comprises the GFC and the high and persis-

tent disagreement that followed it, and it comprises many FOMC meetings where markets updated in the opposite direction of what suggested by the observed high-frequency shocks ("Non-Bayesian"); again, as it was argued before in Figure 7, this period corresponded to a controversial era for monetary efficacy that eventually led to a prolonged period of constrained policy at the zero lower bound, and is compatible with a version of the model where the wedge increases in response to the shock (right panel of Figure 15).

E.10 Euler Wedges [UNDER RESTRUCTURING]

Another way of testing the model's prediction comes directly from the study of the Euler Equation. This exercise uses the derived wedge to test whether monetary surprises raise the wedge on impact, i.e., $\text{Corr}(s_t, W_t) > 0$. In fact, under the simplifying assumption that inflation expectations are aligned, it was shown that:

$$c_t \;=\; c_t^{\text{Agreement}} \;+\; W_t$$

implying

$$W_t \;=\; c_t + \frac{1}{\sigma}\sum_{k=0}^{\infty}\Big\{\mathbb{E}_t^F(i_{t+k}) - \mathbb{E}_t^F(\pi_{t+k+1}) - r_{t+k}^N\Big\} \tag{E.58}$$

Note that all variables on the right-hand side are observed or estimated. In fact, the Tealbooks provide not only interest rates predictions, but a whole range of expectations, including on inflation, $\mathbb{E}_t^F(\pi)$; r_t^N, on the other hand, is among the most heavily researched object in macroeconomics, and a number of influential estimates are routinely used.

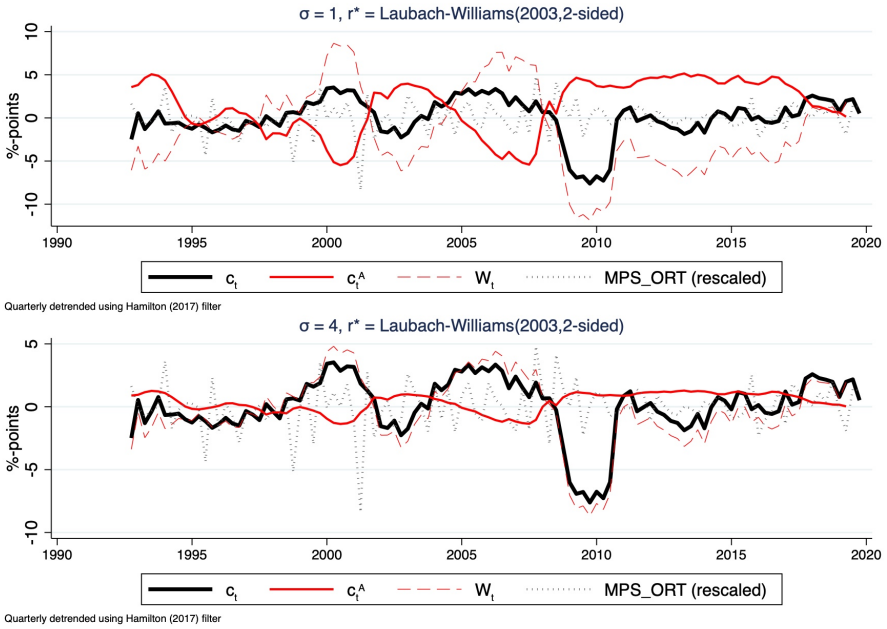
Operationally, I measure c_t using the Personal Consumption Expenditure index and extract its monthly cycle via the Hamilton (1994) filter. I use several of the natural rate of interest estimates available from the Atlanta Fed website,[71] including the widely used Holston et al. (2023) and Lubik and Matthes (2011).

While the unconditional correlation is very close to zero and leaning on the negative side for most shock series, when conditioning on positive responses of industrial production to monetary shocks, I find the expected positive correlation. Across all high-frequency surprise series, the correlation with $W_t\big|_{\Psi_y>0}$ increases. The correlation is economically small but consistently positive for benchmark high-frequency shocks.

[71] Available at this link.

Statistic	High-frequency surprise series s_t				
	$\mathbf{s}^{\perp}_{BS,A}$	$\mathbf{s}_{BS}$	$\mathbf{s}_{GK}$	$\mathbf{s}_{JK}$	$\mathbf{s}_{J}$
$\mathrm{Corr}(s_t, W_t)$	-0.02	0.02	-0.03	0.03	-0.16
N	(107)	(107)	(79)	(107)	(107)
$\mathrm{Corr}\big(s_t, W_t \mid \Psi_y > 0\big)$	0.14	0.07	0.42	0.08	0.20
N	(18)	(18)	(5)	(18)	(18)

Table E1: Pearson correlations between surprises s_t and wedge W_t.



It is well documented that standard consumption Euler equations fit the data poorly (see Hall (1978), Flavin (1981), and Campbell and Mankiw (1989)). Given the deliberately parsimonious specification, a sizable gap between c_t and $c_t^{\text{AGREEMENT}}$—hence a large role for W_t—is unsurprising and not our focus. The test is narrower: we examine whether monetary surprises raise the wedge on impact, i.e., $\mathrm{Corr}(s_t, W_t) > 0$, evaluating this both unconditionally and conditional on expansionary output responses, without taking a stand on the overall fit of the Euler equation.

F A Closer Look at Ψ_w

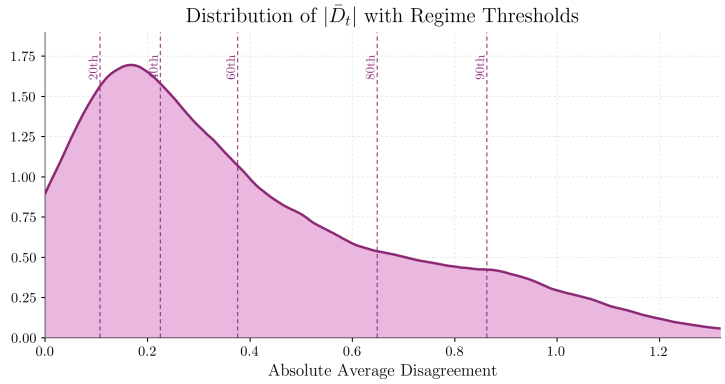


Figure F1: Histogram of interacting variable $|D_t|$ with percentile markers.

Below, several alternative specifications follow:

$$\text{1, INCREMENTAL:} \quad W_{t+h} = c_h + \beta_h^L \, s_t + \beta_h^H \, s_t \mathbf{1}_{\{|\bar{D}_t| \geq d\}} + \boldsymbol{\theta}_h \text{LAGS}_{t-1}^L + \varepsilon_{t+h} \quad \text{(F.1)}$$

with $\text{LAGS}_{t-1}^L \equiv \left[W_{t-1:t-L}, \; s_{t-1:t-L} \right]'$ and $\boldsymbol{\theta}_h \equiv [\, \phi_h^{1:L}, \; \gamma_h^{1:L}]$.

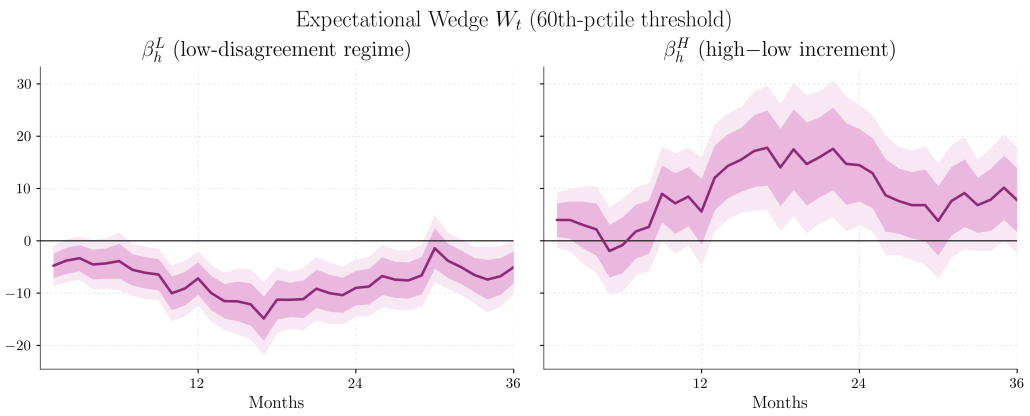


Figure F2: IRFs of $W_t = \sum_{h=1}^{8} D_{t+h}$ | `min(3 quarters)` to a Orthogonalized Monetary Policy Surprise for low/high disagreement (60^{th} percentile threshold).

$$\text{2, BINARY:} \quad W_{t+h} = c_h + \beta_h^{\text{L}} \, s_t \mathbf{1}_{\{|\bar{D}_t| < d\}} + \beta_h^{\text{H}} \, s_t \mathbf{1}_{\{|\bar{D}_t| \geq d\}} + \boldsymbol{\theta}_h \text{LAGS}_{t-1}^L + \varepsilon_{t+h} \quad \text{(F.2)}$$

with $\text{LAGS}_{t-1}^L \equiv \left[W_{t-1:t-L}, \; s_{t-1:t-L} \right]'$ and $\boldsymbol{\theta}_h \equiv [\, \phi_h^{1:L}, \; \gamma_h^{1:L}]$.

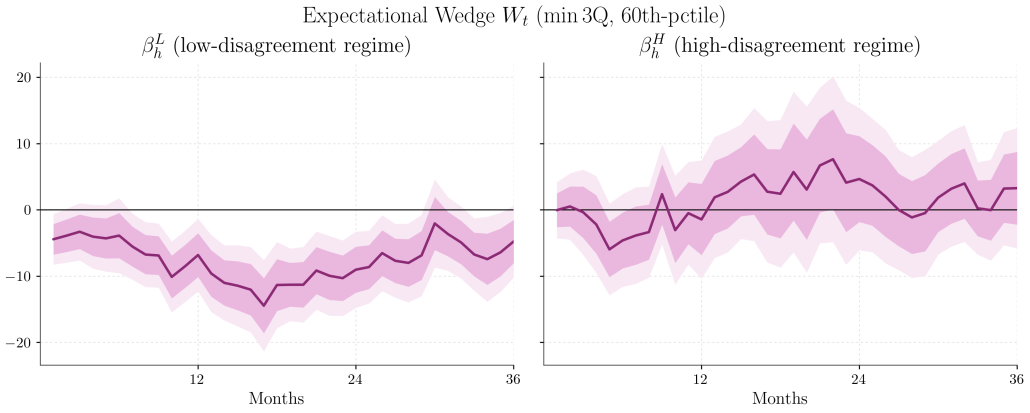


Figure F3: IRF of $W_t = \sum_{h=1}^{8} D_{t+h}$ | `min(3 quarters)` to a Orthogonalized Monetary Policy Surprise for low/high disagreement (60^{th} percentile threshold).



Figure F4: IRF of $W_t = \sum_{h=1}^{8} D_{t+h}$ | `min(8 quarters)` to a Orthogonalized Monetary Policy Surprise for low/high disagreement (60^{th} percentile threshold).

More Thresholds

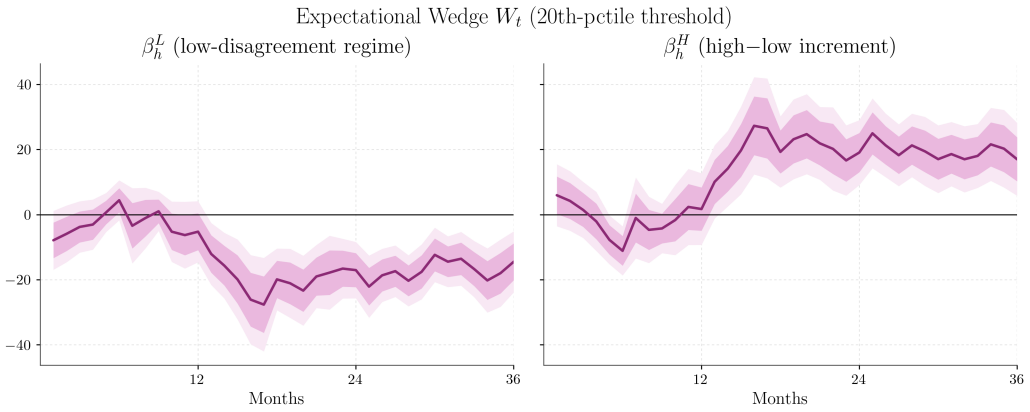


Figure F5: IRF of W_t to a Orthogonalized Monetary Policy Surprise for low/high disagreement (20^{th} percentile threashold).

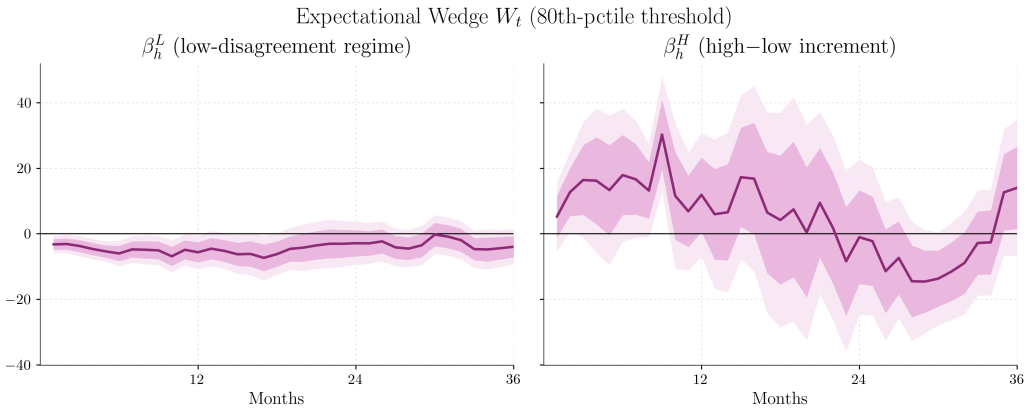


Figure F6: IRF of W_t to a Orthogonalized Monetary Policy Surprise for low/high disagreement (80^{th} percentile threashold).

Unconditional Wedge Response

The unconditional response of the wedge to a monetary policy shock captures an average response across different disagreement states, and can be studied by estimating the following linear projections:

$$W_{t+h} = \mathbf{c}_h + \boldsymbol{\gamma}_h\, s_t + \boldsymbol{\theta}_h \textsc{Lags}_{t-1}^L + \boldsymbol{\varepsilon}_{t+h} \tag{F.3}$$

with $\textsc{Lags}_{t-1}^L \equiv \left[W_{t-1:t-L},\ s_{t-1:t-L} \right]'$ and $\boldsymbol{\theta}_h \equiv [\, \phi_h^{1:L},\ \gamma_h^{1:L}]$.

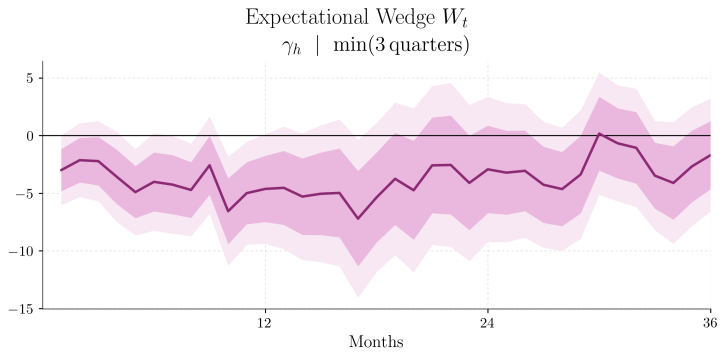


Figure F7: $W_t = \sum_{h=1}^{8} D_{t+h}$ | `min(3 quarters)` to a Orthogonalized Monetary Policy Surprise.

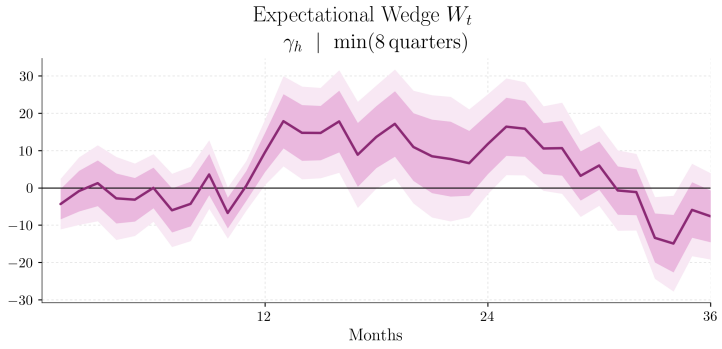


Figure F8: $W_t = \sum_{h=1}^{8} D_{t+h}$ | `min(8 quarters)` to a Orthogonalized Monetary Policy Surprise.

G Testing $\mathbb{E}_t^M(i_{t+h})$'s FOMC-elasticity

I assess whether FOMC announcements induce a systematic revision of the market's expected rate path within 24 hours. Let $\Delta_{24h}\,\mathbb{E}_t^M(i_{t+h}) \equiv \mathbb{E}_t^M(i_{t+h}) - \mathbb{E}_{\underline{t}}^M(i_{t+h})$. First, I test $H_0 : \mathbb{E}[\Delta_{24h}\mathbf{E}_t^M] = \mathbf{0}$ with Hotelling's T^2 on the $(H+1)$-dimensional revision vector; all $p > 0.05$, hence I fail to reject no joint mean shift. Second, I can estimate directly a Bayesian-updating restriction,

$$\mathbb{E}_t^M(i_{t+h}) = \mathbb{E}_{\underline{t}}^M(i_{t+h}) + \lambda_t^{\text{FOMC}}\big(i_t - \mathbb{E}_{\underline{t}}^M(i_t)\big),$$

by regressing the stacked revisions on the monetary policy surprise MPS_t:

$$\Delta_{24h}\,\mathbb{E}_t^M = \beta_0 + \lambda^{\text{FOMC}}\text{MPS}_t + \varepsilon_t.$$

Analogously, the update weight is statistically indistinguishable from zero. Taken together, the tests imply an unconditional slope near zero: within 24 hours the average market revision is not systematically related to the surprise component.

Table G1: Hotelling's T^2 test for a joint mean shift in 24-hour revisions

H	N	T^2	p
8	90	7.98	0.5051
7	109	9.76	0.2499
6	120	3.05	0.8170
5	121	2.26	0.8220
4	133	2.45	0.6641
3	179	4.54	0.2175

Notes: $\Delta_{24h}\mathbf{E}_t^M$ stacks revisions for horizons $h = 0, \ldots, H$. N is the number of meetings with complete data. Null: $H_0 : \mathbb{E}[\Delta_{24h}\mathbf{E}_t^M] = \mathbf{0}$.

Table G2: Bayesian update test of 24-hour market revisions to FOMC surprises

	λ^{FOMC}	p
MPS_t	-0.0067	0.716
Intercept β_0	-0.0043	0.170

Notes: Estimates from $\Delta_{24h}\,\mathbb{E}_t^M = \beta_0 + \lambda^{\text{FOMC}}\text{MPS}_t + \varepsilon_t$, stacking $h = 0, \ldots, H$. p-values use robust standard errors.

Aside This null average is consistent with the learning model. In the EKF, each surprise is mapped into a meeting-specific update weight that depends on the linearization point and the prior covariances. As a result, positive and negative revision episodes offset in the unconditional sample, yielding a slope near zero.