Divided Views, Clearer Signals: FOMC Disagreement and Monetary Policy Uncertainty*

Patrik Räty, Sharath Sonti

November 4, 2026

Abstract

We document that changes in disagreement amongst the members of the Federal Open Market Committee (FOMC), as measured by dispersion in the dot plots, are negatively associated with changes in market-based monetary policy uncertainty around FOMC meetings. Our preferred explanation for this negative relationship comes from a Bayesian learning model in which the market learns about the committee balance of power between hawks and doves by observing the interest rate decision and the dot plot. The model has two countervailing forces: (1) more disagreement makes the learning process stronger, as when the views of hawks and doves are further apart the current rate decision better reveals whose views shape policy, decreasing monetary policy uncertainty, and (2) more disagreement, particularly when highly persistent, makes the future range of policy actions wider, increasing monetary policy uncertainty. Finally, we use heteroskedasticity-based identification exploiting variation in the Federal Reserve's communication strategy with the introduction of the dot plot to explore the effect of monetary policy uncertainty shocks. Our monetary policy uncertainty shocks load negatively on dot plot disagreement, depress investment consistent with a real-options channel, and do not seem to affect credit spreads.

Keywords: monetary policy, disagreement, uncertainty, dot plot

*We thank Jon Steinsson, Yuriy Gorodnichenko, Emi Nakamura, Pierre-Olivier Gourinchas, Rohan Kekre, Eben Lazarus, Peter Maxted, Tim McQuade, Chen Lian, Andreas Schaab, Zheng Liu, Johan Walden, Michael McMahon, Laura Waring, David Sraer, and Ulrike Malmendier, as well as seminar participants at the UC Berkeley Macro Colloquium, UC Berkeley Finance Lunch, and UC Berkeley Macro Café, for helpful comments and suggestions.
Correspondence: Patrik Räty (patrik_raty@berkeley.edu); Sharath Sonti (sharath_sonti@berkeley.edu).

1 Introduction

When monetary policymakers disagree, financial markets must assess what their differing views imply for future policy. The Federal Reserve's "dot plot" makes these differences visible by reporting individual policymakers' assessments of the appropriate policy rate at different future dates. Yet the value of this information is contested. Some worry that the dots can confuse markets by displaying policymakers' preferred rates without clarifying which path the committee is likely to follow (Slok, 2026). Others argue that revealing the range of views helps markets understand how the committee is weighing its policy options (Sahm, 2026). Given recent debate over the dot plot's role in Fed communication, a key empirical question is whether markets become more or less uncertain about future policy rates when the dots reveal greater disagreement.

This paper documents that increases in dot-plot disagreement within the Federal Open Market Committee (FOMC) are associated with declines in market uncertainty about future policy rates. We explain this relationship using a Bayesian learning model featuring two policymakers with different preferred interest rates, and the committee sets the policy rate as a weighted average of the two. Markets do not observe the weights but infer them by comparing the policymakers' prescriptions with the rate the committee implements. When these weights persist, greater disagreement makes the current rate decision more informative about whose views will shape future policy and can therefore reduce uncertainty about future rates.

We then study the macroeconomic consequences of increased monetary policy uncertainty by separating shocks to uncertainty about future policy rates from shocks to the current rate and its expected path. Our heteroskedasticity-based identification exploits shifts in the relative volatility of these shocks across communication regimes marked by the introduction of forward guidance and the dot plot. The identified uncertainty shocks load negatively on dot-plot disagreement and depress investment.

Our baseline analysis compares changes in disagreement across successive releases with changes in market uncertainty over a two-day window around each FOMC announcement. We measure disagreement using the variance of the two-year-ahead dots and market uncertainty using the option-implied variance of future short-term interest rates from Bauer et al. (2022). A one-standard-deviation larger increase in dot-plot variance is associated with a roughly 0.3-standard-deviation lower change in uncertainty about the short rate twelve months ahead. The relationship remains negative after accounting for seasonality, periods when the policy rate is at the zero lower bound, the median policy projection, monetary policy surprises, and the pre-announcement level of uncertainty, as well as across alternative

specifications and uncertainty horizons.

Markets may anticipate some of the changes in dot-plot dispersion between FOMC meetings. To distinguish these anticipated changes from unexpected disagreement, we use mock dot plots published by NatWest before FOMC meetings as a proxy for expected dispersion. These forecasts describe the distribution of dots, allowing us to measure disagreement relative to what was anticipated before the announcement. Measuring realized dispersion relative to these forecasts produces negative estimates of similar magnitude, although they are less precisely estimated in the smaller sample.

We explain this finding using a Bayesian learning model with two policymakers, a hawk and a dove, whose preferred interest rates follow Taylor rules with different coefficients. They share the same macroeconomic forecasts but differ in how those forecasts translate into appropriate policy. The committee sets the policy rate according to their relative influence, which the market does not directly observe.

The market learns about this relative influence by observing the dot plot and the committee's policy decision. When the hawk and the dove recommend similar rates, the decision provides little information about whose views carry more weight. When their recommendations are farther apart, the decision more clearly reveals whose views shape policy. If their relative influence persists over time, learning about it today also helps the market predict future decisions. Greater disagreement can therefore reduce uncertainty about future interest rates by making the current decision more informative.

This learning channel competes with a second force. When disagreement is persistent, wider differences in current policy views imply wider differences in future views. Uncertainty about whose preferences will prevail then translates into greater uncertainty about future interest rates. The model consequently admits either a positive or a negative relationship between disagreement and monetary policy uncertainty. A negative relationship arises when the information revealed about the committee's balance of influence outweighs the wider range of future policy choices. The model thus provides an explanation for our empirical finding while clarifying why disagreement need not reduce uncertainty in every setting. Appendix B presents theoretical extensions and examines the conditions under which the model's predictions continue to hold.

The relationship between committee disagreement and uncertainty about future policy rates raises a further question: how does this uncertainty affect economic activity? We focus on investment, where uncertainty can increase the value of postponing decisions that are costly to reverse. To estimate these effects, we use heteroskedasticity-based identification to separate uncertainty news from news about the current policy rate and the future path of rates. Our identification exploits changes in the relative variances of these news components

across periods distinguished by the introduction of forward guidance and the dot plot, while assuming that their effects on the observed financial variables remain stable across regimes. The resulting uncertainty shocks are negatively related to changes in dot-plot disagreement, in line with our earlier results. Local projections show that a one-standard-deviation quarterly uncertainty shock lowers real nonresidential fixed investment by about 1.5 percent within one year. We do not find clear evidence of a response in the excess bond premium. The decline in investment is consistent with this real-options channel.

Our analysis connects research on disagreement within monetary policy committees with research on monetary policy uncertainty and its economic effects. We document a negative relationship between committee disagreement and market uncertainty, and develop a learning mechanism that explains why greater disagreement can make policy more predictable. We also estimate the investment effects of monetary policy uncertainty shocks, connecting our evidence on dot-plot disagreement to the real economic consequences of uncertainty about future policy rates.

Section 2 reviews the related literature, and Section 3 describes the institutional setting and data. Section 4 presents the evidence on disagreement and market uncertainty. Section 5 develops a learning model to explain the negative relationship. Section 6 estimates the investment effects of monetary policy uncertainty shocks. Section 7 concludes.

2 Literature Review

Disagreement amongst the FOMC has often been measured using dissents (e.g., Chappell et al. (1993) and Belden (1989)), minutes, transcripts, and speeches (e.g., Bennani et al. (2018), Bordo and Istrefi (2023), Meade and Stasavage (2008), Apel et al. (2022)).

We instead use dispersion in the dot plots as our measure of disagreement amongst the FOMC. This is not a strictly better measure of disagreement for all purposes, rather it is a complementary measure that is useful for certain tasks. Dot plot dispersion is a quantitative, continuous, forward-looking measure covering all FOMC members which can be used to research real-time effects as it is released in conjunction with the policy decision. This means that there is a lesser concern of inferring information from other sources or economic conditions having changed as one would have with minutes- or transcript-based measures of FOMC disagreement as the minutes and transcripts are released with a lag.

Additionally, it likely has less strategic voting concerns relative to dissent-based or transcript-based measures. Hansen et al. (2018) show that when FOMC members learned that their deliberations would eventually become public, they exhibited conformism (particularly among more junior members). The dot plots are anonymous to the public forever and

hence such conformism is not to be expected in the dot plot or at least should be expected to a lesser degree.

Howes et al. (2026) find that formal dissents are rare and mask a substantial amount of disagreement amongst the FOMC as well as that dissent is costly to the dissenter. The dot plot is an anonymous (to the public) vehicle through which FOMC members reveal their assessments of appropriate monetary policy. This should not suffer, at least not as strongly, from strategic concerns and can highlight smaller disagreement amongst the FOMC. Tsang and Yang (2025) also find that formal votes fail to capture relevant dissent and construct a text-based hidden dissent measure, which is positively correlated with disagreement in the dot plot.

Substantial literature has focused on why FOMC members disagree. Potential reasons for disagreement include: (1) disagreement about policy goals, (2) disagreement about the economic outlook, (3) disagreement about the pass-through of monetary policy, and (4) strategic concerns (including career concerns, costs of dissenting, chair influence). Hansen et al. (2014) find that private assessments of the state of the economy and individual policy preferences are quantitatively important in explaining votes among the Monetary Policy Committee of the Bank of England. Foerster and Martinez (2023) find that disagreement about economic outlooks explain some of, but fails to explain most, of the disagreement in the dot plots. Howes et al. (2026) find that disagreement about the pass-through of monetary policy is an important driver of disagreement amongst the FOMC with those perceiving stronger output effects being more dovish and those perceiving stronger price effects being more hawkish. Malmendier et al. (2021) find that the personal inflation experiences of FOMC members influence their beliefs, voting behaviours, speeches, inflation forecasts, and policy stances. Madeira et al. (2026) find that in the Federal Reserve, which has a dual mandate, supply shocks generate disagreement as they create a genuine tradeoff for the dual mandate while demand shocks do not. In contrast, they find that supply shocks do not generate disagreement at the Bank of England, which only has a single mandate.

The above highlights several reasons for why disagreement about appropriate monetary policy amongst monetary policymakers exists. However, less work has been able to research what the effects disagreement has. We attempt to tackle this side of the problem by researching the relationship between FOMC disagreement and monetary policy uncertainty.

Caballero and Simsek (2022) and Sastry (2026) document the existence of and model and estimate the effect of disagreement between the Fed and market participants. We differ by focusing on disagreement amongst the FOMC. Couture (2021) is the paper most closely related to our analysis. In one of his results, he uses dispersion measures from the dot plots and regresses the change in a measure of monetary policy uncertainty on the change in these

measures of dispersion. His estimate for the change in the variance of the dots is negative and statistically insignificant (in contrast to our statistically significant negative estimate). Several differences exist between his exercise and what we do, most importantly that his sample runs only from 2012 to 2015.

3 Institutional Setting

Rudebusch and Williams (2008) provide a theoretical justification for conditional interest rate forecasts as a part of the monetary policy toolkit to help agents learn about the central bank's reaction function and thus future policy action. The dot plot is close to such a forward guidance tool. Engstrom (2026) finds that the dot plot provides useful information about subsequent policy actions. However, the focus is on the median dot and the first moment of policy. We focus on the disagreement in the dot plots and the uncertainty surrounding monetary policy.

The Summary of Economic Projections (SEP) was first released in October 2007 and is released quarterly at every other scheduled FOMC meeting. In January 2012, the SEP first included the dot plot. Given that the dot plot is released at every other FOMC meeting in conjunction with the rest of the SEP, whether or not there is a dot plot at a scheduled FOMC meeting is not a function of economic conditions or any other concurrent variables.

The dot plot shows all FOMC members', voting or nonvoting, anonymized responses to the prompt: "FOMC participants' assessments of appropriate monetary policy: Midpoint of target range or target level for the federal funds rate." The prompt is answered from current calendar year-end to up to three years ahead calendar year-end horizons as well as for the long run. Each dot plot contains assessments with horizons up to two years ahead assessments and dot plots near the end of the calendar year contain assessments with horizons up to three years ahead.

Of note is that the dot plot does not contain forecasts of where the Federal Funds Rate will be, rather it is asking each member their opinion on appropriate policy. As such, dispersion in the dot plot reflects differences in assessments of appropriate monetary policy. This we take as disagreement about appropriate monetary policy as they reflect differences in the opinions of FOMC members about what is appropriate monetary policy.

The assessments for the dot plots are set the Friday before the FOMC meeting. By this time, they have received the staff economic forecasts. Members can adjust their assessments during the meeting, however this is uncommon according to Bernanke (2016) and Powell (2024). Conceptually, we interpret dot plots dispersion reflecting pre-meeting disagreement amongst FOMC members.

Our disagreement is distinct from uncertainty about monetary policy. The monetary policy uncertainty we consider is among market participants. We use the market-based short-rate uncertainty from Bauer et al. (2022). It is calculated using out-of-the-money Eurodollar derivatives until September 2022 and SOFR derivatives from there on. The uncertainty object of interest is the implied conditional standard deviation of the 3-month LIBOR or SOFR at time t for h periods into the future.[1] Changes in the short-rate uncertainty around FOMC meetings are likely to be dominated by news coming from FOMC meetings. We consider changes in the short-rate uncertainty around FOMC meetings as changes in monetary policy uncertainty.

Our empirical analysis compares observed dispersion in the released dot plot with the dispersion expected before the announcement to capture the surprise component of disagreement. This comparison complements the measure based on changes since the previous release. Engstrom (2026) shows that dots from the previous dot plot can become stale by the time the next dot plot is released.

NatWest produces FOMC previews prior to FOMC meetings. The exact timing of when these previews are released varies, but they are all released within the week preceding the FOMC meeting with many being released the day before the FOMC meeting. Starting from December 2015[2], these FOMC previews include a mock dot plot. We use these mock dot plots to calculate expected dispersion in an upcoming dot plot. This lets us use a measure of realized dispersion relative to expected dispersion in a single dot plot in our empirical analysis for most of the sample, in addition to realized dispersion in one dot plot relative to realized in the previous dot plot.

4 Committee Disagreement and Market Uncertainty

4.1 Measuring Disagreement and Uncertainty

We measure disagreement among FOMC participants using the variance of their interest-rate assessments in the dot plot. We measure market uncertainty using the option-implied variance of future short-term interest rates from Bauer et al. (2022).

The empirical dependent variable we use is the change in short-rate uncertainty around FOMC meetings at a date t, looking h periods ahead. We use a 2-day window around the

[1] See Appendix A.1 in Bauer et al. (2022) for a detailed explanation of the calculation of the short-rate uncertainty measure.

[2] To our knowledge mock dot plots were not constructed before December 2015. With the exceptions of the June 2020 and March 2023 dot plots, each dot plot from December 2015 onwards has a corresponding mock dot plot.

FOMC meeting going from the close of the day before the FOMC announcement until the close of the next trading day.

The empirical independent variable we use is the change in variance of the dots in the dot plot. In Figure 1 we plot the variance of the dots in the dot plots on average as well as in the time series by the assessment horizon. We see that there is on average much more variance in assessments for horizons two and three years into the future. The three year ahead horizon is only available in half of the dot plots. There is also much more variation in the time series in the two year ahead horizon variance relative to the current year and one year ahead variances. As such, we use the two year ahead horizon in our main specifications.

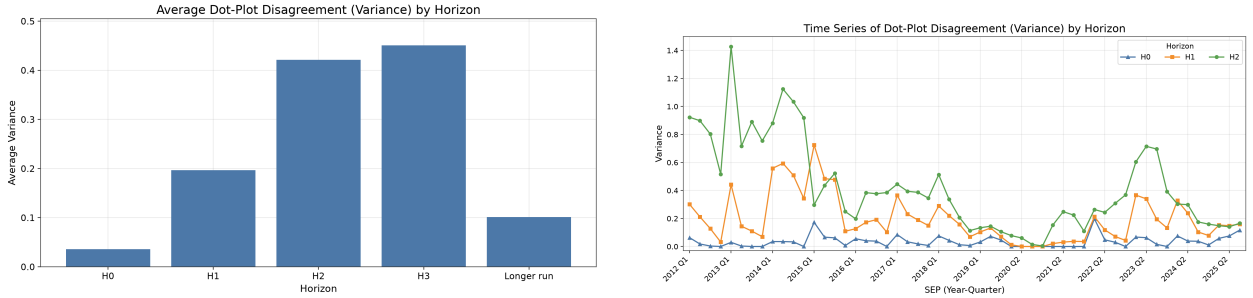


Figure 1: Variance in the dot plots on average and in the time series by horizon

4.2 Baseline Specification and Results

The question we want to answer is how does disagreement amongst the FOMC affect monetary policy uncertainty. The first step is to regress the change in monetary policy uncertainty around an FOMC meeting with a dot plot on the change in the variance of dots in the dot plot. The specification we use is shown below:

$$\Delta SRU_t = \alpha + \beta \Delta \text{Disagreement}_t + X_t' \Gamma + \epsilon_t,$$

where ΔSRU denotes the change in the market implied level of uncertainty about the future short rate in a 2-day window around an FOMC meeting with a dot plot, ΔDisagreement denotes the change in the variance of the dots in the dot plot, X denotes a set of controls, and t indexes time.

The set of controls includes: quarter dummies, the median dot, indicators for being at the ZLB, a post-Covid indicator, the pre-meeting level of short-rate uncertainty, and first moment monetary policy shocks.

We report the results in terms of standardized units, i.e. how many standard deviations does the change in the variance of the short-term interest rate change for a unit standard

deviation change in the variance of the dots in the dot plot. This makes it easier to interpret the economic magnitude of the coefficient of interest, β.

Table 1: Change in Short Rate Uncertainty on FOMC Disagreement

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Standardized effect	-0.330^{***}	-0.326^{**}	-0.297^{**}	-0.323^{**}	-0.294^{*}	-0.318^{*}	-0.304^{**}	-0.302^{*}
HC3 standard error	(0.124)	(0.150)	(0.140)	(0.142)	(0.167)	(0.163)	(0.147)	(0.171)
p-value	0.007988	0.029892	0.033765	0.022822	0.077537	0.051727	0.039543	0.078171
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	53	53	53	53	53	53	53	53
R^2	0.1090	0.1421	0.1447	0.1109	0.1765	0.1445	0.1581	0.1879

The dependent variable is the two-day change in short-rate uncertainty. Reported is the standardized coefficient for the variance of the current H2 dots minus the variance of the rollover-matched prior dots regressor. HC3 standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

The coefficient on the change in dot plot disagreement is statistically significantly negative. The magnitude is such that for a one standard deviation increase in the change of dot plot dispersion, the change in short-rate uncertainty around an FOMC meeting decreases by about 0.3 standard deviations. The coefficient magnitude is remarkably stable across specifications. Table 1 shows the estimated β without controlling for first moment monetary policy shocks while Table 2 adds them as controls. Neither Table 1 nor Table 2 control for the pre-meeting level of short-rate uncertainty. In Table 3 we add the pre-meeting level of short-rate uncertainty as a control and in Table 4 we change the dependent variable in the regressions to be the log change in short-rate uncertainty rather than the raw change. These specifications are both done to address the concern of there not being room for short-rate uncertainty to fall either by controlling for the level of short-rate uncertainty or shifting some of the coefficient weight more to meetings with lower initial short-rate uncertainty.

As a whole, we see that when there is an increase in the observed disagreement in a dot plot, we also see a larger decrease or a smaller increase in short-rate uncertainty around the FOMC meeting.

4.3 Expected and Unexpected Disagreement

For the results using the realized change in dot plot disagreement shown in Table 1 to Table 4, one is concerned about how much of the change is anticipated by the market and how much is genuinely news. To address this, we turn to the mock dot plots produced by NatWest.

The mock dot plots produced by NatWest have the same structure as the actual dot plots produced by the Fed in the SEP. As such they contain the entire distribution of dots, not

Table 2: Change in Short Rate Uncertainty on FOMC Disagreement With NS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Standardized effect	-0.346^{***}	-0.341^{**}	-0.318^{***}	-0.346^{***}	-0.321^{*}	-0.343^{**}	-0.327^{**}	-0.331^{*}
HC3 standard error	(0.105)	(0.155)	(0.121)	(0.119)	(0.170)	(0.160)	(0.128)	(0.172)
p-value	0.000972	0.027450	0.008749	0.003733	0.058383	0.032413	0.010607	0.054006
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	46	46	46	46	46	46	46	46
R^2	0.2554	0.3483	0.2842	0.2554	0.3663	0.3484	0.3058	0.3832

The dependent variable is the two-day change in short-rate uncertainty. Reported is the standardized coefficient for the variance of the current H2 dots minus the variance of the rollover-matched prior dots regressor. HC3 standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table 3: Change in Short Rate Uncertainty on FOMC Disagreement With SRU Control

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: No monetary-policy shock control								
Standardized effect	-0.291^{**}	-0.291^{*}	-0.291^{**}	-0.305^{**}	-0.292^{*}	-0.306^{*}	-0.299^{**}	-0.301^{*}
HC3 standard error	(0.141)	(0.169)	(0.147)	(0.149)	(0.174)	(0.172)	(0.152)	(0.175)
p-value	0.039543	0.086088	0.047976	0.040768	0.092397	0.074639	0.048876	0.085587
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	53	53	53	53	53	53	53	53
R^2	0.2383	0.2739	0.2384	0.2516	0.2740	0.2864	0.2579	0.2916

The dependent variable is the two-day change in short-rate uncertainty. Reported is the standardized coefficient for the variance of the current H2 dots minus the variance of the rollover-matched prior dots regressor. HC3 standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table 4: Log Change in Short Rate Uncertainty on FOMC Disagreement

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: No monetary-policy shock control								
Standardized effect	-0.445^{*}	-0.460^{**}	-0.445^{*}	-0.449^{*}	-0.459^{**}	-0.463^{**}	-0.447^{*}	-0.461^{**}
HC3 standard error	(0.243)	(0.209)	(0.247)	(0.245)	(0.213)	(0.212)	(0.248)	(0.214)
p-value	0.067111	0.027869	0.071430	0.067468	0.031697	0.028724	0.071910	0.031229
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	53	53	53	53	53	53	53	53
R^2	0.1980	0.2709	0.1980	0.1986	0.2710	0.2712	0.1990	0.2721

The dependent variable is the two-day log change in short-rate uncertainty. Reported is the standardized coefficient for the variance of the current H2 dots minus the variance of the rollover-matched prior dots regressor. HC3 standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

just the median dot. They are released prior to the FOMC meeting and as such are only based on information known prior to the meeting. Using these mock dot plots, we calculate expected variance of the dots.

Before using the mock dot plots as expected dot plots, it is sensible to evaluate how accurate the mock dot plots are as forecast for the dot plots. A natural benchmark for evaluation is using the previous dot plot and seeing whether the errors of the mock dot plot are smaller than the errors of using the previous dot plot as the forecast for the next dot plot.

Let A_t denote the realized dot plot distribution and M_t denote the mock dot plot distribution at meeting t. The absolute median dot forecast error (AMFE) and absolute dot variance forecast error (AVFE) when using the mock dot plots as the forecast are:

$$AMFE_t = \frac{|\text{Med}(A_t) - \text{Med}(M_t)|}{\text{Med}(A_{t-1})}$$
$$AVFE_t = \frac{|\text{Var}(A_t) - \text{Var}(M_t)|}{\text{Var}(A_{t-1})}$$

and when using the previous dot plot as the forecast, they are:

$$AMFE_t = \frac{|\text{Med}(A_t) - \text{Med}(A_{t-1})|}{\text{Med}(A_{t-1})}$$
$$AVFE_t = \frac{|\text{Var}(A_t) - \text{Var}(A_{t-1})|}{\text{Var}(A_{t-1})}$$

In Table 5 we show the average absolute median and variance forecast errors calculated for the 38 periods for which both can be calculated. Of most interest to our context is the average absolute variance forecast error which is about 14.5% smaller for the NatWest mock dots than for the previous dot plots. This suggests that while the NatWest mock dots contain some information above and beyond the previous dots, but not much. Hence the realized change used in our main specification is largely unexpected by the market.

Table 5: Average Absolute Forecast Errors for the H2 Dot Distribution

Forecast	Median	Variance
NatWest mock	0.1427	0.2361
Previous dot plot	0.2178	0.2760

In Table 6 we show the analogous as in Table 1, except that the independent variable of interest is the surprise portion of the disagreement in the dot plots measured using the NatWest mock dot plots. Specifically, it is the variance of the H2 dots in a dot plot minus

the expected variance of the H2 dots of the same dot plot.

The sample for which we have the mock dot plots is substantially smaller than full dot plot sample, however we see qualitatively similar results. The sign and magnitude are similar to that seen in Table 1. However, the statistical significance is lessened as expected from having less statistical power.

Table 6: Change in Short Rate Uncertainty on FOMC Disagreement Relative to Expectation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Standardized effect	−0.287	−0.356*	−0.280	−0.314	−0.353	−0.376*	−0.374*	−0.442*
HC3 standard error	(0.183)	(0.203)	(0.205)	(0.191)	(0.221)	(0.205)	(0.210)	(0.230)
p-value	0.117372	0.079185	0.172478	0.100692	0.109749	0.066387	0.074597	0.054426
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	38	38	38	38	38	38	38	38
R^2	0.0824	0.1379	0.1094	0.1049	0.1683	0.1543	0.2909	0.3347

The dependent variable is the two-day change in short-rate uncertainty. Reported is the standardized coefficient for the variance of the current H2 dots minus the variance of the NatWest mock H2 dots for the same dot plot regressor. HC3 standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

4.4 Results Across Uncertainty Horizons

The negative relationship also holds when uncertainty is measured at the twenty-four-month horizon. Appendix Tables A.1 and A.2 report standardized coefficients ranging from −0.34 to −0.45 across specifications, all statistically significant at the 10 percent level. Several specifications controlling for monetary policy shocks are significant at the 5 percent level. These results indicate that the negative association extends beyond the twelve-month horizon used in the baseline analysis.

5 Model

We now develop a Bayesian learning model to explain how greater disagreement can reduce uncertainty about future policy rates.

5.1 Narrative Evidence on Policymakers' Influence

Commentary surrounding the September 2024 FOMC meeting illustrates how analysts use the dots and the policy decision to infer policymakers' relative influence. After the committee

cut rates by 50 basis points, ABN AMRO interpreted the divided dots as evidence that the camp favoring faster easing had prevailed (Quaedvlieg, 2024). Goldman Sachs similarly interpreted some year-end projections as indicating a preference for a smaller initial cut and concluded that the leadership had secured the larger move (Mericle, 2024).

These inferences also entered assessments of future policy. Salman Ahmed of Fidelity International concluded that the "doves are in control" and expected further labor-market weakness to lead to faster and deeper cuts (quoted in Zúñiga, 2024). This interpretation connects perceived influence within the committee to expectations of its future decisions, consistent with the persistence underlying our model.

Earlier accounts provide supporting evidence. Discussing the September 2019 dots, Brown inferred that the balance of power favored the doves despite disagreement among participants (Brown, 2019). Evans likewise described how individual projections convey a range of policy views, while the committee's decision and subsequent communication reveal where its collective weight lies (Evans, 2016). Together, these accounts provide narrative support for the model's premise that markets learn how individual policy preferences translate into committee decisions and use this information to assess future policy.

5.2 Hawkish and Dovish Taylor Rules

For each SEP projection horizon l, let the common date-t macroeconomic forecast be

$$X_{t,l} \equiv \begin{bmatrix} 1 \\ \pi_{t+l|t} - \pi^* \\ x_{t+l|t} \end{bmatrix},$$

where $\pi_{t+l|t}$ denotes the date-t forecast of inflation at projection horizon l, π^* denotes the inflation target, and $x_{t+l|t}$ denotes the date-t forecast of the output gap at projection horizon l.

We assume that the dovish and hawkish policymakers observe the same macroeconomic state. This means that their disagreement stems from different reaction-function coefficients, not from different macroeconomic forecasts.

Let the dovish and hawkish endpoint Taylor-rule coefficient vectors be

$$\theta_s \equiv \begin{bmatrix} \alpha_s \\ \phi_{\pi,s} \\ \phi_{x,s} \end{bmatrix}, \qquad s \in \{D, H\}.$$

The endpoint Taylor rules are given by

$$m_s(X_{t,l}) = \theta_s' X_{t,l} = \alpha_s + \phi_{\pi,s}(\pi_{t+l|t} - \pi^*) + \phi_{x,s} x_{t+l|t}, \qquad s \in \{D, H\}.$$

The realized endpoint dots are given by

$$d_{s,t,l} = \theta_s' X_{t,l} + \xi_{s,t,l}, \qquad s \in \{D, H\},$$

where $\xi_{s,t,l}$ denotes endpoint-specific dot noise capturing factors such as judgment, risk-management considerations, measurement error, and other information not summarized by $X_{t,l}$.

Defining the coefficient gap as

$$\delta \equiv \theta_H - \theta_D = \begin{bmatrix} \alpha_H - \alpha_D \\ \phi_{\pi,H} - \phi_{\pi,D} \\ \phi_{x,H} - \phi_{x,D} \end{bmatrix},$$

then yields that horizon-l dot disagreement is given by

$$\begin{aligned} \Delta_{t,l} &\equiv d_{H,t,l} - d_{D,t,l} \\ &= \delta' X_{t,l} + \zeta_{t,l} \\ &= (\alpha_H - \alpha_D) + (\phi_{\pi,H} - \phi_{\pi,D})(\pi_{t+l|t} - \pi^*) + (\phi_{x,H} - \phi_{x,D}) x_{t+l|t} + \zeta_{t,l}. \end{aligned}$$

where $\zeta_{t,l} \equiv \xi_{H,t,l} - \xi_{D,t,l}$.

5.3 Effective Committee Taylor Rule

We consider an FOMC which aggregates differing views into a single policy action.

Let there be a latent committee center-of-gravity denoted by w_t. The aggregation equation for the horizon-l policy path is given by

$$\begin{aligned} i_{t,l} &= d_{D,t,l} + w_t \Delta_{t,l} + \eta_{t,l} = (1 - w_t) d_{D,t,l} + w_t d_{H,t,l} + \eta_{t,l} \\ &= [\theta_D + w_t(\theta_H - \theta_D)]' X_{t,l} + \varepsilon_{t,l}, \end{aligned}$$

where $\varepsilon_{t,l} \equiv (1 - w_t)\xi_{D,t,l} + w_t \xi_{H,t,l} + \eta_{t,l}$.

Then we define the effective committee-level Taylor-rule coefficient vector as

$$\theta_t^{\text{eff}} \equiv \theta_D + w_t(\theta_H - \theta_D) = \theta_D + w_t\delta,$$

leading to

$$i_{t,l} = \left(\theta_t^{\text{eff}}\right)' X_{t,l} + \varepsilon_{t,l}.$$

5.4 Learning From the Current Policy Decision

The current policy decision is observed at the same meeting as the dot plot. The current policy decision equation is given by

$$i_t = d_{D,t,0} + w_t\Delta_{t,0} + \eta_t, \qquad \eta_t \sim N(0, \sigma_\eta^2).$$

Conditional on the realized dots, $d_{D,t,0}$ and $\Delta_{t,0}$ are observed. The market forms the signal

$$\widetilde{i}_t \equiv i_t - d_{D,t,0} = w_t\Delta_{t,0} + \eta_t.$$

We assume that the center-of-gravity weight follows

$$w_t = \rho w_{t-1} + (1-\rho)\bar{w} + \nu_t, \qquad \nu_t \sim N(0, \sigma_\nu^2), \qquad |\rho| < 1.$$

The market enters the meeting with prior

$$w_t \mid \mathcal{F}_{t-} \sim N(E_{t^-}(w_t), P_{t^-}),$$

where $\mathcal{F}_{t-}$ denotes the market's information set immediately before the policy announcement and SEP release, let $\mathcal{F}_t$ denote the information set immediately after, and $P_{t^-} \equiv Var_{t^-}(w_t)$.

Let

$$Z_t \equiv \Delta_{t,0}^2.$$

Then the posterior variance after observing the dots and the rate decision is

$$P_t(Z_t) = \frac{\sigma_\eta^2 P_{t^-}}{Z_t P_{t^-} + \sigma_\eta^2}.$$

Its derivative with respect to the squared near-term Taylor-rule spread is

$$P_t'(Z_t) = -\frac{\sigma_\eta^2 P_{t^-}^2}{\left(Z_t P_{t^-} + \sigma_\eta^2\right)^2} < 0.$$

Hence, the current decision is more informative about w_t when the hawkish and dovish Taylor-rule prescriptions are farther apart at the near-term macro forecast. This is the information channel.

5.5 Uncertainty About Future Policy Rates

The empirical SRU object is forward-looking, it is the uncertainty about a future rate decision, not the uncertainty about the current stance. At a future meeting at time $t + h$, the near-term decision can be written as

$$i_{t+h} = d_{D,t+h,0} + w_{t+h}\Delta_{t+h,0} + \eta_{t+h}.$$

Under the Taylor-rule interpretation,

$$d_{D,t+h,0} = \theta_D' X_{t+h,0} + \xi_{D,t+h,0},$$

$$\Delta_{t+h,0} = \delta' X_{t+h,0} + \zeta_{t+h,0}.$$

Equivalently,

$$i_{t+h} = \left(\theta_D + w_{t+h}\delta\right)' X_{t+h,0} + \varepsilon_{t+h}.$$

The forward variance of the committee weight is

$$Var_t(w_{t+h}) = \rho^{2h} P_t(Z_t) + \sigma_\nu^2 \frac{1-\rho^{2h}}{1-\rho^2}.$$

Then we can write the variance of the future rate decision as

$$Var_t(i_{t+h}) = E_t[\Delta^2_{t+h,0}]\,\mathrm{Var}_t(w_{t+h}) + \big(E_t[w_{t+h}]\big)^2\,\mathrm{Var}_t(\Delta_{t+h,0}) + \mathrm{Var}_t(d_{D,t+h,0}) + 2E_t[w_{t+h}]\,\mathrm{Cov}_t(d_{D,t+h,0},\Delta_{t+h,0}) + \sigma^2_\eta.$$

5.6 Persistent Disagreement

The disagreement object whose persistence enters the main result is

$$Z_t \equiv \Delta^2_{t,0}.$$

We assume that Z_t follows an AR(1) process

$$Z_{t+1} = \mu_0 + \phi_0(Z_t - \mu_0) + u^Z_{t+1}, \qquad u^Z_{t+1} \sim N(0,\sigma^2_{u,0}), \qquad \phi_0 \in [0,1), \quad \mu_0 > 0.$$

The h-step forecast is

$$E_t[Z_{t+h}] = \mu_0 + \phi_0^h(Z_t - \mu_0).$$

5.7 Disagreement Surprises and Changes in Uncertainty

Define the forward-looking SRU innovation

$$\Delta SRU_{t,h} \equiv Var_t(i_{t+h}) - Var_{t-}(i_{t+h}).$$

We want to get at the event-date slope. Let

$$z_{0,t} \equiv E_{t-}[Z_t], \qquad S_{t,0} \equiv Z_t - z_{0,t}.$$

The object $S_{t,0}$ is the surprise component of near-term Taylor-rule disagreement.

Define

$$G_h(z) \equiv \rho^{2h}\left[\mu_0 + \phi_0^h(z-\mu_0)\right] P_t(z).$$

Then the main-channel contribution to the event-date innovation is

$$G_h(Z_t) - E_{t-}[G_h(Z_t)].$$

A first-order expansion around $z_{0,t}$ gives

$$G_h(Z_t) - E_{t-}[G_h(Z_t)] \approx (G_h)'(z_{0,t}) S_{t,0}.$$

Differentiating,

$$G_h'(z) = \rho^{2h} \left\{ \phi_0^h P_t(z) + \left[\mu_0 + \phi_0^h (z - \mu_0) \right] P_t'(z) \right\}.$$

Evaluating at the unconditional mean $z_{0,t} = \mu_0$, define

$$\bar{P}_0 \equiv P_t(\mu_0).$$

The main-channel Taylor-rule slope is

$$\beta_{h,0} \equiv G_h'(\mu_0) = \rho^{2h} \left[\underbrace{\phi_0^h \bar{P}_0}_{\text{persistent Taylor-rule disagreement } (+)} + \underbrace{\mu_0 P_t'(\mu_0)}_{\text{committee-weight learning } (-)} \right].$$

Because $P_t'(\mu_0) < 0$, the slope is the sum of a positive persistence channel and a negative learning channel.

The corresponding linearized event-date regression is

$$\Delta SRU_{t,h} = \alpha_h + \beta_{h,0} S_{t,0} + \varepsilon_{t,h},$$

where $\varepsilon_{t,h}$ collects higher-order terms, meeting-day news orthogonal to $S_{t,0}$, and the omitted terms in the full variance decomposition.

5.8 When Disagreement Reduces Uncertainty

Proposition 1 (Ambiguous sign at a fixed forward horizon)**.** *Fix $h \geq 1$. The main-channel slope satisfies*

1. $\beta_{h,0} > 0$ *if* $\phi_0 > \phi_{h,0}^*$;
2. $\beta_{h,0} < 0$ *if* $\phi_0 < \phi_{h,0}^*$;
3. $\beta_{h,0} = 0$ *if* $\phi_0 = \phi_{h,0}^*$,

where

$$\phi_{h,0}^* \equiv \left(\frac{\mu_0 \left| P_t'(\mu_0) \right|}{\bar{P}_0} \right)^{1/h} = \left(\frac{\mu_0 P_{t^-}}{\mu_0 P_{t^-} + \sigma_\eta^2} \right)^{1/h} \in (0, 1).$$

Proof. From,

$$\beta_{h,0} = \rho^{2h} \left[\phi_0^h \bar{P}_0 - \mu_0 \left| P_t'(\mu_0) \right| \right].$$

The sign is the sign of $\phi_0^h \bar{P}_0 - \mu_0 |P_t'(\mu_0)|$. Setting this expression equal to zero gives

$$\phi_0^h = \frac{\mu_0 |P_t'(\mu_0)|}{\bar{P}_0}.$$

Then,

$$\frac{\mu_0 \left| P_t'(\mu_0) \right|}{\bar{P}_0} = \frac{\mu_0 P_{t^-}}{\mu_0 P_{t^-} + \sigma_\eta^2} \in (0,1).$$

□

If disagreement is persistent, wider current dots forecast wider future hawk-dove rule gaps, raising forward ΔSRU. If disagreement is not persistent, wider current dots mainly make the policy decision more informative about the committee's effective rule, lowering forward MPU. The sign of the slope depends on which force dominates.

Proposition 2 (Sign monotonicity across ΔSRU horizons)**.** *Fix $\phi_0 \in (0,1)$. The threshold $\phi_{h,0}^*$ is strictly increasing in h and satisfies $\phi_{h,0}^* \to 1$ as $h \to \infty$. Therefore, the sign of $\beta_{h,0}$ is weakly decreasing in h: if $\beta_{h_0,0} \leq 0$ for some h_0, then $\beta_{h,0} \leq 0$ for all $h \geq h_0$. Hence the slope can move from positive to negative as the MPU horizon lengthens, but it cannot move from negative to positive.*

Proof. Let

$$c \equiv \frac{\mu_0 P_{t^-}}{\mu_0 P_{t^-} + \sigma_\eta^2} \in (0,1).$$

Then $\phi_{h,0}^* = c^{1/h}$ so $\phi_h^* = e^{\ln(c)/h}$. As $c \in (0,1)$ we have that ϕ_h^* increases with h. We have that $\lim_{h\to\infty} e^{\ln(c)/h} = 1$ as $\lim_{h\to\infty} \ln(c)/h = 0$ so $\lim_{h\to\infty} \phi_h^* = 1$.

We have that $Sign(\beta_h) = Sign(\phi^h - c)$. Fix $\phi \in (0,1)$, then the sequence $\{\phi^1, \phi^2, \phi^3, ...\}$ is strictly decreasing. So $\phi^h > c$ can only hold for an initial set of horizons $h = 1, 2, ..., h^* - 1$. Once $\phi^h \leq c$, it remains below c forever. □

5.9 The Full Effect of Disagreement on Uncertainty

The main result above isolates the term

$$E_t[\Delta^2_{t+h,0}]\rho^{2h}P_t(Z_t).$$

The full variance decomposition adds two first-order corrections.

First, the full variance of w_{t+h} includes accumulated future committee weight innovations,

$$C_{1,h} \equiv \sigma^2_\nu \frac{1-\rho^{2h}}{1-\rho^2}.$$

Differentiating the term $E_t[Z_{t+h}]C_{1,h}$ with respect to Z_t adds the positive contribution $\phi_0^h C_{1,h}$.

Second, the posterior mean term $\big(E_t[w_{t+h}]\big)^2 Var_t(\Delta_{t+h,0})$ has a positive first-order effect. Taking the conditional expectation $E\big(\big(E_t[w_{t+h}]\big)^2|Z_t=z\big)Var(\Delta_{t+h,0})$ as the disagreement process is homoskedastic. With these terms added in, the linearized full slope can be written as

$$\beta^{full}_{h,0} = \underbrace{\phi_0^h\left(\rho^{2h}\bar{P}_0 + C_{1,h}\right)}_{\text{persistent Taylor-rule disagreement }(+)} + \underbrace{(\mu_0 - V_{\Delta,h})\rho^{2h}P_t'(\mu_0)}_{\text{committee-weight learning }(-)}.$$

Because

$$\mu_0 = E_t[\Delta^2_{t+h,0} \mid Z_t = \mu_0] \geq Var_t(\Delta_{t+h,0} \mid Z_t = \mu_0) = V_{\Delta,h}, \tag{1}$$

and $P_t'(\mu_0) < 0$, the second term is weakly negative. Thus, the full slope preserves the same two-channel structure: a positive persistent-disagreement channel and a negative information channel.

A full-slope threshold is

$$\phi^{*,full}_{h,0} = \left(\frac{(\mu_0 - V_{\Delta,h})\rho^{2h}|P_t'(\mu_0)|}{\rho^{2h}\bar{P}_0 + C_{1,h}}\right)^{1/h}, \tag{2}$$

when the expression inside the root lies in $(0,1)$. The corrections generally lower the persistence threshold relative to the main-channel threshold because they strengthen the positive persistence side and weaken the negative information side. However, they do not change the qualitative ambiguity.

In Section 4, we find negative coefficients at both the twelve- and twenty-four-month uncertainty horizons. Proposition 2 establishes sign monotonicity for the main-channel slope, but this restriction does not automatically extend to the full slope.

6 Monetary Policy Uncertainty and Investment

We now turn to the effects of monetary policy uncertainty on investment.

More monetary policy uncertainty plausibly increases the value of postponing economic decisions for which there is some irreversibility or costly reversibility until some of that uncertainty has been realized. Accordingly an increase in uncertainty can act to depress such economic decisions. This mechanism has been highly explored before by Bernanke (1983), McDonald and Siegel (1986), and Pindyck (1990). We look at what effect our monetary policy uncertainty shocks have on investment. The investment measure we use is real non-residential fixed investment. Monetary policy is highly relevant for investment decisions and recovering investment expenditure after investing is not trivial. Hence investment is a natural object to look at.

Current announcements simultaneously convey information about the current policy and the future path of interest rates in addition to the uncertainty surrounding monetary policy. Our goal is to extract an uncertainty component from announcements and investigate the effect of that component on investment. This exercise closely relates to work by Bloom (2009) and Baker et al. (2016). We focus on monetary policy uncertainty and, in particular, changes in monetary policy uncertainty arising from FOMC announcements.

In order to extract the monetary policy uncertainty news component at an announcement, we use heteroskedasticity-based identification (HBI). Rigobon (2003) introduces HBI and Rigobon and Sack (2004) and Acosta (2026) have used HBI in the monetary policy shock context. We follow in their footsteps by applying HBI to monetary policy uncertainty.

6.1 Regime Selection

HBI requires selecting regimes in which the relative variances of different news coming out of the FOMC meetings changes. Our regime selection is motivated by changes in the communication of the Federal Reserve, in particular the introduction of forward guidance and the introduction of the dot plot. Our first regime covers a period before forward guidance was first used in 2003. This regime runs from 1995 to 2002. Our second regime covers the subsequent period from 2003 to 2006, up until the first SEP in 2007. Here the relative importance of the future path of interest rates increases with the introduction of forward guidance. Our third regime runs from 2012 until the end of 2023 motivated by the introduction of the dot plot in 2012.

6.2 Identification Through Changes in Volatility

HBI can separate the different news components of an FOMC announcement by exploiting relative changes in the variances of the news components. This is because when a particular component becomes relatively more important in a period, the observed comovement among variables will be more strongly driven by the influence of that component. Comparing covariance patterns across regimes then helps distinguish the effects of the different news components.

6.3 Identifying Assumptions

The central identifying assumptions are that each news component has the same effects on observable variables across regimes, while the relative variances of the components change across regimes. We also assume that the news components are contemporaneously uncorrelated with each other within each regime.

6.4 Estimation

Our regimes are indexed by $r(t) \in \{1, 2, 3\}$. Our vector of observables is given by

$$z_t = \begin{pmatrix} T_t \\ P_t \\ \Delta SRU_t \end{pmatrix}$$

where T denotes the target factor, P denotes the path factor and ΔSRU denotes the change in short-rate uncertainty around an FOMC announcement.

We use $z_t = \mu_{r(t)} + A\varepsilon_t$, $E(\varepsilon_t | r(t) = r) = 0$, and $Var(\varepsilon_t | r(t) = r) = D_r$ as our specification of the observables. D_r is diagonal for all regimes and A is nonsingular. $\Sigma_r \equiv Var(z_t | r(t) = r) = AD_rA'$. We normalize $D_1 = I_3$. For our identification, we need relative changes in the entries of D_r.

Let $B = A^{-1}$, $\widehat{\mu}_r$ denote the sample mean vectors, and $\widehat{\Sigma}_r$ denote the sample covariance matrices. We impose the first regime normalization by imposing that $B\widehat{\Sigma}_1 B' = I_3$. Our goal is to find a B such that the three components are uncorrelated within each regime. Our optimization problem then is $\widehat{B} = \underset{B:\, B\widehat{\Sigma}_1 B' = I_3}{\arg\min}\; g(B)'\widehat{W}g(B)$, where $g(B) = \begin{pmatrix} \text{off}\left(B\widehat{\Sigma}_2 B'\right) \\ \text{off}\left(B\widehat{\Sigma}_3 B'\right) \end{pmatrix}$, where $\text{off}(\cdot)$ collects the three distinct off-diagonal elements, so $g(B)$ contains six elements that we assume are zero in population. We estimate B in two steps. In the first step we use the identity matrix as the weighting matrix to recover our shocks. In the second step

the weighting matrix we use is the inverse of a block-diagonal estimate of the covariance matrix of the moment conditions. We label shocks similarly as Acosta (2026) using relative variance changes across regimes.[3] The shocks are signed so that an increase in the shock increases the LHS variable. Importantly for us, a positive uncertainty shock increases short-rate uncertainty on impact.

6.5 Diagnostics

The relative variance patterns which identify the components can be gauged from the vector of the diagonal elements of $\widehat{\Sigma}_1$, $\widehat{\Sigma}_2$, and $\widehat{\Sigma}_3$. Given that we have normalized $\widehat{\Sigma}_1$ to the identity matrix, we would want to see the vector of the diagonal elements in $\widehat{\Sigma}_2$ and $\widehat{\Sigma}_3$ being relatively different from each other and from the vector of ones. The three vectors are

$$\begin{aligned}
\operatorname{diag}\left(\widehat{\Sigma}_1\right) &= \begin{pmatrix} 1 & 1 & 1 \end{pmatrix} \\
\operatorname{diag}\left(\widehat{\Sigma}_2\right) &= \begin{pmatrix} 0.248 & 1.419 & 0.469 \end{pmatrix} \\
\operatorname{diag}\left(\widehat{\Sigma}_3\right) &= \begin{pmatrix} 0.082 & 0.603 & 1.013 \end{pmatrix},
\end{aligned}$$

which show the substantial relative variation required for HBI.

We conduct a test of the overidentifying restrictions. The null hypothesis is that the population covariance matrices can be diagonalized as AD_rA' with a single matrix A. This null would be rejected if either the impact loadings differ by regime or if the shocks are not uncorrelated within each regime. We construct bootstrap samples assuming the fitted covariance restrictions. For each sample, we re-estimate the entire HBI procedure. This test yields a p-value of 0.445.

6.6 Investment Responses to Uncertainty Shocks

We use our monetary policy uncertainty shocks in local projections to estimate their effect on investment. Before we do that, we residualize our shocks with respect to the 6 variables used by Bauer and Swanson (2023). The variables are the surprise in nonfarm payroll employment, nonfarm payroll employment growth over the preceding year, the 3-month S&P 500 return, the 3-month change in the Treasury yield-curve slope, the 3-month change in the Bloomberg Commodity Spot Price Index, and the option-implied skewness of the ten-year Treasury yield.

[3] Using largest loadings of the shocks on the LHS variable to label the shocks yields the same labeling.

Let $y_t = 100 \log I_t$, where I_t denotes real nonresidential fixed investment. Let s_t^T, s_t^P, and s_t^U denote the quarterly sums of the information-adjusted target, path, and uncertainty components, respectively. At each horizon $h = 0, \ldots, 16$, we estimate

$$\begin{aligned} y_{t+h} - y_{t-1} = \alpha_h + \beta_h^T s_t^T + \beta_h^P s_t^P + \beta_h^U s_t^U \\ + \sum_{\ell=1}^{4} \phi_{h,\ell} y_{t-\ell} + \sum_{\ell=1}^{4} \sum_{j \in \{T,P,U\}} \gamma_{h,\ell}^j s_{t-\ell}^j + u_{t+h}^{(h)}. \end{aligned} \tag{3}$$

The dependent variable measures the cumulative change in investment relative to the quarter preceding the shock. The regression includes all three contemporaneous components, four lags of each component, and four lags of the outcome. The coefficient β_h^U measures the response to the uncertainty component, conditional on these controls. We scale the reported responses to a one-standard-deviation quarterly uncertainty shock.

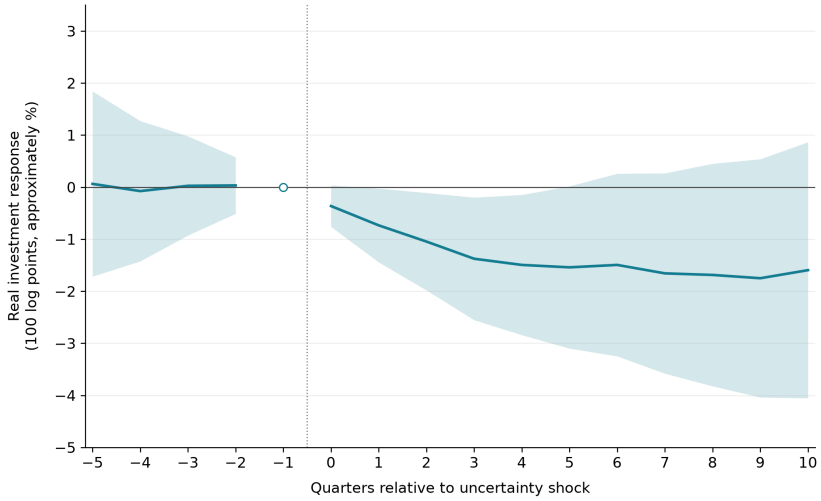


Figure 2: IRF Nonresidential Investment Residualized Uncertainty Shock HC3

7 Conclusion

We have documented that changes in dot plot disagreement are negatively associated with changes in short-rate uncertainty around FOMC meetings. This result also holds when using expected dot plots from NatWest as the reference value rather than the previous dot plot.

This negative finding can be explained using a Bayesian learning model where the market learns about an unobserved persistent balance of power term between hawks and doves. This learning works via the market observing the interest rate decision and dot plot dispersion.

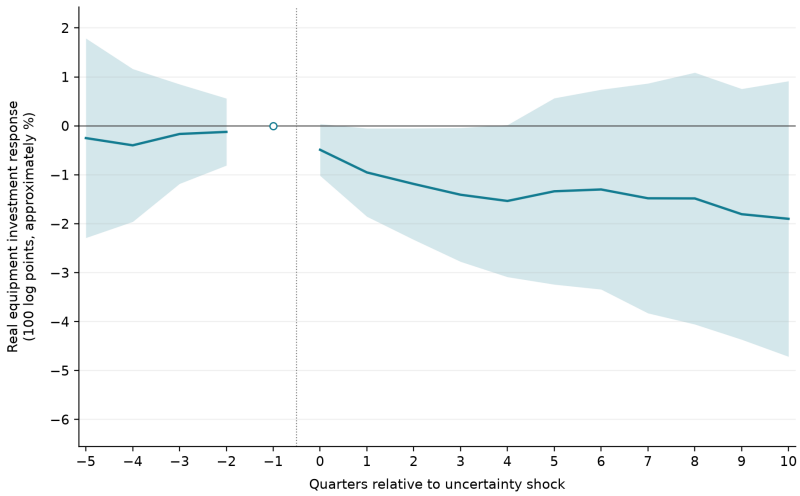


Figure 3: IRF Investment Equipment Residualized Uncertainty Shock HC3

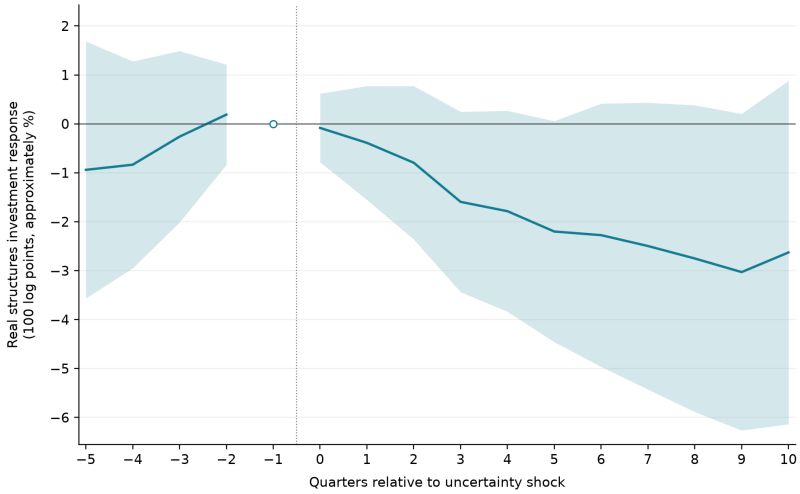


Figure 4: IRF Investment Structures Residualized Uncertainty Shock HC3

Table 7: Residualized uncertainty shocks and changes in dot-plot variance

	Residualized uncertainty shock (standardized)			
	(1)	(2)	(3)	(4)
H2 dot-plot variance change	-0.372^{**}	-0.362^{**}	-0.402^{**}	-0.390^{**}
	(0.152)	(0.166)	(0.159)	(0.175)
p-value	0.0145	0.0295	0.0111	0.0255
H2 median-level control	No	Yes	No	Yes
Quarter-of-year controls	No	No	Yes	Yes
Observations	45	45	45	45
R^2	0.138	0.142	0.202	0.208

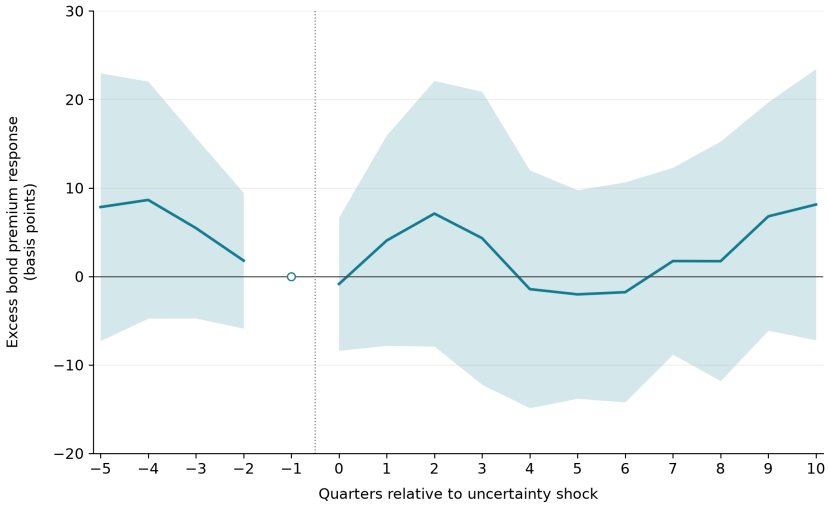


Figure 5: IRF Excess Bond Premium Residualized Uncertainty Shock HC3

The more dispersion there is, the more the market can infer whose opinions are driving policy and as such whose opinions are likely to drive policy in the future. The model admits a positive force as well via persistent disagreement widening the range of future possible actions so the model does not predict a negative slope, but rather can explain a negative slope.

The dot plots have been criticized for several reasons, one of which is that they highlight points of disagreement rather than consensus. Perhaps this is a feature, not a bug as observing this disagreement may lower monetary policy uncertainty.

We also use HBI to generate monetary policy uncertainty shocks which depress investment in the short-run consistent with a real options channel. We do not observe these shocks as having an effect on the excess bond premium. Additionally these shocks load negatively on dot plot disagreement.

References

Acosta, M. (2026). The perceived causes of monetary policy surprises. Working paper, University of Wisconsin–Madison. Version dated March 9, 2026.

Apel, M., Blix Grimaldi, M., and Hull, I. (2022). How much information do monetary policy committees disclose? evidence from the FOMC's minutes and transcripts. *Journal of Money, Credit and Banking*, 54(5):1459–1490.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Bauer, M. D., Lakdawala, A. K., and Mueller, P. (2022). Market-based monetary policy uncertainty. *The Economic Journal*, 132(644):1290–1308.

Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37:87–155.

Belden, S. (1989). Policy preferences of FOMC members as revealed by dissenting votes. *Journal of Money, Credit and Banking*, 21(4):432–441.

Bennani, H., Farvaque, E., and Stanek, P. (2018). Influence of regional cycles and personal background on FOMC members' preferences and disagreement. *Economic Modelling*, 68:416–424.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, 98(1):85–106.

Bernanke, B. S. (2016). Federal reserve economic projections: What are they good for? Ben Bernanke's Blog.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bordo, M. D. and Istrefi, K. (2023). Perceived FOMC: The making of hawks, doves and swingers. *Journal of Monetary Economics*, 136:125–143.

Brown, S. (2019). Insuring against recession? Raymond James; republished in Advisor Perspectives. October 23.

Caballero, R. J. and Simsek, A. (2022). Monetary policy with opinionated markets. *American Economic Review*, 112(7):2353–92.

Chappell, Jr., H. W., Havrilesky, T. M., and McGregor, R. R. (1993). Partisan monetary policies: Presidential influence through the power of appointment. *The Quarterly Journal of Economics*, 108(1):185–218.

Couture, C. (2021). Financial market effects of FOMC projections. *Journal of Macroeconomics*, 67:103279.

Engstrom, E. (2026). Anchored to the dot plot: Central bank projections and interest rate expectations. Technical Report 2026-026, Board of Governors of the Federal Reserve System.

Evans, C. L. (2016). Reflections on the current monetary policy environment. Federal Reserve Bank of Chicago. Speech delivered at the Global Interdependence Center, London, June 3. Quoted passage on p. 10 of the print version.

Foerster, A. and Martinez, Z. (2023). The evolution of disagreement in the dot plot. *FRBSF Economic Letter*, 2023(21):1–6.

Hansen, S., McMahon, M., and Prat, A. (2018). Transparency and deliberation within the FOMC: A computational linguistics approach. *The Quarterly Journal of Economics*, 133(2):801–870.

Hansen, S., McMahon, M., and Velasco Rivera, C. (2014). Preferences or private assessments on a monetary policy committee? *Journal of Monetary Economics*, 67:16–32.

Howes, C., Dordal i Carreras, M., Coibion, O., and Gorodnichenko, Y. (2026). How monetary policy is made: Lessons from historical FOMC discussions. NBER Working Paper 34632, National Bureau of Economic Research.

Madeira, C., Madeira, J., and Monteiro, P. S. (2026). The origins of monetary policy disagreement: The role of supply and demand shocks. *The Review of Economics and Statistics*, 108(2):355–371.

Malmendier, U., Nagel, S., and Yan, Z. (2021). The making of hawks and doves. *Journal of Monetary Economics*, 117:19–42.

McDonald, R. and Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–727.

Meade, E. E. and Stasavage, D. (2008). Publicity of debate and the incentive to dissent: Evidence from the US federal reserve. *The Economic Journal*, 118(528):695–717.

Mericle, D. (2024). US Daily: September FOMC recap: Rate cuts begin with a 50. Technical report, Goldman Sachs. September 18. Relevant discussion on pp. 1–2.

Pindyck, R. S. (1990). Irreversibility, uncertainty, and investment. Working Paper 3307, National Bureau of Economic Research.

Powell, J. H. (2024). Transcript of chair powell's press conference.

Quaedvlieg, R. (2024). Fed Watch: Fed cuts by 50 bps to orchestrate soft landing. ABN AMRO. September 19.

Rigobon, R. (2003). Identification through heteroskedasticity. *The Review of Economics and Statistics*, 85(4):777–792.

Rigobon, R. and Sack, B. (2004). The impact of monetary policy on asset prices. *Journal of Monetary Economics*, 51(8):1553–1575.

Rudebusch, G. D. and Williams, J. C. (2008). Revealing the secrets of the temple: The value of publishing central bank interest rate projections. In Campbell, J. Y., editor, *Asset Prices and Monetary Policy*, pages 247–289. University of Chicago Press, Chicago.

Sahm, C. (2026). A good family fight. Stay-At-Home Macro (SAHM), June 16. https://stayathomemacro.substack.com/p/a-good-family-fight.

Sastry, K. A. (2026). Disagreement about monetary policy. *American Economic Journal: Macroeconomics*, 18(2):486–522.

Slok, T. (2026). The future of forward guidance under Fed chair Warsh. Bloomberg Television interview, June 16. Discussion at 3:39–4:12. https://www.youtube.com/watch?v=YPHSeySBv3g.

Tsang, K. P. and Yang, Z. (2025). Agree to disagree: Measuring hidden dissent in FOMC meetings. *Journal of Economic Dynamics and Control*, 180:105197.

Zúñiga, B. (2024). Why does the Fed's 50 basis point cut make sense? Funds Society. September 23. Quotes Salman Ahmed of Fidelity International.

A Appendix

A.1 Example of a Low and High Disagreement Dot Plot

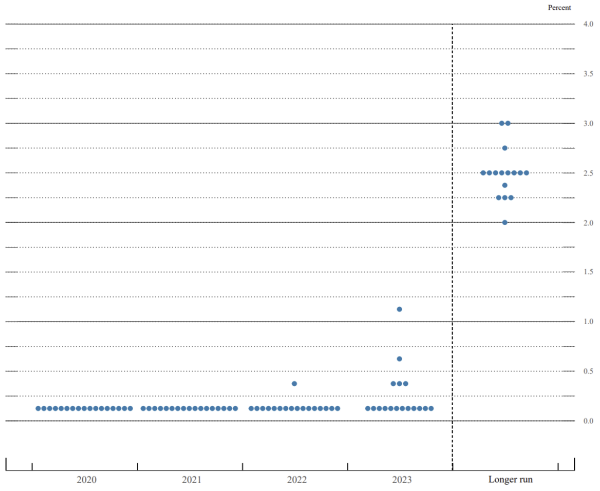


Figure A.1: Low Disagreement Dot Plot

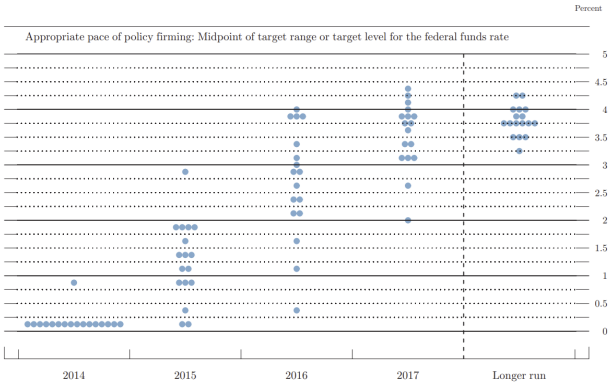


Figure A.2: High Disagreement Dot Plot

A.2 Additional Regression Results

Table A.1: Change in Short Rate Uncertainty on FOMC Disagreement, 24m SRU

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: No monetary-policy shock control								
Standardized effect	-0.385^{*}	-0.410^{*}	-0.342^{*}	-0.364^{*}	-0.368^{*}	-0.389^{*}	-0.345^{*}	-0.372^{*}
HC3 standard error	(0.213)	(0.209)	(0.204)	(0.212)	(0.192)	(0.204)	(0.208)	(0.195)
p-value	0.070731	0.050015	0.094190	0.085372	0.054856	0.056805	0.097353	0.056780
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	53	53	53	53	53	53	53	53
R^2	0.1481	0.1681	0.2083	0.1643	0.2282	0.1844	0.2112	0.2312

The dependent variable is the two-day change in short-rate uncertainty. Reported is the standardized coefficient for the variance of the current H2 dots minus the variance of the rollover-matched prior dots regressor. HC3 standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table A.2: Change in Short Rate Uncertainty on FOMC Disagreement With NS, 24m SRU

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel B: Nakamura–Steinsson (NS) shock control								
Standardized effect	-0.405^{*}	-0.446^{**}	-0.360^{*}	-0.383^{*}	-0.406^{**}	-0.424^{**}	-0.363^{*}	-0.409^{**}
HC3 standard error	(0.213)	(0.197)	(0.203)	(0.213)	(0.179)	(0.193)	(0.207)	(0.183)
p-value	0.056961	0.023712	0.076289	0.071672	0.023635	0.028359	0.080058	0.025245
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	46	46	46	46	46	46	46	46
R^2	0.1643	0.2036	0.2389	0.1850	0.2746	0.2226	0.2409	0.2769

The dependent variable is the two-day change in short-rate uncertainty. Reported is the standardized coefficient for the variance of the current H2 dots minus the variance of the rollover-matched prior dots regressor. HC3 standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table A.3: Log Change in Short Rate Uncertainty on FOMC Disagreement With SRU Control

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: No monetary-policy shock control								
Standardized effect	-0.444^{*}	-0.458^{**}	-0.446^{*}	-0.448^{*}	-0.459^{**}	-0.462^{**}	-0.447^{*}	-0.462^{**}
HC3 standard error	(0.244)	(0.208)	(0.248)	(0.246)	(0.215)	(0.211)	(0.250)	(0.216)
p-value	0.068409	0.027919	0.072575	0.068446	0.033056	0.028686	0.073063	0.032660
Quarter fixed effects	No	Yes	No	No	Yes	Yes	No	Yes
Median-dot control	No	No	Yes	No	Yes	No	Yes	Yes
ZLB indicator	No	No	No	Yes	No	Yes	Yes	Yes
Observations	53	53	53	53	53	53	53	53
R^2	0.1981	0.2717	0.1983	0.1993	0.2718	0.2729	0.1993	0.2730

Notes: The dependent variable is the two-day change in log short-rate uncertainty variance. The main regressor is the variance of the current H2 dots minus the variance of the rollover-matched prior dots. Columns vary quarter fixed effects, the median-dot control, and the ZLB indicator. The pre-meeting log sru variance control is included. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Table A.4: Forecasting H2 dot-plot variance and median: separate level regressions

	H2 dot variance		H2 median dot	
Forecast	Previous dots	NatWest mock	Previous dots	NatWest mock
Forecast coefficient	0.785	0.790	0.931	0.973
	(0.139)	(0.118)	(0.047)	(0.038)
Intercept	0.041	0.063	0.223	0.179
	(0.033)	(0.026)	(0.133)	(0.105)
R^2	0.6990	0.6994	0.9017	0.9196
Observations	38	38	38	38

Notes: Each column estimates a separate regression $A_t = \alpha + \beta F_t + u_t$, where A_t is the realized H2 dot-plot variance or median and F_t is either the previous horizon-matched actual-dot statistic or the statistic implied by the corrected NatWest mock dot plot. All columns use the same 38 meetings between 2015Q4 and 2025Q4, include an intercept, and contain no additional controls. HC3 standard errors are in parentheses.

A.3 Additional Impulse Responses

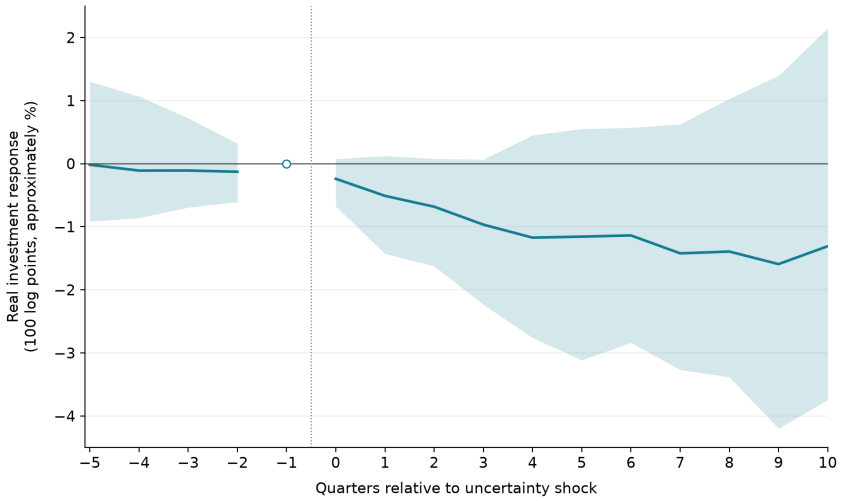


Figure A.3: IRF Nonresidential Investment Uncertainty Shock 95% Bootstrap CI

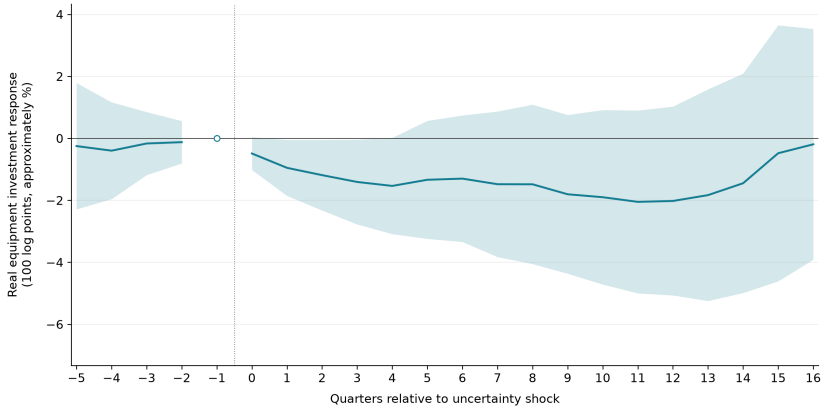


Figure A.4: IRF Investment Equipment Extended Residualized Uncertainty Shock HC3

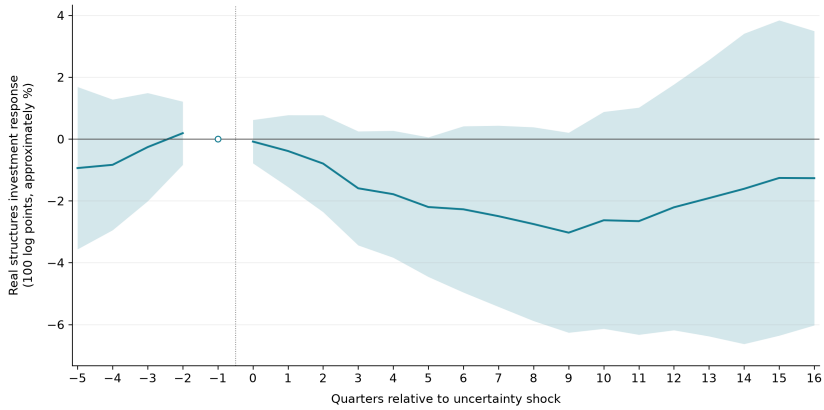


Figure A.5: IRF Investment Structures Extended Residualized Uncertainty Shock HC3

B Disagreement, learning, and the horizon of policy uncertainty

Wider disagreement can make policy more predictable by revealing whose views shape the committee's decisions. It can also make policy less predictable by widening the range of future choices. This appendix extends that tradeoff to disagreement about the neutral rate, learning from the communicated policy path, and different horizons for the dots and uncertainty. The two mechanisms survive these extensions. A negative disagreement–uncertainty relationship, and especially a particular ordering of its sign across horizons, require additional conditions.

B.1 Neutral-rate disagreement and the economic environment

Let t index SEP meetings, l the horizon of a dot, and $h \geq 1$ the horizon of the future policy decision whose uncertainty we measure. Horizon zero denotes the current decision. The market observes the released endpoint dots, the SEP forecast path, and the committee's decision and guidance. Its uncertainty concerns the influence of hawkish relative to dovish views. Subscripts $t-$ and t on expectations and variances refer to information just before and after the release; $\mathcal{F}_{t-}$ denotes the former information set.

The dovish and hawkish endpoints, $s \in \{D, H\}$, share forecasts of inflation and output gaps but have different, known reaction coefficients. Write $\tilde{X}_{t,l}$ for the two forecast gaps and $\tilde{\phi}_s$ for the coefficients. Their dots and disagreement satisfy

$$\begin{aligned} d_{s,t,l} &= \alpha_{s,t} + \tilde{\phi}_s' \tilde{X}_{t,l} + \xi_{s,t,l}, \\ \Delta_{t,l} &\equiv d_{H,t,l} - d_{D,t,l} = a_t + g_{t,l} + \zeta_{t,l}. \end{aligned} \tag{B.1}$$

Here $a_t = \alpha_{H,t} - \alpha_{D,t}$ is disagreement about the neutral nominal rate, $g_{t,l} = (\tilde{\phi}_H - \tilde{\phi}_D)' \tilde{X}_{t,l}$ is disagreement about the response to forecast gaps, and $\zeta_{t,l} = \xi_{H,t,l} - \xi_{D,t,l}$ is transitory endpoint noise. The neutral-rate gap affects every horizon. The gap-driven component varies along the forecast path. Noiseless longer-run dots reveal the intercepts, $d_{s,t,\mathrm{LR}} = \alpha_{s,t}$.

The neutral-rate gap follows $a_{t+1} - \bar{a} = \phi_a(a_t - \bar{a}) + u^a_{t+1}$, with $0 \leq \phi_a < 1$. Within the projection window $1 \leq h \leq L$, the market adopts the SEP forecasts, so $E_t[g_{t+h,0}] = g_{t,h}$. Beyond L, it extrapolates the last forecast toward its mean at rate $\phi_x \in [0, 1)$. Expected future disagreement is therefore

$$m_{t,h} \equiv E_t[\Delta_{t+h,0}] = \bar{a} + \phi_a^h(a_t - \bar{a}) + E_t[g_{t+h,0}]. \tag{B.2}$$

The horizon-h dot gap forecasts future near-term disagreement after removing transitory

noise and allowing the neutral-rate gap to mean-revert. Persistence is thus a property of the sources of disagreement and the forecast path; a single autoregressive coefficient for squared dispersion is a reduced-form approximation.

Assumptions and scope. The Gaussian specification is a local approximation around a positive endpoint gap. Endpoint views and their forecast errors are independent of the latent committee weight and of communication noise, conditional on pre-meeting information. Neutral-rate innovations, gap-forecast revisions, and endpoint noise are mutually uncorrelated; endpoint noise is independent across horizons and meetings. Communication errors are independent across the observed horizons. Forecast-error variances and covariances are deterministic and do not vary with the released dots. These restrictions let the endpoints determine signal precision without directly revealing the weight, and isolate learning from changes in fundamental volatility. The results below are conditional on these restrictions.

B.2 Learning from the communicated path

The committee's decision and communicated path aggregate endpoint views using a common weight:

$$\begin{aligned} i_{t,l} &= d_{D,t,l} + w_t \Delta_{t,l} + \eta_{t,l}, \\ w_t - \bar{w} &= \rho(w_{t-1} - \bar{w}) + \nu_t, \qquad 0 < \rho < 1. \end{aligned} \tag{B.3}$$

The Gaussian innovations have variances $\sigma^2_{\eta,l} > 0$ and $\sigma^2_\nu \geq 0$. For $l = 0$, $i_{t,0}$ is the current decision. At longer horizons, it represents a communicated rate path. A median dot can proxy this path if its position among the endpoints reflects committee influence; this is an aggregation assumption, not a consequence of observing the endpoints.

Let P_t^- denote the pre-meeting variance of w_t, and let $\mathcal{L}_t$ contain the horizons at which the market observes an informative path signal. Given the endpoints, each signal is $i_{t,l} - d_{D,t,l} = w_t \Delta_{t,l} + \eta_{t,l}$. The posterior variance is

$$P_t = \left[(P_t^-)^{-1} + Z_t^{\text{eff}}\right]^{-1}, \qquad Z_t^{\text{eff}} = \sum_{l \in \mathcal{L}_t} \frac{\Delta^2_{t,l}}{\sigma^2_{\eta,l}}. \tag{B.4}$$

At each observed horizon, the relative position $(i_{t,l} - d_{D,t,l})/\Delta_{t,l}$ estimates the weight with error variance $\sigma^2_{\eta,l}/\Delta^2_{t,l}$. A wider range makes a given communication error less consequential. The posterior variance falls with Z_t^{eff}; the realized position of the path determines the posterior mean.

Path signals allow learning when the current decision conveys little information. At a binding lower bound, we omit the censored current decision from $\mathcal{L}_t$ and retain informative

future horizons. This preserves the linear filter. A censored decision can still convey information when some endpoint prescriptions exceed the floor, so omission understates learning in that case. A weaker information channel at the lower bound requires less informative remaining signals, for example because their communication noise is larger. The lower bound alone does not determine the sign of the total dispersion slope.

Forecastable disagreement. Let $U_{t,h} = \mathrm{Var}_t(i_{t+h})$ denote future policy uncertainty. Set $V_h = \mathrm{Var}_t(\Delta_{t+h,0})$. Future weight uncertainty is $\rho^{2h}P_t + C_h$, where $C_h = \sigma_\nu^2(1-\rho^{2h})/(1-\rho^2)$ accumulates innovations after today's meeting. Averaging over the decision and path realizations conditional on the released endpoint information gives

$$E[U_{t,h} \mid \text{endpoints and forecasts}, \mathcal{F}_{t-}] = K_{t,h} + \underbrace{m_{t,h}^2(\rho^{2h}P_t + C_h) + V_h C_h}_{G_{t,h}}. \tag{B.5}$$

Here $K_{t,h}$ is known before the release and independent of the dots. Increasing $m_{t,h}^2$, holding precision fixed, raises uncertainty. Increasing Z_t^{eff}, holding the forecast fixed, lowers it at rate $-\rho^{2h}m_{t,h}^2P_t^2$. These are the persistence and information channels.

The learning term depends on $(E_t\Delta_{t+h,0})^2$, not $E_t[\Delta_{t+h,0}^2]$. Learning reduces the variance of the weight but increases the dispersion of its posterior mean by the same amount in expectation. The contribution of uncertain future disagreement is therefore independent of signal precision after averaging over signals. This is the posterior-mean correction in the main paper's full slope.

At a deterministic steady state with zero forecast gaps, common neutral rates, and no endpoint noise, all signals have zero loadings: the weight neither matters for policy nor can be learned. Neutral-rate disagreement restores both learning and forecastable disagreement even with the economy at target. Transitory endpoint noise alone can reveal the weight, but when $m_{t,h} = 0$ that learning has no expected effect on forward uncertainty. Future macroeconomic shocks can make the weight matter for future policy without making today's decision informative.

B.3 Dispersion surprises at different horizons

The event-date outcome is $\Delta U_{t,h} = U_{t,h} - \mathrm{Var}_{t-}(i_{t+h})$. The regressor in the model is the surprise in squared endpoint disagreement, $S_{t,l} = \Delta_{t,l}^2 - E_{t-}[\Delta_{t,l}^2]$. The following slopes are first-order approximations around a common positive gap $\Delta_{t,l} = m_{t,h} = \mu_\Delta > 0$, conditional on the prior. They are not global sign restrictions for arbitrarily large dispersion changes.

A level surprise $e_{t,l} = \Delta_{t,l} - E_{t-}\Delta_{t,l}$ combines a neutral-rate innovation, a gap-forecast

revision, and endpoint noise. Let ω_l^a, ω_l^g, and ω_l^ζ denote their shares of its variance. Define $b_{h,l}$ as the projection coefficient of the horizon-h gap-forecast revision on the horizon-l revision, with $b_{l,l} = 1$. The loading of forecastable future disagreement on the observed surprise is

$$\kappa_{h,l} = \frac{\mathrm{Cov}_{t-}(m_{t,h} - E_{t-}m_{t,h}, e_{t,l})}{\mathrm{Var}_{t-}(e_{t,l})} = \omega_l^a \phi_a^h + \omega_l^g b_{h,l}. \tag{B.6}$$

Endpoint noise contributes no forecast content. For $h > L$, the gap-revision loading is $b_{h,l} = \phi_x^{h-L} b_{L,l}$. A nonnegative loading that declines beyond the dot horizon is an additional restriction on forecast revisions, not a property of Gaussian learning alone.

The surprise also changes the precision of path signals. Let $b_{j,l}^\Delta$ be the projection coefficient of $e_{t,j}$ on $e_{t,l}$, describing its co-movement with the gap at observed horizon j. Let $\bar{P}$ be the posterior variance at the evaluation point. Define

$$r_l = \mu_\Delta^2 \bar{P} \sum_{j \in \mathcal{L}_t} \frac{b_{j,l}^\Delta}{\sigma_{\eta,j}^2}, \qquad W_h = \rho^{2h}\bar{P} + C_h, \qquad \theta_h = \frac{\rho^{2h}\bar{P}}{W_h}. \tag{B.7}$$

The learning exposure r_l measures how the horizon-l surprise changes information across all observed signals. With only a signal at l, it equals the fraction of prior weight variance resolved at the evaluation point. With several signals it is a co-movement-weighted exposure, lying between zero and one if all relevant $b_{j,l}^\Delta$ lie in $[0, 1]$. The share θ_h is the fraction of forward weight variance attributable to today's unresolved weight, rather than future innovations.

The local slope of $\Delta U_{t,h}$ on $S_{t,l}$ is

$$\beta_{h,l} = \underbrace{W_h \kappa_{h,l}}_{\text{persistence}} - \underbrace{\rho^{2h}\bar{P} r_l}_{\text{information}} = W_h(\kappa_{h,l} - \theta_h r_l). \tag{B.8}$$

Thus a negative slope requires $\kappa_{h,l} < \theta_h r_l$: the surprise must improve learning enough to offset its forecast content. The two signed contributions require $\kappa_{h,l}, r_l \geq 0$. If cross-horizon revisions are negatively related, the formula still applies but the usual positive-persistence, negative-information interpretation need not apply to that regressor.

When learning uses only the current decision and the regressor uses distant dots, information passes through $b_{0,l}^\Delta$. A distant dispersion surprise unrelated to current disagreement has no information effect through that decision. Observing the communicated path at the dot horizon adds a direct signal. The empirical horizon with the largest raw dispersion variation need not supply the most information, since precision and cross-horizon co-movement also matter.

What the source of disagreement changes. A neutral-rate surprise moves every endpoint gap and carries forecast loading ϕ_a^h. Its slope is $W_h(\phi_a^h - \theta_h r)$, where $r = \bar{P}\bar{Z}^{\text{eff}} \in [0, 1)$ is the resolved share at the evaluation point. Its sign remains ambiguous. For any fixed horizon, sufficiently persistent neutral-rate disagreement makes the slope positive.

A gap-forecast revision confined to an observed horizon l raises uncertainty at the matched horizon $h = l$ and lowers it at other horizons within the projection window, provided learning about today's weight matters there. At the matched horizon its forecast loading is one, exceeding its discounted information exposure. Revisions that co-move across horizons instead use the full $b_{h,l}$ profile; the confined-revision result is not a general ranking of regression slopes.

Transitory endpoint noise lowers expected uncertainty if it widens an observed signal's loading without changing expected future disagreement. This conclusion concerns noise in the views entering the aggregation equation. Pure measurement error added by the econometrician does not create that information. Under the source-independence assumptions, the total local slope is the variance-share-weighted average of the three source-specific slopes. Its sign alone does not identify those shares or establish that neutral-rate slopes exceed all other slopes.

B.4 When sign monotonicity survives

Fix the dot horizon l, the observed signals, and the prior. If there are no future weight innovations, $C_h = 0$ and $\theta_h = 1$. Equation (B.8) then has the sign of $\kappa_{h,l} - r_l$. Provided $b_{h,l}$ is nonnegative and nonincreasing for $h \geq l$, the forecast loading declines there. The sign can move from positive to negative, but cannot move back. This is the counterpart of the main paper's main-channel result. It is a restriction on signs, not on the magnitudes of negative coefficients.

Horizon alignment can break the prediction. Below the dot horizon, forecast revisions can become more relevant as h approaches l. A gap revision confined to an observed dot horizon gives a direct counterexample: its uncertainty slope is negative before that horizon, positive at it, and negative afterward within the projection window. With correlated revisions, a reversal occurs whenever $b_{h,l}$ crosses the information threshold from below. A hump-shaped loading permits this pattern; it does not guarantee a crossing.

Consequently, comparing twelve- and twenty-four-month uncertainty slopes on two-year-ahead dots does not by itself test the main-channel monotonicity prediction. The shorter uncertainty horizon lies before that dot horizon. Calendar year-end dots must be matched to the actual maturity of the uncertainty measure. A suitable comparison uses a dot horizon

at or below every uncertainty horizon considered, together with the required restrictions on forecast loadings.

Future weight innovations impose a further condition. When $\sigma_\nu^2 > 0$, today's learning becomes a smaller fraction of future weight uncertainty. For positive loadings, a sufficient condition for sign monotonicity is

$$\frac{\kappa_{h+1,l}}{\kappa_{h,l}} \leq \frac{\theta_{h+1}}{\theta_h} = \frac{\rho^2 W_h}{\rho^2 W_h + \sigma_\nu^2} \quad \text{at every horizon being compared.} \tag{B.9}$$

This condition makes $\kappa_{h,l}/\theta_h$ nonincreasing, so it cannot recross the fixed threshold r_l from below. It is sufficient, not necessary: a slope may keep the same sign even if the ratio increases. A sufficient primitive restriction bounds both active component decay rates, ϕ_a and $b_{h+1,l}/b_{h,l}$, by the right-hand side. Merely requiring disagreement to mean-revert does not suffice.

The neutral-rate component gives a second counterexample. If $\sigma_\nu^2 > 0$ and $\phi_a > \rho^2$, its persistence term eventually dominates: W_h approaches a positive constant while the learning term decays as ρ^{2h}. Its slope is therefore positive at sufficiently long horizons, even if negative initially. If $\phi_a < \rho^2$ and $r > 0$, the pure neutral-rate slope is eventually negative; equality depends on the coefficients. These are statements about the eventual sign, not its size. The two channels survive future weight innovations, but the one-direction sign prediction does not survive without the relative-decay restriction.

B.5 Separate committee weights

A committee may put different weights on neutral-rate views and on responses to macroeconomic gaps. Let $w_t = (w_t^a, w_t^g)'$ and $v_{t,l} = (a_t + \zeta_{t,l}, g_{t,l})'$, so the signal becomes $i_{t,l} - d_{D,t,l} = v_{t,l}' w_t + \eta_{t,l}$. Each weight has its own autoregressive persistence, collected in $R = \mathrm{diag}(\rho_a, \rho_g)$; their innovations may be correlated. The posterior covariance satisfies

$$P_t^{-1} = (P_t^-)^{-1} + \mathcal{Z}_t, \qquad \mathcal{Z}_t = \sum_{l \in \mathcal{L}_t} \frac{v_{t,l} v_{t,l}'}{\sigma_{\eta,l}^2}. \tag{B.10}$$

A single signal reveals one combination of the weights. Separate information about both dimensions within a meeting requires noncollinear loading vectors across observed horizons. Prior correlation can also transmit learning from one weight to the other.

Let $m_{t,h}$ now denote the vector of expected future component disagreements. The forecastable part of uncertainty becomes $m_{t,h}'[R^h P_t R^h + C_h] m_{t,h}$, where C_h is the covariance

accumulated from future weight innovations. Holding the forecast vector fixed, an increase in information in the positive-semidefinite matrix sense lowers this uncertainty. Scaling a signal's loading vector along its existing direction supplies such an increase. Raising total dot dispersion while changing the mix of neutral-rate and gap disagreement need not do so. Likewise, the positive quadratic form rises with a proportional expansion of forecastable disagreement, but need not rise with every change in an individual component when covariances matter.

The two-channel intuition therefore carries over as a comparison between forecastable policy differences and information in relevant directions. The scalar threshold and source or horizon rankings do not automatically carry over. At a deterministic steady state with zero gap forecasts, only the neutral-rate weight directly affects policy and is directly observed through the signals; the reaction weight can still be updated indirectly through prior correlation.

B.6 Implications for the empirical interpretation

The model concerns conditional variances and surprises in squared endpoint gaps. Cross-sectional dot variance equals the squared range times the variance of dots' normalized positions inside that range. A stable distributional shape therefore preserves the sign mapping; changes in shape need not. Range-based estimates offer the most direct counterpart. Mock dot plots proxy pre-meeting expected dispersion, while changes from the preceding release also contain a predictable component. That component has no surprise content conditional on pre-meeting information, but its pooled regression coefficient need not vanish if it co-varies with the expected amount of uncertainty resolved at a meeting.

The model also separates an announcement's average effect from the dispersion slope. The law of total variance implies $E_{t-}[\Delta U_{t,h}] = -\operatorname{Var}_{t-}(E_t[i_{t+h}]) \leq 0$. An announcement resolves uncertainty on average even when the slope on wider-than-expected dots is positive. The slope describes variation around that expected resolution, not the sign of every meeting's change.

A negative estimated slope is consistent with learning dominating forecastable disagreement. Interpreting differences across horizons requires aligned maturities and restrictions on the persistence and co-movement of the underlying sources. Interpreting differences across states requires attention to both signal precision and forecastable disagreement: a larger prior variance, a larger macroeconomic gap, or the lower-bound regime can change both channels. None alone guarantees a more negative total coefficient. These distinctions preserve the model's economic explanation while limiting its empirical predictions to the

conditions under which they follow.