Hawks, Doves, and the Dot Plot: Heterogeneous FOMC Reaction Functions*

Vladislav Abramov†

August 12, 2026

Abstract

The dots in the Federal Reserve's Summary of Economic Projections (SEP) reveal that FOMC members disagree about the appropriate policy rate, but not what they disagree about. Using members' individual projections from 2012 to 2019, I estimate each member's policy reaction function and decompose the cross-member dispersion in the dot plot into differences in economic outlook and differences in reaction functions. The composition of this disagreement shifts systematically with the forecast horizon: at the nearest horizons it reflects mainly differences in how members would respond to a given outlook, at intermediate horizons differences in the outlook itself, and at the longest horizons differences in members' perceived long-run neutral rate. Reaction-function differences therefore dominate at both ends of the horizon spectrum, including the near-term dots most relevant for the coming decision. These differences are large enough to move the implied rate path by 25 to 75 basis points, are systematically related to observable characteristics such as experienced inflation and educational background, and help predict dissenting votes. The individual dots thus carry behaviorally meaningful information about how members would set policy that the median projection discards.

Keywords: monetary policy rules, FOMC members' preferences, dissent voting, dot plot

JEL Classification: E52, E58, E43

*I am grateful to Tatevik Sekhposyan, Joon Jo, Pedro Bento, Sarah Zubairy, and all participants of the Texas A&M Macro Reading Group, the North Texas Economic Conference, and the Southern Economic Association for helpful suggestions and comments.

†Department of Economics, Texas A&M University, 2935 Research Parkway Suite 200, College Station, TX 77843, United States. E-mail: vabramov@tamu.edu.

1 Introduction

Monetary policy in the United States is conducted by the Federal Open Market Committee (FOMC), a body composed of Board of Governors members and regional Federal Reserve Bank presidents. There are eight regular policy meetings per year, and twelve voting members who make policy decisions. While the FOMC operates under a shared mandate to pursue price stability and maximum employment, each member brings unique experiences, educational and professional backgrounds that can influence how aggressively—or cautiously—they respond to changes in inflation or labor market conditions. Since 2012, these differences have been on public display: four times a year, each participant's projected path for the federal funds rate is published as a dot in the Summary of Economic Projections (SEP), and the spread of those dots is widely read as a measure of disagreement within the Committee.

The dot plot reveals that members disagree, but it is silent on what they disagree about. Two participants can project different rates for two very different reasons: they may foresee different economies, or they may respond differently to the same economy. The first is a difference in beliefs about the outlook; the second, a difference in reaction functions. The distinction is not merely academic—the value of the dot plot is now contested at the highest level of the Federal Reserve. Chair Warsh, a longstanding skeptic of forward guidance, declined to submit his own projection at his first meeting in 2026 and opened a review of the Committee's communications, amid reports that the dot plot could be pared back or dropped altogether.[1] A prominent strand of the communications literature argues the opposite: the individual projections are valuable precisely because they embed each participant's reaction function, and the SEP's real weakness is that it reveals too little of it—the rate each member expects, but not the responses to inflation and unemployment that lie behind it (Cecchetti and Schoenholtz, 2026). Whether the dispersion in the dots reflects genuine differences in how members would set policy, or merely different readings of the same economy, is the question this paper takes up.

To answer it, I estimate each member's policy reaction function directly from their own SEP projections, in the spirit of Taylor (1993), relating each participant's projected policy rate to the inflation and unemployment gaps they themselves forecast.[2] Because the coefficients are identified from a member's own forecasts, they isolate differences

[1]On Warsh's skepticism toward forward guidance, see his testimony at the U.S. Senate Committee on Banking, Housing, and Urban Affairs nomination hearing, April 21, 2026 (C-SPAN). Warsh did not submit a federal funds rate projection at the June 2026 FOMC meeting; see CNBC, June 17, 2026.

[2]I assume that FOMC members assign different relative weights to inflation versus unemployment stabilization, so that there exists a trade-off between the two objectives, and that how each member perceives and navigates this trade-off reflects their individual policy stance.

in reaction functions from differences in beliefs about the economy. Building on these estimates, my central contribution is to decompose the cross-member dispersion in the dot plot into two channels: a *forecast channel*, reflecting that members project different outlooks, and a *reaction-function channel*, reflecting that they respond differently to a given outlook. The composition of disagreement shifts sharply with the forecast horizon. At the nearest horizons, where members largely agree on the near-term economy, most of the dispersion in the dots reflects differences in how they would respond to it; at intermediate horizons the outlook itself becomes contested and forecast differences dominate; and at the longest horizons disagreement is again about the reaction function, now operating through differences in the perceived long-run neutral rate. This structure is invisible in the median dot, and it implies that the near-term dots—those most relevant for the upcoming decision—carry the most information about how members differ in setting policy.

The individual projections cover headline and core inflation, real GDP growth, the unemployment rate, and the federal funds rate, at horizons from the current year to the long run; the long-run projections let me construct each member's forecasted inflation and unemployment gaps. The first SEP projections were published after the Fed's October 2007 meeting, and the Fed has since treated them as an important component of its forward guidance.[3] Because the Fed releases the key linking projections to named participants only with a five-year lag, I merge the projections with these keys to build a panel of individual forecasts spanning 2007–2019.

The decomposition rests on the estimated reaction functions, which differ substantially across members. Estimating them for the 24 members with sufficiently long tenure in the sample,[4] I find that the average marginal effects (AMEs) of the inflation and unemployment gaps vary widely—including within the same Reserve Bank district and among Governors. A one-percentage-point increase in the forecasted unemployment gap is associated with a change in the projected federal funds rate ranging from about -1.8 to -0.2 percentage points across members, and a comparable increase in the inflation gap with a change from about 0 to 3.1 points. This heterogeneity is robust across estimation methods and does not hinge on any single member.

These differences are consequential rather than cosmetic. In counterfactual exercises, had all voting members shared the Chair's reaction function while keeping their own forecasts, the implied consensus rate would differ from the actual one by 25–75 basis points in most periods; replacing Jerome Powell's reaction function with Janet Yellen's, and vice versa, while holding forecasts fixed, moves the implied rate by 25–50

[3] See `https://www.federalreserve.gov/newsevents/speech/bernanke20071114a.htm`.

[4] A total of 43 FOMC members served between 2007 and 2019, but many served only brief terms.

basis points. Variation in reaction functions, not just in the macroeconomic outlook, materially shapes the projected policy path.

To characterize the reaction-function differences along a single dimension, I construct a continuous index—the *Hawk* measure—equal to the sum of a member's inflation- and unemployment-gap AMEs, capturing the relative weight they place on inflation versus unemployment. Unlike the classification in Istrefi (2019), which sorts members into hawks, doves, and swingers from textual analysis of news coverage, this measure is built entirely from members' own economic projections and places them on a continuous spectrum rather than in discrete groups.[5] The measure may also reflect the economic environment, the influence of other members, and factors unrelated to intrinsic preferences (Bobrov, Kamdar, and Ulate 2025, Fos and Xu 2025); I discuss the extent to which it can be read as preferences in Section 4, and a full causal separation of intrinsic preferences from institutional and regional influences is beyond the paper's scope.

If the reaction-function differences were mostly estimation noise, they would not align with anything observable. I show that they do. Whether members' characteristics explain their policy stances has a substantial prior literature: Malmendier, Nagel, and Yan (2021) find that lifetime inflation experiences shape hawkishness; Bordo and Istrefi (2023) link stances to birth cohort and education, with hawks more often trained at "freshwater" universities and doves at "saltwater" ones;[6] Smales and Apergis (2016) emphasize tenure in the Federal Reserve System and on the FOMC; and Bennani, Farvaque, and Stanek (2018), Bobrov, Kamdar, and Ulate (2025) document heterogeneity tied to regional conditions and personal background, inferring each member's desired rate from meeting transcripts, whereas I estimate the full reaction function from published projections. Relating the estimated reaction functions to a set of these characteristics—institutional position, voting status, Fed experience, education, and research background—I find that a small number of them account for a substantial share of the cross-member variation. The reaction-function differences therefore reflect systematic, structural differences among members, not idiosyncratic noise.

The differences are also behaviorally meaningful. The rationale for FOMC dissent

[5] Istrefi (2019) classifies FOMC members into hawks (who emphasize inflation), doves (who emphasize economic activity), and swingers (who lack a strong inclination), and show that this classification is related to members' preferred rates and helps predict dissenting votes. My work is also complementary to, but distinct from, approaches that infer objectives from policymakers' sentiment (Shapiro and Wilson, 2022) and that examine the perceived monetary policy rule of the Fed using micro data on professional forecasters (Bauer, Pflueger, and Sunderam, 2024).

[6] "Freshwater" and "saltwater" schools refer to differing macroeconomic traditions, with the former emphasizing market efficiency and policy rules (e.g., University of Chicago, University of Minnesota) and the latter favoring Keynesian approaches and active stabilization policy (e.g., Massachusetts Institute of Technology, Harvard University).

has been studied since the early 1990s (Belden 1991, Havrilesky and Gildea 1991), with dissent often attributed to inherent preferences for tighter or looser policy, though it may also stem from strategic, political, or regional considerations (Meade and Stasavage 2008, Bobrov, Kamdar, and Ulate 2025, Tillmann 2011, Kocherlakota 2025). I find that dissent is systematically related to the estimated reaction functions: controlling for individual characteristics and time fixed effects, members with more inflation-oriented reaction functions are significantly more likely to cast a hawkish dissent and less likely to cast a dovish one. I also find limited evidence of strategic forecasting—members do not systematically distort their projections when their voting status changes—so the estimated differences are unlikely to be artifacts of members gaming the projections.

Taken together, these results answer the question with which I began. The dispersion in the dot plot is not merely different readings of the economy: a large and systematic part of it reflects genuine differences in how members would respond to a shared outlook, differences that are tied to observable characteristics, move the projected policy path, and help predict dissent. The cross-section of dots therefore carries structured, behaviorally meaningful information that the median discards—direct evidence, in precisely the dimension the reform debate turns on, that setting the individual projections aside would sacrifice signal, not merely trim noise.

The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 estimates individual reaction functions, documents their heterogeneity and its consequences for the projected rate path, and decomposes the disagreement in the dot plot into forecast and reaction-function channels. Section 4 constructs the *Hawk* measure, examines which characteristics explain the reaction-function differences, and relates them to members' forecasts and dissenting votes. Section 5 concludes.

2 Data

The Summary of Economic Projections (SEP) is a relatively recent data source, with the first release dating back to October 2007. It is the only comprehensive dataset that provides individual projections from Federal Reserve policymakers, covering multiple economic variables and forecast horizons. Because of this, the SEP is widely followed by analysts and financial markets as a key indicator of the FOMC's policy outlook and macroeconomic expectations.

The most recent SEP publications do not identify individual contributors; instead, they present the distribution of forecasts across participants without revealing names or affiliations. The Federal Reserve releases participant identifiers with a considerable lag. At the time of writing, individual-level SEP data are available for the periods 2007–2020. Data after 2020 remain unavailable due to the Federal Reserve's five years publication lag policy.

The data structure includes the FOMC member's name along with forecasts of the federal funds rate, headline and core inflation, real GDP growth, and the unemployment rate. With the exception of long-term projections, all forecasts are made for the fourth quarter of the corresponding year. Typically, forecasts are provided for five horizons: the fourth quarter of the current year t, and the fourth quarter of years $t+1$, $t+2$, $t+3$, and the long run. Note that SEP data can be partially extended to the past, as some of the individual projections are available before 2007, for example, the Monetary Policy Reports in Romer (2010). However, as far as the author knows, no datasets exist with individual-level projections of the policy rate before 2007. For the purposes of this paper, it is essential to have individual policy rate forecasts, thus I focus exclusively on the period where SEP forecasts were available. In particular, although individual inflation and unemployment projections are available back to the early 1990s (Romer, 2010), the individual federal funds rate projections that serve as my dependent variable were first published only in 2012, which sets the earliest possible start of the estimation sample.

The SEP dataset is an unbalanced panel. Although the sample began relatively recently, in 2007, only 2 out of 43 FOMC members were present throughout the entire sample period. Each year, some members leave the committee as their terms end, and new members replace them. Additionally, the annual rotation of voting rights among regional Federal Reserve Bank presidents contributes further to the panel's unbalanced nature. As a result, many policymakers have relatively few observations. Because estimating FOMC reaction functions relies on individual-level data, these limitations may lead to imprecise or unstable estimates.

To address this issue, I require each forecaster to have at least 40 time-horizon pairs of forecasts, see Figure 1. This corresponds to roughly two years of forecasts per member. Applying this requirement reduces the sample to 24 members. Some members were present in the FOMC for just a couple of meetings, and thus, I do not have a sufficient number of forecasts for them. To ensure stable and interpretable estimates of individual policy rules, I impose a minimum threshold of 40 observations per member. This balances the trade-off between sample size and estimation precision, retaining most members while maintaining a sufficient number of data points for robust inference.[7]

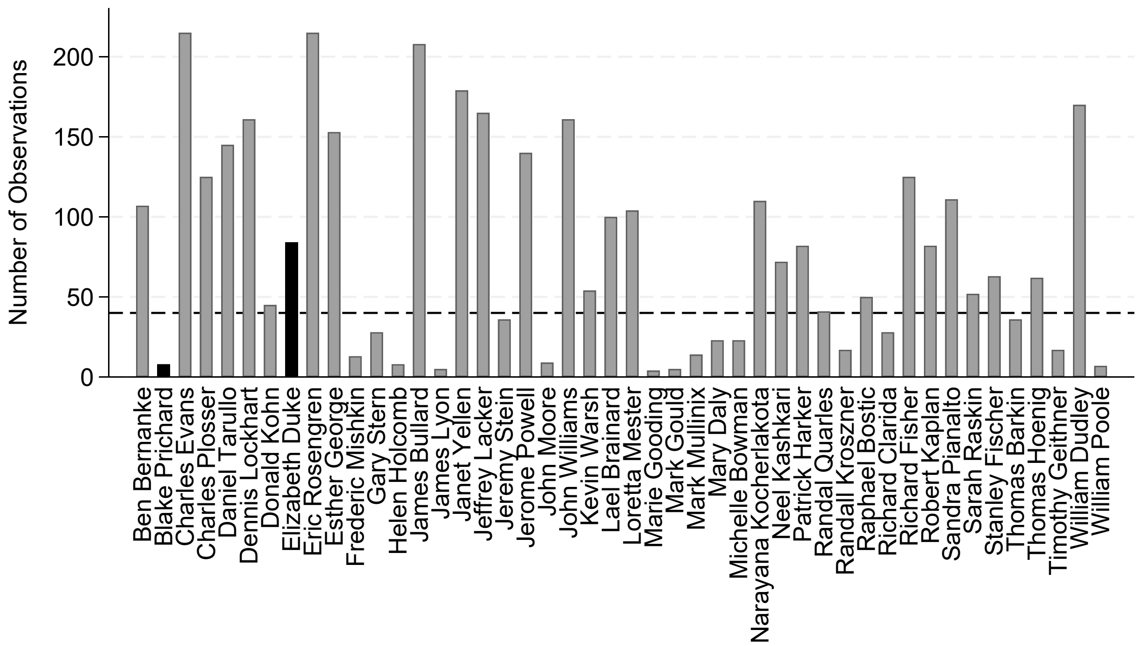

Figure 1: Sample Coverage: Number of Observations for Each FOMC Member

Figure 2 displays individual projections of the federal funds rate — the interest rate on overnight lending between depository institutions and the Fed's primary tool for managing short-term interest rates in the economy. Popular among financial analysts, interest rate forecasts — commonly referred to as "dot plots" due to their visual presentation — were first included in SEP publications in 2012 as part of the Fed's broader effort to increase transparency and provide more forward guidance to the public and financial markets.[8] For much of the early sample the federal funds rate was held at its effective lower bound (ELB), a target range of 0–0.25 percent whose midpoint of 0.13% participants report in the SEP.

[7]Thresholds of 35, 45, and 50 observations produce similar results.

[8]However, individual forecasts of other variables are available before 2012.

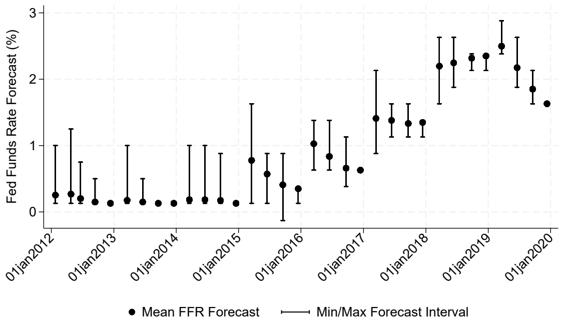

Figure 2: Federal funds Rate Forecasts

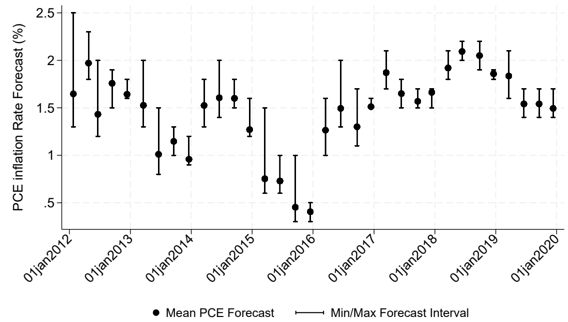

Figure 3: PCE Inflation Rate Forecasts

Many projections sit exactly at this bound, but they are not spread evenly through the data. Censoring falls almost entirely on the near horizons: 39.9% of current-year projections are at the ELB, against 24.1% one year out, 8.9% two years out, and essentially none at three years or in the long run. Overall, 21.0% of the estimation sample is censored [9]. Censoring is also concentrated in the early years of the sample: between 89% and 97% of current-year projections are at the bound in 2012–2014, while no projection at any horizon is censored in 2016–2019 (Table 1). This clustering at the lower bound calls for a censored modeling approach, which I implement in the next section. It also has a consequence for identification, which I return to in Section 4: members who served only after liftoff contribute no censored observations, while those serving in 2012–2014 contribute a policy rate forecast that barely moves.

3 Reaction functions of FOMC members

Answering whether the disagreement in the dot plot reflects differences in beliefs or in reaction functions requires, first, a measure of each member's reaction function. This section estimates those functions from members' own projections, documents how much they differ and what that implies for the projected policy path, and then decomposes the dispersion in the dots into a forecast and a reaction-function component.

3.1 Empirical specification

Individual reaction functions can be described by a policy rule that links the policy rate (federal funds rate) to macroeconomic variables in the Fed's mandate, Taylor (1993). In this paper, I take a similar approach and estimate individual-level policy

[9] In the September 2015 SEP, Narayana Kocherlakota (Federal Reserve Bank of Minneapolis) submitted a federal funds rate projection of −0.13% — the midpoint of a −0.25 to 0 percent target range — for both end-2015 and end-2016. These are the only sub-zero projections in the sample.

Table 1: SEP forecasts: Effective Lower Bound

Year	Meetings	Obs.	Members	At ELB (% of obs.)	Unemp. forecast Mean	Unemp. forecast Max
2012	5	311	19	62.1	7.44	8.60
2013	4	250	19	54.0	6.49	7.60
2014	4	232	19	31.0	5.50	6.50
2015	4	238	20	4.6	4.93	5.50
2016	4	238	17	0.0	4.62	5.00
2017	4	227	20	0.0	4.19	4.80
2018	4	222	19	0.0	3.64	4.40
2019	4	238	17	0.0	3.71	4.20
2012–2019	**33**	**1,956**	**37**	**21.0**	**5.18**	**8.60**
2020 *(excluded)*	3	187	17	**93.0**	5.81	**14.00**

Notes: Each cell is computed over individual projections, pooling all forecast horizons in the year (target years t through $t+3$; long-run projections are excluded). *Obs.* counts these projections, and *Meetings* and *Members* the distinct meetings and participants they come from; the total row therefore counts distinct participants over 2012–2019 rather than summing the annual figures. *At ELB* is the share of these projections at the effective (zero) lower bound. *Mean* and *Max* unemployment are computed over the unemployment forecasts attached to the same projections.

rules using SEP forecasts.[10]

The ELB imposes a natural censoring constraint on the policy rate. Short-term forecasts of the federal funds rate were near the zero level throughout more than a third of the sample, as can be seen in Figure 2. To address this constraint, I estimate individual policy rules using a left-censored Tobit model — a standard approach for modeling monetary policy during periods when the policy rate is constrained by the ELB (e.g., Huang 2015, Kim and Pruitt 2017, Mavroeidis 2021, Arai 2023).[11] Ultimately, the Tobit model is preferred because it provides consistent estimates by properly handling the censored nature of the dependent variable.

[10]To my knowledge, this is the first paper to estimate and analyze individual policy reaction functions using FOMC forecast data. Fendel and Rülke (2012), Arai (2023), and González-Astudillo and Tanvir (2024) use similar data (Summary of Economic Projections, or the Monetary Policy Reports), but estimate a single reaction function for the specified group of members, and therefore lack any insights on heterogeneity of reaction functions inside Fed.

[11]The Tobit model accounts for censoring in the dependent variable. While OLS estimates differ somewhat, the unemployment gap coefficient seems to be biased toward zero for most of the FOMC members, the qualitative conclusions of the paper remain robust, see Figure 1 in the Online Appendix A.

I estimate the following Tobit model:

$$\widehat{FFR}_{it}^{h} = \begin{cases} \widetilde{FFR}_{it}^{h}, & \text{if } \widetilde{FFR}_{it}^{h} > \tau_{\text{lower}} \\ \tau_{\text{lower}}, & \text{if } \widetilde{FFR}_{it}^{h} \leq \tau_{\text{lower}} \end{cases}, \tag{1}$$

where the latent policy rate forecast is given by:

$$\widetilde{FFR}_{it}^{h} = X_{it}^{h}\beta_i + \epsilon_{it}^{h} \equiv \alpha_i + \alpha_t + \gamma_i \cdot (\widehat{UN}_{it}^{h} - \widehat{UN}_{it}^{LR}) + \omega_i \cdot (\widehat{\pi}_{it}^{h} - \widehat{\pi}_{it}^{LR}) + Z_{it}\delta + \epsilon_{it}^{h}. \tag{2}$$

Throughout, i indexes FOMC members, t indexes SEP release date, and the superscript h denotes the forecast target year. I consider horizons $h \in \{0, 1, 2, 3\}$ for current, next, two and three years ahead forecasts. The latent policy rate forecast is modeled as a linear function of the inflation and unemployment gaps, reflecting the Federal Reserve's dual mandate. The observed federal funds rate forecast ($\widehat{FFR}_{it}^{h}$) equals the latent forecast ($\widetilde{FFR}_{it}^{h}$) if it is greater than the censoring threshold (τ_{lower}) imposed by the ELB, and equals the threshold otherwise. I use the long-term forecasts of each policymaker to proxy their target values. This allows me to utilize their beliefs about targets and avoid imposing fixed target values, which might not suit all policymakers. The coefficients γ_i and ω_i represent the weights on unemployment and inflation gaps, respectively, in the latent policy rule and are allowed to vary across FOMC members. Z_{it} includes a voting status dummy and its interactions with the forecasted inflation and unemployment gaps, allowing the responsiveness of forecasts to vary by voting status. The error term is assumed to be i.i.d. Normal, $\epsilon_{it}^{h} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$.

Time fixed effects α_t and member fixed effects are also included in the model, allowing the natural rate of interest to vary across both time and individuals. This contrasts with the standard formulation of the Taylor rule, which assumes a constant natural rate, see Taylor (1993). My approach is motivated by a growing body of evidence suggesting that the natural rate, while relatively stable in the short run, is subject to substantial fluctuations over time, see Holston, Laubach, and Williams 2017. This variation is also evident in the SEP forecasts of the long-run interest rate within my sample, see Figure 4. These projections can be interpreted as forecasts of the natural rate, under the assumption that inflation and unemployment gaps close in the long run. The data show that projected long-run interest rates vary significantly over time and, to some extent, across members, with the median forecast ranging from a minimum of 2.7% to a maximum of 4.25%. Assuming a constant natural rate can lead to biased estimates of the coefficients of inflation and economic activity, Trehan and Wu (2007). To address this concern, I exploit the panel structure of the data by incorporating time fixed effects, which control for shifts in the natural rate and/or policy persistence. Additionally, Figures 5 and 6 show long-run projections of the unemployment rate and inflation, which are used to construct the gaps in the reaction functions. While inflation rate long-run projections are almost always equal to the Fed's target of 2% for all members and time periods, unemployment rate long-run projections exhibit somewhat greater cross-sectional dispersion as well as greater variation over time.

The contemporaneous formulation of the policy rule is chosen because it preserves all available observations while yielding a similar pseudo R^2 to alternative forward-looking or recursive specifications.[12] Conceptually, this formulation treats forecasted inflation and real economic activity as exogenous, outside the central bank's control, reflecting the lagged effects of monetary policy. In contrast, using the current or short-run interest rate alongside long- or medium-run forecasts of inflation and unemployment likely introduces endogeneity, Shapiro and Wilson (2022). I treat inflation and unemployment expectations as predetermined inputs to the policy rule, consistent with the SEP data structure. Modeling expectation formation is beyond the scope of this paper.

The model is estimated on the sample of both voting and non-voting members. One potential concern here is that there might be systematic differences in individual forecasts submitted by a member whose voting status changes over time. For example, Tillmann (2011) found that non-voters systematically overpredict inflation relative to the consensus forecast if they favor tighter policy and underpredict inflation if they favor looser policy. Forecasts from voting members, in contrast, show no systematic relationship with policy preferences. This is evidence on strategic forecasting of FOMC members. To account for potential strategic forecasting, I include Z_{it} — a set of interactions between voting status and the explanatory variables.[13] Finally, I assume that the coefficients in the individual reaction functions remain fixed throughout a policymaker's term. This assumption may not hold strictly in reality, as there is evidence that policymakers change their views about certain policy goals and priorities over time. However, estimating individual reaction functions with time-varying parameters is implausible given sample size limitations.

[12] The structure of the SEP data complicates estimation of forward-looking policy rules, as forecasts are made in different quarters but always target the fourth quarter values.

[13] I formally test whether the deviations in forecasts from the consensus depend on the voting status in what follows, and find no evidence that voting status systematically affects the forecasts of regional Federal Reserve Bank presidents. This is further confirmed by the fact that all effects of interaction terms in Z_{it} in the latent policy rate equation are statistically insignificant and economically small.

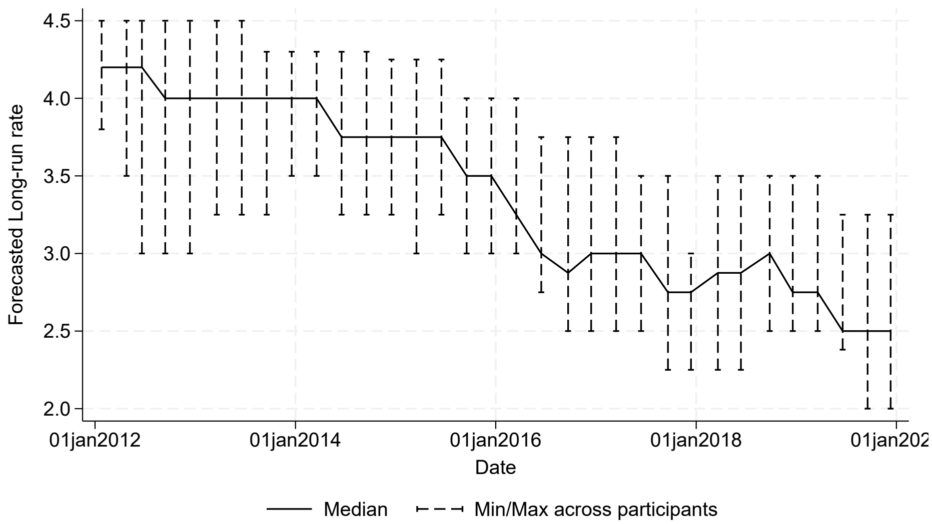

Figure 4: Long-run interest rate projections over time

Notes: Dashed intervals indicate the range between the minimum and maximum forecasts.

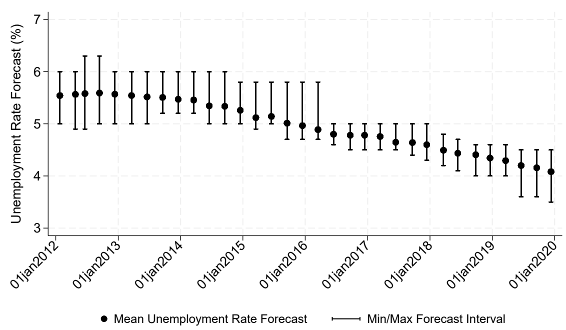

Figure 5: Long-run Unemployment Rate Forecasts

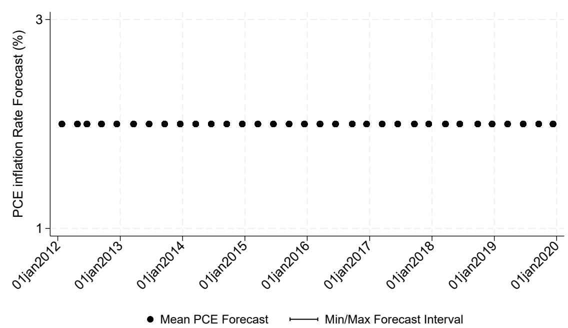

Figure 6: Long-run PCE Inflation Rate Forecasts

The model is estimated using the Maximum Likelihood Estimator. Because the regression errors may be correlated across horizons within a member–date cell (i, t), we compute standard errors clustered at the member×date level, which allows for arbitrary correlation across horizons within each cell.[14] First, define the indicator function:

$$\mathcal{I}(\widehat{FFR}_{it}^{h}) = \begin{cases} 1, & if \quad \widehat{FFR}_{it}^{h} > \tau_{\text{lower}} \\ 0, & if \quad \widehat{FFR}_{it}^{h} = \tau_{\text{lower}}. \end{cases} \tag{3}$$

[14]Almost identical confidence intervals are obtained when clustering at the meeting level. A more serious concern that could potentially lead to inconsistent estimates is heteroskedasticity. However, the standard errors with and without heteroskedasticity-robust adjustments (Eicker–Huber–White) are very similar (see Figure 5 in Online Appendix A), suggesting that the error terms are close to homoscedastic.

The likelihood function is given by:

$$\mathcal{L} = \prod_{i=1}^{N}\prod_{t=1}^{T}\prod_{h=1}^{H}\left(1-\Phi\left[\frac{X_{it}^{h}\beta_i-\tau_{\text{lower}}}{\sigma}\right]\right)^{1-\mathcal{I}(\widehat{FFR}_{it}^{h})}\cdot\left(\frac{1}{\sigma}\varphi\left[\frac{\widehat{FFR}_{it}^{h}-X_{it}^{h}\beta_i}{\sigma}\right]\right)^{\mathcal{I}(\widehat{FFR}_{it}^{h})}, \tag{4}$$

where $\Phi(\cdot)$ and $\varphi(\cdot)$ denote the standard normal cumulative distribution function and probability density function, respectively. τ_{lower} is the censoring threshold, which is treated as fixed and exogenous. In the data, the federal funds rate forecasts are measured as the midpoint of the target range. For example, if the forecasted range is [0, 0.25], the reported value of the forecast is 0.13 — the midpoint of the [0, 0.25] range, rounded to two decimal places. Accordingly, I set $\tau_{\text{lower}} = 0.13$.

3.2 Average marginal effects

I focus on the average marginal effects of the forecasted federal funds rate with respect to a unit change in either the unemployment or inflation gap. This captures how each member adjusts their policy rate forecast in response to deviations in these key inputs. These response coefficients are assumed to be fixed for each FOMC member over their tenure. Figure 7 shows estimated inflation and unemployment gap AMEs across FOMC members. As a benchmark, the dashed lines show AMEs from a pooled linear regression estimated with consensus (mean) SEP forecasts. First, for the sample of 24 FOMC members, all of the unemployment gap coefficients are negative, and all inflation gap coefficients are positive. This aligns well with the theoretical predictions, and mean AMEs are consistent with recent empirical results (on aggregated data) from the literature, see for example Carvalho, Nechio, and Tristao (2021).

Second, the magnitude of the estimated coefficients varies substantially across members. The unemployment gap coefficient ranges from -0.2 to -1.8, with an average of -1.2. This implies that a 1 percentage point increase in the unemployment gap, *ceteris paribus*, is associated with a reduction in the forecasted federal funds rate of between 0.2 and 1.8 percentage points, depending on the member, suggesting substantial variation in the perceived costs of labor market slack. For the inflation gap, the effect ranges from 0 to 3.1 percentage points, with an average of 1.6. Notably, for many FOMC members the estimates of the AMEs are outside of the consensus-based OLS estimate 95% confidence interval.

Partly, this heterogeneity could be attributed to a limited sample size - it is hard to estimate coefficients precisely if the number of observations per member is low. However, with only a few exceptions, all coefficients are statistically significant, with many coefficients also statistically different from each other. Moreover, the alternative regression with fixed weights provides a significantly worse fit to the data.[15] The goodness-of-fit measures are statistically higher for a more flexible specification.[16]. Finally, this heterogeneity is not driven by one or two members with extreme values of marginal effects. Taken together, these findings suggest meaningful intrinsic differences across FOMC members—including those within the same institutional structure, such as the Board or a specific regional Federal Reserve Bank.

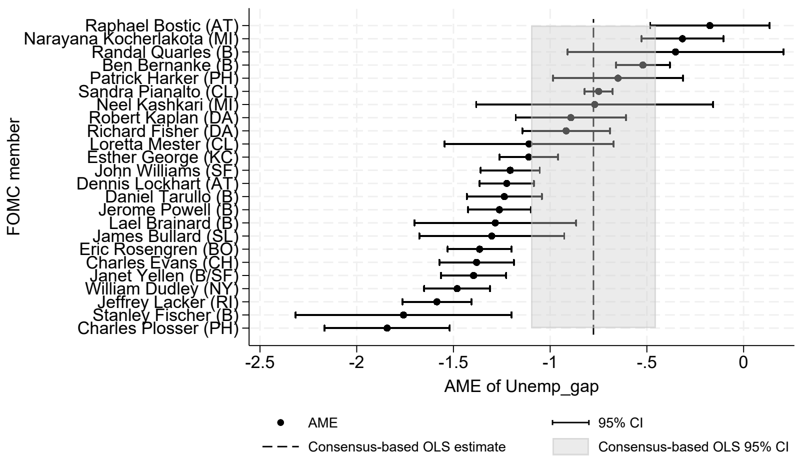

(a) Unemployment gap AMEs

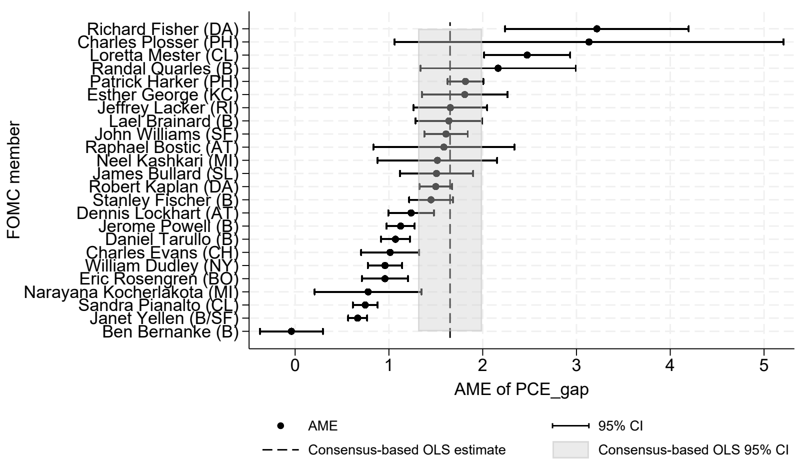

(b) Inflation gap AMEs

Figure 7: Individual policy rule AMEs

[15]The fixed weights Tobit model produces AMEs of -1.22 and 1.38 for the unemployment gap and inflation, respectively.

[16]Based on the likelihood ratio test and the Akaike Information Criterion, the model with individual coefficients fits the data significantly better than the simpler model with fixed coefficients for all members. For example, AIC is 30% lower for the more flexible specification, indicating superior performance even after penalizing for additional parameters. Both the LR test and the AIC criterion suggest that this more flexible specification is preferable over the fixed weights model specification. This means that the cross-sectional heterogeneity of AMEs is statistically large in the data.

3.3 Robustness exercises

To test the robustness of these results, I perform several exercises. First, I change the threshold on the number of observations required per FOMC member (I use 35, 40, 45, 50 values); this has a limited effect on the results. Second, I estimate the policy rules with OLS instead of the Tobit model; see Figure 1 in the Online Appendix A for the comparison. Estimated coefficients are not substantially different across the two estimation methods, and I still observe substantial heterogeneity in the estimates. Third, policymakers can form their forecasts of the federal funds rate based on the core inflation rate, which is arguably better suited for the purpose of analyzing the inflation trends. Substituting the PCE inflation gap with core PCE inflation yields similar results, with inflation point AMEs being somewhat larger for core measure for all members; see Figure 2 in the Online Appendix A. Part of it is likely mechanical: the headline gap contains a volatile food and energy component that raises its variance to more than twice that of the core gap, attenuating the estimated coefficient. Same result would be expected if participants would respond to the persistent component of inflation and look through transitory food and energy movements, so that the core gap is the variable they actually react to. Using $\pi^* = 2$ inflation target as stated in the Statement on Longer-Run Goals and Strategy[17] instead of long-run inflation forecasts and the estimate of average non-cyclical unemployment rate from U.S. Congressional Budget Office for the considered period (4.7%) instead of individual long-run unemployment forecasts also produces similar results[18]. Another robustness exercise estimates the baseline Tobit model specification where time fixed effects are replaced with the member-specific projections of the long-run real interest rate. This produces almost identical results in terms of AMEs.

Finally, SEP projections are for fixed calendar endpoints rather than fixed forecast horizons; the effective horizon shortens over the year, expanding participants' information sets. I nonetheless retain meeting fixed effects without adding horizon fixed effects. Each submission is a schedule: the rate a participant assigns to a given horizon is the one her own rule prescribes for the gaps she projects there, so the two move together along her path. A horizon fixed effect would attribute that movement to the horizon rather than to the gaps, and cannot distinguish the two — absorbing variation that may equally reflect the participant's reaction function. Since participants at the same meeting and horizon project similar outlooks, too little disagreement remains to identify individual rules once that variation is gone. The estimates should therefore be

[17] https://www.federalreserve.gov/monetarypolicy/files/fomc_longerrungoals.pdf

[18] Most members forecast around 2% inflation rate in the long run in all time periods, while variability of long run unemployment rate forecasts is also limited.

read as horizon-weighted averages of each participant's reaction function.

An interesting and closely related question concerns the degree of policy-rate persistence in individual FOMC members' forecasts. To see that, I estimate an alternative specification in which time fixed effects are replaced by the lagged federal funds rate forecast, allowing for interest-rate smoothing and gradual adjustment in projected policy paths. This model allows to capture interest rate smoothing and policy rate persistence, which might also differ across FOMC members:

$$\begin{aligned}\widetilde{FFR}_{it}^{h} &= \alpha_i + \rho_i \widehat{FFR}_{t-1}^{h} \\ &\quad + \gamma_i \cdot (\widehat{UN}_{it}^{h} - \widehat{UN}_{it}^{LR}) + \omega_i \cdot (\widehat{\pi}_{it}^{h} - \widehat{\pi}_{it}^{LR}) + Z_{it}\delta + \epsilon_{it}^{h}.\end{aligned}$$

The estimated values for ρ_i are of interest. These coefficients quantify the degree of interest-rate smoothing in each member's forecasts: a higher ρ_i implies greater policy-rate persistence, meaning that member i's forecasted rate remains more tightly anchored to the previous period's forecasted federal funds rate, conditional on current macroeconomic gaps and other controls. While we know that monetary policy is highly persistent at the aggregate level, Coibion and Gorodnichenko (2012), there are no empirical estimates at the FOMC member level in the literature. The estimated values of this coefficient are present in Figure 8, which reports the estimated average marginal effects of the lagged policy rate forecasts, revealing substantial heterogeneity across members, with point estimates ranging from approximately 0.2 to 0.9. This dispersion indicates meaningful differences in how quickly policymakers adjust their projected policy paths in response to new information, with higher values reflecting more gradual, inertial updating of forecasts. Notably, Ben Bernanke and Janet Yellen are estimated among the least inertial FOMC members. Allowing for policy-rate persistence does improve the model's overall explanatory power, relative to the baseline specification with time fixed effects. However, this improvement comes at the cost of a markedly different identifying structure, as lagged policy forecasts absorb much of the time-series variation that is summarized by meeting fixed effects in the baseline specification. Since the focus of this paper is on cross-sectional heterogeneity in policy reaction functions—rather than on modeling forecast dynamics per se — the baseline specification with time fixed effects remains the preferred framework.

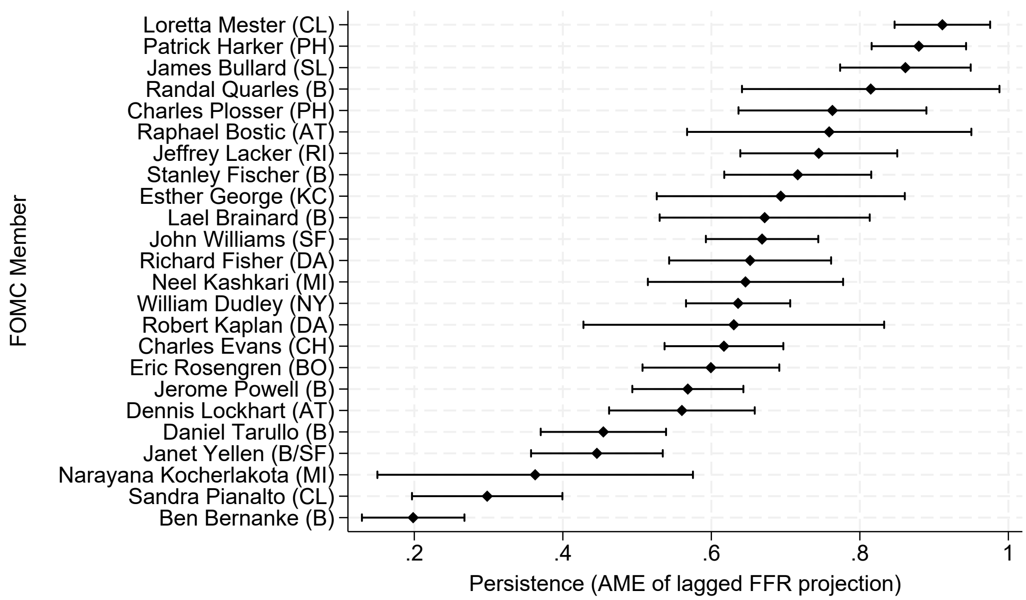

Figure 8: Coefficient ρ_i across members.

Notes: Coefficients are given with 95% confidence intervals.

Next, I also estimate a version of the policy rules using GDP gap forecasts in place of unemployment gaps. The results differ markedly: for many members, the estimated coefficients on the GDP gap are negative; see Figure 3 in the Online Appendix A. This would imply that output above potential is associated with a looser policy stance, which contradicts standard monetary policy intuition. At the same time, the inflation gap coefficients become much larger. Similar patterns are documented in Arai (2023). The GDP gap is inherently harder to measure than the unemployment gap because it relies on imprecise estimates of potential output, which are subject to substantial real-time uncertainty and revision. During most of 2011–2019, FOMC members' GDP gap forecasts were systematically larger than GDP gaps constructed using the CBO's estimates of potential output, likely leading to measurement error that biases the estimated policy response relative to specifications based on the unemployment gap. Given this evidence, as well as the fact that unemployment is a part of the Fed mandate, I use an unemployment-based policy rule as a baseline.

Another potential concern in estimating a Tobit model with member-specific intercepts and slopes is the incidental parameters problem, Greene (2002). The incidental parameters problem arises when the number of parameters to be estimated increases with the number of individuals, potentially leading to biased and inconsistent estimates, especially in short panels with limited observations per unit. In our case, three parameters are estimated per member: a fixed effect, the unemployment gap slope, and the inflation gap slope. However, I also require each FOMC member to contribute at least 40 forecasts, with the average number of forecasts per member being around 71 over the sample. This relatively large number of observations per member mitigates the severity of the incidental parameters problem. As a result, while some finite-sample bias may remain, it is unlikely to materially affect the estimated coefficients or the overall conclusions.

Finally, because SEP forecasts are formed simultaneously, there is a potential for reverse causality: changes in expected macro variables affect the expected policy rate, and changes in the expected policy rate might alter the projections of macro variables. However, the contemporaneous specification alleviates this concern, as it is unlikely that changes in the policy rate forecast would affect same-period projections of macroeconomic conditions. Also, there is evidence that the endogeneity in various forms of policy rules, estimated on the aggregate data, is likely to be limited; see Carvalho, Nechio, and Tristao (2021). To reduce the effect of this potential simultaneity, I provide results where the forecasted inflation and unemployment gaps are replaced with their first lags. The resulting AMEs are quantitatively similar to the baseline, confirming the robustness of the estimates; see Figure 4 in the Online Appendix A.

3.4 Implications of Committee Structure for Forecasted Policy Rates

In this part, I explore how heterogeneity in individual reaction functions translates into differences in the forecasted policy rate. I evaluate this by considering several counterfactual short-term predictions of the federal funds rate from the Tobit model, estimated previously. Figure 9 shows the following set of predictions for the end of current year in a subsample of voting members: (1) average federal funds rate prediction from the estimated Tobit model, together with 95% confidence bands, (2) counterfactual predicted federal funds rate if I use Chair's AMEs for all other voting members in the committee, (3) counterfactual prediction if committee consists of only a Chair, (4) actual consensus (mean) federal funds rate forecast across members[19].

To construct counterfactual (2) I use the previously estimated Tobit model and generate predicted values of the fed rate such that I use forecasts of all FOMC members in the period, but restrict AMEs of these members to be exactly equal to the Chair's estimated AMEs. Then I take the average across these forecasts. Counterfactual (3) uses the estimated equation corresponding to the Chair. In other words, I ignore all other members' predictions and estimated AMEs and assume that the committee

[19] All four projections are calculated on the identical sample of members to ensure comparability.

consists of only the Chair, and calculate the predicted values for this scenario.

The two counterfactual scenarios imply different forecasts, compared with both actual consensus forecasts (4) and with the average predicted forecast from the Tobit model (1). In some time periods, a given counterfactual prediction (2 or 3) is lower/greater than the average federal funds rate prediction from the estimated Tobit model (1) by 50 or more basis points, while also being outside the 95% interval of the latter. Similarly, these counterfactual predictions, though not shown on the graph, frequently lie outside of the range of actual forecasts of members. All this indicates that heterogeneity in reaction functions results in different forecasts of the policy rate. In section 4.2, I investigate whether forecast differences across FOMC members are indeed systematically related to their position on the estimated loadings, as captured by the baseline Tobit specification.

Next, we can see that predictions from the Tobit model (1) exceed actual predictions in (4) for 2016-2018. Such deviations likely occur because the actual forecasts reflect additional judgment and risk-management considerations (including uncertainty and asymmetric tail risks), not captured by the model, which can shift the actual policy path above or below the model prediction. Finally, if we compare counterfactuals (2) and (3) with each other, the only difference between them would come from the forecasts (in one case, all members' forecasts are considered, in another, only the Chair's forecasts are considered). For example, during 2016, counterfactual (2) was above counterfactual (3), implying that forecasts of the other committee members favored tighter policy, compared with forecasts of Janet Yellen, who was the Chair at the time.

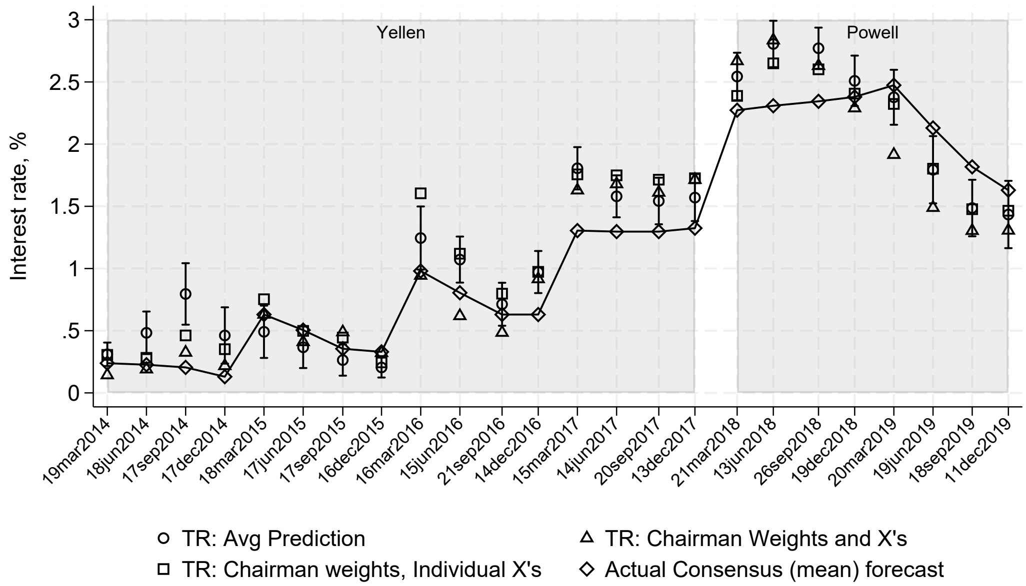

Figure 9: Counterfactual Forecasts of the Federal Funds Rate.

Notes: Actual/counterfactual forecasts made in year T for the fourth quarter of year T. 95% confidence intervals are given for the average federal funds rate prediction from the estimated Tobit model.

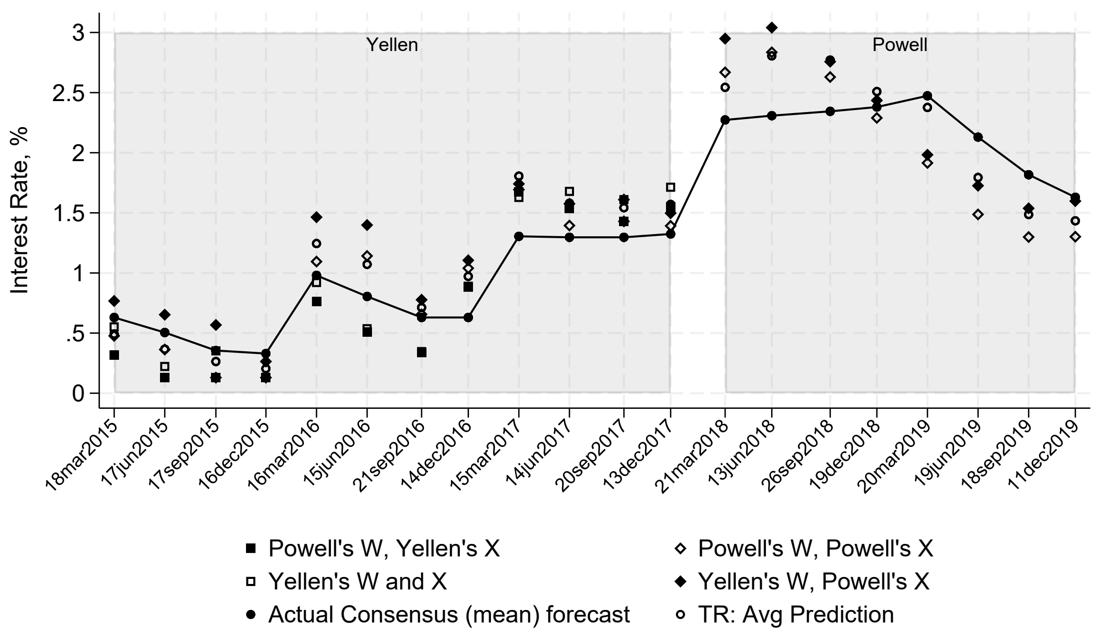

Figure 10: Counterfactual Forecasts of the Federal Funds Rate.

Notes: Actual/counterfactual forecasts made in year T for the fourth quarter of year T. W = reaction-function weights; X = input forecasts (inflation & unemployment gaps).

Figure 10 focuses on the period 2015–2019, when two individuals served as Chair:

Janet Yellen and Jerome Powell. In addition to baseline Chair-only scenarios, I introduce two additional counterfactuals that "swap" reaction functions across Chairs while holding their macroeconomic projections fixed. In the "Powell's W, Yellen's X" case, I apply estimated Powell's AMEs to Yellen's forecasts; in the reverse case, Yellen's AMEs are applied to Powell's forecasts. These swaps yield non-trivial differences in projected interest rates—reinforcing the point that differences in reaction functions, not just macro outlooks, may drive policy variation. For instance, if J. Powell "uses" J. Yellen's reaction function, while keeping his views on the economy intact, the change in the forecasted policy rate by him would be on average 25-50 basis points for most of the sample, even though the estimated AMEs are not very different between the two (see Figures 7a,7b).

3.5 Decomposing disagreement in the dot plot

The spread of the dots at a given meeting is read as disagreement among participants about the appropriate policy rate. That spread can come from two distinct sources. Participants may project different outlooks for inflation and unemployment, and so prefer different rates even if their reaction functions are identical; or they may apply different reaction functions to a common outlook. The first is disagreement about *beliefs* about the economy, the second about the *reaction function*—how projected conditions map into a preferred rate. Because I estimate each participant's reaction function from their own projections, I can separate the two and measure how much each contributes to the observed spread in the dots. This speaks to the current debate over what the dot plot conveys: whether its dispersion reflects genuine differences in how members would set policy, or simply different readings of the same economy.[20]

Let $\widehat{m}_{it}^{h}$ be participant i's fitted preferred rate at meeting t and horizon h, from the reaction function estimated in Section 3.2. This rate depends on two participant-specific inputs: the coefficients of i's reaction function (its intercept and its slopes on the gaps), and i's own projected gaps. I hold one input at its cross-participant average while letting the other vary. Working within meeting×horizon cells over the 2012–2019 sample—so that anything common to all participants at a given meeting and horizon is removed and only cross-participant disagreement remains—I build two counterfactual rates.[21] In the first, $\widehat{m}_{it}^{F}$, every participant is given the cell-average reaction function but keeps their own forecasts, so its variation across participants isolates the *forecast channel*. In the second, $\widehat{m}_{it}^{R}$, every participant keeps their own reaction function but is given the

[20] See, for e.g., Cecchetti and Schoenholtz (2026), Bernanke (2016).

[21] The decomposition is applied to the fitted latent preferred rate rather than the censored dot; the two coincide away from the effective lower bound, and excluding the ELB period leaves the shares below essentially unchanged.

cell-average forecasts, so its variation isolates the *reaction-function channel*.

I summarize the two channels by their relative contribution to disagreement,

$$S^F = \frac{\mathrm{Var}(\widehat{m}^F)}{\mathrm{Var}(\widehat{m}^F) + \mathrm{Var}(\widehat{m}^R)}, \qquad S^R = \frac{\mathrm{Var}(\widehat{m}^R)}{\mathrm{Var}(\widehat{m}^F) + \mathrm{Var}(\widehat{m}^R)},$$

so that S^F (forecast) and S^R (reaction function) express each channel as a share of the variation the two generate together and sum to one.[22]

Alongside these shares, I report the size of the disagreement using its root-mean-square (RMS) spread: the typical gap, in percentage points, between a participant's preferred rate and the average rate across participants at the same meeting and horizon. Formally, for a set of cells c with n_c participants each,

$$\mathrm{RMS} \;=\; \sqrt{\frac{1}{N}\sum_{c}\sum_{i\in c}\big(\widehat{m}_i - \bar{\bar{m}}_c\big)^2}, \qquad N = \sum_{c} n_c,$$

where $\bar{\bar{m}}_c$ is the cell mean. The shares say what the disagreement is made of; the RMS spread says how large it is, so that a channel with a given share can be judged in economic terms and not only in proportional ones.

Pooling across horizons, the forecast channel accounts for most of the disagreement the two channels generate (68%) and the reaction-function channel for the rest (32%). In levels, the total RMS spread across participants is about 0.78 percentage points, of which the reaction-function channel alone accounts for roughly 0.46—so differences in members' reaction functions generate on the order of 45 basis points of disagreement in the preferred rate, an economically meaningful amount.

These pooled numbers, however, hide a striking non-monotonicity across the forecast horizon (Table 2). The reaction-function channel is the larger of the two at *both ends* of the horizon spectrum and the smaller in the middle, and it means something different at each end. At the nearest horizon it accounts for about 59% of the two channels' variation: participants largely agree on the near-term economy, so what disagreement remains reflects differences in how they would respond to those gaps—that is, differences in the *slopes* of their reaction functions. Over the intermediate horizons, roughly half a year to two-and-a-half years out, the forecast channel dominates (70–77%): the outlook itself becomes contested and drives most of the spread. At

[22] The two channels are not independent, so the fitted-rate variance also contains a covariance term, negative in the data (about -10% of total variance, pooled), meaning the channels partly offset and the observed spread is somewhat smaller than they separately imply; I focus on the relative size of the two channels throughout.

the longest horizons the reaction-function channel again becomes the larger, but now through its intercept rather than its slopes. Three years out, participants expect the inflation and unemployment gaps to have largely closed, so their preferred rates fall back toward their individual intercepts, and the remaining spread reflects differences in the *perceived long-run (neutral) rate* rather than differences in how they respond to current gaps.[23] The two ends of the spectrum therefore reflect distinct forces: near-term disagreement is about how members would *respond* to a shared outlook, while long-horizon disagreement is about their beliefs about where rates settle in the long run. This structure is invisible in the median dot, and it is the near-term dots—those most relevant for the coming policy decision—that carry the most information about differences in how members would respond.

Table 2: Decomposition of dot-plot dispersion by forecast horizon

Horizon		Share (%)		RMS spread
(quarters)	N	Forecast	Reaction function	(pp)
1	135	41.2	58.8	0.68
2	120	60.2	39.8	0.71
3	117	59.4	40.6	0.76
4	118	71.5	28.5	0.91
5	135	73.8	26.2	0.85
6	120	75.0	25.0	0.88
7	117	75.4	24.6	0.87
8	144	76.7	23.3	0.88
9	135	73.0	27.0	0.77
10	94	70.5	29.5	0.88
11	117	66.0	34.0	0.72
12	144	68.6	31.4	0.72
13	135	59.7	40.3	0.55
14	78	45.7	54.3	0.56
Pooled	1,709	67.5	32.5	0.78

Notes: The table decomposes the cross-participant spread in the fitted preferred rate, within meeting$\times$horizon cells over the 2012–2019 sample. The forecast and reaction-function shares are each channel's variance as a fraction of the two channels' combined variance and sum to 100; a negative covariance between the channels (not shown) means the total spread is somewhat smaller than the two channels imply (see text). The horizon is the number of quarters between the meeting and the projection's target quarter; horizon 1 is the nowcast and horizon 14 is three years ahead. N is the number of individual projections used at each horizon (participant–meeting observations, pooled over the sample). The RMS spread is the total spread on the percentage-point scale.

[23] Splitting the reaction-function channel into its intercept and slope components confirms this: the intercept accounts for about 94% of the intercept-plus-slope variation at the three-year horizon, against roughly 29% at the nearest horizon, see Table 3 in the Online Appendix C. The two components rise and fall monotonically across the horizon, crossing at around two years.

The decomposition shows that a substantial part of the disagreement in the dots is about reaction functions rather than forecasts. The rest of the paper characterizes those reaction-function differences: it summarizes them in a single index, asks which member characteristics account for them, and tests whether they are reflected in members' forecasts and votes.

4 Summary measure of reaction functions heterogeneity

Sections 3.2 and 3.5 established that members' reaction functions differ, and that these differences account for a large part of the disagreement in the dot plot. This section characterizes the differences themselves. I first summarize each member's reaction function in a single index, then ask which observable characteristics explain the variation, and finally test whether the differences are reflected in members' forecasts and dissenting votes. As a starting point, there is clear variation across members in both input dimensions, see Figure 11. Some members, such as Charles Plosser and Richard Fisher, display strong reactions to inflation deviations, while others, including Ben Bernanke and Narayana Kocherlakota, exhibit relatively muted responses to both inflation and unemployment. Members like Janet Yellen, William Dudley, and John Williams exhibit more moderate and balanced policy reaction functions.

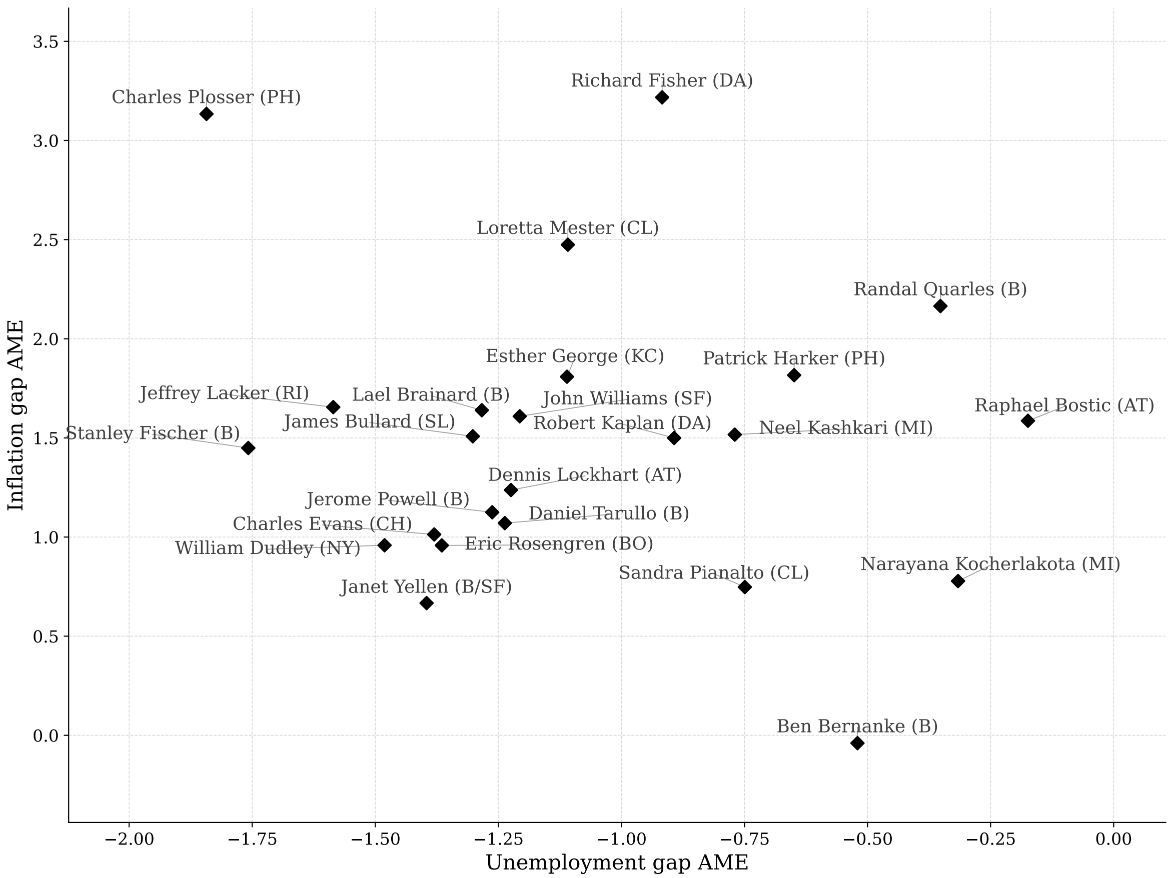

Figure 11: Scatterplot of average marginal effects.

Variation in marginal effects across members captures differences in the absolute importance of each input—it does not reflect the inflation-unemployment tradeoff per se. For example, it is not clear whether Narayana Kocherlakota is more hawkish or dovish than Neel Kashkari. The former has smaller responses to both objectives than the latter. To summarize each member's policy reaction functions along a single dimension, I construct a continuous, parsimonious measure of the perceived tradeoff between inflation and unemployment based on the estimated policy rules.

There are several ways to construct such a measure. A common approach defines a member as hawkish if they place greater emphasis on controlling inflation, and as dovish if they prioritize reducing unemployment or supporting economic activity.[24] Given estimated individual policy rules, a natural way to define such a measure is to take a sum of AMEs of inflation and unemployment gaps.[25] This measure is depicted in

[24]This characterization is regardless of members' views on the effects of monetary policy.

[25]While a simple sum of AMEs captures overall policy stance, it might not capture differences in magnitudes. For example, if a member's A and member's B sum of AME is approximately equal (like in the case of Narayana Kocherlakota and Neel Kashkari), it might not mean, in general, that these two members react equally to the same change in a policy rule input. One input might be emphasized

Figure 12.

I find that the relative importance of inflation vs unemployment differs across members, stemming from substantial variation in reaction functions. Even across the Board members or across the members from the same regional Fed my measure sometimes implies statistically significantly different inflation/unemployment tradeoff. One notable pattern is that Board members tend to be more dovish than regional Fed presidents, consistent with past qualitative classifications, Bordo and Istrefi (2023).

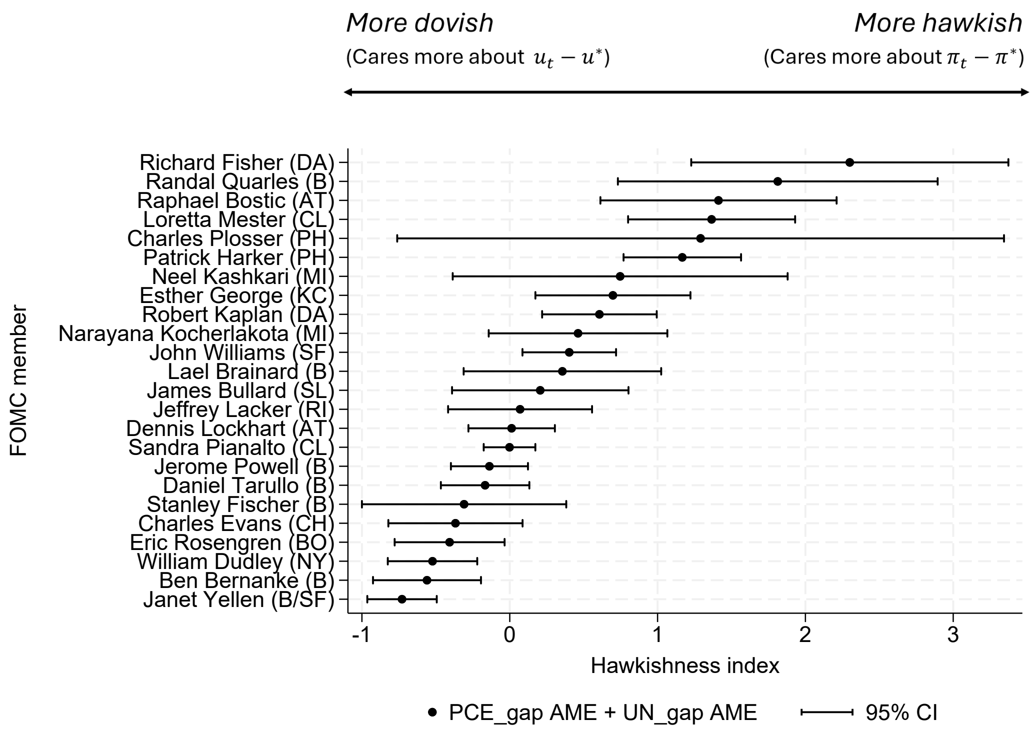

Figure 12: Estimated preferences measure

It is important to emphasize that the constructed Hawk measure reflects perceived (or revealed) tradeoff and not necessarily the true preferences of FOMC members, which are not observable. The constructed *Hawk* measure is thus viewed as relative valuation of inflation versus unemployment stabilization, *regardless of the underlying reasons for this perceived tradeoff*. Members who place greater weight on controlling inflation for whatever reasons, relative to reducing unemployment, are considered more

more strongly by one member than the other, i.e., one member can have a higher AME for this input. One way to take that into account is to standardize estimated AMEs:

$$Hawk_i^{'} = (AME_{UN,i} - mean(AME_{UN,i}))/SD(AME_{UN,i}) \\ + (AME_{\pi,i} - mean(AME_{\pi,i}))/SD(AME_{\pi,i}).$$

This alternative measure is shown in Figure 7 in Online Appendix A and highly correlated with the original one (corr>0.85). All main conclusions remain the same regardless of the choice of measure.

hawkish, while those who prioritize minimizing unemployment are considered to be more dovish.

For example, the level of this measure for a particular member is likely affected by what the member thinks about other members' decisions. Consequently, it cannot be directly viewed as a measure of true preferences. For example, FOMC members do not vote in isolation — they understand that the final policy outcome depends on committee consensus. Therefore, a member's forecasted rate (dependent variable) may reflect not only their own preference, but also anticipation of the median or majority vote, and/or strategic positioning to influence or signal to others. Thus, heterogeneity in estimated reaction functions could reflect differences in expectations about other members, rather than differences in preferences per se. In addition to strategic considerations, differences in estimated functions may also stem from informational asymmetry. Policymakers might have access to different sets of information. For example, though all have access to the same staff forecasts, each FOMC member has their own set of advisors, which can result in different estimated reaction functions.

Next, FOMC members operate in distinct macroeconomic environments—for example, periods of low inflation and low unemployment, or high inflation and low unemployment. This can mechanically lead to differences in estimated reaction functions, even if the underlying preferences over inflation and unemployment are perceived similarly across members. For example, it is not clear whether a member places low weight on inflation because inflation was near target during their tenure, or because of a genuine preference to deprioritize it. For instance, Bernanke exhibits a near-zero inflation response (AME of inflation), but it is unclear whether this reflects his intrinsic preferences or the characteristics of the sample period (low inflation) in which he served. Consequently, observed heterogeneity in reaction functions may partly reflect variation in the economic environments faced by policymakers.[26] That said, the sample period analyzed does not include major crisis episodes (starting after the end of the GFC, and ending before Covid-19) and is relatively homogeneous. This should limit the extent to which environmental factors drive differences in reaction functions. Extending the sample to include the 2020 meetings, in which the projected unemployment gap reached 10.5 percentage points — more than three times its maximum in the baseline sample — leaves the estimated marginal effects essentially unchanged (see Figure 6 in Online Appendix A).

[26] If that is indeed the case, it would suggest that preferences are not structural but instead state-dependent, implying that policymakers' responses vary non-linearly with the economic environment. Modeling such non-linearities would require a more complex specification of the policy reaction function, which is difficult to estimate given the limited data available.

On the other hand, federal funds rate forecasts can also represent what the FOMC member *would like* the interest rate to be. The Fed provides the following interpretation of SEP forecasts: *"[Interest rate forecasts are] based on policymakers' assessments of appropriate monetary policy, which, by definition, is the future path of policy that each participant deems most likely to foster outcomes for economic activity and inflation that best satisfy his or her interpretation of the Federal Reserve's dual objectives of maximum employment and stable prices"*. This supports the interpretation that forecasts represent the views on appropriate policy and preferred interest rates of FOMC members. Similarly, Bernanke (2016) argues that SEP provides indirect information about the FOMC reaction's function.

However, in general, it is not clear how exactly FOMC members form their forecasts. For example, from Blinder (2022): *"Notably, in January 2012 it began offering market participants its now-famous 'dot plot', showing where FOMC members* ***expected (or was it wanted?)***[27] *the funds rate to be over the next few years."* Thus, an important question is whether these estimated preferences accurately reflect the underlying policy inclinations of FOMC members[28].

I argue that the differences in estimated reaction functions might indeed suggest differences in preferences of FOMC members. I provide several empirical pieces of evidence supporting this view. First, as I discuss in the following Subsection 4.1, individual characteristics help explain a substantial part of the heterogeneity in reaction functions. Fundamental and exogenous individual characteristics (such as education and working experience) are indeed known determinants of preferences (Bordo and Istrefi 2023, Malmendier, Nagel, and Yan 2021). Individual characteristics would be mostly irrelevant if members primarily cared only about other members' views or if estimated AMEs were affected by the economic environment. Second, in Subsection 4.2 I show that more hawkish FOMC members, according to my measure, perceive more inflation risk, which motivates a higher policy rate—i.e., they have hawkish expectations. Third, in Subsection 4.3, I also find that voting decisions are associated with the reaction function implied inflation-unemployment tradeoff measure. Finally, in Subsection 4.4 I show that evidence of strategic forecasting is limited in the data and unlikely to be a primary driver of observed differences.

How does my measure correspond with the literature? One prominent example is Istrefi (2019), which characterizes members as hawks, doves, and swingers based on

[27]Emphasis added by the author to highlight the interpretive ambiguity.

[28]Over the last couple of years, there has been a discussion on how to improve SEP design, including, but not limited to making it better suited to inform of members reaction functions.

textual analysis of more than 20,000 Fed documents.[29] Although the full classification is not publicly available, several similarities can be observed. First, I find that board members are clustered on the dovish side of the spectrum, which coincides with the findings Bordo and Istrefi (2023) that board members are likely to be doves. Second, as we will see later, more hawkish members tend to dissent in favor of more restrictive policy more often, while also providing higher inflation forecasts. Finally, my measure echoes findings from Bordo and Istrefi (2023), including the importance of educational background in shaping preferences.

4.1 Explaining the Heterogeneity in Estimated AMEs

A natural concern is that the estimated reaction-function differences are largely estimation noise. If they were, they would not be systematically related to observable, predetermined characteristics of members. This subsection shows that they are—which both indicates where the heterogeneity comes from and confirms that the reaction-function channel identified in the decomposition reflects genuine differences across members.

Heterogeneity in the reaction functions among Federal Reserve (Fed) members likely reflects a range of structural, institutional, and personal influences. Prior research points to differences in academic training, economic philosophy (Bordo and Istrefi 2023, Chappell Jr., McGregor, and Vermilyea 2008), regional economic conditions (Bobrov, Kamdar, and Ulate 2025), career incentives (Abrams and Iossifov 2006), and political pressures (Schnakenberg, Turner, and Uribe-McGuire 2017) as key contributors to preference heterogeneity among FOMC members.

Lifetime experiences and memories play an important role in determining the formation of FOMC preferences. For instance, Malmendier, Nagel, and Yan (2021) find that past experiences of high inflation make FOMC members more hawkish, with the effect diminishing as time passes since the event. Bordo and Istrefi (2023) point out that not only inflation but also unemployment experiences are shaping the views of committee members. To capture experienced inflation directly, I construct for each member the average rate of CPI inflation over their adult life (from age 21 onward), following the experience-weighting logic of Malmendier, Nagel, and Yan (2021). A member's age is closely related to this measure—in the cross-section the two are correlated at roughly 0.86—and also proxies for other lifetime exposures, such as unemployment experience and cohort effects; I therefore use experienced inflation as the primary experience vari-

[29] This classification is mainly based on media coverage, thus reflecting their perceptions of the policymakers' intentions.

able and additionally include the number of years the member has spent in the Federal Reserve (*Tenure*) to capture institutional experience.

As shown in Bordo and Istrefi (2023), the preferences of FOMC members are most likely formed during their university studies; the education they obtain may be an important determinant of their professional views and policy choices. To capture that, I use the type of university degree (PhD in economics or related fields, JD degree, other degrees), as well as the "freshwater"/"saltwater" type distinction. I draw on the classification of economic ideology based on educational background, specifically referencing the "freshwater" and "saltwater" schools of thought, as discussed in Bordo and Istrefi (2023). This terminology reflects the geographical distribution of universities with differing macroeconomic perspectives. "Freshwater" institutions, primarily located near the Great Lakes in the United States, are contrasted with "saltwater" universities, which are situated closer to the coasts. This distinction is naturally associated with the prevalence of monetarist or Keynesian schools of thought.

Regional economic conditions may also shape a member's stance: Reserve Bank presidents represent districts whose labor markets can diverge from the national picture, and Bobrov, Kamdar, and Ulate (2025) and Chappell Jr., McGregor, and Vermilyea (2008) document that district conditions are related to policy positions and dissent. To capture this channel, I include the gap between the member's own district unemployment rate and the national rate in the month preceding each meeting; this variable is zero by construction for Board members, who represent no single district.

In this part, I assess the importance of individual characteristics of FOMC members in explaining the observed heterogeneity in AMEs of policy rules across Fed members. I consider characteristics emphasized in the literature—experienced inflation, tenure, institutional position (Board versus Reserve Bank president), educational background (degree type and "freshwater"/"saltwater" affiliation), and district economic conditions—together with a small number of additional biographical and institutional factors that may plausibly matter, namely political affiliation, gender, whether the member has published in a leading economics journal, and whether the member serves as Committee Chair. Descriptive statistics of the members are shown in Table 1 in the Online Appendix B.

Of course, these characteristics do not fully capture the nuances of each member's preferences. For instance, two members may value unemployment stabilization equally but differ in their beliefs about the effectiveness of monetary policy in achieving it. Such differences—rooted in views about transmission mechanisms—are not directly observable and are difficult to proxy without a separate structural model. However, the characteristics used in this paper mainly come from the literature, and evidently, most of them correlate with the preferences of FOMC members.

Given these constraints in mind, I estimate the following non-linear model to test whether individual characteristics can explain heterogeneity in estimated individual policy rule AMEs:

$$\widetilde{FFR}^h_{it} = \alpha_i + \alpha_t + (\gamma_1 Exp\pi + \gamma_2 PhD + \dots) \cdot (\widehat{UN}^h_{it} - \widehat{UN}^{LR}_{it}) + \\ (\omega_1 Exp\pi + \omega_2 PhD + \dots) \cdot (\widehat{\pi}^h_{it} - \widehat{\pi}^{LR}_{it}) + Z_{it}\delta + \epsilon^h_{it}. \qquad (5)$$

I follow the approach in Section 3 and estimate a left-censored Tobit model, in which the member-specific slopes on the inflation and unemployment gaps are replaced by a common slope shifted along the characteristics above[30]. To determine which characteristics are relevant, I select the combination that minimizes the Akaike Information Criterion (AIC).[31]

I assess each characteristic's contribution to fit by the increase in the model's AIC when it is removed from the selected specification; a larger increase indicates a characteristic that contributes more to explaining the cross-member variation in reaction functions. Table 2 in the Online Appendix C reports these contributions. Experienced inflation contributes most, followed by educational background—both graduate-school tradition and degree type—and district economic conditions; institutional position and research record follow, with tenure and the Chair indicator contributing least among the retained characteristics. Several of these characteristics overlap, so their individual contributions are not cleanly separable: the district unemployment gap is zero for all Board members, and every Committee Chair in the sample is a member of the Board, so district conditions, Board membership, and the Chair indicator share explanatory content. I therefore emphasize the ranking rather than the precise contribution of any single characteristic, and I do not interpret the direction of individual effects.

Consistent with education being a central determinant of members' views, the "freshwater"/"saltwater" distinction emphasized by Bordo and Istrefi (2023) remains among the more important characteristics for explaining the heterogeneity. Because school-of-thought is correlated with institutional position in this sample—none of the Board

[30]An alternative would be to regress the estimated *Hawk* measure directly on these characteristics. I embed the characteristics in the Tobit slopes instead because that specification is estimated on the full panel of forecasts rather than on 24 collapsed member-level points, which preserves the power needed to separate correlated regressors. Regressing a generated *Hawk* measure on the characteristics would additionally underestimate uncertainty of the coefficients

[31]This lets model fit, rather than a priori choices, determine the set of characteristics, while the AIC penalty guards against including factors that do not improve fit by more than the cost of the parameters they add.

members graduated from a "freshwater" university—its standalone contribution should be read alongside that of Board membership rather than in isolation.

To evaluate how well observable characteristics account for heterogeneity in reaction functions, I compare the AMEs from the flexible, member-specific specification with those implied by the model above. The selected characteristics recover almost all of the explainable cross-member variation: measured by the AIC, the characteristics model closes roughly 99% of the gap between a specification with a single common slope for all members and one that estimates a separate, unrestricted slope for each member.[32] Figures 14 and 13 show the comparison directly. For both the unemployment gap and inflation the two sets of AMEs are closely aligned across most of the members. Overall, the fact that a small number of observable characteristics reproduces most of the heterogeneity suggests that differences in reaction functions reflect systematic differences among FOMC members rather than noise.

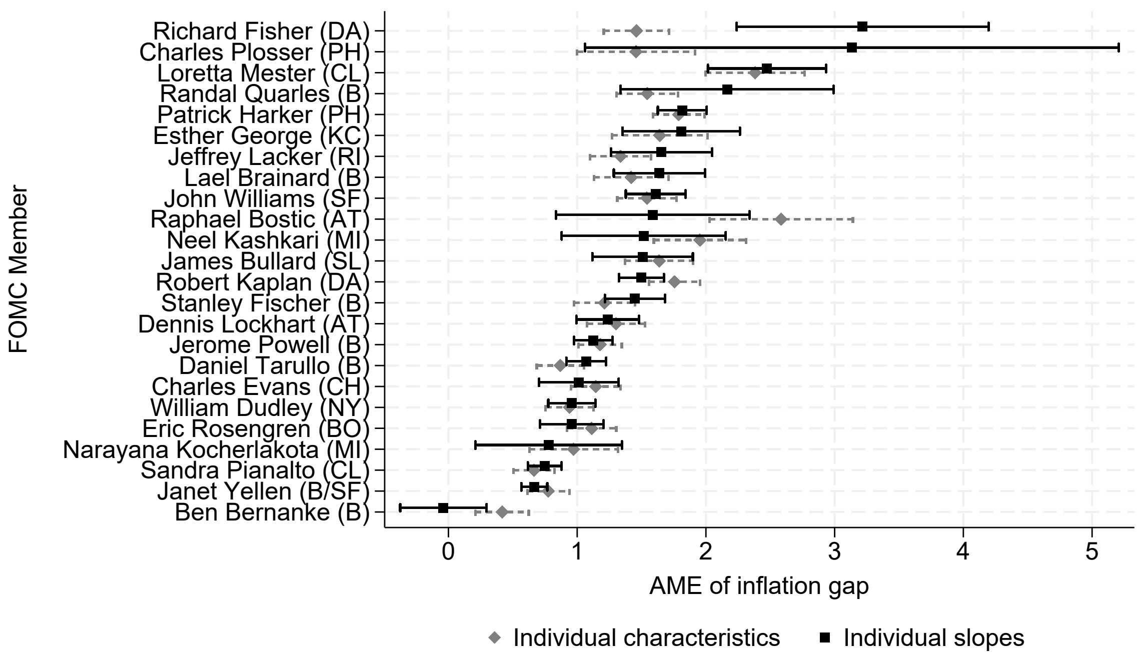

Figure 13: Inflation gap coefficients: individual slope vs individual characteristics.

Notes: confidence intervals correspond to 95% level

[32]The AIC is 2942.7 for the common-slope specification, 2778.6 for the fully flexible member-specific specification, and 2780.7 for the selected characteristics model, so the latter has the same overall fit to data as the fully flexible specification.

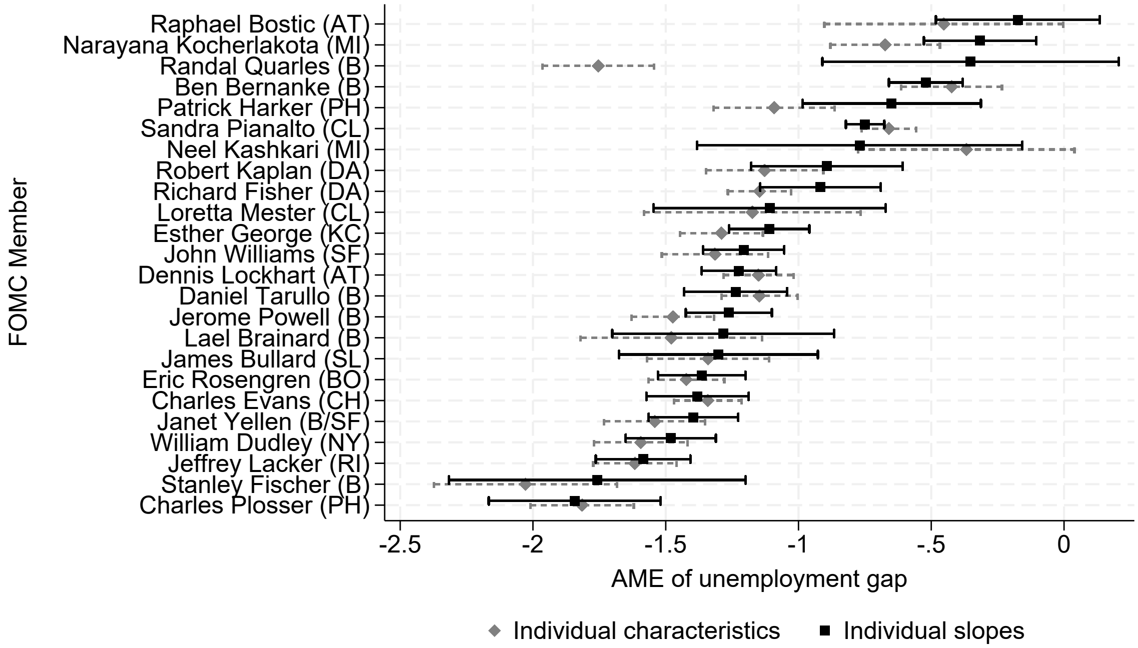

Figure 14: Unemployment gap coefficients: individual slope vs individual characteristics.

Notes: confidence intervals correspond to 95% level

4.2 FOMC forecasts and the *Hawk* measure

A further test of whether the reaction-function differences are behaviorally meaningful is whether they line up with what members actually forecast and how they vote. In this subsection I test whether forecast differences across FOMC members are systematically related to their position on the estimated preference spectrum. Table 3 reports OLS estimates in which each participant's projection of i) inflation, ii) unemployment, and iii) GDP growth is regressed on their *Hawk* measure, separately for the current-year, next-year, and two-year-ahead horizons. Time fixed effects absorb features of the economic environment common to all participants at a given meeting, so the coefficient is identified from differences across participants within a meeting. Standard errors are clustered by participant.

The results support the commonly held view that hawkish policymakers anticipate greater inflationary pressure. A higher *Hawk* measure is associated with a higher projected inflation rate at all three horizons, and the estimates are statistically significant. The magnitudes are economically meaningful: moving from the most dovish to the most hawkish participant in the sample raises the projected inflation rate by roughly 0.2 percentage points for the current year, 0.4 percentage points for the following year, and 0.3 percentage points two years ahead — substantial against a 2% target.

No such relationship appears for unemployment or GDP growth. The coefficients are small, change sign across horizons, and are not significantly different from zero at any horizon. Hawkishness therefore travels with beliefs about inflation specifically, rather than with a generally more optimistic or more pessimistic view of the economy.

One possible explanation is that more hawkish participants project a higher inflation rate — perhaps higher than their "true" forecast — in order to justify a rate hike. Two pieces of evidence weigh against this. First, a participant building a case for tightening would have reason to project a tighter labour market as well, yet hawkishness is unrelated to unemployment projections at any horizon. Second, if participants at the extremes of the spectrum were strategically distorting their forecasts, their projections would be less accurate than those of more neutral participants; measuring the forecast accuracy of each participant's nowcast, I find no evidence of systematic over- or under-prediction across participants with different values of the *Hawk* measure. I return to strategic forecasting in Subsection 4.4.

Table 3: FOMC forecasts and *Hawk*

	Inflation	Unemployment	GDP
Current Year Forecasts			
β estimate	0.065***	-0.032	0.018
SE	(0.022)	(0.025)	(0.036)
Adjusted R^2	0.908	0.995	0.951
Next Year Forecasts			
β estimate	0.129***	-0.036	0.001
SE	(0.039)	(0.058)	(0.027)
Adjusted R^2	0.434	0.983	0.755
Two-Year-Ahead Forecasts			
β estimate	0.088**	0.018	-0.102
SE	(0.032)	(0.060)	(0.071)
Adjusted R^2	0.377	0.957	0.812
Time FEs	Yes	Yes	Yes

Notes: Standard errors clustered by participant in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Significance is confirmed by a wild cluster bootstrap with 9,999 replications, reported because the number of clusters is small. The *Hawk* measure is estimated in a first stage; the standard errors treat it as known and therefore understate uncertainty. N = 709 in each column.

4.3 Dissent voting and *Hawk* measure

Dissent data from the Federal Reserve systematically records instances in which individual members of the Federal Open Market Committee (FOMC) express disagreement with the prevailing consensus on monetary policy decisions. This information is captured through formal statements and detailed voting records during FOMC meetings. Analyzing dissent data provides critical insights into the heterogeneity of perspectives within the committee and elucidates the underlying factors that influence policy outcomes.

However, preferences are not the only factor explaining why certain members deviate from the consensus. Other influences, such as strategic considerations, career incentives, institutional roles, or misinterpretations of economic conditions, can also shape dissenting behavior. Dissenting votes are typically deliberate and reflect deeper policy disagreements. Sometimes they serve a strategic purpose, allowing members to signal their hawkish or dovish leanings, shape future discussions, or build their reputations within the institution. Another reason for dissent might be alternative interpretations of economic conditions. In some cases, dissenting members may even be ahead of the consensus, identifying risks or opportunities that others overlook.

Table 4 summarizes the frequency of dissent at the individual member level. Dissent is not evenly distributed across policymakers: ten of the twenty-four participants account for every dissenting vote in the sample, and fourteen never dissent at all. Dissent is also persistent. The transition matrix in Table 5 shows that, conditional on having dissented at the previous meeting, the probability of dissenting again is 32.0 percent, against 4.7 percent for a participant who voted with the consensus — a sevenfold difference, suggesting that dissent reflects stable individual dispositions rather than transient reactions to circumstance. At the same time, dissent is not the mechanical behavior of a fixed group: dissenters return to the consensus in roughly two-thirds of cases, and dissent remains rare overall, occurring in 6.9 percent of participant-meeting observations. Together these patterns indicate that dissent is concentrated among a subset of policymakers with persistent leanings, but is expressed episodically rather than at every opportunity.

Table 4: Dissent Frequency and Switching by FOMC Member

Member	Meetings	Dissents	Share dissent	Switches
Neel Kashkari	4	3	0.750	2
Esther George	13	6	0.462	7
Jeffrey Lacker	13	4	0.308	4
Richard Fisher	12	3	0.250	1
James Bullard	16	3	0.188	4
Eric Rosengren	17	3	0.176	6
Narayana Kocherlakota	8	1	0.125	1
Loretta Mester	11	1	0.091	2
Charles Plosser	12	1	0.083	2
Charles Evans	25	2	0.080	4
Lael Brainard	22	0	0.000	0
Patrick Harker	4	0	0.000	0
Janet Yellen	34	0	0.000	0
William Dudley	39	0	0.000	0
Jerome Powell	30	0	0.000	0
Daniel Tarullo	33	0	0.000	0
Stanley Fischer	14	0	0.000	0
Robert Kaplan	4	0	0.000	0
Raphael Bostic	4	0	0.000	0
Dennis Lockhart	12	0	0.000	0
Randal Quarles	9	0	0.000	0
Sandra Pianalto	13	0	0.000	0
John Williams	17	0	0.000	0
Ben Bernanke	26	0	0.000	0

Notes: The table reports the members used in the estimation of the model in Section 3. Only FOMC members with voting rights at a given meeting are counted. *Switches* is the number of times a member moves between consensus and dissent across consecutive meetings at which they held a vote.

Table 5: Transition Matrix of Dissent

Previous meeting	Current meeting		
Dissent	No dissent (0)	Dissent (1)	Total
No dissent (0)	327 (95.34)	16 (4.66)	343 (100.00)
Dissent (1)	17 (68.00)	8 (32.00)	25 (100.00)
Total	344 (93.48)	24 (6.52)	368 (100.00)

Notes: Entries report counts, with row percentages in parentheses. Dissent is defined as any non-consensus vote. The sample is smaller than in Table 6 because each member's first meeting has no lagged observation.

In order to demonstrate the role of heterogeneity of reaction functions, I show how the estimated tradeoff measure (*Hawk*) is related to the voting behavior of FOMC members. To do that, I estimate an ordered logit model in which the dependent variable is *Dissent* $\in \{-1, 0, +1\}$ (-1 = dovish dissent, 0 = consensus, +1 = hawkish dissent), and the regressor of interest is the estimated tradeoff measure (*Hawk*), the sum of the inflation and unemployment gap AMEs. Formally, I assume an ordered-logit structure with latent index:

$$z_{it}^* = X_{it}\eta + \eta_{\mathrm{H}}\,\mathrm{Hawk}_i + \nu_{it}, \qquad \nu_{it} \sim \mathrm{Logistic}(0,1),$$

where X_{it} collects the individual characteristics used earlier in equation (5) (education, experienced inflation, working experience, district conditions, etc.).

Neither meeting nor participant fixed effects can be included. In any meeting with no dissenting vote, the meeting indicator perfectly predicts the outcome, so those observations are dropped; because most meetings have no dissent at all, including meeting effects would discard almost the entire sample. Participant fixed effects are ruled out by construction, since $Hawk_i$ is time-invariant and hence collinear with them.

Observed outcomes follow threshold crossings:

$$\mathrm{Dissent}_{it} = \begin{cases} -1, & \text{if } z_{it}^* \le \kappa_1, \\ \;\;0, & \text{if } \kappa_1 < z_{it}^* \le \kappa_2, \\ +1, & \text{if } z_{it}^* > \kappa_2, \end{cases} \qquad \kappa_1 < \kappa_2,$$

with the usual ordered-logit normalization. The implied probabilities are:

$$\begin{aligned} \Pr(-1 \mid \cdot) &= \Lambda\big(\kappa_1 - \eta' X_{it} - \eta_{\mathrm{H}}\mathrm{Hawk}_i\big), \\ \Pr(0 \mid \cdot) &= \Lambda\big(\kappa_2 - \eta' X_{it} - \eta_{\mathrm{H}}\mathrm{Hawk}_i\big) - \Lambda\big(\kappa_1 - \eta' X_{it} - \eta_{\mathrm{H}}\mathrm{Hawk}_i\big), \\ \Pr(+1 \mid \cdot) &= 1 - \Lambda\big(\kappa_2 - \eta' X_{it} - \eta_{\mathrm{H}}\mathrm{Hawk}_i\big), \end{aligned} \tag{6}$$

where $\Lambda(\cdot)$ is the standard logistic CDF.

The first column of Table 6 reports the marginal effects of the *Hawk* measure without controls; the second adds the full set of individual characteristics, including district economic conditions. I interpret the specification with controls. There, a one-unit increase in the *Hawk* measure (approximately one standard deviation) is associated with a statistically significant increase in the probability of a hawkish dissent of about 7.6 percentage points and a decrease in the probability of a dovish dissent of about 4.5 percentage points (both significant at the 1 percent level). The estimates are smaller without controls, but the hawkish effect is significant at the 10 percent level and the dovish effect at the 5 percent level.[33] The results are robust to the link function: an ordered probit yields similar marginal effects.

Table 6: Average marginal effects of *Hawk* measure on the dissent types

	Hawk AMEs	
Dovish Dissent	-1.7 **	-4.5 ***
Consensus	-1.1	-3.1
Hawkish Dissent	2.8 *	7.6 ***
Pseudo R2	0.03	0.26
AIC	230.8	194.1
Controls	No	Yes
N	392	392
Participants	24	24

Notes: Ordered logit; average marginal effects reported in percentage points. Standard errors clustered by participant (24 clusters). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Though these estimates might appear small, they are economically meaningful given how rarely dissent occurs. In the sample, dissent is recorded in only 6.9 percent of participant-meeting observations — 4.3 percent hawkish and 2.6 percent dovish. A marginal effect of 7.6 percentage points on hawkish dissent therefore exceeds the unconditional hawkish-dissent rate itself (4.3 percent): a one-unit increase in Hawk raises the likelihood of a hawkish dissent by more than its base rate — an increase of over 75 percent relative to that base. The estimated effects of *Hawk* are thus substantial, demonstrating the measure's relevance for explaining dissenting votes. These results are qualitatively similar to the findings of Istrefi (2019), though obtained with completely different data and framework.

4.4 Strategic forecasting in FOMC

If members were engaged in strategic forecasting, the change in voting status would change the incentives to predict according to consensus. For example, a non-voting member might be willing to steer the consensus by submitting abnormally low or high

[33] Every Board member in the sample — including all Chairs and Vice Chairs — dissent in no meeting, reflecting the convention that the Board votes with the Chair. The results are similar when the sample is restricted to Reserve Bank presidents, so the *Hawk* measure is not merely proxying for institutional position.

forecasts. When a member gains voting status, such incentives should become weaker, as voting members can influence the policy outcome directly through their vote. This is a potential reason why the forecasted federal funds rates would not imply the preferred rates, which could result in failing to capture the true preferences of FOMC members.

Voting status changes frequently in the sample due to the FOMC's rotation system. Focusing on regional Federal Reserve Bank presidents, the majority of members experience multiple transitions between voting and non-voting status over the sample period, providing substantial within-member variation for identifying the effect of voting status on forecasting behavior. To test the extent of strategic forecasting in the sample, I look at how voting status changes affect the forecasts, relative to consensus. I estimate the following regression on the sample of regional Fed presidents:

$$\left|\mathbf{E}_t^i X_{t+h} - \mathbf{E}_t^{mean} X_{t+h}\right| = \alpha_i + \alpha_t + \beta_1 NonV_t^i + \epsilon_t,$$

where $\mathbf{E}_t^i X_{t+h}$ denotes the forecast of one of the three: $\{$*inflation*, *unemployment rate*, *federal funds rate*$\}$, $\mathbf{E}_t^{mean} X_{t+h}$ is the cross-sectional mean of forecasts across all voting members except i, and $NonV_t^i = 1$ if member i is a non-voter in period t, and zero otherwise. I restrict the sample to members who change their voting status over the sample period. By examining the coefficient β_1, we can assess whether members alter their forecasting behavior when their voting status changes.[34]

Tables 7 and 8 report the estimated coefficients for inflation and unemployment, respectively. In all cases, the Non-Vote dummy variable coefficients are relatively close to zero. The effects are not statistically significant at the 5% level for all except one case. Thus, non-voting status does not appear to result in economically meaningful deviations from the consensus forecast inputs. Importantly, Table 9 similarly shows that non-voting members do not systematically deviate from the consensus federal funds rate forecast relative to voting members. Therefore, voting status should not substantially affect the estimates of average marginal effects. Thus, I do not observe systematic strategic forecasting in the sample.[35]

[34] Note that the use of absolute deviations focuses on the magnitude of strategic forecasts rather than their direction.

[35] Alternative approaches exist to test for strategic forecasting behavior (e.g., using narrative data or survey metadata), but these are beyond the scope of this paper.

Table 7: Effect of Non-Voting Status on **Inflation Rate** Forecasts

	$h = 1$	$h = 2$	$h = 3$	$h = LR$
β_1	0.009	0.002	0.006	-0.001
	(0.017)	(0.026)	(0.017)	(0.004)
Member FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
$Adj.R^2$	0.35	0.36	0.37	0.79
N	600	600	600	495

Notes: Standard errors clustered by participant in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Effect of Non-Voting Status on **Unemployment Rate** Forecasts

	$h = 1$	$h = 2$	$h = 3$	$h = LR$
β_1	0.009	0.034	0.031	0.033**
	(0.011)	(0.029)	(0.025)	(0.016)
Member FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
$Adj.R^2$	0.34	0.37	0.30	0.47
N	600	600	600	491

Notes: Standard errors clustered by participant in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Effect of Non-Voting Status on **Federal Funds Rate** Forecasts

	$h = 1$	$h = 2$	$h = 3$	$h = LR$
β_1	0.021	0.042	0.035	-0.005
	(0.021)	(0.029)	(0.072)	(0.015)
Member FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
$Adj.R^2$	0.51	0.57	0.56	0.37
N	396	396	396	359

Notes: Standard errors clustered by participant in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Conclusion

In this paper, I use Summary of Economic Projections (SEP) data to estimate FOMC member-level policy reaction functions for 2012–2019, in order to ask what the disagreement displayed in the dot plot actually reflects. I merge individual SEP forecasts with participant keys (released with a five-year lag) to construct a panel of policymakers' forecasts, and I estimate a Tobit model for expectations-based reaction functions of 24 members.

My central finding is that the dispersion in the dot plot reflects two distinct sources—differences in members' economic forecasts and differences in their reaction functions—and that their relative importance shifts systematically with the forecast horizon. Near-term disagreement is driven mainly by differences in how members would respond to a shared outlook, whereas at longer horizons it increasingly reflects differences in their beliefs about the economy and about the long-run policy rate. Disagreement in the near-term dots is thus largely about policy responses, not merely about differing readings of the economy.

These reaction-function differences are substantial, consequential, and structured. A one-percentage-point increase in the forecasted unemployment gap is associated with a change in the forecasted federal funds rate ranging from -1.8 to -0.2 percentage points across members, and a comparable increase in the inflation gap with a change from 0 to 3.1 points; these differences move the implied policy path by tens of basis points, and they are systematically related to observable characteristics—most importantly experienced inflation in adulthood, followed by educational background, both graduate-school tradition and degree type.

To summarize the reaction-function differences, I construct a parsimonious measure, *Hawk*, based on the inflation- and unemployment-gap average marginal effects, capturing how the inflation–unemployment trade-off is perceived across members. This measure is behaviorally meaningful: dissent voting is systematically associated with it, with members estimated to be more hawkish (dovish) much more likely to vote for tighter (looser) policy than the consensus.

A key limitation of the analysis is structural: participants' identities are released with a five-year lag, so individual reaction functions cannot be recovered in real time. If anything, the prevailing direction of Federal Reserve communication now runs toward less disclosure, with the dot plot itself under review and possibly to be scaled back—which makes what the individual projections contain, rather than how quickly they arrive, the more pressing question.

These results speak to that question. The individual dots are not redundant forecasts: they embed systematic, characteristic-linked differences in how members weigh inflation against unemployment, and these differences shift the implied policy path and help predict dissent. The cross-section of projections thus carries behaviorally meaningful information that a shorter publication lag would surface in real time and that paring back the dot plot would discard.

The analysis suggests that this heterogeneity might stem from different preferences or different beliefs of FOMC members about data-generating processes in the economy. There are, however, other potential sources of variation in reaction functions. Clear identification of these sources appears to be an important avenue for future research.

References

Abrams, Burton A. and Plamen Iossifov. (2006). "Does the Fed contribute to a political business cycle?" *Public Choice*, 129, 249–262.

Arai, Natsuki. (2023). "The FOMCs new individual economic projections and macroeconomic theories." *Journal of Banking and Finance*, 151, 106845.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam. (2024). "Perceptions about monetary policy." *The Quarterly Journal of Economics*, 139, 2227–2278.

Belden, Susan. (1991). "Rationale for dissent: the case of FOMC members." *Contemporary Economic Policy*, 9, 59–70.

Bennani, Hamza, Etienne Farvaque, and Piotr Stanek. (2018). "Influence of Regional Cycles and Personal Background on FOMC Members' Preferences and Disagreement." *Economic Modelling*.

Bernanke, Ben S. (2016). "Federal Reserve economic projections: What are they good for?" *Brookings blog, November*, 28.

Blinder, Alan S. (2022). "A Monetary and Fiscal History of the United States, 1961-2021."

Bobrov, Anton, Rupal Kamdar, and Mauricio Ulate. (June 2025). "Regional Dissent: Do Local Economic Conditions Influence FOMC Votes?" *American Economic Review: Insights*, 7, 268–84.

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

Carvalho, Carlos, Fernanda Nechio, and Tiago Tristao. (2021). "Taylor rule estimation by OLS." *Journal of Monetary Economics*, 124, 140–154.

Cecchetti, Stephen G. and Kermit L. Schoenholtz. (June 14, 2026). *Warsh's Communications Dilemma: Talk Less, or Say More?* Money, Banking and Financial Markets. URL: https://www.moneyandbanking.com/commentary/2026/6/14/warshs-communications-dilemma-talk-less-or-say-more (visited on 07/14/2026).

Chappell Jr., Henry W., Rob Roy McGregor, and Todd A. Vermilyea. (2008). "Regional economic conditions and monetary policy." *European Journal of Political Economy*, 24, 283–293.

Coibion, Olivier and Yuriy Gorodnichenko. (2012). "Why are target interest rate changes so persistent?" *American Economic Journal: Macroeconomics*, 4, 126–162.

Fendel, Ralf and Jan-Christoph Rülke. (2012). "Are heterogeneous FOMC forecasts consistent with the Fed's monetary policy?" *Economics Letters*, 116, 5–7.

Fos, Vyacheslav and Nancy R. Xu. (2025). "When Do FOMC Voting Rights Affect Monetary Policy?"

González-Astudillo, Manuel and Rakeen Tanvir. (2024). "Hawkish or Dovish? Using the Fed's Dot Plot to Gauge Intended Policy Responses to Inflation and Economic Activity in Real Time." *Working Paper*.

Greene, William. (2002). "The behavior of the fixed effects estimator in nonlinear models."

Havrilesky, Thomas and John A. Gildea. (1991). "The policy preferences of FOMC members as revealed by dissenting votes: comment." *Journal of Money, Credit and Banking*, 23, 130–138.

Holston, Kathryn, Thomas Laubach, and John C. Williams. (2017). "Measuring the natural rate of interest: International trends and determinants." *Journal of International Economics*, 108, S59–S75.

Huang, Yu-Fan. (2015). "Time variation in US monetary policy and credit spreads." *Journal of Macroeconomics*, 43, 205–215.

Istrefi, Klodiana. (2019). "In Fed watchers' eyes: Hawks, doves and monetary policy."

Kim, Jinill and Seth Pruitt. (2017). "Estimating monetary policy rules when nominal interest rates are stuck at zero." *Journal of Money, Credit and Banking*, 49, 585–602.

Kocherlakota, Narayana R. (2025). "Asymmetries in federal reserve objectives." *Economic Policy*, 40, 371–400.

Malmendier, Ulrike, Stefan Nagel, and Zhen Yan. (2021). "The making of hawks and doves." *Journal of Monetary Economics*, 117, 19–42.

Mavroeidis, Sophocles. (2021). "Identification at the Zero Lower Bound." *Econometrica*, 89, 2855–2885.

Meade, Ellen E. and David Stasavage. (2008). "Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve." *The Economic Journal*, 118, 695–717.

Romer, David. (2010). "A New Data Set on Monetary Policy: The Economic Forecasts of Individual Members of the FOMC." *Journal of Money, Credit and Banking*, 42, 951–957.

Schnakenberg, Keith E., Ian R. Turner, and Alicia Uribe-McGuire. (2017). "Allies or commitment devices? A model of appointments to the Federal Reserve." *Economics & Politics*, 29, 118–132.

Shapiro, Adam Hale and Daniel J. Wilson. (2022). "Taking the fed at its word: A new approach to estimating central bank objectives using text analysis." *The Review of Economic Studies*, 89, 2768–2805.

Smales, Lee A. and Nick Apergis. (2016). "The influence of FOMC member characteristics on the monetary policy decision-making process." *Journal of Banking and Finance*, 64, 216–231.

Taylor, John B. (1993). "Discretion versus policy rules in practice." *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Tillmann, Peter. (2011). "Strategic forecasting on the FOMC." *European Journal of Political Economy*, 27, 547–553.

Trehan, Bharat and Tao Wu. (2007). "Time-varying equilibrium real rates and monetary policy analysis." *Journal of Economic dynamics and Control*, 31, 1584–1609.