Two Sides of the Same Coin: Perceived Reaction Functions and Monetary Policy Expectations Shocks

Christopher S. Sutherland*
Monetary Policy Studies Division, Financial Markets Department, Bank of Canada
Department of Economics, University of Oxford

Abstract

To identify the causal effects of monetary policy, researchers often use either high-frequency or narrative monetary policy shock measures as instruments. By their nature, these instruments tend towards being either weak or endogenous, respectively. This paper proposes a third approach. First, we develop a framework that specifies the links between true and perceived FOMC reaction functions. Second, we use this framework to develop an econometric model to estimate the perceived FOMC reaction functions implied by monetary policy expectations. This allows us to extract both expected reaction function coefficient estimates and two monetary policy expectations shocks. We provide evidence supporting the relevance and exogeneity of these instruments, demonstrate that the impulse response function estimates are consistent with New Keynesian theory, and improve statistical inference relative to the literature.

September 15, 2026

Click here for the latest draft.

Keywords: monetary policy shocks, reaction function, forward guidance, expectations.

JEL Classification: D83, D84, E37, E52, E58.

*csutherland@bankofcanada.ca, christopher.sutherland@exeter.ox.ac.uk. I am very grateful to Steve Bond for ongoing discussion and comments. I also appreciate discussion and comments from Philippe Andrade, Michael Bauer, Francesco Bianchi, Giancarlo Corsetti, Antonio Diez de los Rios, Jan Ditzen, Michael Ehrmann, Martin Eichenbaum, Filippo Ferroni, Bruno Feunou, Luca Gambetti, Robert Hill, Oscar Jordà, Alexandre Kohlhas, Gary Koop, Ingomar Krohn, Daniel Lewis, Michael McMahon, Pascal Meichtry, Luba Petersen, Carolin Pflueger, Giovanni Ricco, Neil Shephard, James Stock, Adi Sunderam, & Jon Witmer. Thank you also to seminar participants at The Barcelona School of Economics *Advances in Structural Shocks Identification* Workshop, The 2026 IAAE Annual Conference, The 2026 IIF International Symposium on Forecasting, The Bank of Canada, and The University of Oxford. I thank the authors of Bauer and Swanson (2023) and Aruoba and Drechsel (2025) for kindly sharing their data. Finally, I am very grateful to the Harvard University Department of Economics for hosting me in Autumn 2024, during which the initial phases of this paper were drafted.

1 Introduction

Monetary policy shock measures are an indispensable tool used by macroeconometricians to identify the causal effect of monetary policy on the economy. The idea is to construct a relevant monetary policy shock measure to be used as an exogenous instrument for an endogenous measure of monetary policy. If successful, these instruments can facilitate valid estimation of impulse response functions that trace out the dynamic effects of monetary policy on macroeconomic aggregates.

The most popular approach to constructing such monetary policy shock measures is the high-frequency approach: to construct a monetary shock measure based on the changes in prices of futures contracts tied to the policy rate during short time windows that encompass monetary policy events, during which time the only (or main) plausible cause of those changes is monetary news. The approach was introduced by Kuttner (2001), but has been refined ever since (e.g. Gürkaynak, Sack, and Swanson (2005), Nakamura and Steinsson (2018)). Perhaps the current state-of-the-art measures were introduced by Miranda-Agrippino and Ricco (2021) and Acosta, Ajello, Bauer, Loria, and Miranda-Agrippino (2025).

Another popular approach is the narrative approach, which began by using narrative evidence from FOMC meeting transcripts and the Fed's real-time Greenbook forecasts to construct a series of intended federal funds rate changes, then regressing these intentions on the Fed's forecasts to remove systematic policy responses (Romer and Romer (2004)). This approach has also been updated and refined ever since, with perhaps the current state-of-the-art measure introduced by Aruoba and Drechsel (2025).

By their nature, high-frequency shock measures will tend to be weak instruments because they capture only a small fraction of the updates to market participants' interest rate expectations (see Ramey (2016), Bauer and Swanson (2023), Bianchi, Ludvigson, and Ma (2025)). Conversely, narrative measures will tend to be endogenous because, although they capture a larger fraction of such changes, they may struggle to completely disentangle the endogenous and exogenous components. Also, to the extent that narrative measures focus on the federal funds rate, they are not an ideal summary statistic for the monetary policy stance, which is better captured by a term structure of interest rate expectations, not just one point.

So, this paper proposes a complementary approach to estimating monetary policy instruments that attempts to better balance relevance and exogeneity. We develop a framework designed for the structural estimation of time-varying expected FOMC reaction function parameters implied by matching sets of macroeconomic forecasts. The framework allows us to delineate separate information (macroeconomic projections), macroeconomic response ("Fed response to news"), smoothing/gradualism (Coibion and Gorodnichenko (2012)), and, crucially, FOMC judgment channels. The channels are of interest per se, but also allow researchers to more precisely define and analyze concepts such as forward guidance and information effects. We use private forecast data to do so because the interest rate forecasts are likely to have been constructed using simple rules and heuristics (unlike central bank projections), which makes disentangling these channels more straightforward. We then extract the residuals implied by the expected reaction function parameter estimates, and show that they have statistical properties consistent with valid econometric instruments.

This paper builds on Bauer, Pflueger, and Sunderam (2024), which also estimates the time-varying, expected FOMC reaction function parameters implied by (Blue Chip) interest rate forecasts. By contrast, they do not extract monetary shocks. The authors argue that their error term "*contains forecaster j's expectation of the future monetary policy shock at t+h*," "*as well as possible measurement and specification errors affecting the funds rate forecasts*."

Based on the reaction function framework (§2.1, §2.2), we propose an econometric model (§2.3, §2.4) designed to mitigate these many sources of bias. We outline eight main identification challenges and propose solutions to mitigate each. For instance, using Pesaran (2015) test statistics, we detect strong cross-sectional dependence between forecasts. Failure to account for such cross-sectional dependence causes the error term to be autocorrelated and the expected reaction function coefficient estimates to be biased. So, we adapt the Chudik and Pesaran (2015) dynamic panel estimator, which adjusts for such cross-sectional dependence. To account for multiple documented sources of forecast data heterogeneity (Patton and Timmermann (2010)), we parse the data by forecaster, forecast month, and forecast horizon. We argue that our approach allows us to treat the residuals as shocks to monetary policy expectations.

We apply this econometric model to the *Consensus Economics* U.S. forecast data to estimate forecaster-expected, forward-looking FOMC reaction functions and their implied monetary policy rate path shocks (also applicable to the Blue Chip forecast data). Our method could be adapted to many other countries, not least because *Consensus Economics* provides standardized forecast data for over 100 countries.

We make a number of contributions to the empirical monetary economics literature. First, we provide estimates of the expected (or, perceived) FOMC reaction functions implied by monetary policy expectations purged of several sources of bias that had not yet been explored. Second, we provide estimates of shocks to monetary policy expectations at the three-month and twelve-month forecast horizons, which conceptually, are related to the target and path factors from Gürkaynak, Sack, and Swanson (2005). Crucially, test statistics imply that these shock measures are less likely to be weak instruments than the high-frequency or narrative monetary shock measures available to date.

Third, we provide a framework that demonstrates that the reaction function estimates and the perceived policy rate path shocks are two sides of the same coin. Having both allows us to disentangle movements in interest rate forecasts into changes in: macroeconomic forecasts (inflation, growth, unemployment), expected responses to macroeconomic forecasts (coefficient estimates), expected smoothing, expected judgment (e.g. due to forward guidance), and finally, expected monetary policy rate path shocks.

Fourth, this paper makes clear, easy-to-implement proposals to improve statistical inference in empirical monetary economics. (i) Macroeconometricians should use Montiel Olea and Pflueger (2013) weak instruments tests when using monetary policy instruments (as in Eberly, Stock, and Wright (2020))—or, better yet, Lewis and Mertens (2025) tests. Doing so, we reject the hypothesis that our instruments are weak instruments at reasonable bias levels, while we cannot do the same for either high-frequency or narrative instruments (as alluded to in Nakamura and Steinsson (2018)). (ii) Macroeconometricians should combine multiple monetary instruments containing distinct information, which not only allows us to test the hypothesis that our instruments are exogenous in the first place (which do not reject), but also to sharpen identification (as investigated in Huang and Neri (2026)). (iii) Macroeconometricians can improve inference by using weak-instrument-robust confidence intervals, which are very easy to implement (Lee, McCrary, Moreira, and Porter (2022)).

The paper is structured as follows. Section (§) 2 presents the econometric model. §3 presents expected reaction function and perceived monetary policy shock estimates. §4 discusses instrument validity, and compares impulse response functions constructed using our instrument to those from the literature. §5 concludes. §A details the reaction function framework. We use the $\mathcal{A}$ symbol throughout to denote areas for which further details can be found in the online $\mathcal{A}$ppendix, which can be accessed by clicking the $\mathcal{A}$.

2 Econometric Model

2.1 Central Bank Reaction Function Framework

The foundation of the econometric model is a central bank reaction function framework, which is derived more formally in Appendix A, §A.1. Here, we present a less formal summary. First, we begin with a general specification of a *Federal Open Market Committee* (FOMC) reaction function, using this to develop our estimator (Eq. 47):

$$\underbrace{\tilde{r}_t}_{\substack{\text{Fed Funds}\\ \text{Rate}}} = \underbrace{\tilde{\rho}_t \tilde{r}_{t-\ell}}_{\substack{\text{Monetary policy}\\ \text{gradualism}}} + \underbrace{\tilde{\phi}_t^{r^*} r_t^*}_{\text{r-star}} + \underbrace{\tilde{\phi}_t^{\pi^*} \pi_t^*}_{\substack{\text{Inflation}\\ \text{target}}} + \underbrace{\left(\tilde{\mathbb{E}}_{t^-}\left[\mathbf{x_{t+h_x}} - \mathbf{x^*_{t+h_x}} \;\middle|\; \tilde{\Omega}_{t^-}\right]\right)'}_{\substack{\text{Macroeconomic}\\ \text{projections}}} \underbrace{\tilde{\boldsymbol{\beta}}_{\mathbf{t}}}_{\substack{\text{Macro}\\ \text{response}}} + \underbrace{\tilde{\varepsilon}_t}_{\substack{\text{Monetary policy}\\ \text{shock}}} \tag{1}$$

Notation. $\sim$ denotes objects attributable to the FOMC. t denotes the time period. ℓ denotes lag length. $*$ denotes long-run equilibrium variables. t^- is some time shortly before the monetary policy shock, $\tilde{\varepsilon}_t$, is realized at t. t^+ is some time shortly after the shock.

Gradualism. Following Clarida, Gali, and Gertler (2000), Woodford (2003), and particularly Coibion and Gorodnichenko (2012), we take the stance throughout this paper that our monetary policy rule should allow for gradualism (smoothing, partial adjustment) i.e. a lagged dependent variable $\tilde{\rho}_t \tilde{r}_{t-\ell}$, where we assume (but do not impose) that $|\tilde{\rho}_t| < 1$. As such, $\tilde{\phi}_t^{\pi^*}, \tilde{\phi}_t^{r^*}, \tilde{\boldsymbol{\beta}}_{\mathbf{t}} \propto (1 - \tilde{\rho}_t)$, so estimates of $\tilde{\phi}_t^{\pi^*}$, $\tilde{\phi}_t^{r^*}$, and $\tilde{\boldsymbol{\beta}}_{\mathbf{t}}$ are implicitly scaled by $(1 - \tilde{\rho}_t)$.

Variables. π_t^* denotes the inflation target. r_t^* denotes the natural rate of interest. $\mathbf{x_t}$ is a k x 1 vector of macroeconomic variables—the inflation rate, real GDP growth rate, and unemployment rate ($\mathcal{A}$). $\mathbf{x_t^*}$ is a matching vector that denotes the long-run equilibrium rates of inflation, real GDP growth, and the unemployment rate, respectively. h_x is a forecast horizon that determines how forward-looking the reaction function is.

Parameters. $\tilde{\boldsymbol{\beta}}_{\mathbf{t}}$ is the FOMC's time-varying response to its macroeconomic projections. $\tilde{\varepsilon}_t$ is a traditional monetary policy shock, which is orthogonal to the FOMC's information set at t^-, $\tilde{\Omega}_{t^-}$, but is $\tilde{\Omega}_t$–measurable (at time t).[1]

Yet, at each FOMC meeting, members do not decide only on the Fed Funds Rate (FFR), but also on *a set* of FFR projections—one for each projection horizon, h.[2] The FOMC updates its policy rate and policy rate projections in a recursive manner, which is detailed in Appendix A. Shifting Equation 1 forward h periods, and then taking the conditional expectation at time t:

$$\underbrace{\tilde{\mathbb{E}}_t\left[\tilde{r}_{t+h} \;\middle|\; \tilde{\Omega}_t\right]}_{\substack{\text{FOMC Fed Funds}\\ \text{Rate projection}}} = \tilde{\mathbb{E}}_t\left[\tilde{\rho}_{t+h}\tilde{r}_{t-\ell+h} + \tilde{\phi}^{r^*}_{t+h} r^*_{t+h} + \tilde{\phi}^{\pi^*}_{t+h}\pi^*_{t+h} + \left(\tilde{\mathbb{E}}_{t^-+h}\left[\mathbf{x_{t+H}} - \mathbf{x^*_{t+H}} \;\middle|\; \tilde{\Omega}_{t^-+h}\right]\right)' \tilde{\boldsymbol{\beta}}_{\mathbf{t+h}} \;\middle|\; \underbrace{\tilde{\Omega}_t}_{\substack{\text{Information}\\ \text{set (FOMC)}}}\right] \tag{2}$$

[1] That is, $\tilde{\varepsilon}_t \perp\!\!\!\perp \tilde{\Omega}_{t^-}$ and $\tilde{\mathbb{E}}_{t^-}\left[\tilde{\varepsilon}_t \;\middle|\; \tilde{\Omega}_{t^-}\right] = 0$ (i.e. unpredictability of structural shocks).

[2] González-Astudillo and Tanvir (2023) also use time-varying, forward-looking FOMC reaction functions.

where $H = h_x + h$, and where $\tilde{\mathbb{E}}_t\left[\tilde{\varepsilon}_{t+h} \;\middle|\; \tilde{\Omega}_t\right] = 0$, and so drops out. Finally, observe that FOMC members are free to, and often do, add judgment to their FFR projections:

$$\underbrace{\tilde{\mathbb{E}}_{t^+}\left[\tilde{r}_{t+h}\right]}_{\substack{\text{FOMC Fed Funds}\\ \text{Rate projection}}} = \tilde{\mathbb{E}}_t\left[\tilde{\rho}_{t+h}\tilde{r}_{t-\ell+h} + \tilde{\phi}^{r^*}_{t+h}r^*_{t+h} + \tilde{\phi}^{\pi^*}_{t+h}\pi^*_{t+h} + \left(\tilde{\mathbb{E}}_{t^-+h}\left[\mathbf{x_{t+H}} - \mathbf{x^*_{t+H}}\right]\right)'\tilde{\boldsymbol{\beta}}_{\mathbf{t+h}}\right] + \underbrace{\tilde{\lambda}^h_{t^+}}_{\substack{\text{FOMC}\\ \text{judgment}}} \tag{3}$$

where, henceforth, we suppress the conditioning sets for brevity only. $\tilde{\lambda}^h_{t^+}$ is known when the FFR projection is finalized (or, $\tilde{\lambda}^h_{t^+}$ is $\tilde{\Omega}_{t^+}$–measurable). This is important because we embed **a monetary policy expectations shock**, $\tilde{\varepsilon}^h_{t^+}$, in judgment ($\tilde{\varepsilon}^h_{t^+}$ is also $\tilde{\Omega}_{t^+}$–measurable). We postulate that judgment is persistent and follows an AR(1) process:

Assumption 1. ***FOMC-expected judgment****, $\tilde{\lambda}^h_t$, evolves according to the following process:*

$$\tilde{\lambda}^h_t = \tilde{\theta}_t\tilde{\lambda}^h_{t-\ell} + \tilde{\varepsilon}^h_t \qquad \tilde{\varepsilon}^h_t \sim i.i.d.\left(0, \sigma^2_{\tilde{\varepsilon}_t}\right) \tag{4}$$

where $\tilde{\varepsilon}^h_t$ is a shock to expected judgment, and can equivalently be interpreted as a monetary policy *expectations* shock. It is assumed that $|\tilde{\theta}_t| < 1$, but this is not imposed. The AR(1) structure captures the idea that judgment that is persistent over time (due to, for example, ongoing forward guidance).[3] If new judgment enters through $\tilde{\varepsilon}^h_t$, $\tilde{\lambda}^h_{t-1}$ allows us to capture its persistent-yet-decaying nature:

$$\tilde{\lambda}^h_t = \tilde{\varepsilon}^h_t + \tilde{\theta}_t\tilde{\varepsilon}^h_{t-1} + \tilde{\theta}_t\tilde{\theta}_{t-1}\tilde{\varepsilon}^h_{t-2} + \tilde{\theta}_t\tilde{\theta}_{t-1}\tilde{\theta}_{t-2}\tilde{\varepsilon}^h_{t-3} + \cdots \tag{5}$$

What is FOMC-expected judgment? $\tilde{\lambda}^h_t$ represents other influences on the FFR projection not captured by the systematic component of the reaction function:

1. **Timing/rounding**. The FFR changes in (sometimes multiple) 25-basis-point increments, but a reaction function output will not neatly coincide with 25-basis point increments from zero. This can be conceptualised as timing/rounding choice. For example, if FOMC projections suggest that the target range for the FFR should be 3.375% to 3.625% (where there prevailing range is 3.5% to 3.75%), then the FOMC must decide whether to round up (hold the FFR now, and likely cut at the next meeting), or to round down (cut at the current meeting and likely hold at the next meeting). This shows up as positive or negative judgment, respectively, in the $\tilde{\lambda}^h_t$ term.
2. **Forward guidance** from the past and/or newly introduced.[4]
3. **Risk management** due to elevated projection uncertainty (e.g. anticipated insurance rate cuts due to asymmetric risks to the outlook).
4. **Model uncertainty** (e.g. about the monetary policy rule itself, r-star, etc.).
5. **Dissent** to make a point (to the market and/or to other FOMC participants).

[3] Campbell, Evans, Fisher, Justiniano, Calomiris, and Woodford (2012) capture a variant of this idea when they describe the central bank reaction function error as the sum of a contemporaneous policy rate shock and a series of past, predictable FG shocks (Eq.1, p.22). Miranda-Agrippino and Ricco (2021) also capture a similar idea (Eq.3, p.79): "*Also, the central bank can announce—or leak—at time $t-1$ a deviation from the monetary policy rule, $w_{t|t-1}$, that takes places at time t. This announced deviation—that can be thought of as due to either implicit or explicit forward guidance—creates a wedge between monetary policy rule residuals and monetary policy shocks*."

[4] The following forward guidance may contain this type of judgment "*In particular, the Committee also decided today to keep the target range for the federal funds rate at 0 to 1/4 percent and currently anticipates that exceptionally low levels for the federal funds rate are likely to be warranted at least through mid-2015*." (source).

6. **Persuasion** by judgment expressed in The Tealbook, staff presentations, or other FOMC participants during deliberations.

7. **An unanticipated event** not reflected in the macroeconomic projections or macro response ($\tilde{\beta}_{\mathbf{t}}$).

Given the sources of judgment, we hypothesize that the law of motion for judgment is a persistent process (Eq. 4). $\tilde{\lambda}_t^h$ is also intended to account for the role of persistent shocks in monetary policy rules discussed in Rudebusch (2002) and Coibion and Gorodnichenko (2012), and the serial correlation in shock measures discussed in Miranda-Agrippino and Ricco (2021) and Bauer and Swanson (2023).

2.2 The Expected Central Bank Reaction Function

Second, at all times t other than the four FOMC meetings per year at which projections are released, agents must rely on their *perception* of the FOMC's FFR projection, that is, how the projection may have changed in response to data and events. Because a FFR path projection is intrinsically forward-looking, it implies the existence of an *expected* reaction function. To derive this, we take expectations of Eq. 3 with respect to the information set of forecaster i at time t.

$$\ddot{\mathbb{E}}_{it^+}\left[\tilde{\mathbb{E}}_{t^+}\left[\tilde{r}_{t+h}\right]\right] = \ddot{\mathbb{E}}_{it^+}\left[\tilde{\mathbb{E}}_t\left[\tilde{\rho}_{t+h}\tilde{r}_{t-\ell+h} + \tilde{\phi}^{r^*}_{t+h}r^*_{t+h} + \tilde{\phi}^{\pi^*}_{t+h}\pi^*_{t+h} + \left(\tilde{\mathbb{E}}_{t^-+h}\left[\mathbf{x_{t+H}} - \mathbf{x^*_{t+H}}\right]\right)'\tilde{\beta}_{\mathbf{t+h}}\right] + \tilde{\lambda}^h_{t^+}\right] \quad (6)$$

where $\ddot{\mathbb{E}}_{it^+}\left[. \mid \ddot{\Omega}_{it^+}\right]$ is forecaster i's ($i = 1, ..., N$) expectation conditional on their information set, $\ddot{\Omega}_{it^+}$ at time t^+. We assume that both the FOMC and forecasters know their own expected coefficients ($\tilde{\Omega}_{t^+}$–measurable and $\ddot{\Omega}_{it^+}$–measurable, respectively—Assumption 10 in §A.2). This assumption implies that we can move those coefficients outside the the conditional expectations. Once outside, we denote those coefficients as, for example, ρ^h_{it}. The analogous definitions apply for all other coefficients in Eq. 8 below. Further, we denote:

$$\ddot{\mathbb{E}}_{it^+}\left[\tilde{\lambda}^h_{t^+}\right] := \lambda^h_{it^+} = \theta_{it^+}\lambda^h_{i,t^+-\ell} + \varepsilon^h_{it^+} \quad (7)$$

This is **forecaster-expected FOMC judgment**, which is the forecaster's expectation of FOMC-expected judgment (a more precise and technical definition can be found in §A.2, Equations 77 to 87). Finally, for notational simplicity, moving forward, we abstract from the notation for slight timing differences in expectation formation (t^-, t, t^+) without loss of generality. So, we can rewrite Eq. 6 as follows:[5]

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t+h}\right]\right] = \rho^h_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-\ell+h}\right]\right] + \phi^{r^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[r^*_{t+h}\right]\right] + \phi^{\pi^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\pi^*_{t+h}\right]\right] + \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\left(\mathbf{x_{t+H}} - \mathbf{x^*_{t+H}}\right)'\right]\right]\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}} + \lambda^h_{it} \quad (8)$$

2.3 Mapping The Expected Reaction Function to the Data

Third, we take the forecaster-expected reaction function (Eq. 8), the parameters of which we seek to estimate, to the monthly (t) professional forecast panel data. So, henceforth, let t refer to monthly forecasting periods and set ℓ=1. Here, we use the *Consensus Economics* data, but the *Blue Chip Financial Forecasts* data have similar attributes, and are also suited to our framework. We begin by defining the relationship between the unobservable dependent variable from above, $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t+h}\right]\right]$, and an observable-yet-noisy measure of the dependent variable, y^h_{it}. y^h_{it} is individual forecaster i's forecast of the risk-free, short-term interest rate (the three-month Treasury bill rate) at time t and forecast horizon h.

[5] In general, if $\mathcal{H} \subseteq \mathcal{G}$ (filtrations), and Z is integrable $\Rightarrow \mathbb{E}\left[\mathbb{E}\left[Z \mid \mathcal{G}\right] \mid \mathcal{H}\right] = \mathbb{E}\left[Z \mid \mathcal{H}\right]$, by the Law of Iterated Expectations. Since we can assume $\tilde{\Omega}_t \subseteq \tilde{\Omega}_{t^-+h}$ (filtrations) $\Rightarrow \tilde{\mathbb{E}}_t\left[\tilde{\mathbb{E}}_{t^-+h}\left[Z\right]\right] = \tilde{\mathbb{E}}_t\left[Z\right]$. Then, if we are willing to assume that $\tilde{\beta}_{\mathbf{t+h}}$ is $\tilde{\Omega}_t$–measurable (deterministic conditional on $\tilde{\Omega}_t$, or, that the FOMC knows its own macro responsiveness), then it follows that $\ddot{\mathbb{E}}_{it^+}\left[\tilde{\mathbb{E}}_t\left[\left(\tilde{\mathbb{E}}_{t^-+h}\left[\mathbf{x_{t+H}} - \mathbf{x^*_{t+H}}\right]\right)'\tilde{\beta}_{t+h}\right]\right] = \ddot{\mathbb{E}}_{it^+}\left[\left(\tilde{\mathbb{E}}_t\left[\mathbf{x_{t+H}} - \mathbf{x^*_{t+H}}\right]\right)'\tilde{\beta}_{t+h}\right]$.

$$\underbrace{y_{it}^{h}}_{\substack{\text{Private interest}\\ \text{rate forecast}}} = \underbrace{\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\tilde{r}_{t+h}\right]\right]}_{\substack{\text{Expected FOMC}\\ \text{FFR projection}}} + \underbrace{\ddot{\mathbb{E}}_{it}\left[e_{t+h}^{y}\right]}_{\substack{\text{Expected FOMC}\\ \text{FFR projection } \textit{error}}} + \underbrace{\ddot{\mathbb{E}}_{it}\left[\tau_{t+h}\right]}_{\substack{\text{Expected}\\ \text{spread}}} \tag{9}$$

$\ddot{\mathbb{E}}_{it}\left[e_{t+h}^{y}\right]$ is the expected FOMC FFR projection error. If $\ddot{\mathbb{E}}_{it}\left[e_{t+h}^{y}\right] \neq 0$, the forecaster thinks the FOMC is wrong about future policy rates as in Caballero and Simsek (2022). $\ddot{\mathbb{E}}_{it}\left[\tau_{t+h}\right]$ is the expected spread between the three-month Treasury bill rate and the FFR. It is an amalgam of the convenience yield, liquidity premium, and short-horizon term premium (although evidence suggests that this component is small at such short forecast horizons, e.g. Schmeling, Schrimpf, and Steffensen (2022)). Next, we substitute Eq. 8 into $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\tilde{r}_{t+h}\right]\right]$ in Eq. 9:

$$y_{it}^{h} = \underbrace{\rho_{it}^{h}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\tilde{r}_{t-1+h}\right]\right] + \phi_{it}^{r^{*}}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[r_{t+h}^{*}\right]\right] + \phi_{it}^{\pi^{*}}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\pi_{t+h}^{*}\right]\right] + \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)'\right]\right]\ddot{\boldsymbol{\beta}}_{\mathbf{it}}^{\mathbf{h}} + \lambda_{it}^{h}}_{\text{Expected FOMC FFR projection}} + \ddot{\mathbb{E}}_{it}\left[e_{t+h}^{y}\right] + \ddot{\mathbb{E}}_{it}\left[\tau_{t+h}\right] \tag{10}$$

We now need to replace each expectation above with *observable* forecasts to facilitate estimation. We can, similar to Eq. 9, define a relationship between the unobservable main independent variables in Eq. 10, $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)'\right]\right]$, and an observable-yet-noisy measure of the main independent variables, $\ddot{\mathbb{E}}_{it}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)\right]$, or, $\mathbf{x_{it}^{H}}$ for short. So, $\mathbf{x_{it}^{H}}$ is a vector of private forecaster i's macroeconomic forecasts (inflation, growth, unemployment rate) at time t and forecast horizon $H = h_x + h$.

$$\underbrace{\mathbf{x_{it}^{H}}}_{\substack{\text{Private macro}\\ \text{forecasts}}} := \underbrace{\ddot{\mathbb{E}}_{it}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)\right]}_{\substack{\text{Private macro}\\ \text{forecasts}}} = \underbrace{\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)\right]\right]}_{\substack{\text{Expected FOMC}\\ \text{macro projections}}} + \underbrace{\ddot{\mathbb{E}}_{it}\left[\mathbf{e_{t+H}^{x}}\right]}_{\substack{\text{Expected FOMC}\\ \text{macro projection } \textit{errors}}} \tag{11}$$

However, $\ddot{\mathbb{E}}_{it}\left[\mathbf{e_{t+H}^{x}}\right]$ is also unobservable, so we must otherwise isolate $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)\right]\right]$.

Assumption 2. *Given that both the FOMC and our sample group are sophisticated forecasters forecasting the same common macro aggregates, and that expected FOMC projections are anchored by the publicly-released Summary of Economic Projections every three months (while private forecasts are comparatively unanchored), we can infer that:*

$$span\left\{\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)\right]\right]\right\} \subseteq span\left\{\ddot{\mathbb{E}}_{it}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)\right]\right\} \tag{12}$$

That is, the private-sector macroeconomic forecasts span those same forecasters' expected FOMC macroeconomic projections (that the space generated by private forecasts contains the space generated by the perceived FOMC projections). This implies that (where $\mathbf{A_{it}^{H}}$ is a k x k nonsingular matrix):

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\left(\mathbf{x_{t+H}} - \mathbf{x_{t+H}^{*}}\right)\right]\right] = \mathbf{A_{it}^{H}}\mathbf{x_{it}^{H}} + \upsilon_{\mathbf{it}}^{\mathbf{H}} \quad s.t. \quad \ddot{\mathbb{E}}_{it}\left[\ \upsilon_{\mathbf{it}}^{\mathbf{H}} \ \middle|\ \mathbf{x_{it}^{H}}\ \right] = \mathbf{0} \tag{13}$$

Assumption 3. *Third, we treat the inflation target, neutral rate, and T-bill spread in Eq. 10 as common-knowledge Martingales:*

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t}\left[\pi_{t+h}^{*}\right]\right] = \ddot{\mathbb{E}}_{it}\left[\pi_{t+h}^{*}\right] = \tilde{\mathbb{E}}_{t}\left[\pi_{t+h}^{*}\right] := \pi_{t}^{*} \tag{14}$$

which is the prevailing inflation target, is observable, and can be substituted directly into Eq. 10. Similarly:

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[r^*_{t+h}\right]\right] = \ddot{\mathbb{E}}_{it}\left[r^*_{t+h}\right] = \tilde{\mathbb{E}}_t\left[r^*_{t+h}\right] := r^*_t \tag{15}$$

where r^*_t could, in principle, be the last-published estimates from the FOMC Summary of Economic Projections (SEP). Yet, our data frequency is monthly, and the SEP data are only available from 2012. So, instead, we use the widely-used Laubach and Williams (2003) r-star estimates (updated versions of which are provided Holston, Laubach, and Williams (2023) and available on the FRBNY website). These are observable, and so can be substituted directly into Eq. 10. Consistent with the literature on monetary policy rules, the following is implied in Eq. 10, $\phi^{\pi^*}_{it} \propto \left(1-\rho^h_{it}\right)$, $\phi^{r^*}_{it} \propto \left(1-\rho^h_{it}\right)$.

$$\ddot{\mathbb{E}}_{it}\left[\tau_{t+h}\right] := \tau_t \tag{16}$$

We can now substitute Eq. 13, 14, 15, 16 into Eq. 10, which moves us towards an equation we can estimate with observable data:[6]

$$y^h_{it} = \rho^h_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] + \phi^{r^*}_{it} r^*_t + \phi^{\pi^*}_{it}\pi^*_t + \phi^{\tau}_{it}\tau_t + \mathbf{x}^{\mathbf{H}\prime}_{\mathbf{it}}\beta^{\mathbf{h}}_{\mathbf{it}} + \ddot{\mathbb{E}}_{it}\left[e^y_{t+h}\right] + \upsilon^{\mathbf{H}\prime}_{\mathbf{it}}\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}} + \lambda^h_{it} \tag{17}$$

$$s.t. \qquad \beta^{\mathbf{h}}_{\mathbf{it}} := \mathbf{A}^{\mathbf{H}\prime}_{\mathbf{it}}\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}} \tag{18}$$

Assumption 4. *Drawing on Chudik and Pesaran (2015), Stock and Watson (2016), and Giacomini, Lu, and Smetanina (2024), the vector of private macroeconomic forecasts,* $\mathbf{x}^{\mathbf{H}}_{\mathbf{it}}$*, can be described by the following linear covariance stationary dynamic heterogeneous panel data model:*

$$\mathbf{x}^{\mathbf{H}}_{\mathbf{it}} := \ddot{\mathbb{E}}_{it}\left[\left(\mathbf{x}_{\mathbf{t+H}} - \mathbf{x}^*_{\mathbf{t}}\right)\right] = \mathbf{c}_{xi} + \mathbf{\Gamma}^{\mathbf{H}}_{\mathbf{it}}\mathbf{f}_{\mathbf{t}} + \mathbf{v}^{\mathbf{H}}_{\mathbf{it}} \qquad s.t. \tag{19}$$

$$\mathbf{\Gamma}^{\mathbf{H}}_{\mathbf{it}} = \bar{\mathbf{\Gamma}}^{\mathbf{H}}_{\mathbf{t}} + \ddot{\mathbf{\Gamma}}^{\mathbf{H}}_{\mathbf{it}}, \qquad \frac{1}{N}\sum_i \ddot{\mathbf{\Gamma}}^{\mathbf{H}}_{\mathbf{it}} = \mathbf{0}, \qquad rank\left(\frac{1}{N}\sum_i \ddot{\mathbf{\Gamma}}^{\mathbf{H}}_{\mathbf{it}}\ddot{\mathbf{\Gamma}}^{\mathbf{H}\prime}_{\mathbf{it}}\right) = m \tag{20}$$

$\mathbf{c}_{xi}$ is an individual fixed effect. $\mathbf{f}_{\mathbf{t}}$ is an m x 1 vector of latent factors. $\mathbf{\Gamma}^{\mathbf{H}}_{\mathbf{it}}$ is an k *x* m matrix of factor loadings. The factor structure reflects the effect of contemporaneous structural shocks on expectations of future macroeconomic gaps. $\mathbf{v}^{\mathbf{H}}_{\mathbf{it}}$ follows a general linear covariance stationary process distributed independently of the idiosyncratic errors in Eq. 7, ε^h_{it}.

Assumption 5. *The vector* $\mathbf{f}_{\mathbf{t}}$ *is composed of two structural shocks:* $\mathbf{f}_{\mathbf{t}}{}' = \left[f^s_t, f^d_t\right]$ *where* f^s_t *is a supply factor, and* f^d_t *is a demand factor (where a monetary policy expectations shock,* ε^h_{it}*, was introduced in Eq. 7 earlier).*

Giacomini, Lu, and Smetanina (2024), among others, explain how we can represent macroeconomic forecast survey data using time-varying (dynamic) factor models. Stock and Watson (2016) provide strong empirical support that such factor models capture macroeconomic time series well—that is, an m *x* 1 vector of latent common factors, which are contemporaneously orthogonal and possibly serially correlated. $\mathbf{f}_{\mathbf{t}}$ can be interpreted as perceived structural shocks—the interpretation of which requires additional assumptions, such as those provided in our reaction function framework here, or those based on impulse response function interpretation.[7] Accordingly, we postulate that we can approximate a number

[6] Where we have applied the notation from Eq. 11, $\mathbf{x}^{\mathbf{H}}_{\mathbf{it}}$, to Eq. 13.

[7] See Giacomini, Lu, and Smetanina (2024), p.10. The elements of $\mathbf{f}_{\mathbf{t}}$ *can* be serially correlated (see p.15).

of the elements of the expected reaction functions using time-varying dynamic factor models.

Next, we address a remaining unobservable term in Eq. 17, the expected (forward-looking) monetary policy gradualism term, $\rho^h_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right]$. Observe that we can shift Eq. 9 back one period:

$$y^h_{i,t-1} = \ddot{\mathbb{E}}_{i,t-1}\left[\tilde{\mathbb{E}}_{t-1}\left[\tilde{r}_{t-1+h}\right]\right] + \ddot{\mathbb{E}}_{i,t-1}\left[e^y_{t-1+h}\right] + \ddot{\mathbb{E}}_{i,t-1}\left[\tau_{t-1+h}\right] \tag{21}$$

which corresponds to the h-period-ahead forecast formed at time $t-1$. Decomposing the first term on the right-hand-side using the following identity:

$$y^h_{i,t-1} = \underbrace{\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] - \Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right]}_{=\ddot{\mathbb{E}}_{i,t-1}\left[\tilde{\mathbb{E}}_{t-1}\left[\tilde{r}_{t-1+h}\right]\right]} + \ddot{\mathbb{E}}_{i,t-1}\left[e^y_{t-1+h}\right] + \ddot{\mathbb{E}}_{i,t-1}\left[\tau_{t-1+h}\right] \tag{22}$$

Rearranging to isolate $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right]$:

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] = y^h_{i,t-1} + \Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] - \ddot{\mathbb{E}}_{i,t-1}\left[e^y_{t-1+h}\right] - \ddot{\mathbb{E}}_{i,t-1}\left[\tau_{t-1+h}\right] \tag{23}$$

Using Lemma A.1:

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] = y^h_{i,t-1} + \underbrace{\left(\varphi^{r^*}_{it} r^*_t + \varphi^{\pi^*}_{it}\pi^*_t + \mathbf{\Psi^h_{it}}\mathbf{f_t} + \mathbf{\check{\Psi}^h_{it}}\mathbf{\Delta f_t} + \Xi^h_{it}\right)}_{=\Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] \text{ by Lemma A.1}} - \ddot{\mathbb{E}}_{i,t-1}\left[e^y_{t-1+h}\right] - \ddot{\mathbb{E}}_{i,t-1}\left[\tau_{t-1+h}\right] \tag{24}$$

We then substitute this into Eq. 17 as desired:

$$\begin{aligned} y^h_{it} = \rho^h_{it}\Big(y^h_{i,t-1} + \left(\varphi^{r^*}_{it} r^*_t + \varphi^{\pi^*}_{it}\pi^*_t + \mathbf{\Psi^h_{it}}\mathbf{f_t} + \mathbf{\check{\Psi}^h_{it}}\mathbf{\Delta f_t} + \Xi^h_{it}\right) - \ddot{\mathbb{E}}_{i,t-1}\left[e^y_{t-1+h}\right] - \ddot{\mathbb{E}}_{i,t-1}\left[\tau_{t-1+h}\right]\Big) + \\ \phi^{r^*}_{it} r^*_t + \phi^{\pi^*}_{it}\pi^*_t + \phi^{\tau}_{it}\tau_t + \mathbf{x^{H'}_{it}}\beta^{\mathbf{h}}_{\mathbf{it}} + \ddot{\mathbb{E}}_{it}\left[e^y_{t+h}\right] + \upsilon^{\mathbf{H'}}_{\mathbf{it}}\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}} + \lambda^h_{it} \end{aligned} \tag{25}$$

We collect like variables and parameters in vectors to compress the notation, and apply Eq. 16 again:

$$\begin{aligned} y^h_{it} = \rho^h_{it} y^h_{i,t-1} + \underbrace{\rho^h_{it}\varphi^{r^*}_{it} r^*_t + \rho^h_{it}\varphi^{\pi^*}_{it}\pi^*_t + \phi^{r^*}_{it} r^*_t + \phi^{\pi^*}_{it}\pi^*_t}_{:=\mathbf{z_t'}\mathbf{\Phi_{it}}} \quad + \quad \mathbf{x^{H'}_{it}}\beta^{\mathbf{h}}_{\mathbf{it}} \quad + \\ \underbrace{\rho^h_{it}\mathbf{\Psi^h_{it}}\mathbf{f_t} + \rho^h_{it}\mathbf{\check{\Psi}^h_{it}}\mathbf{\Delta f_t} + \ddot{\mathbb{E}}_{it}\left[e^y_{t+h}\right] - \rho^h_{it}\ddot{\mathbb{E}}_{i,t-1}\left[e^y_{t-1+h}\right]}_{:=\sum_{\ell=0}^{L}\mathbf{f'_{t-\ell}}\gamma_{\mathbf{it}\ell}} + \underbrace{\upsilon^{\mathbf{H'}}_{\mathbf{it}}\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}} + \rho^h_{it}\Xi^h_{it} + \phi^{\tau}_{it}\tau_t - \rho^h_{it}\phi^{\tau}_{i,t-1}\tau_{t-1}}_{:=\zeta_{it}} + \lambda^h_{it} \end{aligned} \tag{26}$$

where $\mathbf{\Phi'_{it}} = \left[\rho^h_{it}\varphi^{r^*}_{it} + \phi^{r^*}_{it}, \rho^h_{it}\varphi^{\pi^*}_{it} + \phi^{\pi^*}_{it}\right]$, $\mathbf{z_t}' = [r^*_t, \pi^*_t]$. Using the identity $\mathbf{\Delta f_t = f_t - f_{t-1}}$, the factor innovations can be expressed as a distributed lag in $\mathbf{f_t}$. $\ddot{\mathbb{E}}_{it}\left[e^y_{t+h}\right]$ is assumed to be spanned by the factor space. ζ_{it} is expected to be small under the conditions of Lemma A.1.[8]

[8]Forecasts of Treasury bill rates differ from forecasts of the federal funds rate because of expected liquidity, convenience-yield, and other market-structure effects. To assess the empirical relevance of this component, we construct an observable proxy based on the spread between Blue Chip forecasts of Treasury bill rates and federal funds rates. This proxy exhibits only a weak correlation with the recovered monetary policy shock series (approximately 0.04–0.07 depending on the sample). Further, including the spread in the DCCE specification materially destabilizes coefficient estimates, consistent with the fact that the proxy is mechanically linked to the dependent variable and therefore difficult to separately identify from the forecast object itself. Consequently, we treat this component as small and largely absorbed by the common-factor structure and residual component of the model.

2.4 Identification Strategy

In this section, we develop our estimator for Eq. 26, which, based on the derivations flowing through §2.1, §2.2, and §2.3, begins as follows:

$$y_{it} = \rho_{it} y_{i,t-1} + \mathbf{z_t}'\mathbf{\Phi_{it}} + \mathbf{x_{it}}'\beta_{\mathbf{it}} + \sum_{\ell=0}^{L} \mathbf{f}'_{\mathbf{t}-\ell}\gamma_{\mathbf{it}\ell} + u_{it} \tag{27}$$

$$u_{it} = \zeta_{it} + \lambda_{it} \tag{28}$$

$$\lambda_{it} = \theta_{it}\lambda_{i,t-1} + \varepsilon_{it} \tag{29}$$

$$\varepsilon_t := (\varepsilon_{1t}, \ldots, \varepsilon_{Nt})' \qquad \mathbb{E}[\varepsilon_t] = \mathbf{0}, \qquad \text{Var}[\varepsilon_t] = \mathbf{\Sigma}_\varepsilon \tag{30}$$

$$\mathbb{E}[\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_s'] = \mathbf{0}, \qquad t \neq s \tag{31}$$

$\mathbf{\Sigma}_\varepsilon$ is a positive definite covariance matrix allowing for contemporaneous cross-sectional dependence.

Assumption 6. *The other parameters of interest follow a random coefficient model:*

$$\rho_{it} = \rho + \eta_{\rho i} + \nu_{\rho t}, \quad \eta_{\rho i} \sim iid\left(0, \sigma^2_{\eta_\rho}\right), \quad \nu_{\rho t} \sim iid\left(0, \sigma^2_{\nu_\rho}\right) \tag{32}$$

$$\boldsymbol{\beta_{it}} = \boldsymbol{\beta} + \boldsymbol{\eta_{\beta i}} + \boldsymbol{\nu_{\beta t}}, \quad \eta_{\beta \mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\eta_\beta}\right), \quad \nu_{\beta \mathbf{t}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\nu_\beta}\right) \tag{33}$$

$$\boldsymbol{\gamma_{it\ell}} = \boldsymbol{\gamma_\ell} + \boldsymbol{\eta_{\gamma i\ell}} + \boldsymbol{\nu_{\gamma t\ell}}, \quad \eta_{\gamma \mathbf{i}\ell} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\eta_\gamma}\right), \quad \nu_{\gamma \mathbf{t}\ell} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\nu_\gamma}\right) \tag{34}$$

$$\theta_{it} = \theta + \eta_{\theta i} + \nu_{\theta t}, \quad \eta_{\theta i} \sim iid\left(0, \sigma^2_{\eta_\theta}\right), \quad \nu_{\theta t} \sim iid\left(0, \sigma^2_{\nu_\theta}\right) \tag{35}$$

where each parameter is centered around a common population mean and varies through additive unit-specific and time-specific random effects. This helps clarify that, in the model, there are two main sources of perceived monetary policy shocks: variation in the perceived reaction function parameters—captured through the unit-specific and time-specific random coefficient components ($\boldsymbol{\eta}_i$ and $\boldsymbol{\nu}_t$), and shocks to perceived FOMC judgment (λ_{it})—captured by the idiosyncratic disturbances (ε_{it}).

Identification challenge 1. Overparameterization. This generalized forecaster-expected reaction function is plausibly an accurate representation of the true expected reaction function, but as such, it is overparameterized because we know that parameters vary by forecaster, i (Altavilla and Giannone (2017)), by country, c (Sutherland (2023)), time, t (Bauer, Pflueger, and Sunderam (2024)), and forecast horizon, h (Hamilton, Pruitt, and Borger (2011)). So, it cannot be estimated without further assumptions. Specifically, we will need to pool data across one of the three main dimensions of the data: i, t, or h.

Identification challenge 2. Forecaster (i) heterogeneity. We cannot pool data across i (as in Bauer, Pflueger, and Sunderam (2024)). If the assumption of parameter homogeneity across forecasters does not hold in a *dynamic* model such as ours, then Pesaran and Smith (1995) shows that a pooled OLS estimator leads to inconsistent and potentially highly misleading estimates of the coefficients. To account for forecaster parameter heterogeneity, we estimate the model forecaster by forecaster and compute the cross-sectional average of the estimated coefficients, which is a novel approach. This is the mean-group estimator, which Pesaran and Smith (1995) show is a consistent estimator of the mean coefficient. By contrast, the pooled estimator imposes parameter homogeneity and instead converges to a variance-weighted average of the individual coefficients, where the weights depend on the regressor variation for each forecaster (and, so, implicitly, the sample size of each forecaster in each window). And, as this paper shows, the two approaches can yield very different estimates, especially with respect to monetary policy inertia (see Figures 2 and 3).

Identification challenge 3. Forecast horizon (h) heterogeneity. We opt not to pool data across h. As such, we have abstracted from the h superscript in Eq. 27 because we will estimate the model separately for each h. Arguably, imposing parameter homogeneity across forecast horizons is a fairly innocuous assumption (as used in Bauer, Pflueger, and Sunderam (2024)), yet in our structural setting, it would also be a missed opportunity: we can extract two separate monetary policy rate path shocks at two forecast horizons (3-months and 12-months) that gives us two separate econometric instruments—each with a slightly different interpretation (§3.4), akin to the target and path factors from Gürkaynak, Sack, and Swanson (2005). In turn, when estimating impulse response functions, having two instruments for a single endogenous regressor yields an overidentified model (see Eq. 55), which allows us to test the exogeneity of the instruments, and which leads to new insights (see §4.1 and §4.2).

Identification challenge 4. Time-varying (t) parameters. We opt to pool data across t, and use seven-year rolling windows. This forces us to average away some of the time variation, yet still allows for a high degree of parameter variation. Bauer, Pflueger, and Sunderam (2024) also use seven-year rolling windows (T=84) following the same reasoning.[9] Hence, we can further simplify Eq. 27 by focusing on the equation for each rolling window. So, coefficients no longer vary by t *within* a given window, but are time-varying because they can vary *by* rolling window:

$$y_{it} = \rho_i y_{i,t-1} + \mathbf{z_t}'\mathbf{\Phi_i} + \mathbf{x_{it}}'\beta_\mathbf{i} + \sum_{\ell=0}^{L} \mathbf{f_{t-\ell}}'\gamma_{\mathbf{i}\ell} + u_{it} \tag{36}$$

Accordingly, the heterogeneous coefficients now follow a simpler random coefficients model within each window:[10]

$$\rho_\mathbf{i} = \rho + \eta_{\rho\mathbf{i}}, \quad \eta_{\rho\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_\rho\right) \tag{37}$$

$$\beta_\mathbf{i} = \beta + \eta_{\beta\mathbf{i}}, \quad \eta_{\beta\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_\beta\right) \tag{38}$$

$$\gamma_{\mathbf{i}\ell} = \gamma_\ell + \eta_{\gamma\mathbf{i}\ell}, \quad \eta_{\gamma\mathbf{i}\ell} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_\gamma\right) \tag{39}$$

$$\theta_\mathbf{i} = \theta + \eta_{\theta\mathbf{i}}, \quad \eta_{\theta\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_\theta\right) \tag{40}$$

Identification challenge 5. Recovering the structural beta parameters. Comparison of Eq. 17, 18, and 36 implies that these estimated reduced-form coefficients satisfy:[11]

$$\beta_\mathbf{i} = (\mathbf{A_i})' \ddot{\beta}_\mathbf{i} \tag{41}$$

Provided that $\mathbf{A_i}$ is nonsingular, the structural macro-response coefficients can therefore be recovered as

$$\widehat{\ddot{\beta}}_i = \left[\left(\widehat{\mathbf{A}}_i\right)'\right]^{-1} \widehat{\beta}_i \tag{42}$$

where both $\widehat{\beta}_i$ (blue line) and $\widehat{\ddot{\beta}}_i$ (black line) are plotted in Figures 2 and 3.

Identification challenge 6. Cross-sectional dependence. A priori, the presence of $\mathbf{f_t}$ in Eq. 36 will tend to introduce cross-sectional dependence between forecasts; this leads to substantial bias and size distortions if ignored (Pesaran (2006), p.970). Unsurprisingly, in rolling window regressions, we invariably

[9] Bauer, Pflueger, and Sunderam (2024) also estimate a series of cross-sectional regressions with $N \approx 100$.

[10] Where each $\mathbf{\Sigma}$ is a symmetric nonnegative definite matrix, and where each of the $\eta_\mathbf{i}$ are independently distributed of $\gamma_\mathbf{j}, \mathbf{\Gamma_j}, \mathbf{v_{jt}}, \mathbf{f_t}, \varepsilon_{jt} \;\; \forall i, j, t$. Further, ρ_i lies strictly inside the unit circle.

[11] $\beta^\mathbf{h}_\mathbf{it}$ is identified by the variation in $\mathbf{v}^\mathbf{h}_\mathbf{it}$, which is the idiosyncratic movement in $\mathbf{x}^\mathbf{H}_\mathbf{it}$. See below for the rationale.

detect cross-sectional dependence across forecasts using the test from Pesaran (2015).[12]

To account for cross-sectional dependence across forecasters, that is, exposure to common shocks ($\mathbf{f_t}$), we adapt the Common Correlated Effects (CCE) estimator developed by Pesaran (2006). This approach controls for unobserved common factors by augmenting each forecaster's regression with cross-sectional averages (across all forecasters that period) of the dependent and independent variables.

The *Dynamic* Common Correlated Effects (Chudik and Pesaran (2015)) estimator extends the CCE framework to allow for lagged dependent variables (in addition to weakly exogenous regressors), which is crucial here. The DCCE estimator has excellent small-sample properties given even substantial heterogeneity, dynamics, and relatively small values of N and T (e.g. $N = 20$, $T = 20$). DCCE allows one to account for both common effects and parameter heterogeneity.

In the DCCE framework, the $\mathbf{f_t}$ can exhibit an arbitrary degree of correlation among themselves and with the individual-specific regressors, $\mathbf{x_{it}}$. The factors (or cross-sectional averages) are assumed to span the common component, although later, we allow for additional factors ($\mathbf{g_{it}}$) to be certain. The ε_{it} in λ_{it} can be spatially correlated. So, the $\mathbf{x}_{it}$, which are weakly exogenous regressors within the broader linear covariance stationary dynamic heterogeneous panel data model, are assumed to be generated by the following:

$$\omega_{\mathbf{it}} = \begin{bmatrix} \mathbf{x_{it}} \\ \mathbf{g_{it}} \end{bmatrix} = \mathbf{c}_{\omega i} + \mathbf{\Gamma_i'} \mathbf{f_t} + \mathbf{v_{it}} \tag{43}$$

The $mx1$ vector of unobserved common factors, $\mathbf{f_t}$, is assumed to be covariance stationary, and distributed independently of ε_{it} and $\mathbf{v_{it}}$ $\forall i, t, t'$. $\mathbf{\Gamma_i}$ is an m x k matrix of factor loadings. $\gamma_\mathbf{i}$ and $\mathbf{\Gamma_i}$ are independently and identically distributed across i, and of $\mathbf{f_t}$, for all i and t, with mean γ and $\mathbf{\Gamma}$, respectively.[13] $\mathbf{g_{it}}$ is a vector of variables, such as consumption expenditure and business investment forecasts, that do not enter Eq. 36 directly, but that can be used to proxy for additional factors if necessary. $\mathbf{v_{it}}$ follows general linear covariance stationary process distributed independently of the idiosyncratic errors, ε_{it}.

If the errors, u_{it}, are cross-sectionally weakly dependent, and $\mathbb{E}[\gamma_\mathbf{i}] = \mathbb{E}[\gamma]$, the lagged dependent variable is no longer strictly exogenous. Yet, Chudik and Pesaran (2015) show that the estimator can become consistent if it is augmented with cross-sectional (weighted) averages of the dependent variable and the individual-specific regressors. Specifically, the floor of $\sqrt[3]{T}$ lags of the cross-section averages of the regressors are added to Eq. 36.[14] We denote the number of lags needed by $p_T = \sqrt[3]{T}$.[15] We also need at least as many (independent) cross-sectional averages as we have latent factors. That condition is satisfied here because we use the cross-sectional averages of the very variables from which the factors originate as well as two others.

$$y_{it} = \rho_i y_{i,t-1} + \mathbf{z_t'} \mathbf{\Phi_i} + \mathbf{x_{it}}' \beta_\mathbf{i} + \sum_{\ell=0}^{p_T} \overline{\mathbf{x}}_{t-\ell}' \delta_{\mathbf{i}\ell} + \sum_{\ell=0}^{p_T} \overline{\mathbf{g}}_{t-\ell}' \kappa_{\mathbf{i}\ell} + e_{it} \tag{44}$$

$$e_{it} = u_{it} + \sum_{\ell=p_T+1}^{\infty} \overline{\mathbf{x}}_{t-\ell}' \delta_{\mathbf{i}\ell} + \sum_{\ell=p_T+1}^{\infty} \overline{\mathbf{g}}_{t-\ell}' \kappa_{\mathbf{i}\ell} + O_p\left(N^{-1/2}\right) \tag{45}$$

[12] That is, we invariably reject weak cross-sectional dependence during estimation.

[13] Where $vec(\mathbf{\Gamma_i}) = vec(\mathbf{\Gamma}) + \eta_{\mathbf{\Gamma i}}, \quad \eta_{\mathbf{\Gamma i}} \sim iid\,(\mathbf{0}, \mathbf{\Omega_\Gamma})$, for $i = 1, 2, \ldots, N$, and where $\mathbf{0}$ is km x 1.

[14] See Ditzen (2018), p. 587 for more details.

[15] In practice, Chudik and Pesaran (2015) recommend restricting the number of cross-sectional lags to maintain a sufficient number of degrees of freedom, which we do.

In doing so, we filter the individual-specific regressors by using cross-section averages (denoted with an overbar) such that, asymptotically, the differential effects of unobserved common factors are eliminated.[16] That is, the cross-sectional averages and their lags are used as proxies for the unobserved common factors. The estimator relies on $(N, T) \to \infty$ jointly for consistency, with the number of lags of cross-sectional averages growing slowly with the time dimension.[17] Crucially, Chudik and Pesaran (2015) showed that this CCE estimator has good small-sample properties despite substantial heterogeneity and dynamics, and for relatively small values of N and T.

We also allow for the possibility that there may be other, outside latent factors present in the data, which could be proxied by other forecasts made by forecaster i at time t, such as consumption expenditure and business investment forecasts.

Identification challenge 7. Serial correlation.

Thus far, the key difference from Chudik and Pesaran (2015) has been that our error term, λ_{it}, is serially correlated by definition (Eq. 5). Yet, if we reparameterize Eq. 44 as follows:[18]

$$y_{it} - \theta_i y_{i,t-1} = \left(\rho_i y_{i,t-1} + \mathbf{z_t}' \mathbf{\Phi_i} + \mathbf{x_{it}}' \beta_\mathbf{i} + \sum_{\ell=0}^{p_T} \overline{\mathbf{x}}_{t-\ell}' \delta_{\mathbf{i}\ell} + \sum_{\ell=0}^{p_T} \overline{\mathbf{g}}_{t-\ell}' \kappa_{\mathbf{i}\ell} + e_{it} \right) - \theta_i \left(\rho_i y_{i,t-2} + \mathbf{z_{t-1}}' \mathbf{\Phi_i} + \mathbf{x_{i,t-1}}' \beta_\mathbf{i} + \sum_{\ell=0}^{p_T} \overline{\mathbf{x}}_{t-1-\ell}' \delta_{\mathbf{i}\ell} + \sum_{\ell=0}^{p_T} \overline{\mathbf{g}}_{t-1-\ell}' \kappa_{\mathbf{i}\ell} + e_{i,t-1} \right) \quad (46)$$

Rearranging gives the final form of our estimator:

$$y_{it} = (\theta_i + \rho_i) y_{i,t-1} - \theta_i \rho_i y_{i,t-2} + \mathbf{z_t}' \mathbf{\Phi_i} - \theta_i \mathbf{z_{t-1}}' \mathbf{\Phi_i} + \mathbf{x_{it}}' \beta_\mathbf{i} - \theta_i \mathbf{x_{i,t-1}}' \beta_\mathbf{i} + \left(\sum_{\ell=0}^{p_T} \overline{\mathbf{x}}_{t-\ell}' \delta_{\mathbf{i}\ell} - \theta_i \sum_{\ell=0}^{p_T} \overline{\mathbf{x}}_{t-1-\ell}' \delta_{\mathbf{i}\ell} \right) + \left(\sum_{\ell=0}^{p_T} \overline{\mathbf{g}}_{t-\ell}' \kappa_{\mathbf{i}\ell} - \theta_i \sum_{\ell=0}^{p_T} \overline{\mathbf{g}}_{t-1-\ell}' \kappa_{\mathbf{i}\ell} \right) + (e_{it} - \theta_i e_{i,t-1}) \quad (47)$$

Expanding the error term gives:

$$(e_{it} - \theta_i e_{i,t-1}) = (u_{it} - \theta_i u_{i,t-1}) \quad + \left(\sum_{\ell=p_T+1}^{\infty} \overline{\mathbf{x}}_{t-\ell}' \delta_{\mathbf{i}\ell} + \sum_{\ell=p_T+1}^{\infty} \overline{\mathbf{g}}_{t-\ell}' \kappa_{\mathbf{i}\ell} + O_p\left(N^{-1/2}\right) \right) - \theta_i \left(\sum_{\ell=p_T+1}^{\infty} \overline{\mathbf{x}}_{t-1-\ell}' \delta_{\mathbf{i}\ell} + \sum_{\ell=p_T+1}^{\infty} \overline{\mathbf{g}}_{t-1-\ell}' \kappa_{\mathbf{i}\ell} + O_p\left(N^{-1/2}\right) \right) \quad (48)$$

[16] We omit cross-sectional averages of lagged dependent variables. The dependent variable is an expectation, not a realized macro outcome. Any common persistent component in expectations operates through observable macroeconomic fundamentals, which are controlled for using cross-sectional averages of regressors and, therefore, factor structures. Including lagged cross-sectional averages of expectations would over-partial common belief dynamics and obscure the structural interpretation of the lagged dependent variable coefficients. We also could not include $\overline{y}_{it}$ because we would purge the very shock measure we seek to extract.

[17] This requires the matrix $\mathbf{C} = (\gamma, \mathbf{\Gamma})'$ to have full column rank (Chudik and Pesaran (2015), p. 397).

[18] This reparameterization is reminiscent of quasi-differencing, which is often used as a statistical device to eliminate serial correlation in the error process by estimating and filtering out an autoregressive component (Prais and Winsten (1954)). Here the use of quasi-differencing is theory-driven: we impose a structural persistence parameter implied by the model rather than estimating it from the data. This transformation should therefore be understood not as a statistical correction but rather as part of the identifying structure of the model, rendering the residuals interpretable as structural shocks.

where we can assume that the terms in the second line (essentially, the differenced effect of distant lagged regressors) above are vanishingly small.[19] So, then:

$$e_{it} - \theta_i e_{i,t-1} \approx u_{it} - \theta_i u_{i,t-1} \Rightarrow \tag{49}$$

And, by Eq. 28:

$$u_{it} - \theta_i u_{i,t-1} = (\zeta_{it} + \lambda_{it}) - \theta_i (\zeta_{i,t-1} + \lambda_{i,t-1}) \Rightarrow \tag{50}$$

And, similarly, because $\zeta_{it} - \theta_i \zeta_{i,t-1}$ should also be vanishingly small (by Lemma A.1), then (rearranging):

$$\lambda_{it} - \theta_i \lambda_{i,t-1} + (\zeta_{it} - \theta \zeta_{i,t-1}) \approx \lambda_{it} - \theta_i \lambda_{i,t-1} \Rightarrow \tag{51}$$

Now, substituting in Eq. 29 (while suppressing the t parameter subscripts) into the first term only gives:

$$\lambda_{it} - \theta_i \lambda_{i,t-1} = (\theta_i \lambda_{i,t-1} + \varepsilon_{it}) - \theta_i \lambda_{i,t-1} = \varepsilon_{it} \Rightarrow \tag{52}$$

$$e_{it} - \theta_i e_{i,t-1} \approx u_{it} - \theta_i u_{i,t-1} \approx \lambda_{it} - \theta_i \lambda_{i,t-1} = \varepsilon_{it} \tag{53}$$

and so, as a close approximation, ε_{it} becomes the new error term, which is not serially correlated. As such, we can estimate Eq. 47 with the adapted DCCE estimator. Then, the forecast horizon-specific (h) mean perceived shock estimates are:

$$\widehat{\varepsilon_t^h} = \frac{1}{N} \sum_i^N \widehat{\varepsilon_{it}^h} \tag{54}$$

where the forecaster-specific parameter and ε_{it}^h estimates come from the rolling-window estimate whose last observation is t (and first observation is $t-W+1$), where W is window length in months. The mean-group coefficients and ε_{it}^h estimates are averaged in the same manner.

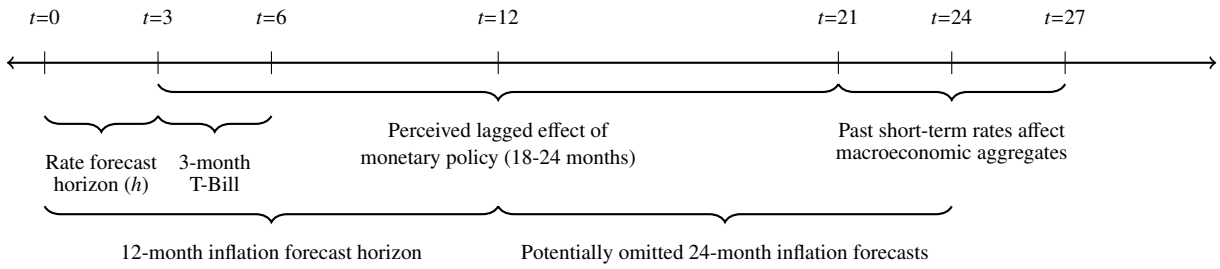
The $\widehat{\varepsilon_{it}^h}$ have little-to-no serial correlation *within each rolling window* (as hypothesized), which reassures us that our DCCE approach is well motivated. However, the subsequent rolling-window collation approach induces a new source of serial correlation (as demonstrated in §A.4). To counteract this, we regress the $\widehat{\varepsilon_t^h}$ on its own lags (twelve), which purges $\widehat{\varepsilon_t^h}$ of this new source of serial correlation. It is the residual from this final step that we show below in Fig. 4, 5.

[19] See Chudik and Pesaran (2015) p.397.

Identification challenge 8. Macro forecast horizons: simultaneity vs. omitted variable bias.

Monetary policy is forward-looking (as reflected in Eq. 1 onwards). This presents an identification challenge for our estimator (Eq. 47). The dependent variable is a forecast of an interest rate either three or twelve months from now ($h \in \{3, 12\}$). The key independent variable is a forecast of a macroeconomic variable, such as inflation, 12 months from now (H=12). Depending on how the forecast horizons are calibrated, endogeneity may be introduced to the estimator in the form of either simultaneity bias or omitted variable bias.

Figure 1: **Forecast Horizons (t months) with h=3, H=12.**



Aruoba and Drechsel (2024) find that, after an exogenous monetary policy contraction, the headline *Personal Consumption Expenditures Price Index* displays a gradual decline, with a significantly negative response at conventional levels only after 43 months and the peak reduction occurs after 54 months (consistent with Christiano, Eichenbaum, and Evans (1999), and tangentially, Nakamura and Steinsson (2018)). However, because we are using expectations data, what matters is not the true long and variable lags of monetary policy but the *perception* thereof. A more conventional period would be 18 to 24 months.

As long as the *12-month inflation forecast horizon* window in Figure 1 does not overlap the *Past short-term rates affect macroeconomic aggregates* window in Figure 1, then we avoid endogeneity due to simultaneity. In Figure 1 with h=3, there is a buffer of nine months; with h=12, there is a slight overlap of three months, but only due to our conservative assumption about perceived monetary policy lags.

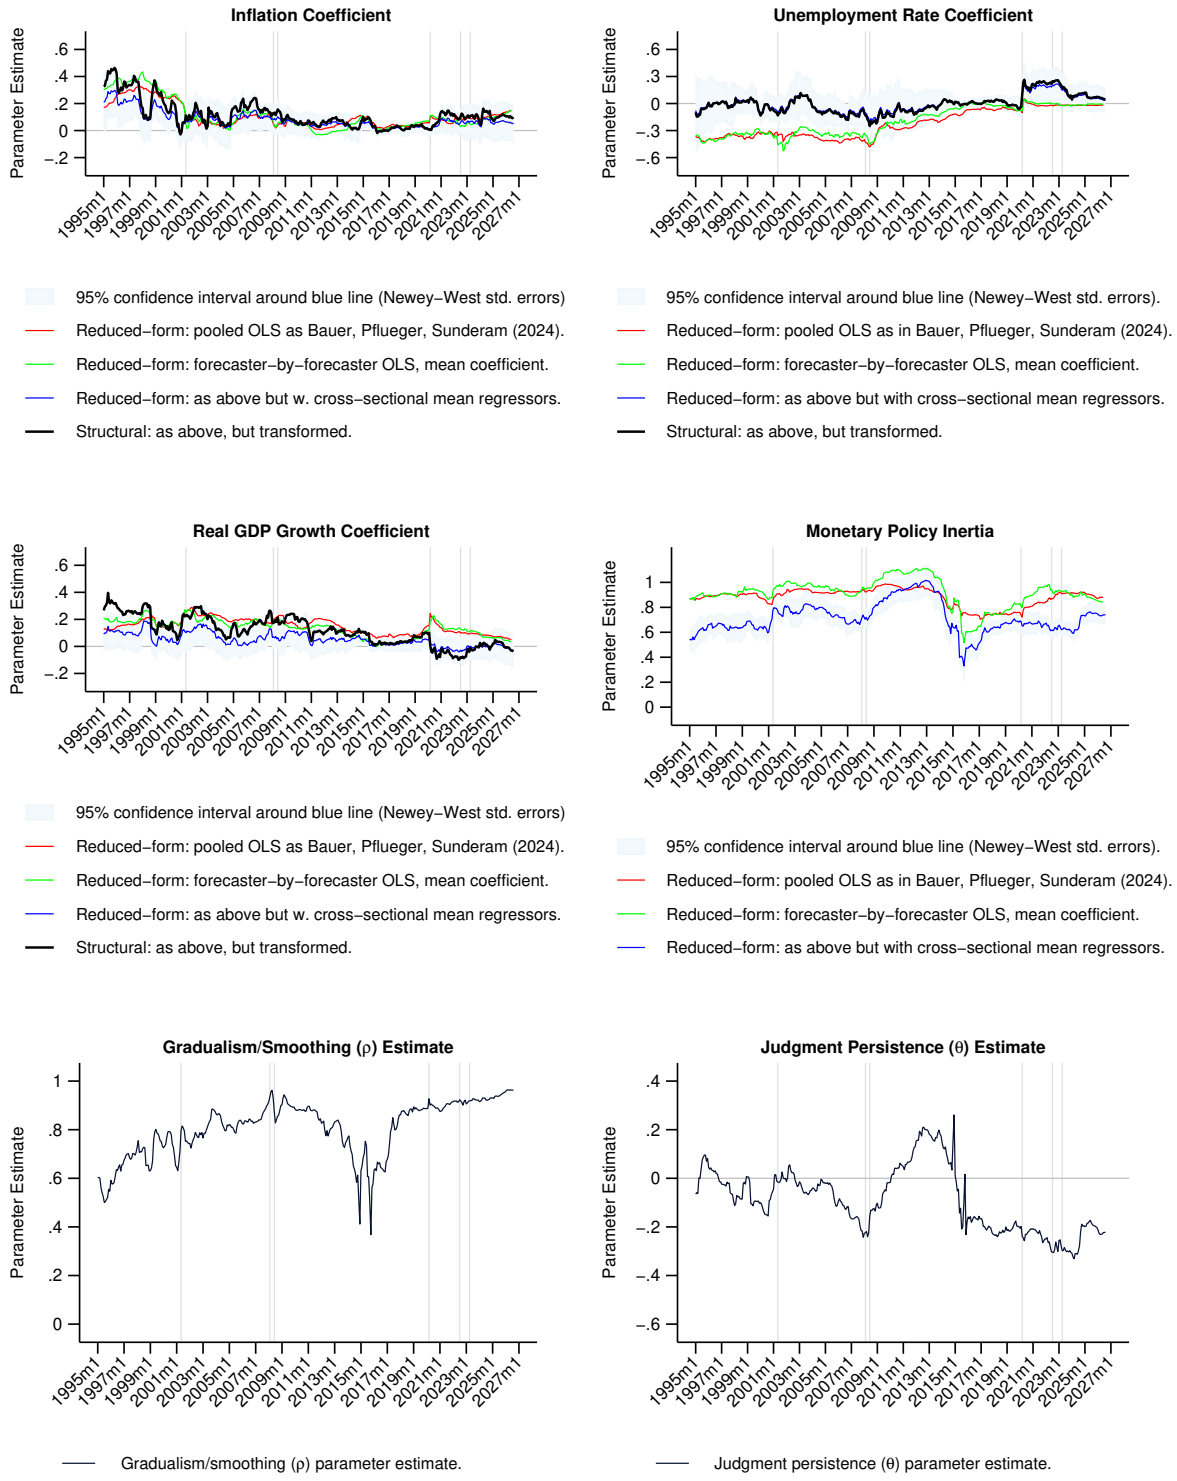
Conversely, if the inflation forecasts are not forward-looking enough, then we may have some omitted variable bias. This is captured in Figure 1 by considering the possibility that forecasters use inflation forecasts as far as 24 months ahead to form their policy rate forecasts. We can check for this by including both current-year and next-year inflation forecasts, which at their peak in January, look forward 24 months. When we do so, our results are not materially different.

3 Estimation

In this section, we present the forecast data, the forecaster-expected reaction function parameter estimates, and the monetary policy expectation shock estimates. We then consider three case studies to help interpret the results. Estimation is performed with the support of the xtdcce2 package in Stata (Ditzen (2018)), which should facilitate straightforward replication and adaptation of our empirical approach and code to one of the over 100 countries for which data are available from *Consensus Economics*.

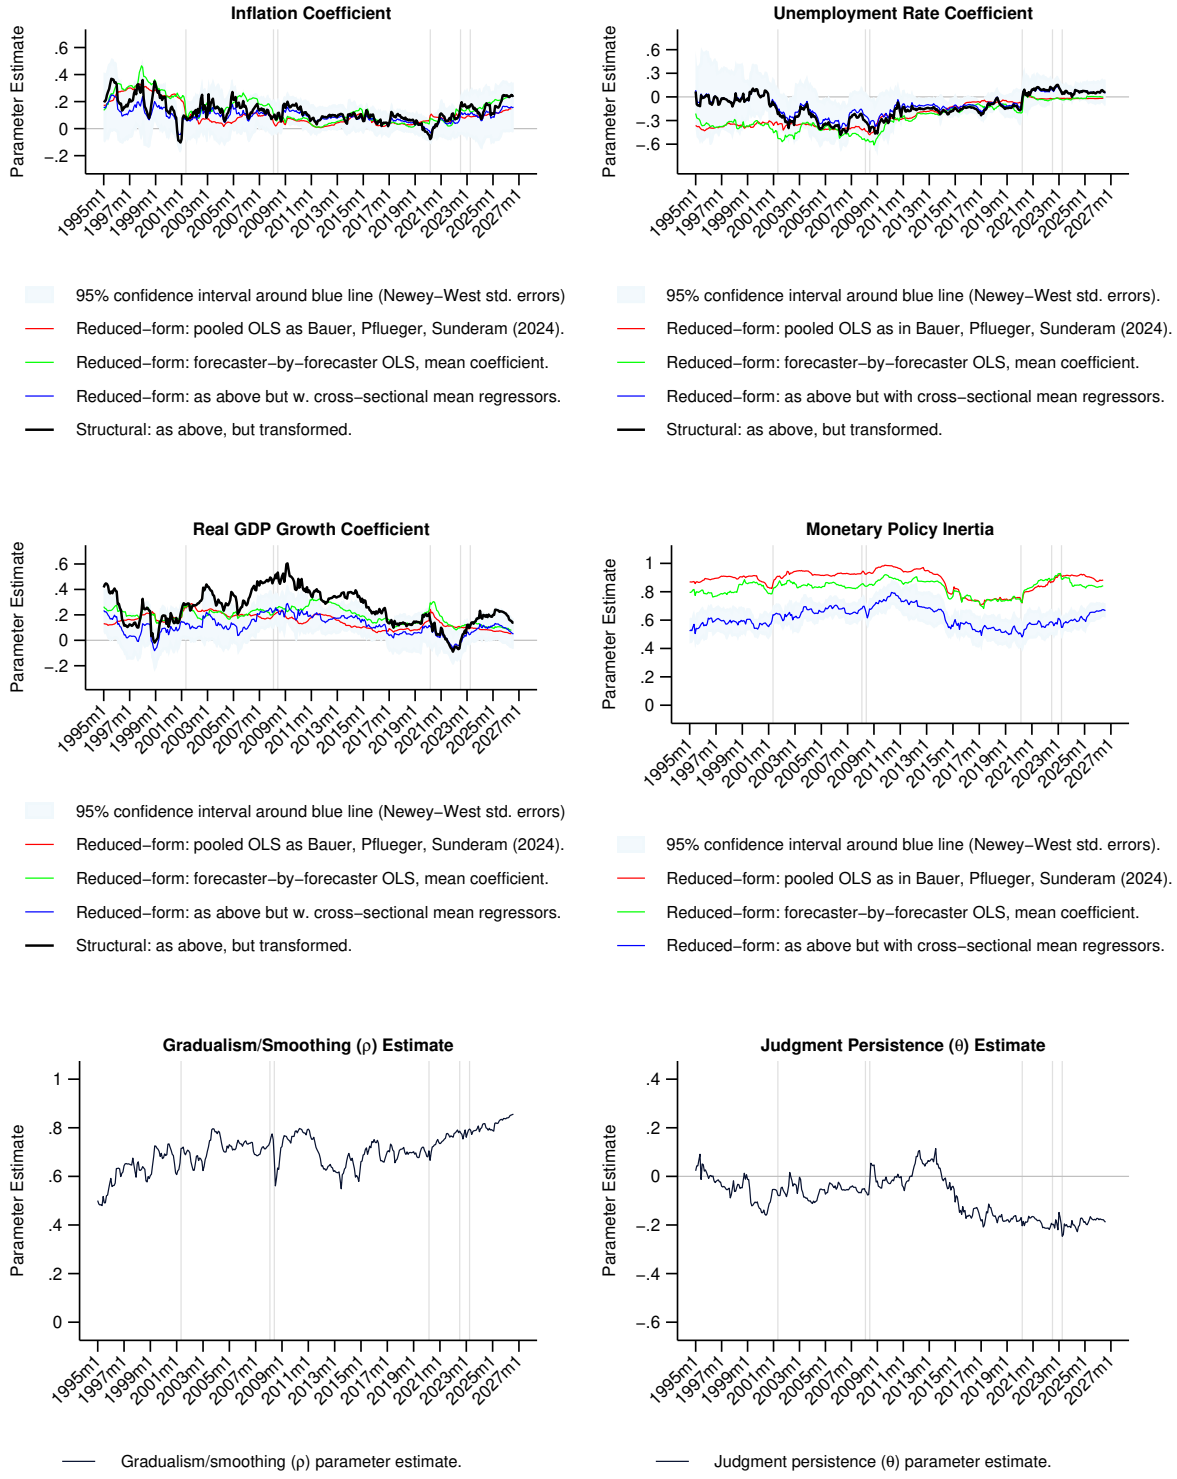
3.1 Expected FOMC Reaction Function Estimates

Figure 2: **Expected FOMC Reaction Function Coefficient Estimates (3-month forecast horizon)**



Note 1: The thin, grey, vertical lines denote the case studies described below.
Note 2: Results available for months 1995m1 to 2026m8. Forecasts made on second or third Monday of each month.

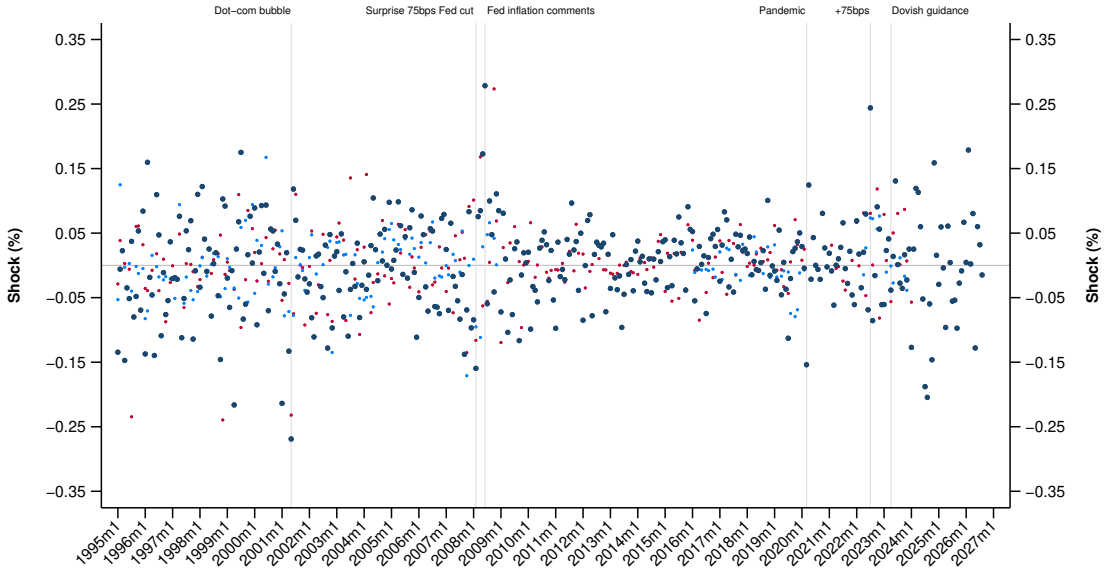
Figure 3: **Expected FOMC Reaction Function Coefficient Estimates (12-month forecast horizon)**



Note 1: The thin, grey, vertical lines denote the case studies described below.
Note 2: Results available for months 1995m1 to 2026m8. Forecasts made on second or third Monday of each month.

3.2 Expected FOMC Reaction Function-Implied Judgment Shocks[20]

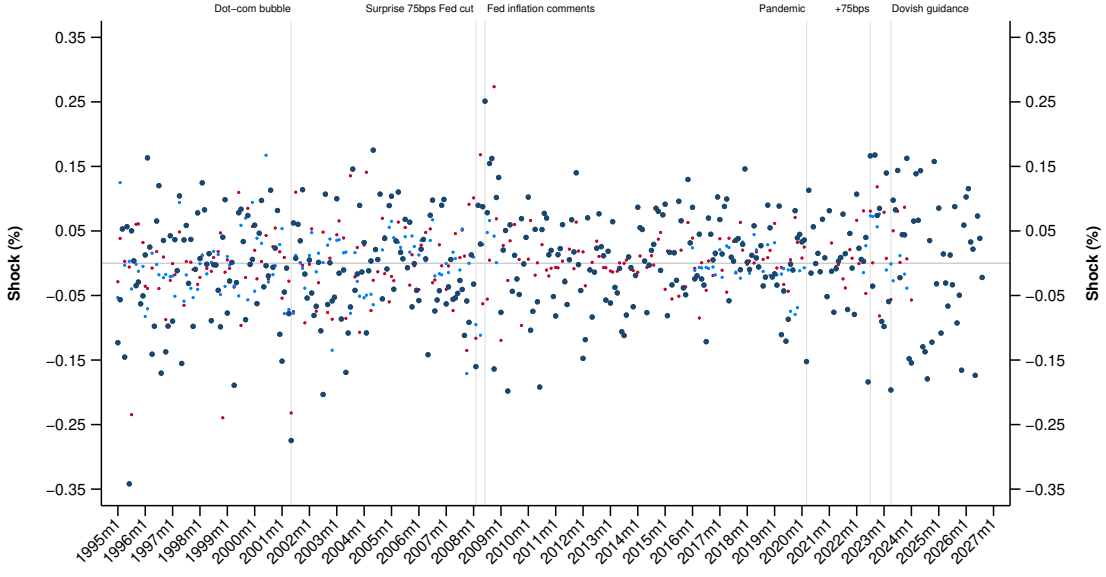
Figure 4: **Monetary Policy Expectations Shocks (3-month Forecast Horizon)**



- **High−frequency monetary policy surprises** (non−zero values only). Bauer & Swanson (2023). Method: principal components analysis.
- **Narrative policy rate shocks** (non−zero values only). Aruoba & Dreschel (2026). Method: machine learning.
- **Forecaster−perceived monetary policy expectations shocks**. Method: dynamic common correlated effects.

Sample period: 1995m1 to 2026m8. The thin, grey, vertical lines denote the case studies described below.

Figure 5: **Monetary Policy Expectations Shocks (12-month Forecast Horizon)**



- **High−frequency monetary policy surprises** (non−zero values only). Bauer & Swanson (2023). Method: principal components analysis.
- **Narrative policy rate shocks** (non−zero values only). Aruoba & Dreschel (2026). Method: machine learning.
- **Forecaster−perceived monetary policy expectations shocks**. Method: dynamic common correlated effects.

Sample period: 1995m1 to 2026m8. The thin, grey, vertical lines denote the case studies described below.

[20] **Pairwise shock correlations**: perceived expectations shocks (3-month horizon with 12-month horizon): 0.68; 3-month horizon with Bauer and Swanson (2023): 0.13; 3-month horizon with Aruoba and Drechsel (2025): 0.15; 12-month horizon with Bauer and Swanson (2023): 0.08; 12-month horizon with Aruoba and Drechsel (2025): 0.05; Bauer and Swanson (2023) with Aruoba and Drechsel (2025): 0.11.

3.3 Why are the Estimated Inflation Coefficients Less Than One?

First, the dynamic nature of the monetary policy rule implies that the macro response coefficient estimates in Figures 2 and 3 are **short-run** in nature—that is, they are proportional to $(1-\rho_{it})$.[21] Now, since the estimates of expected $\bar{\rho}_{it}$ in Figures 2 and 3 oscillate around 0.8 most of the time, that implies that the long-run coefficients are approximately 5 times the magnitude of the short-run expected coefficient estimates at a given time. So, the long-run coefficients are, sometimes, likely close to or even above one.

The idea that The Fed should raise nominal rates more than one-for-one with inflation dates back to Taylor (1993), and reflected the need to establish inflation-fighting credibility. However, the expected macro response coefficient estimates in Figures 2 and 3, which are significantly lower than one, corroborate several points made by Nakamura, Riblier, and Steinsson (2025): "While the Taylor rule (generally) fits well during the Greenspan period, it (generally) fits poorly in the early 1980s and after the early 2000s."

Nakamura, Riblier, and Steinsson (2025) offer three reasons why central banks may optimally vary the nominal interest rate less than one-for-one with inflation in certain circumstances.[22] First, in response to cost-push and sectoral shocks, optimal policy calls for a substantial amount of forward guidance, which weakens the contemporaneous relationship between the nominal interest rate and inflation (i.e. central banks respond in a shock-contingent manner). Second, a combination of inflationary cost push shocks and negative demand shocks calls for a muted response of interest rates (i.e. shocks to the Phillips curve and the dynamic IS equation may be correlated). Third, "long and variable" lags imply that there may be little monetary policy can do in response to transitory shocks.

Two additional, related reasons include anchored forecaster inflation expectations due to the success of inflation-targeting. That is, when inflation spends most of the time (in the sample data anyway) around 2%, inflation news may become less important with respect to forecaster-expected reaction functions. As well, the effective lower bound (ELB) limits The Fed's ability to respond to below-target inflation, which could also drive down the expected inflation response coefficient during periods in which the policy rate is forecasted to be at the ELB.[23]

3.4 Case Studies

Case 1. Dot-com bubble shock (2001m5). On April 18, 2001, the FOMC unexpectedly reduced the federal funds rate target by 50 basis points (bps) between scheduled meetings. The Committee cited weakening capital investment, declining corporate profitability, falling equity wealth, and slower foreign growth as threats to already-fragile U.S. growth. By this point, the Fed had lowered the target rate by 200 bps within six months. Accordingly, perceived FOMC responsiveness to the growth outlook increased steadily, and continued to do so (Fig. 2, 3). The Bauer and Swanson (2023) measure includes a negative surprise comparable in magnitude to the forecaster-perceived monetary policy expectations shock—intuitive given the surprise nature of the meeting; the Aruoba and Drechsel (2025) measure is also negative but much smaller—also intuitive given that this measure would have adjusted heavily for the points mentioned by the FOMC above (Fig. 4, 5).

Case 2a. Financial crisis shocks (2008m2). Amidst a weakening economic outlook, increasing downside risks to growth, strains in short-term funding markets, deteriorating financial market conditions,

[21] As Coibion and Gorodnichenko (2012), Bauer, Pflueger, and Sunderam (2024), and others have also observed.
[22] This assumes anchored long-run inflation expectations.
[23] See Debortoli, Galí, and Gambetti (2020) for an analysis of the empirical (ir) relevance of the ELB.

tightening credit conditions, deepening housing contraction, and a softening in labor markets, the FOMC cut the federal funds rate at an unscheduled meeting on Jan. 22, 2008 by 75 bps—the largest single-day reduction in over two decades. Immediately after the first cut, futures contracts implied a 72 percent chance of another 50 bps rate cut at the scheduled January 30 meeting (source). Accordingly, Fig. 4, 5 show a large, negative monetary policy expectations shocks at both horizons. Meanwhile, perceived FOMC unemployment rate responsiveness increased (Fig. 2, 3). Soon after, on Jan. 30 (also in 2008m2), the FOMC cut the target rate another 50 bps. Both the Bauer and Swanson (2023) and Aruoba and Drechsel (2025) shock measures include a negative surprise/shock comparable in magnitude to the forecaster-perceived monetary policy expectations shock—albeit somewhat smaller (Fig. 4, 5).

Case 2b. Financial crisis shocks (2008m6). Given the rapid pace of target rate cuts to this point, the FOMC was no longer perceived as applying its usual level of gradualism ($\widehat{\rho_t} \approx 0.7$), and $\widehat{\rho_t}$ began dropping toward about 0.5 (Fig. 3), which is what we might expect to see in a crisis. At this time, the monetary policy stance had been loosened substantially, the price of oil had risen significantly, and inflation fears were emerging. Following comments by Richard Fisher and Ben Bernanke signaling that interest rates may need to be raised in the near future should inflation expectations increase, markets began expecting that the Federal Reserve may need to raise interest rates later in the year. This led to a large, positive monetary policy expectations shock at the three-month forecast horizon ($\approx$ 36 bps), and a substantial shock at the twelve-month horizon ($\approx$ 25 bps). Neither the Bauer and Swanson (2023) surprise measure nor the Aruoba and Drechsel (2025) shock measure include any surprise/shock as there was no FOMC meeting between May 13 and June 9, 2008 (the sixth forecast month ending in 2008) (Fig. 4, 5).

Case 3. Pandemic shock (2020m3). At an unscheduled meeting on March 3, 2020, the FFR was unexpectedly lowered by 50 bps to the effective lower bound (ELB) due to the economic shock caused by the COVID-19 pandemic. The FOMC also used forward guidance that indicated a willingness to act further. In March 2020 (2020m3), we observe large, negative monetary policy expectations shocks ($\approx$ 16 bps) at both the three-month and twelve-month forecast horizons (Fig. 4, 5). Neither the Bauer and Swanson (2023) measure nor the Aruoba and Drechsel (2025) measure include a shock corresponding to this announcement (Fig. 4, 5).

Case 4. Post-pandemic monetary policy reversal (2022m7). At this time, the FOMC was in the middle of tightening monetary policy at a rapid pace due to the surge in inflationary pressures. On June 15, 2022, the FOMC raised the target rate by 75 bps despite indicating at the May 2022 meeting that a 75 basis point rate increase was "not something that the committee is actively considering". Despite the large increase, the FOMC also added hawkish forward guidance: "[the Committee] anticipates that ongoing increases in the target range will be appropriate". Accordingly, we see large positive shocks to forecaster-perceived monetary policy expectations at both the three-month horizon ($\approx$24 bps) and the twelve-month horizon ($\approx$17 bps); the Bauer and Swanson (2023) and Aruoba and Drechsel (2025) shocks were the same sign but smaller in magnitude ($\approx$8 bps each) (Fig. 4, 5). Having both horizons helps us understand the nature of the shock to monetary policy expectations: this shock provided news about both the timing and the magnitude of monetary policy tightening. Accordingly, compared to the March 2022 Summary of Economic Projections, FOMC participants projected substantially more target rate increases and sooner than expected in March. Further, the FOMC was clearly signaling its strong commitment to achieving its inflation target. Intuitively, we can see that forecasters perceived FOMC inflation responsiveness as having increased substantially around this time (Fig. 3)—an increase that is most clearly glimpsed with the estimator proposed in this paper, as compared to the others in Fig. 3.

Case 5. A hawkish shift in forward guidance (2023m4). On March 22, 2023, the Federal Reserve

raised its target rate by 25 bps. However, the FOMC also signaled that the recent precipitous rise in the target rate (425 bps in one year) may soon be ending by softening its longstanding forward guidance from "The Committee anticipates that ongoing increases in the target range will be appropriate in order to attain a stance of monetary policy that is sufficiently restrictive to return inflation to 2 percent over time", to "The Committee anticipates that some additional policy firming may be appropriate in order to attain a stance of monetary policy that is sufficiently restrictive to return inflation to 2 percent over time". This episode demonstrates the value of estimating shocks at multiple forecast horizons: the forecaster-perceived monetary policy expectations shock measure at the three-month forecast horizon decreased by only 2 bps, but decreased by 21 bps at the twelve-month forecast horizon. The Bauer and Swanson (2023) measure includes a negative surprise of 6 bps; the Aruoba and Drechsel (2025) shock measure is almost zero for this period (Fig. 4, 5).

3.5 A Parsimonious Coibion and Gorodnichenko (2012) Test

Coibion and Gorodnichenko (2012) ask why are target interest rate changes so persistent? They acknowledge that inertia in central banks' reaction functions is now standard in theoretical and empirical monetary economics, but argue that the source of this inertia is controversial. The authors propose a number of tests to discriminate between competing explanations of inertia, and conclude that the evidence strongly favors the interest-smoothing explanation over persistent monetary policy shocks.

The perceived reaction function framework and the corresponding estimator in this paper provide a structural test of the competing hypotheses from Coibion and Gorodnichenko (2012). Figures 2 and 3 provide estimates of the mean gradualism/smoothing ($\widehat{\rho_{it}}$) and judgment persistence ($\widehat{\theta_{it}}$) structural parameters, respectively, over time, providing a straightforward way to visualize the extent to which each channel is dominating. As we seek to explain positive persistence, we can see that, consistent with Coibion and Gorodnichenko (2012), gradualism/smoothing explains that positive persistence the great majority of the time, with positively persistent judgment dominating only occasionally—particularly during the post-financial crisis ELB/forward guidance period, which accords with economic intuition. Interestingly, smoothing still dominated in the wake of the pandemic, which is also consistent with the idea that central banks were behind the curve on inflation, responding perhaps too slowly to inflationary pressures.

4 Macroeconomic Effects

In this section, we use our shock measures as instrumental variables to estimate impulse response functions by OLS in a local projections framework (LP-IV as in Jordà, Schularick, and Taylor (2015)). We favour LP over VAR to contain extrapolation bias, as LP dominate VAR with respect to bias. Given the inherent bias-variance tradeoff, this comes at the cost of higher variance—(much) larger standard errors (compared to the common, mis-specified, small-to-moderate-lag VARs). The two sets of estimates do coincide, however, as increasingly many lags are (appropriately) added to VARs (Montiel Olea, Plagborg-Møller, Qian, and Wolf (2024)).

For comparison to the relevant literature, the system of equations is based on Gertler and Karadi (2015).[24]

$$y_{n,t+h} = \mu_n^h + \psi_n^h r_t + \sum_{\ell=0}^{3} \mathbf{y}_{\mathbf{t}-\ell}^{'} \kappa_{\mathbf{l}}^{\mathbf{h}} + \xi_{nt}^h \tag{55}$$

[24] Bauer and Swanson (2023) also based their model on Gertler and Karadi (2015), although they use the two-year Treasury yield. We favour the one-year yield as this endogenous monetary policy variable is more closely aligned with the structure of all candidate monetary policy instruments. The results are not sensitive to this choice.

where $\mathbf{y_t}$ is a vector of n (n=1, 2, ...) endogenous monthly (t) variables: the natural logarithm (ln) of the consumer price index, ln of industrial production, the one-year Treasury yield, and the (updated) Gilchrist and Zakrajšek (2012) Excess Bond Premium, respectively.[25] r_t is the policy variable (the one-year Treasury yield). The lag length is three. Whereas the bias and standard deviation of VAR estimates are heavily dependent on the choice of lag length, this choice matters much less for the Local Projections estimator (which was tested and confirmed). Let $\mathbf{z_t'} = \left[\widehat{\varepsilon_t^{h=3}}, \widehat{\varepsilon_t^{h=12}}\right]$ be the vector of instrumental variables. The conditions for valid estimation of Eq. 55 are:[26]

1. $\mathbb{E}\left[\xi_{1t}^{h,\perp}\mathbf{z_t^{\perp}}\right] = \alpha > \alpha^*$ (instrument **relevance** for monetary shock, $\xi_{1t}^{h,\perp}$).
2. $\mathbb{E}\left[\xi_{nt}^{h,\perp}\mathbf{z_t^{\perp}}\right] = 0$ (**contemporaneous exogeneity** of instrument to other shocks, $\xi_{nt}^{h,\perp}, n \neq 1$).
3. $\mathbb{E}\left[\xi_{n,t+j}^{h,\perp}\mathbf{z_t^{\perp}}\right] = 0 \quad \forall j{\neq}0$ (exogeneity of instrument to **leads and lags** of all shocks).

where the $\perp$ superscript denotes orthogonalization with respect to control variables, α^* is a vector of critical values, such as those derived in Montiel Olea and Pflueger (2013), When Eq. 55 is estimated by two-stage least squares, as we do in §4.3, the unit effect normalization is imposed automatically (Stock and Watson (2018), p. 924). Condition 3 is testable, and indeed holds, as neither the $\widehat{\varepsilon_t^h}$ nor the shocks from the literature are forecastable using leads or lags of y_{nt}.

4.1 Instrument Relevance

In this section, we provide weak instrument test statistics based on our shock measures and those from the literature. When using monetary instruments to estimate the effects of monetary policy, researchers should, at a minimum, provide Montiel Olea and Pflueger (2013) test statistics and critical values, which are robust to serial correlation, heteroskedasticity, and clustering (*better still* would be to use the Lewis and Mertens (2025) test[27]). We provide three main results here. First, we find evidence that the forecast-derived perceived monetary policy shocks are not weak instruments (Figure 6).[28] Second, we find clear evidence that high-frequency monetary policy shocks are weak instruments (Figure 6). Third, narrative-based monetary policy shocks (particularly the state-of-the-art Aruoba and Drechsel (2025) variant) appear to lie somewhere in between (Figure 6).

We suggest this is because monetary policy expectations evolve gradually and continuously. Bianchi, Ludvigson, and Ma (2025) "*finds that much of the causal impact of monetary policy on markets occurs outside of tight windows around policy announcements*." Using monthly forecast data that allows for this slow-moving expectation formation process to be reflected in the data allows us to capture more of both the endogenous and exogenous movements in monetary policy expectations.

[25] We could include more controls in an effort to increase efficiency (by soaking up residual variation in the outcome). Yet, once we control for the variables that are *strong* predictors of either the outcome or the impulse variable, LP estimation is relatively robust to dynamic misspecification (Montiel Olea, Plagborg-Møller, Qian, and Wolf (2025)), so we favour preserving degrees of freedom, and comparability to the literature.

[26] This is an augmented version of the list from Stock and Watson (2018), which is discussed on p. 924.

[27] This test generalizes the tests of Stock and Yogo (2005) and Sanderson and Windmeijer (2016) to be robust to heteroskedasticity and autocorrelation, and extends the Montiel Olea and Pflueger (2013) robust test for models with a single endogenous regressor to multiple endogenous regressors.

[28] Nor, interestingly, are the residuals from Bauer, Pflueger, and Sunderam (2024) rolling regressions, even if they are unlikely to be exogenous.

The weak instruments tests are based on our Local Projections Instrumental Variables estimator (Eq. 55). During estimation, we take steps to adjust for potential bias from weak instruments. Eq. 55 is estimated by limited information maximum likelihood (LIML). LIML is a linear combination of OLS and Two-Stage Least Squares (TSLS) estimates. The weights are chosen so that they (approximately) eliminate any TSLS bias. With weak instruments, LIML has better small-sample properties than TSLS. The just-identified estimates below are approximately (median-) unbiased (or less biased) when using weak instruments.

In the monetary policy shock literature, most authors appear to use the Stock and Yogo (2005) tests for weak instruments, if any. Yet, these tests assume conditionally homoskedastic and serially uncorrelated model errors. Hence, the Stock and Yogo (2005) tests are not robust to heteroskedasticity, serial correlation, and/or clustering, but Montiel Olea and Pflueger (2013) develops tests that are.

That is, most studies test *the null hypothesis of weak instruments* by comparing the commonly-reported yet not heteroskedasticity-robust Cragg-Donald Wald F statistic or the heteroskedasticity-robust Kleibergen-Paap Wald F statistic to the Stock, Wright, and Yogo (2002), Stock and Yogo (2005) critical values (e.g. 8.68, the 5% critical value for testing the null hypothesis that the TSLS bias exceeds 10% of the OLS bias). Despite the popularity of this approach, there is no theoretical justification for its use (Andrews, Stock, and Sun (2019), p. 738). This is because these statistics are not computed using the appropriate estimator of the covariance matrix. That is, the Stock, Wright, and Yogo (2002) Stock and Yogo (2005) critical values are based on the assumption of conditionally homoskedastic, and serially uncorrelated model errors. Montiel Olea and Pflueger (2013) show that even the heteroskedasticity-robust (Kleibergen-Paap) F statistic may be too high when instruments are weak.

By contrast, the Montiel Olea and Pflueger (2013) effective F statistic and critical values reported below in Figure 6 are based on Newey and West (1987) errors. This more robust *effective F statistic* is a scaled, corrected version of the non-robust first-stage F statistic; it depends on the robust covariance matrix estimate. The robust *critical values* depend on an estimate of the covariance matrix of the OLS reduced form regression coefficients, and on the covariance matrix of the reduced form errors, which can be estimated using standard procedures (Montiel Olea and Pflueger (2013), p. 358).

The weak instruments test we need to perform differs depending on whether the model is overidentified (i.e. using the two forecaster-perceived reaction function shocks, or combining the Bauer and Swanson (2023) shock, Bauer, Pflueger, and Sunderam (2024) residual, and the Aruoba and Drechsel (2025) shock), or just identified (i.e. using each shock measure one at a time). **In the overidentified case**, we must use the effective F statistic (solid light blue line, Figure 6) and compare this to the Stock and Yogo (2005) critical values (e.g. 8.68, or the rule of thumb of 10). **In the just identified case**, we must again use the effective F statistic,[29] but in this case, compare it to the Montiel Olea and Pflueger (2013) critical values (two dashed lines, Figure 6), which are LIML-specific.[30] When we use the Montiel Olea and Pflueger (2013) effective F statistic and critical values, we are testing the null hypothesis that the estimator's approximate asymptotic bias exceeds a fraction "τ" of a "worst-case" benchmark.[31]

[29] If just identified, the Kleibergen-Paap Wald F statistic equals the effective F statistic.

[30] See Andrews, Stock, and Sun (2019), p. 739.

[31] In Figure 6, one can compare Montiel Olea and Pflueger (2013) critical values for *the null hypothesis that the TSLS approximate asymptotic bias exceeds 10% of the "worst-case" benchmark with a size of 5%* with Stock, Wright, and Yogo (2002) critical values for *the null hypothesis that the TSLS bias exceeds 10% of the OLS bias* (for now, the dashed grey line equal to ten).

Test Results. Figure 6 shows us that we reject the null hypothesis that the perceived reaction function-implied monetary policy shocks (overidentified model) are weak instruments most of the time, but that the decision is always borderline. So, while our impulse response function estimates in §4.3 are robust to weak instruments problems, inference would still be improved by stronger instruments.

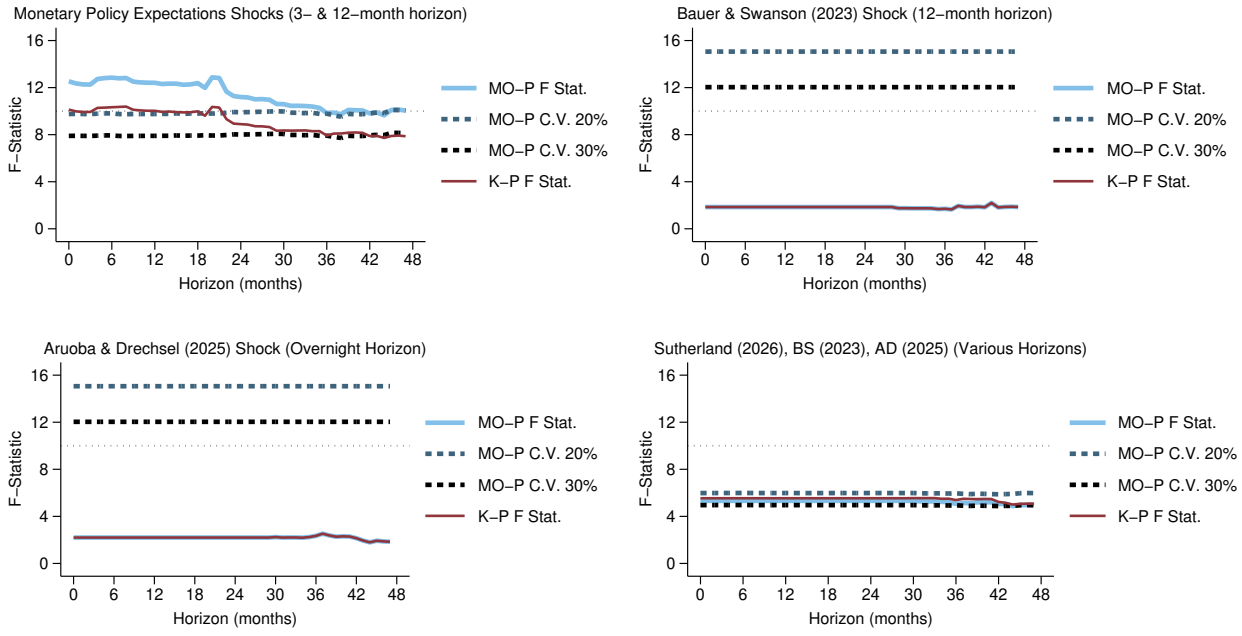
The other three sets of test statistics in Figure 6 pertain to just-identified models. When we use the Bauer, Pflueger, and Sunderam (2024) residuals as an instrument, we always reject the null hypothesis that the TSLS approximate asymptotic bias exceeds 5% of the "worst-case" benchmark with a size of 5%.

For the Bauer and Swanson (2023) shocks, we cannot reject the null (even that bias exceeds 30%). So, when using this instrument, this implies that the distribution of the LP-IV estimator is not centered at ψ_n^h and inference based on conventional IV standard errors is unreliable (Stock and Watson (2018), p. 929), implying the need for weak-instrument-robust inference. That instruments based on high-frequency market surprises may be weak instruments was suggested in Eberly, Stock, and Wright (2020) p. 26. Figure 6 provides further evidence of this. The Aruoba and Drechsel (2025) shocks lie somewhere in between as shown in Figure 6.

Implications for Inference. Should researchers wish to proceed using high-frequency shocks as instruments, then a large and expanding literature shows that special care must be taken when using such potentially weak instruments, particularly with respect to inference (see Stock and Wright (2000), Stock, Wright, and Yogo (2002), Dufour and Taamouti (2005), Mikusheva (2010), Andrews, Stock, and Sun (2019), Ganics, Inoue, and Rossi (2021), Keane and Neal (2023), Barnichon and Mesters (2024), Lewis and Mertens (2025); Gürkaynak, Swanson, and Wright (2026)).

Figure 6: **Weak Instruments Test Statistics**

Each graph corresponds to a different set of instruments.



MO–P C.V.: Montiel Olea–Pflueger critical values.
Confidence level = 5%, tau = 20% or tau = 30%.
Dotted, grey line: Stock & Yogo (2005) F–Statistic threshold of 10, which is a standard threshold used in the literature.
Sample period (monthly data): 1990m1 to 2026m8.

4.2 Instrument Exogeneity

In this section, we test whether the instruments are exogenous using tests of over-identifying restrictions. The evidence suggests that neither the two forecaster-perceived reaction function shocks, nor the combination of the Bauer and Swanson (2023) shocks, Bauer, Pflueger, and Sunderam (2024) residuals, and the Aruoba and Drechsel (2025) shocks show much sign of endogeneity.

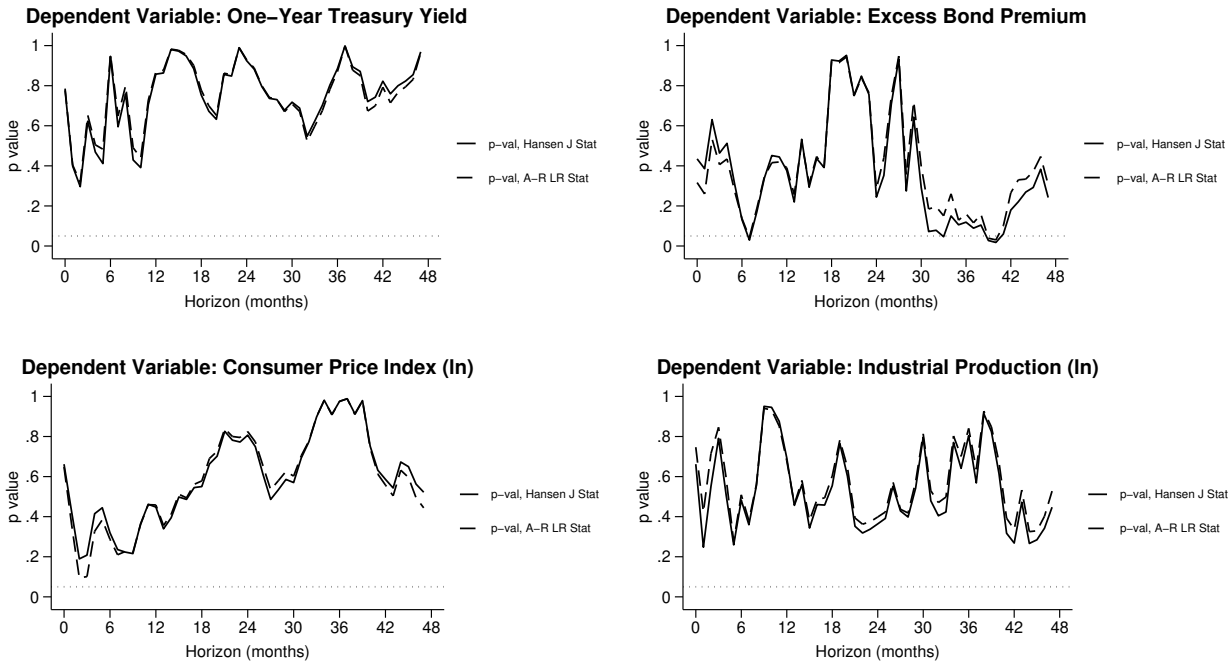
In Figures 7 and 8 below, we plot two sets of statistics.

First, for each horizon, we plot the p values from Hansen J Tests, which test the null hypothesis that all instruments are exogenous (conditional on being relevant)—that they are uncorrelated with the error term and correctly excluded from the structural equation. The Hansen J Test is robust to heteroskedasticity but can be size distorted under weak identification.

Second, for each horizon, we plot the p values from Anderson-Rubin Tests, which (jointly) test the null hypothesis that the instruments are exogenous and that the structural coefficients are identified (whether the instruments are exogenous at a hypothesized value of the structural parameter). The Anderson-Rubin Test maintains proper size under weak identification.

Figure 7: **Tests of Overidentifying Restrictions**

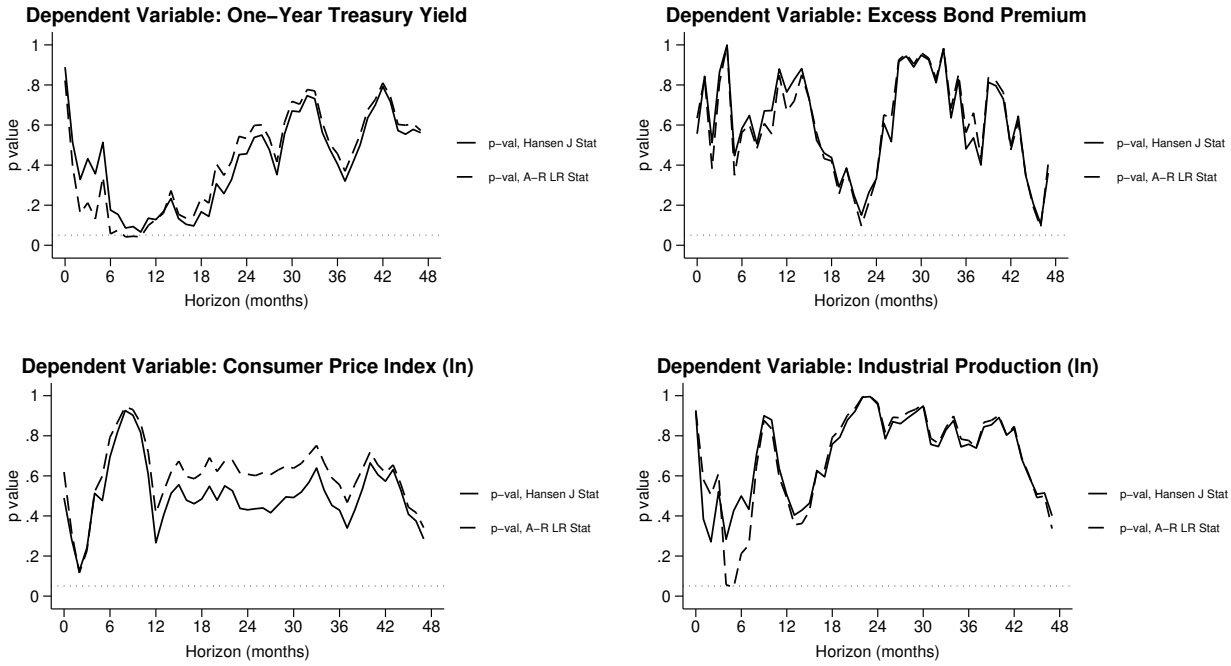
Test statistics based on using forecaster-perceived monetary policy expectations shocks as instruments.



Dotted, straight line denotes p=0.05. Sample period (monthly data): 1990m1 to 2026m8.
Hansen J−test: robust to heteroskedasticity. Anderson−Rubin test: robust to weak instruments.

Figure 8: **Tests of Overidentifying Restrictions**

Test statistics. Instrument set: Bauer and Swanson (2023), Aruoba and Drechsel (2025), forecaster-perceived mon. pol. expectations shocks.



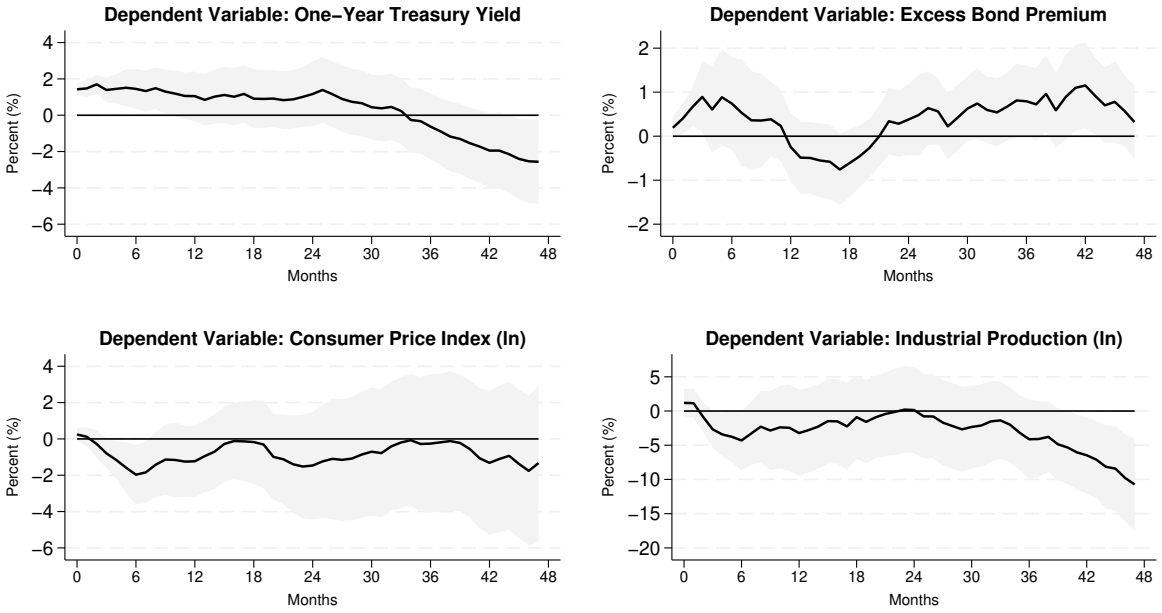
Dotted, straight line denotes p=0.05. Sample period (monthly data): 1990m1 to 2026m8.
Hansen J−test: robust to heteroskedasticity. Anderson−Rubin test: robust to weak instruments.

4.3 Impulse Response Functions

Below, we plot the impulse response function estimates associated with each instrument. The estimates created using the forecaster-perceived reaction function-implied monetary policy shocks (Figure 9) are broadly consistent with New Keynesian theory. First, the one-year Treasury yield rises in response to a monetary policy shock. Second, prices tend to fall, although with wide confidence intervals. Third, industrial production tends to fall. The estimates created using the Bauer and Swanson (2023) shocks (Figure 10) exhibit a price puzzle. The estimates created using the Aruoba and Drechsel (2025) shocks (Figure 11) are also broadly consistent with New Keynesian theory and exhibit the type of long lags of monetary policy consistent with those documented in Aruoba and Drechsel (2024).

Figure 9: **Impulse Response Functions**

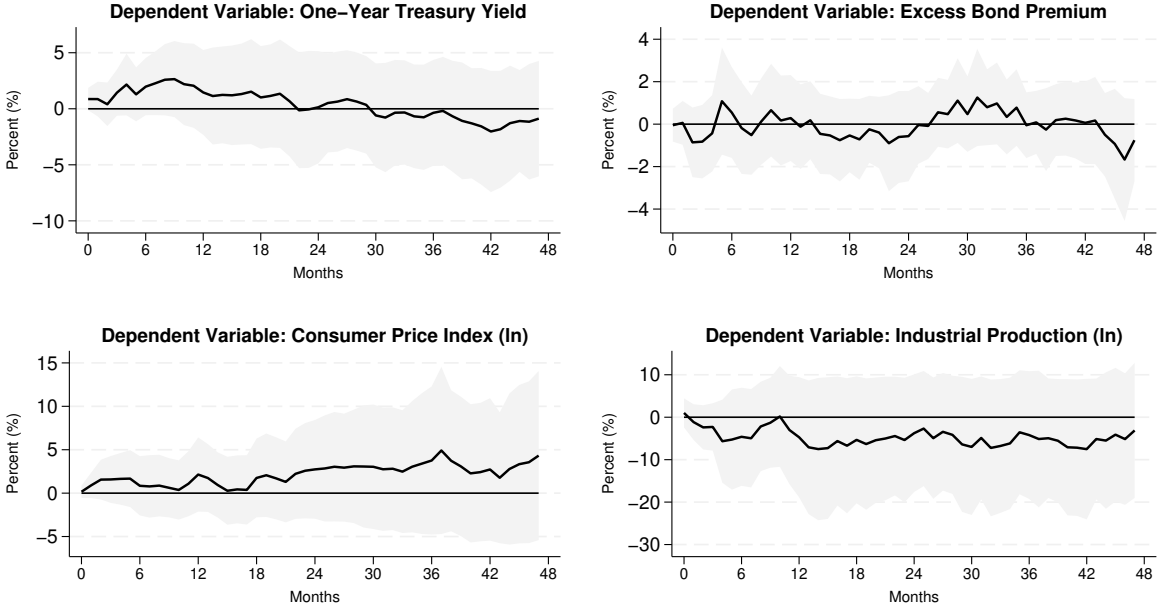
Instruments: forecaster-perceived monetary policy expectations shocks.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2026m8.

Figure 10: **Impulse Response Functions**

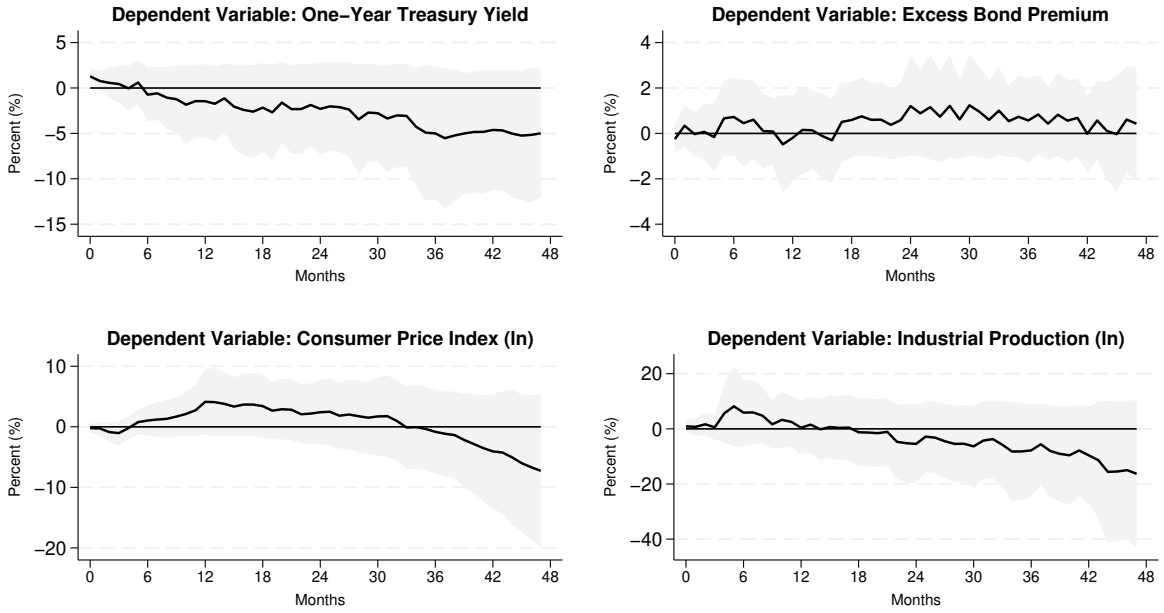
Instrument: Bauer and Swanson (2023) shock measure.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2026m8.

Figure 11: **Impulse Response Functions**

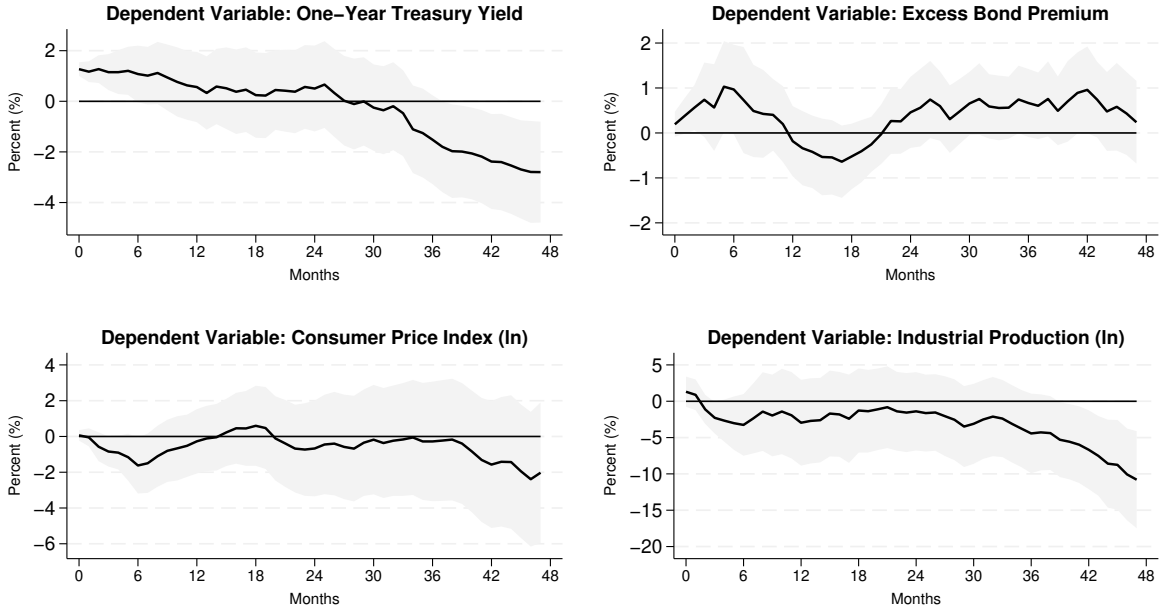
Instrument: Aruoba and Drechsel (2025) monetary policy surprises.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2026m8.

Figure 12: **Impulse Response Functions**

Instrument set: Bauer and Swanson (2023), Aruoba and Drechsel (2025), forecaster-perceived monetary policy expectations shocks.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2026m8.

5 Conclusion

This paper has proposed a number of ideas. First, the reaction function framework incorporates central bank judgment and multidimensional parameter heterogeneity. Crucially, it also delineates smoothing, neutral rate, information, macro response, and FOMC judgment channels.

Second, we have proposed an identification strategy and econometric framework based on Dynamic Common Correlated Effects for the structural estimation of forecaster-expected central bank reaction functions using forecast data. We hope this may be particularly useful as our framework and code can easily be adapted to other countries—not least because *Consensus Economics* offers standardized data for over 100 countries.

Third, a consequence of our framework is that we can extract numerous monetary shock measures—one for each forecast horizon. We propose to use these measures as instruments with a Local Projections Instrumental Variables estimator to measure the effect of monetary policy on the economy. In doing so, we test the validity of the instrument set. The evidence suggests that the instruments are neither endogenous nor weak. We also provide impulse response function estimates that are consistent with New Keynesian models (with shifting steady states). We draw some broader conclusions for the empirical monetary economics literature.

Monetary policy instrument construction approach I: high-frequency market-based surprises. These are shocks to monetary policy expectations (as perceived by markets), and are a function of the underlying, true FOMC shocks. Instruments constructed using this methodology tend to be weak. As such, results based on such instruments and unaccompanied by weak-instrument robust inference should be read through the lens of weak instruments—estimates biased towards OLS with understated standard errors. Fortunately, more recent efforts to incorporate more monetary policy events address this issue (for example the *U.S. Monetary Policy Event-Study Database* (USMPD) of Acosta, Ajello, Bauer, Loria, and Miranda-Agrippino (2025)).

Monetary policy instrument construction approach II: narrative shocks. These represent underlying, true FOMC shocks. It is likely that the original versions of these instruments (Romer and Romer (2004)) did not fully disentangle endogenous components from exogenous components (Aruoba and Drechsel (2025)). Newer versions of these instruments (Aruoba and Drechsel (2025)) seem to have addressed this problem, perhaps even extracting too much information—not leaving enough exogenous variation—potentially leading to a weak instruments problem. In principle, the approach could straightforwardly be adjusted accordingly to solve any such problem.

Monetary policy instrument construction approach III: forecaster-perceived monetary policy shocks. These are shocks to monetary policy expectations (as perceived by forecasters), and are a function of the underlying, true FOMC shocks. They are extracted by estimating the reaction function implied by short-term interest rate forecasts, and the shocks these imply. These measures appear to balance relevance and exogeneity better than its predecessors, but in principle, all three approaches should be able to accomplish such balance, and it may be advantageous to use them in conjunction.

Testing monetary policy instruments for relevance and exogeneity. *Relevance*. To test this, either the Montiel Olea and Pflueger (2013), or better yet, the Lewis and Mertens (2025) robust tests for weak instruments for 2SLS should be used. *Exogeneity*. It is useful to use numerous monetary policy instru-

ments, not least because we can test the exogeneity of the instrument set. Doing so suggests that all three methodologies listed above can lead to exogenous instruments.

In future work, we plan to relax some of the assumptions made in this paper. First, can we increase the frequency of this analysis to increase its monetary policy practicality by applying the method Ghysels and Wright (2009). Second, what if parameters follow a random walk instead of a random coefficients model? Third, what if coefficients are not known at the time forecasts are made? Fourth, can we improve our Local Projections estimates by, for example, using Smooth or Bayesian Local Projections ? Finally, we plan to adapt our approach to The Bank of Canada, Bank of England, European Central Bank, and to publish these estimates.

References

Acosta, M., A. Ajello, M. D. Bauer, F. Loria, and S. Miranda-Agrippino (2025): "Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database," *Federal Reserve Bank of San Francisco Working Paper 2025-30*.

Altavilla, C. and D. Giannone (2017): "The effectiveness of non-standard monetary policy measures: Evidence from survey data," *Journal of Applied Econometrics*, 32, 952–964.

Andrews, I., J. H. Stock, and L. Sun (2019): "Weak instruments in instrumental variables regression: Theory and practice," *Annual Review of Economics*, 11, 727–753.

Arai, N. (2023): "The FOMC's new individual economic projections and macroeconomic theories," *Journal of Banking & Finance*, 151, 106845.

Aruoba, S. B. and T. Drechsel (2024): "The long and variable lags of monetary policy: Evidence from disaggregated price indices," *Journal of Monetary Economics*, 148, 103635.

——— (2025): "Identifying monetary policy shocks: A natural language approach," *American Economic Journal: Macroeconomics (forthcoming)*.

Barnichon, R. and G. Mesters (2024): "Innovation Powered Narrative Inference," *Working Paper*.

Bauer, M. D., C. E. Pflueger, and A. Sunderam (2024): "Perceptions About Monetary Policy," *The Quarterly Journal of Economics*, 139, 2227–2278.

Bauer, M. D. and E. T. Swanson (2023): "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual*, 37, 87–155.

Bianchi, F., C. L. Ilut, and H. Saijo (2024): "Smooth diagnostic expectations," *National Bureau of Economic Research Working Paper No.32152*.

Bianchi, F., S. C. Ludvigson, and S. Ma (2025): "A Structural Approach to High-Frequency Event Studies: The Fed and Markets as Case History," *Working Paper*.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020): "Overreaction in macroeconomic expectations," *American Economic Review*, 110, 2748–82.

Byrne, D., R. Goodhead, M. McMahon, and C. Parle (2023): "The central bank crystal ball: Temporal information in monetary policy communication," *CEPR Discussion Paper DP17930*.

Caballero, R. J. and A. Simsek (2022): "Monetary policy with opinionated markets," *American Economic Review*, 112, 2353–2392.

Campbell, J. R., C. L. Evans, J. D. Fisher, A. Justiniano, C. W. Calomiris, and M. Woodford (2012): "Macroeconomic effects of federal reserve forward guidance [with comments and discussion]," *Brookings Papers on Economic Activity*, 1–80.

Carvalho, C., F. Nechio, and T. Tristao (2021): "Taylor rule estimation by OLS," *Journal of Monetary Economics*, 124, 140–154.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary policy shocks: What have we learned and to what end?" *Handbook of Macroeconomics*, 1, 65–148.

Chudik, A. and M. H. Pesaran (2015): "Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors," *Journal of Econometrics*, 188, 393–420.

Clarida, R., J. Gali, and M. Gertler (2000): "Monetary policy rules and macroeconomic stability: evidence and some theory," *The Quarterly Journal of Economics*, 115, 147–180.

Coibion, O. and Y. Gorodnichenko (2012): "Why are target interest rate changes so persistent?" *American Economic Journal: Macroeconomics*, 4, 126–62.

——— (2015): "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 105, 2644–78.

Debortoli, D., J. Galí, and L. Gambetti (2020): "On the empirical (ir) relevance of the zero lower bound constraint," *NBER Macroeconomics Annual*, 34, 141–170.

Ditzen, J. (2018): "Estimating dynamic common-correlated effects in Stata," *The Stata Journal*, 18, 585–617.

Dufour, J.-M. and M. Taamouti (2005): "Projection-based statistical inference in linear structural models with possibly weak instruments," *Econometrica*, 73, 1351–1365.

Eberly, J. C., J. H. Stock, and J. H. Wright (2020): "The federal reserve's current framework for monetary policy: A review and assessment," *International Journal of Central Banking*, 16, 5–71.

Ganics, G., A. Inoue, and B. Rossi (2021): "Confidence Intervals for Bias and Size Distortion in IV and Local Projections-IV Models," *Journal of Business & Economic Statistics*, 39, 307–324.

Gertler, M. and P. Karadi (2015): "Monetary policy surprises, credit costs, and economic activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

Ghysels, E. and J. H. Wright (2009): "Forecasting professional forecasters," *Journal of Business & Economic Statistics*, 27, 504–516.

Giacomini, R., J. Lu, and K. Smetanina (2024): "Perceived shocks and impulse responses," .

Gilchrist, S. and E. Zakrajšek (2012): "Credit spreads and business cycle fluctuations," *American economic review*, 102, 1692–1720.

González-Astudillo, M. and R. Tanvir (2023): "Hawkish or Dovish Fed? Estimating a Time-Varying Reaction Function of the Federal Open Market Committee's Median Participant," *Finance and Economics Discussion Series 2023-070*.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005): "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements," *International Journal of Central Banking*, 1, 55–93.

Gürkaynak, R. S., E. T. Swanson, and J. H. Wright (2026): "High-Frequency Identification," *Working paper*.

Hack, L., K. Istrefi, and M. Meier (2023): "Identification of Systematic Monetary Policy," *CEPR Discussion Paper 17999*.

Hamilton, J. D., S. Pruitt, and S. Borger (2011): "Estimating the market-perceived monetary policy rule," *American Economic Journal: Macroeconomics*, 3, 1–28.

Hoesch, L., B. Rossi, and T. Sekhposyan (2023): "Has the information channel of monetary policy disappeared? Revisiting the empirical evidence," *American Economic Journal: Macroeconomics, forthcoming*.

Holston, K., T. Laubach, and J. C. Williams (2023): "Measuring the natural rate of interest after COVID-19," *Federal Reserve Bank of New York Staff Reports, no. 1063*.

Jordà, Ò., M. Schularick, and A. M. Taylor (2015): "Betting the house," *Journal of International Economics*, 96, S2–S18.

Keane, M. and T. Neal (2023): "Instrument strength in IV estimation and inference: A guide to theory and practice," *Journal of Econometrics*, 235, 1625–1653.

Kuttner, K. N. (2001): "Monetary policy surprises and interest rates: Evidence from the fed funds futures market," *Journal of Monetary Economics*, 47, 523–544.

Laubach, T. and J. C. Williams (2003): "Measuring the natural rate of interest," *Review of Economics and Statistics*, 85, 1063–1070.

Lee, D. S., J. McCrary, M. J. Moreira, and J. Porter (2022): "Valid t-ratio Inference for IV," *American Economic Review*, 112, 3260–3290.

Lewis, D. J. and K. Mertens (2025): "A Robust Test for Weak Instruments for 2SLS with Multiple Endogenous Regressors," *Review of Economic Studies*.

Mikusheva, A. (2010): "Robust confidence sets in the presence of weak instruments," *Journal of Econometrics*, 157, 236–247.

Miranda-Agrippino, S. and G. Ricco (2021): "The transmission of monetary policy shocks," *American Economic Journal: Macroeconomics*, 13, 74–107.

Montiel Olea, J. L. and C. Pflueger (2013): "A robust test for weak instruments," *Journal of Business & Economic Statistics*, 31, 358–369.

Montiel Olea, J. L., M. Plagborg-Møller, E. Qian, and C. K. Wolf (2024): "Double Robustness of Local Projections and Some Unpleasant VARithmetic," *Working Paper*.

——— (2025): "Local Projections or VARs? A Primer for Macroeconomists," *Working Paper*.

Nakamura, E., V. Riblier, and J. Steinsson (2025): "Beyond the Taylor Rule," *NBER Working Paper No. 34200*.

Nakamura, E. and J. Steinsson (2018): "High-frequency identification of monetary non-neutrality: The information effect," *The Quarterly Journal of Economics*, 133, 1283–1330.

Newey, W. K. and K. D. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica: Journal of the Econometric Society*, 703–708.

Patton, A. J. and A. Timmermann (2010): "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion," *Journal of Monetary Economics*, 57, 803–820.

Pesaran, M. H. (2006): "Estimation and inference in large heterogeneous panels with a multifactor error structure," *Econometrica*, 74, 967–1012.

——— (2015): "Testing weak cross-sectional dependence in large panels," *Econometric reviews*, 34, 1089–1117.

Pesaran, M. H. and R. Smith (1995): "Estimating long-run relationships from dynamic heterogeneous panels," *Journal of Econometrics*, 68, 79–113.

Prais, S. J. and C. B. Winsten (1954): "Trend estimators and serial correlation," *Cowles Commission Discussion Paper*.

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

Romer, C. D. and D. H. Romer (2000): "Federal Reserve information and the behavior of interest rates," *American Economic Review*, 90, 429–457.

——— (2004): "A new measure of monetary shocks: Derivation and implications," *American Economic Review*, 94, 1055–1084.

Rudebusch, G. D. (2002): "Term structure evidence on interest rate smoothing and monetary policy inertia," *Journal of Monetary Economics*, 49, 1161–1187.

——— (2006): "Monetary Policy Inertia: Fact or Fiction?" *International Journal of Central Banking*.

Schmeling, M., A. Schrimpf, and S. A. Steffensen (2022): "Monetary policy expectation errors," *Journal of Financial Economics*, 146, 841–858.

Stein, J. C. and A. Sunderam (2018): "The fed, the bond market, and gradualism in monetary policy," *The Journal of Finance*, 73, 1015–1060.

Stock, J. and M. Yogo (2005): "Asymptotic distributions of instrumental variables statistics with many instruments," *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 6, 109–120.

Stock, J. H. and M. W. Watson (2016): "Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics," 2, 415–525.

——— (2018): "Identification and estimation of dynamic causal effects in macroeconomics using external instruments," *The Economic Journal*, 128, 917–948.

Stock, J. H. and J. H. Wright (2000): "GMM with weak identification," *Econometrica*, 68, 1055–1096.
Stock, J. H., J. H. Wright, and M. Yogo (2002): "A survey of weak instruments and weak identification in generalized method of moments," *Journal of Business & Economic Statistics*, 20, 518–529.
Sutherland, C. S. (2023): "Forward guidance and expectation formation: A narrative approach," *Journal of Applied Econometrics*, 38, 222–241.
Woodford, M. (2003): *Interest and prices: Foundations of a theory of monetary policy*, Princeton University Press.

A Reaction Function Framework

A.1 The Federal Open Market Committee Reaction Function

We begin by defining the information set (Ω) of each FOMC ($\sim$) participant, $\tilde{i}$, at time t_n, $\tilde{\Omega}_{\tilde{i}t_n}$.[32] Now, because our focus will eventually be the *forecaster-expected* reaction function, and because financial markets, forecasters, journalists, and the broader public generally focus on the **median FOMC projection**, we abstract from each FOMC member, $\tilde{i}$, and focus on the median FOMC projection. Because each point projection is anonymous and not linked to its other matching point projections, forecasters must try to "*infer a reaction function from committee-wide median economic and interest rate projections that often come from different individuals*." (Goldman Sachs Economics Research, July 1, 2025).

Yet, even if the projections from each FOMC member were linked, "*this information would likely be useful to investors—**knowing the reaction function of the median participant inferred from their linked projections** would have helped to predict monetary policy surprises in the past*." (ibid, emphasis added). In other words, forecasters would still focus on the median participant. So, for simplicity, and without any substantial loss of generality, we will treat all median projections as part of the same set of projections, i.e. having followed the projection process below. As such, we do not need the $\tilde{i}$ subscript.

Next, we define function $\tilde{f}$, a function that maps the information set, $\tilde{\Omega}_{t_n}$ to a conditional expectation, $\tilde{\mathbb{E}}_{t_n}\left[\,.\;\middle|\;\tilde{\Omega}_{t_n}\right]$.[33] This conditional expectation represents an FOMC macroeconomic projection at time t_n.

$$\tilde{f} : \tilde{\Omega}_{t_n} \rightarrow \tilde{\mathbb{E}}_{t_n}\left[\;\mathbf{x_{t_n+h_{t_n}}}\;\middle|\;\tilde{\Omega}_{t_n}\;\right] \tag{56}$$

$\mathbf{x_{t_n+h_{t_n}}}$ denotes a vector of k=4 macroeconomic variables (real GDP growth, unemployment rate, inflation rate, core inflation rate), at some point in the future, $t+h_{t_n}$, where $h_{t_n} \in \{h^0_{t_n}, h^1_{t_n}, h^2_{t_n}, h^3_{t_n}, h^4_{t_n}\}$ denotes a fixed forecast horizon. In March, June, September, and December, the FOMC projects (only) these macroeconomic variables for the current year ($h^0_{t_n}$), one year from now ($h^1_{t_n}$), two years from now ($h^2_{t_n}$), three years from now ($h^3_{t_n}$) (in September and December only), and longer run ($h^4_{t_n}$). As such, the forecast horizon, h_{t_n}, is t_n-dependent because each horizon becomes shorter as the calendar year progresses.

The median FOMC member's information set, $\tilde{\Omega}_{t_n}$, is composed of available macroeconomic data at t_n, the path of the policy rate, i.e. the federal funds rate (FFR), path up to t_n, past FOMC projections, past FOMC judgment, financial market data, market intelligence, staff judgment, and news.[34]

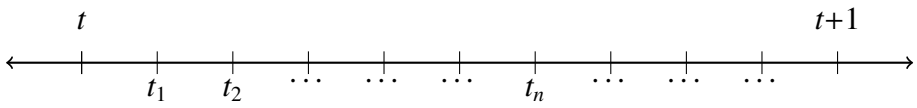
Time. To make sense of evolving information sets (filtrations), shocks, and what is known at the time a projection is made, we must be very careful with the sequence of events. Throughout the paper, t=$\{1, 2, \ldots, T\}$, $t \in \mathbb{N}^+$ is a standard indexing set, which denotes monthly periods, as all data are monthly. Further, let the set t_n=$\{t_1, t_2, \ldots\}$, $n \in \mathbb{N}^+$ denote arbitrarily many sub-periods. The gaps between each t_n correspond to small increments of time, over which information is accumulated, information sets are updated, and new expectations are formed. This would usually happen during the FOMC policy deliberation process, but in principle, it could happen at any time as new information arrives continuously (and forecasters & markets know this). The notation that follows is intricate but necessary—at least until §A.2, at which point we simplify the notation considerably.

[32] Throughout the paper, $\sim$ denotes objects attributable to the FOMC.

[33] The conditional expectation $\tilde{\mathbb{E}}_{t_n}\left[\,.\;\middle|\;\tilde{\Omega}_{t_n}\right]$ is defined as an $\tilde{\Omega}_{t_n}$–measurable function.

[34] We define $\tilde{\Omega}_{t_n}$ as a filtration or sigma-algebra to keep the framework general. This maintains flexibility over possible informational states, which allows us to consider counterfactuals (alternative scenarios) within the model.

Figure 13: **Monthly Period (t) and Sub-Period (t_n) Timeline**



Before the FOMC releases its *Summary of Economic Projections* (SEP), each FOMC participant maps their realized information set ($\tilde{\omega}_{t_n}$) to a vector of point forecasts. For ease of exposition, here FOMC participants start their recursive projection process at t_1:

$$\tilde{f} : \tilde{\omega}_{t_1} \in \tilde{\Omega}_{t_1} \rightarrow \tilde{\mathbb{E}}_{t_1}\left[\mathbf{x}_{\mathbf{t_1}+\mathbf{h}^{\mathbf{0}}_{\mathbf{t_1}}} \;\middle|\; \tilde{\Omega}_{t_1} = \tilde{\omega}_{t_1} \right] \in \mathbb{R}^k \tag{57}$$

This is a recursive process in which the FOMC participant adds these macroeconomic point forecasts to their information set, yielding $\tilde{\Omega}_{t_2}$, where $t_2 > t_1$ and $\tilde{\Omega}_{t_1} \subset \tilde{\Omega}_{t_2}$.[35] Next, as monetary policy is forward-looking (because it operates with a lag), each FOMC participant maps these point macroeconomic projections to a future FFR (i.e. an FFR projection) using $\tilde{g}$, a **monetary policy reaction function**:

$$\tilde{g} : \mathbb{R}^k \rightarrow \mathbb{R} \tag{58}$$

$$\tilde{g} : \tilde{\mathbb{E}}_{t_1}\left[\mathbf{x}_{\mathbf{t_1}+\mathbf{h}^{\mathbf{0}}_{\mathbf{t_1}}} \;\middle|\; \tilde{\Omega}_{t_1} = \tilde{\omega}_{t_1} \right] \rightarrow \tilde{\mathbb{E}}_{t_2}\left[\tilde{r}_{t_2+h^0_{t_2}} \;\middle|\; \tilde{\Omega}_{t_2} = \tilde{\omega}_{t_2} \right] \in \mathbb{R} \tag{59}$$

Or, equivalently:

$$\tilde{g}\left(\tilde{f}\left(\tilde{\omega}_{t_1}\right)\right) = \tilde{g}\left(\tilde{\mathbb{E}}_{t_1}\left[\mathbf{x}_{\mathbf{t_1}+\mathbf{h}^{\mathbf{0}}_{\mathbf{t_1}}} \;\middle|\; \tilde{\Omega}_{t_1} = \tilde{\omega}_{t_1}\right]\right) = \tilde{\mathbb{E}}_{t_2}\left[\tilde{r}_{t_2+h^0_{t_2}} \;\middle|\; \tilde{\Omega}_{t_2} = \tilde{\omega}_{t_2}\right] \in \mathbb{R} \tag{60}$$

Next, before submitting this FFR point projection to The Federal Reserve Board of Governors, the FOMC participant considers whether to add judgment $\tilde{\lambda}_{t_n}$, to the projection.[36] That is, the final projection from Eq. 59 may ultimately differ somewhat from the point projection published as part of the SEP because the FFR projection may be adjusted by **FOMC-expected judgment**, $\tilde{\mathbb{E}}_{t_n}\left[\tilde{\lambda}_{t_n+h_{t_n}} \;\middle|\; \tilde{\Omega}_{t_n} = \tilde{\omega}_{t_n} \right]$.

Given the sources of judgment (discussed in §2), we hypothesize that the law of motion for judgment is a persistent process:

$$\tilde{\mathbb{E}}_{t_n}\left[\tilde{\lambda}_{t_n+h_{t_n}} \;\middle|\; \tilde{\Omega}_{t_n} = \tilde{\omega}_{t_n} \right] = \tilde{\theta}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\left[\tilde{\lambda}_{t_{n-2}+h_{t_{n-2}}} \;\middle|\; \tilde{\Omega}_{t_{n-2}} = \tilde{\omega}_{t_{n-2}} \right] + \tilde{\varepsilon}^{h_{t_{n-1}}}_{t_{n-1}}, \qquad \tilde{\varepsilon}^{h_{t_n}}_{t_n} \overset{iid}{\sim} \left(0, \sigma^2_{\tilde{\varepsilon}}\right) \tag{61}$$

$\tilde{\varepsilon}^{h_{t_n}}_{t_n}$ is a shock to FOMC-expected judgment. When h_{t_n}=0, then $\tilde{\mathbb{E}}_{t_2}\left[\tilde{r}_{t_2+h^0_{t_2}} \;\middle|\; \tilde{\Omega}_{t_2} = \tilde{\omega}_{t_2}\right]$ simply corresponds to the policy rate, and so $\tilde{\varepsilon}^{h_{t_n}}_{t_n}$ is a **traditional monetary policy shock**. When h_{t_n}>0, then $\tilde{\varepsilon}^{h_{t_n}}_{t_n}$ is a **monetary policy rate path projection shock**, which plausibly has comparable influence on markets and the economy to the traditional monetary policy shock. Although not formally stated as such, estimating perceptions of $\tilde{\varepsilon}^{h_{t_n}}_{t_n}$ has been the focus of the literature that estimates perceived monetary shocks based on high-frequency interest rate futures price changes (i.e. from Kuttner (2001) to Bauer and Swanson (2023)). This further underscores the importance of clarifying the definition of $\tilde{\varepsilon}^{h_{t_n}}_{t_n}$.

[35] More formally, we add to the increasing filtration $\{\tilde{\Omega}_{t_n}\}$ that reflects growing information over time.

[36] The FOMC participant could have similarly added judgment to their macroeconomic projections of $\mathbf{x}_{\mathbf{t_n}}$ earlier as well, but that is not our focus here.

So, FOMC-expected judgment, or, $\tilde{\mathbb{E}}_{t_n}\left[\ \tilde{\lambda}_{t_n+h_{t_n}}\ \middle|\ \tilde{\Omega}_{t_n}=\tilde{\omega}_{t_n}\ \right] := \tilde{\lambda}_{t_n}^{h_{t_n}}$, for short, is known when the FFR projection is finalized ($\tilde{\Omega}_{t_4}$–measurable s.t. $\tilde{\Omega}_{t_2}\subset\tilde{\Omega}_{t_3}\subset\tilde{\Omega}_{t_4}$).

$$\tilde{\mathbb{E}}_{t_4}\left[\ \tilde{r}_{t_4+h^0_{t_4}}\ \middle|\ \tilde{\Omega}_{t_4}=\tilde{\omega}_{t_4}\ \right]=\tilde{\mathbb{E}}_{t_2}\left[\ \tilde{r}_{t_2+h^0_{t_2}}\ \middle|\ \tilde{\Omega}_{t_2}=\tilde{\omega}_{t_2}\ \right]+\tilde{\mathbb{E}}_{t_3}\left[\ \tilde{\lambda}_{t_3+h^0_{t_3}}\ \middle|\ \tilde{\Omega}_{t_3}=\tilde{\omega}_{t_3}\ \right] \tag{62}$$

Next, the projection produced using Eq. 62 is added to $\tilde{\Omega}_{t_5}$, and the FOMC participant then begins producing longer-horizon ($h^1_{t_n} > h^0_{t_n}$) projections.

$$\tilde{f}:\tilde{\omega}_{t_5}\in\tilde{\Omega}_{t_5}\rightarrow\tilde{\mathbb{E}}_{t_5}\left[\ \mathbf{x}_{\mathbf{t_5}+\mathbf{h}^{\mathbf{1}}_{\mathbf{t_5}}}\ \middle|\ \tilde{\Omega}_{t_5}=\tilde{\omega}_{t_5}\ \right]\in\mathbb{R}^k \tag{63}$$

The FOMC participant then adds these longer-horizon ($h^1_{t_5}$) macroeconomic projections to their updated information set, $\tilde{\omega}_{t_6}$ ($\tilde{\Omega}_{t_5}\subset\tilde{\Omega}_{t_6}$). These projections are then mapped to a more-distant FFR projection:

$$\tilde{g}:\tilde{\mathbb{E}}_{t_5}\left[\ \mathbf{x}_{\mathbf{t_5}+\mathbf{h}^{\mathbf{1}}_{\mathbf{t_5}}}\ \middle|\ \tilde{\Omega}_{t_5}=\tilde{\omega}_{t_5}\ \right]\rightarrow\tilde{\mathbb{E}}_{t_6}\left[\ \tilde{r}_{t_6+h^1_{t_6}}\ \middle|\ \tilde{\Omega}_{t_6}=\tilde{\omega}_{t_6}\ \right]\in\mathbb{R} \tag{64}$$

Once again, FOMC-expected judgment, $\tilde{\mathbb{E}}_{t_7}\left[\ \tilde{\lambda}_{t_7+h^1_{t_7}}\ \middle|\ \tilde{\Omega}_{t_7}=\tilde{\omega}_{t_7}\ \right]$ may then be added.

$$\tilde{\mathbb{E}}_{t_8}\left[\ \tilde{r}_{t_8+h^1_{t_8}}\ \middle|\ \tilde{\Omega}_{t_8}=\tilde{\omega}_{t_8}\ \right]=\tilde{\mathbb{E}}_{t_6}\left[\ \tilde{r}_{t_6+h^1_{t_6}}\ \middle|\ \tilde{\Omega}_{t_6}=\tilde{\omega}_{t_6}\ \right]+\tilde{\mathbb{E}}_{t_7}\left[\ \tilde{\lambda}_{t_7+h^1_{t_7}}\ \middle|\ \tilde{\Omega}_{t_7}=\tilde{\omega}_{t_7}\ \right] \tag{65}$$

The FOMC participant iterates the process until arriving at the *Long-Run Projection* of The FFR and each macroeconomic variable included in each SEP. Turning to **functional form**, following the detailed studies of the functional form of The Fed's monetary policy rule by Rudebusch (2002), Rudebusch (2006), Coibion and Gorodnichenko (2012) and Stein and Sunderam (2018), as well as historical Tealbooks. To simplify notation, **we will suppress the explicit conditioning event and write conditional expectations as** $\tilde{\mathbb{E}}_{t_n}[\cdot]$, where the subscript t_n denotes conditioning on the realized information set $\tilde{\Omega}_{t_n}=\tilde{\omega}_{t_n}$.

Assumption 7. *We assume the forward-looking monetary policy reaction function takes the form:*

$$\begin{aligned}&\tilde{\mathbb{E}}_{t_n}\left[\ \tilde{r}_{t_n+h_{t_n}}\ \right]=\tilde{\mathbb{E}}_{t_{n-1}}\left[\ \tilde{\lambda}_{t_{n-1}+h_{t_{n-1}}}\ \right]+\\&\tilde{\mathbb{E}}_{t_{n-2}}\left[\ \tilde{\rho}^{h_{t_{n-2}}}_{t_{n-2}}(L)\tilde{r}_{t_{n-2}+h_{t_{n-2}}}+\tilde{\phi}^{\pi^*}_{t_{n-2}}\pi^*_{t_n}+\tilde{\phi}^{r^*}_{t_{n-2}}r^*_{t_{n-2}}+\left(\mathbf{x}_{\mathbf{t_{n-2}}+\mathbf{h}_{\mathbf{t_{n-2}}}}-\mathbf{x}^*_{\mathbf{t_{n-2}}}\right)'\tilde{\boldsymbol{\beta}}^{\mathbf{h}_{\mathbf{t_{n-2}}}}_{\mathbf{t_{n-2}}}\ \right]\end{aligned} \tag{66}$$

$\tilde{\rho}^{h_{t_n}}_{t_n}(L)$ is a lag polynomial such that $\tilde{\rho}^{h_{t_n}}_{t_n}(L)\tilde{r}_{t_n+h_{t_n}}=\sum_{\ell=1}^{L}\tilde{\rho}^{\ell}_{t_n}\tilde{r}_{t_n+h_{t_n}-l}$. $\tilde{\rho}^{\ell}_{t_n}$ is a monetary policy gradualism parameter.[37] Accordingly, $\tilde{\phi}^{\pi^*}_{t_n}\propto\left(1-\sum_{\ell=1}^{L}\tilde{\rho}^{\ell}_{t_n}\right)$, $\tilde{\phi}^{r^*}_{t_n}\propto\left(1-\sum_{\ell=1}^{L}\tilde{\rho}^{\ell}_{t_n}\right)$ (as in Bauer et al. (2024)).

$\pi^*_{t_n}\in\mathbb{R}$ is the inflation target, which is set at two percent, but which could, in theory, be time-varying. $r^*_{t_n}\in\mathbb{R}$ is the time-varying central bank estimate of the real neutral policy rate. $\mathbf{x}^*_{\mathbf{t_n}}\in\mathbb{R}^k$ is a vector that matches $\mathbf{x}_{\mathbf{t_n}+\mathbf{h}_{\mathbf{t_n}}}$, which includes the inflation target, and the time-varying, median-FOMC-projected, long-run real GDP growth and unemployment rates, respectively—each in the appropriate place.

[37] Using Bayesian Information Criterion, Coibion and Gorodnichenko (2012) find that a higher-order autoregressive process best captures historical Fed reaction functions. Woodford (2003) shows that the optimal interest rate rule in a New Keynesian model has AR(2) interest rate inertia. In this general form of the reaction function, we are agnostic about the lag order, L.

As an aside, if we simply set h_{t_n}=0, $\tilde{\lambda}^{h}_{t_n}$=0, and the parameters to be constant, then Eq. 66 nests a standard monetary (Taylor) rule—where $\tilde{\mathbb{E}}_{t_n}\left[\mathbf{x_{t_n}}\right]$ is a nowcast. For clarity, we use t instead of t_n here:

$$\tilde{\mathbb{E}}_t\left[\ \tilde{r}_t\ \right] = \tilde{\mathbb{E}}_t\left[\ \tilde{\rho}(L)\tilde{r}_t + \tilde{\phi}^{\pi^*}\pi^*_t + \tilde{\phi}^{r^*}r^*_t + \left(\mathbf{x_t} - \mathbf{x^*_t}\right)'\tilde{\beta}\ \right] \Rightarrow \tag{67}$$

$$\tilde{r}_t = \tilde{\rho}(L)\tilde{r}_t + \tilde{\phi}^{\pi^*}\pi^*_t + \tilde{\phi}^{r^*}r^*_t + \left(\mathbf{x_t} - \mathbf{x^*_t}\right)'\tilde{\beta} \tag{68}$$

where, invariably, a monetary policy shock is added, which implicitly, must follow the projection timing sequence and definitions outlined above:

$$\tilde{r}_t = \tilde{\rho}(L)\tilde{r}_t + \tilde{\phi}^{\pi^*}\pi^*_t + \tilde{\phi}^{r^*}r^*_t + \left(\mathbf{x_t} - \mathbf{x^*_t}\right)'\tilde{\beta} + \tilde{\varepsilon}_t \tag{69}$$

Assumption 8. *Coefficients are known given* $\tilde{\Omega}_{t_n}$ *($\tilde{\Omega}_{t_n}$–measurable).*

This implies that we can pull the coefficients outside the conditional expectation operator. Combining Assumptions 7, 8, and the linear property of the conditional expectation operator implies:

$$\begin{aligned}\tilde{\mathbb{E}}_{t_n}\left[\ \tilde{r}_{t_n+h_{t_n}}\ \right] = \tilde{\rho}^{h_{t_n}}_{t_n}(L)\tilde{\mathbb{E}}_{t_{n-2}}\left[\tilde{r}_{t_{n-2}+h_{t_{n-2}}}\ \right] + \tilde{\phi}^{\pi^*}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\left[\ \pi^*\ \right] + \tilde{\phi}^{r^*}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\left[\ r^*_{t_{n-2}}\ \right] + \\ \tilde{\mathbb{E}}_{t_{n-2}}\left[\ \left(\mathbf{x_{t_{n-2}+h_{t_{n-2}}}} - \mathbf{x^*_{t_{n-2}}}\right)'\ \right]\tilde{\beta}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} + \tilde{\mathbb{E}}_{t_{n-1}}\left[\ \lambda^{h_{t_{n-1}}}_{t_{n-1}}\ \right]\end{aligned} \tag{70}$$

Assumption 9. *The coefficients in Equation 70 follow a random coefficients model.*

Bauer, Pflueger, and Sunderam (2024) instead assume a random walk model, which implies the need to use Kalman Filtering. Their Kalman Filtering estimation results are very similar to those from their baseline approach, and their baseline approach is similar to our approach below. More specifically, $\tilde{\beta}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} \in \mathbb{R}^k$ is a vector of time-varying (and also projection-horizon-specific), macro-response coefficients:

$$\tilde{\beta}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} = \tilde{\beta} + \tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} \tag{71}$$

where $\tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} \in \mathbb{R}^k$ is a vector of macro-response coefficient shocks, where $\tilde{\upsilon}^{h_{t_n}}_{t_n} \sim i.i.d.\left(0, \sigma^{\mathbf{2}}_{\tilde{\upsilon}}\right)$.[38] The source of $\tilde{\upsilon}^{h_{t_n}}_{t_n}$ could be a new committee participant, changing preferences, uncertainty, or inconsistency.

Forward guidance (FG) provides additional information about the FOMC's current reaction function(s) to the public *beyond* the monetary policy committee's macroeconomic projections. FG is broader than just FOMC-expected judgment, $\tilde{\lambda}^{h_{t_n}}_{t_n}$, only. To understand how FG works in practice, in theory, and to make policy recommendations, it is crucial that we account for its multi-faceted nature, rather than conflating distinct signals. FG takes one of the following four forms:

1. A **quantitative policy rate projection**, $\tilde{\mathbb{E}}_{t_n}\left[\tilde{r}_{t_n+h_{t_n}}\right]$ (left-hand side of Eq. 66), which is time-contingent FG, i.e., the FOMC Fed Funds Rate (FFR) projections released quarterly in The Summary of Economic Projections (SEP).

[38] Time-varying coefficients depart from many models in this literature, such as the noisy information model of Miranda-Agrippino and Ricco (2021), but are consistent with Bauer and Swanson (2023), Bauer, Pflueger, and Sunderam (2024), and Hack, Istrefi, and Meier (2023).

2. **Qualitative judgment** $\left(\tilde{\mathbb{E}}_{t_n}\left[\ \tilde{\lambda}_{t_n+h_{t_n}}\ \right]\right)$ delivered as qualitative FG, e.g., "*The Committee will maintain the target range for the federal funds rate at 0 to 1/4 percent and continues to anticipate that economic conditions are likely to warrant exceptionally low levels of the federal funds rate for an extended period.*" (source).[39]

3. An updated description of the **response to macroeconomic projections** in the reaction function:

$$\Delta_n \tilde{\beta}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} = \tilde{\beta}_{\mathbf{t_n}} - \tilde{\beta}_{\mathbf{t_{n-1}}} = \left(\tilde{\beta} + \tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}}\right) - \left(\tilde{\beta} + \tilde{\upsilon}^{\mathbf{h_{t_{n-1}}}}_{\mathbf{t_{n-1}}}\right) = \Delta_n \tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} \tag{72}$$

   (from Eq. 66 and 71), which is state-contingent FG, e.g. FG indicating a stronger response to inflation post-pandemic, such as "*The Committee is strongly committed to returning inflation to its 2 percent objective.*" (source).[40]

4. New information about the **functional form of the reaction *function***, $\tilde{g}^{h_{t_n}}_{t_n}$ e.g. COVID information being de facto added to the reaction function: "*The Committee will continue to monitor the implications of incoming information for the economic outlook, including **information related to public health**, as well as global developments and muted inflation pressures, and will use its tools and act as appropriate to support the economy.*" (source).[41]

The information channel is separate and is represented by either quantitative or qualitative updates to the central bank's macroeconomic projections (*excluding* policy rate projections), particularly $\Delta_n \tilde{\mathbb{E}}_{t_n}\left[\mathbf{x_{t_n+h_{t_n}}}\right] = \tilde{\mathbb{E}}_{t_n}\left[\mathbf{x_{t_n+h_{t_n}}}\right] - \tilde{\mathbb{E}}_{t_{n-1}}\left[\mathbf{x_{t_{n-1}+h_{t_{n-1}}}}\right]$. The concept of the information channel dates back to Romer and Romer (2000) at least.

> "*The key idea is that information the Federal Reserve has about the economy that is not known to market participants is likely to be reflected in the Federal Reserve's internal forecasts.*" "*[The paper] demonstrates that the Federal Reserve has considerable **information about inflation** beyond what is known to commercial forecasters. It also shows that monetary-policy actions provide signals of the Federal Reserve's information and that commercial forecasters modify their forecasts in response to those signals.*" (Romer and Romer (2000), p.429).

The idea has since been carried forward, "*Following Romer and Romer (2000), we refer to the effect of FOMC announcements on private-sector views of non-monetary economic fundamentals as "Fed information effects.*"" Nakamura and Steinsson (2018), p.1286.[42] Several papers have adjusted interest rate-based measures for information effects. Romer and Romer (2004) use Federal Reserve Greenbook forecasts to control for information effects, as do Miranda-Agrippino and Ricco (2021).

$\Delta_n \tilde{\mathbb{E}}_{t_n}\left[\mathbf{x_{t_n+h_{t_n}}}\right]$ is driven by: new macroeconomic data, new information about how the projections are formed, FOMC judgment added to the macro projections.[43] Although new macroeconomic data

[39] $\tilde{\lambda}^{h_{t_n}}_{t_n}$ may be either qualitative or time-contingent FG, and less likely, state-dependent FG; it may be either Delphic or Odyssean FG (Sutherland (2023) elaborates on types of FG).

[40] So, here, $\Delta_n \tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} > 0$. The update could, in principle, be quantitative also, but monetary policy committees, to date, do not provide real-time quantitative monetary policy reaction functions.

[41] We **exclude** information about the central bank's macroeconomic projections. This is a subtle yet crucial distinction, and one that must be made to differentiate between state-dependent FG and information effects.

[42] Information effects are studied in, e.g., Bauer and Swanson (2023), Hoesch, Rossi, and Sekhposyan (2023).

[43] Byrne, Goodhead, McMahon, and Parle (2023) delve even deeper into the information channel and provide a complementary framework for conceptualizing and analyzing economic assessment updates.

are public information, the central bank does not always explain which of these three factors drove $\Delta_n \tilde{\mathbb{E}}_{t_n}\left[\mathbf{x_{t_n+h_{t_n}}}\right]$, and by how much. So, the public usually cannot disentangle which of the three channels drove $\Delta_n \tilde{\mathbb{E}}_{t_n}\left[\mathbf{x_{t_n+h_{t_n}}}\right]$. Further, some central banks, such as The European Central Bank, Bank of England, and Reserve Bank of Australia use financial market-implied policy rate paths to produce their macroeconomic projections, so the information and signal content of those projections is different.

Ideally, we would estimate Eq. 66 directly. Essentially, Aruoba and Drechsel (2025) present a natural-language approach for doing something similar with h_{t_n}=0, and assuming constant parameters. Yet, at present, we do not have enough data to do so for h_{t_n}>0. We would if we have complete *Summary of Economic Projections* panel data attributable to each FOMC participant, but there are substantial gaps.[44] We *might* if we chose to pool data over forecast horizon h (as in González-Astudillo and Tanvir (2023)), but that would not be consistent with our framework. Instead, we adopt a strategy that affords us more data: we estimate the forecaster-expected reaction function using monthly forecaster data.

A.2 The Expected Central Bank Reaction Function

A crucial distinction is the true versus *perceived* central bank reaction function. It is the perceived reaction function from which most of the monetary shocks series in the literature are derived. That is, most papers derive or use monetary policy shocks from high-frequency shifts in the expected path of monetary policy as implied by interest rate futures and swaps. As such, we could also derive monetary policy shocks from shifts in the expected path of monetary policy as judged by professional forecasters, or even the public. The key reason to use perceived shocks is that we observe monetary expectations data at higher frequencies (monthly, daily, continuously) than FOMC decisions (eight times per year).

Following a large literature on the systematic violations of full information rational expectations shown by professional forecasters (e.g. Coibion and Gorodnichenko (2015), Bordalo, Gennaioli, Ma, and Shleifer (2020), Bianchi, Ilut, and Saijo (2024)), forecaster expectation formation is not assumed to be consistent with rational expectations. Further, if the forecasts were rational, they should have the same time series properties as the underlying data being forecasted, but they do not ($\mathcal{A}$). So, we assume that forecasters use a simple, linear model, which makes the parameters of interest identifiable.

Time. To begin, we take expectations of Eq. 66 with respect to the information set of a forecaster, $\ddot{\Omega}_{it}$, at the next monthly forecasting submission deadline, t+1, which by definition (Figure 13), always occurs after t_n. $\ddot{\mathbb{E}}_{it}\left[. \;\middle|\; \ddot{\Omega}_{it}\right]$ is individual i's ($i = 1, ..., N$) expectation conditional on their information set, $\ddot{\Omega}_{it}$.[45] We use the double dots to refer to a forecaster expectation. So, the forecaster-expected central bank reaction function is:[46]

$$\begin{aligned}\ddot{\mathbb{E}}_{i,t+1}\Big[\;\tilde{\mathbb{E}}_{t_n}\Big[\;\tilde{r}_{t_n+h_{t_n}}\;\Big]\;\Big|\;\ddot{\Omega}_{i,t+1}\Big] = \ddot{\mathbb{E}}_{i,t+1}\Big[\;\tilde{\mathbb{E}}_{t_{n-1}}\Big[\;\tilde{\lambda}_{t_{n-1}+h_{t_{n-1}}}\;\Big]\;\Big|\;\ddot{\Omega}_{i,t+1}\;\Big] + \\ \ddot{\mathbb{E}}_{i,t+1}\Big[\;\tilde{\mathbb{E}}_{t_{n-2}}\Big[\;\tilde{\rho}^{h_{t_n}}_{t_{n-2}}(L)\tilde{r}_{t_{n-2}+h_{t_{n-2}}} + \tilde{\phi}^{\pi^*}_{t_{n-2}}\pi^*_t + \tilde{\phi}^{r^*}_{t_{n-2}}r^*_{t_{n-2}} + \left(\mathbf{x_{t_{n-2}+h_{t_{n-2}}}} - \mathbf{x^*_{t_{n-2}}}\right)'\tilde{\beta}^{\mathbf{h_{t_{n-2}}}}_{\mathbf{t_{n-2}}}\;\Big]\;\Big|\;\ddot{\Omega}_{i,t+1}\;\Big]\end{aligned} \tag{73}$$

[44] Such data are available from 2007 to 2012, and from 2016 to 2020 only (Arai (2023)).

[45] If the function that forms $\ddot{\Omega}_{i,t+1}$ were $f^{\ddot{\Omega}}: \tilde{\Omega}_t \to \ddot{\Omega}_{i,t+1}$, then we could simplify the conditioning set in Eq. 74 and write $\ddot{\mathbb{E}}_{i,t+1}\left[\tilde{r}_{t+h} \;\middle|\; \ddot{\Omega}_{i,t+1}\right]$ by the Law of Iterated Expectations, but it is more likely that $f^{\ddot{\Omega}}: \tilde{\Omega}_t \bigcup \Omega^i_t \to \ddot{\Omega}_{i,t+1}$, where Ω^i_t is a set of private signals. Caballero and Simsek (2022) suggests that central banks partially accommodate market views, so we can further postulate $f^{\tilde{\Omega}}: \ddot{\Omega}_{i,t+1} \bigcup \Omega^c_t \to \tilde{\Omega}_t$, where Ω^c_t is a set of private central bank signals.

[46] We treat the function $\tilde{\mathbb{E}}_{t_n}\left[\;.\;\middle|\tilde{\Omega}_t\right]$ as a measurable random variable with respect to the outer sigma-algebra $\sigma\left(\ddot{\Omega}_{i,t+1}\right)$—consistent with the information asymmetry across agents.

Now, the FOMC releases its projections every three months, but in principle, an FOMC participant could update their projection at any time t_n (formally or informally). For example, if a major U.S. trade agreement were to be agreed some time after the latest SEP was released, then the FOMC participant may begin adjusting their growth projections in response (at least mentally, if not formally by updating a projections model). Forecasters know this.

Further, forecasters may also update their forecasts at any time, but must submit them to *Consensus Economics* and *Blue Chip*, respectively, once per month. Now, an FOMC projection at t_n could happen immediately before the next private forecast submission deadline at t+1—essentially coinciding with t+1. In that case, a FOMC projection at t_n would represent the latest state of the projection at t+1. We take the stance that, while forming their own forecasts, forecasters indeed have a view on the latest state of the FOMC SEP given recent news.

As such, we are interested in precisely the case in which t_{n-2}, t_{n-1}, t_n, and t+1 occur very close together in time. So, we can now abstract from the t_n sub-periods and their notation, and instead focus on monthly forecasting periods t with the knowledge that the projection formation process still allows for a causal effect of FOMC projections on private forecasts, and that our reaction function framework appropriately nests information sets and is consistent with measure theory such that it can be mapped to our econometric model.[47] And so we can simplify the notation in Eq. 73 as follows:

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \tilde{r}_{t+h_t}\ \right]\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \tilde{\rho}_t^{h_t}(L)\tilde{r}_{t+h_t} + \tilde{\phi}_t^{\pi^*}\pi_t^* + \tilde{\phi}_t^{r^*}r_t^* + \left(\mathbf{x_{t+h_t}} - \mathbf{x_t^*}\right)'\tilde{\boldsymbol{\beta}}_{\mathbf{t}}^{\mathbf{h_t}}\ \right]\ \right] + \ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \tilde{\lambda}_{t+h_t}\ \right]\ \right] \tag{74}$$

Forecasters are assumed to form their projections and to update their information sets in a manner consistent with the FOMC (Equations 56 to 71), which will be important to bear in mind as we consider forecaster-expected FOMC judgment later on (Eq. 77).

Assumption 10. *The coefficients are known given* $\ddot{\Omega}_{it}$ *($\ddot{\Omega}_{it}$–measurable). This implies that:*

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \left(\mathbf{x_{t+h_t}} - \mathbf{x_t^*}\right)'\ \right]\ \tilde{\boldsymbol{\beta}}_{\mathbf{t}}^{\mathbf{h_t}}\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \left(\mathbf{x_{t+h_t}} - \mathbf{x_t^*}\right)'\ \right]\ \right]\ \ddot{\boldsymbol{\beta}}_{\mathbf{it}}^{\mathbf{h_t}} \tag{75}$$

This implies that we can pull the coefficients outside the conditional expectation operator. This means that each forecaster knows the coefficients of their own perceived, forward-looking reaction function. To simplify subsequent derivations somewhat, we use the following notation for forecaster's coefficients:

$$\ddot{\boldsymbol{\beta}}_{\mathbf{it}}^{\mathbf{h_t}} := \ddot{\mathbb{E}}_{it}\left[\ \tilde{\boldsymbol{\beta}}_{\mathbf{t}}^{\mathbf{h_t}}\ \right] \tag{76}$$

By the linearity of the expectation operator, individuals can form a separate expectation for each component of Eq. 74. We consider each in turn below.

> "*Do market participants accept at face value what the FOMC declares about future policy, or do they form their own inferences about likely FOMC policy from other clues in the statements? More important, do their forecasts of future funds rate change because the statement changes their beliefs about the FOMC's reaction function, or because it changes their forecast of economic conditions that are expected to determine FOMC policy change, as a result of inferences they make from the statement about information available to the FOMC*?" (Michael Woodford in Campbell et al. (2012), p.66).

[47] This framework could also apply to market participants, who may also update their expectations at any point.

The framework above, below (§A.3), and its accompanying identification strategy (§2), allow us to differentiate between not only these channels, but more. That is, we can identify whether forecasters' "*forecasts of the future funds rate change because the statement changes their beliefs about the FOMC's reaction function*," (the perceived parameters): ρ_{it}^{ℓ}, $\ddot{\beta}_{\mathbf{it}}^{\mathbf{h}}$, and ε_{it}^{h}, "*or because it changes their forecast of economic conditions*" (the information channel): $\ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \left(\mathbf{x_{t+h}} - \mathbf{x_t^*}\right)'\ \right]\ \right]$.

$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\rho}_t^{h_t}(L)\ \tilde{\mathbb{E}}_t\left[\ \tilde{r}_{t+h_t}\ \right]\ \right]$ is **forecaster-expected monetary policy *gradualism***. That is, it is the monetary policy gradualism that is expected to prevail at time $t+h_t$. As above, we use the following notation for brevity: $\rho_{it}^{h_t}(L) := \ddot{\mathbb{E}}_{it}\left[\ \tilde{\rho}_t^{h_t}(L)\ \right]$.

$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ r_{t+h_t}^*\ \right]\ \right]$ is the **forecaster-expected FOMC real neutral rate projection**. If we treat the "Longer-run" FFR projection from The SEP as $\tilde{\mathbb{E}}_t\left[\ r_{t+h_t}^*\ \right]$, then forecasters receive updates quarterly.

$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \left(\mathbf{x_{t+h_t}} - \mathbf{x_t^*}\right)'\ \right]\ \right]$ is a **vector** (bolded term) containing the forecaster-expected set of FOMC macroeconomic projections. Because the FOMC provides macroeconomic projections on a quarterly basis in the SEP, these figures are known at those times. At all other times, forecasters must process information to form private views about how FOMC participants may have revised their projections. $\pi_{\mathbf{t+h_t}}^*$ is a vector also of length k that includes median long-run real GDP growth from the SEP, median long-run unemployment rate from the SEP, and the inflation target (or median long-run inflation growth from the SEP as they coincide).

$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\beta}_{\mathbf{t}}^{\mathbf{h_t}}\ \right] := \ddot{\beta}_{\mathbf{it}}^{\mathbf{h_t}}$ is the vector of forecaster-expected central bank macroeconomic projection response coefficients. Estimates of these coefficients are almost certainly heterogeneous because so little quantifiable information about them is provided by the FOMC. The FOMC occasionally provides implicit information about these coefficients' values in the form of qualitative or state-contingent forward guidance. Yet, central banks seldom articulate their reaction functions precisely enough to provide estimates of these coefficients.[48] Dispersion in forecaster-expected macro response coefficients is hypothesized to be a principal driver of the observed heterogeneity in forecasters' perceived policy rate projections.

Forecaster-expected FOMC judgment, our main focus, by Eq. 61, is:

$$\ddot{\mathbb{E}}_{i,t+1}\left[\ \tilde{\mathbb{E}}_{t_n}\left[\ \tilde{\lambda}_{t_n+h_{t_n}}\ \right]\ \right] = \ddot{\mathbb{E}}_{i,t+1}\left[\ \tilde{\mathbb{E}}_{t_n}\left[\ \tilde{\theta}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\left[\ \tilde{\lambda}_{t_{n-2}+h_{t_{n-2}}}\ \right] + \tilde{\varepsilon}_{t_{n-1}}^{h_{t_{n-1}}}\ \right]\ \right] \tag{77}$$

Using our shorthand notation, $\tilde{\lambda}_{t_n}^{h_{t_n}} := \tilde{\mathbb{E}}_{t_n}\left[\ \tilde{\lambda}_{t_n+h_{t_n}}\ \right]$, and assumptions thus far:

$$\ddot{\mathbb{E}}_{i,t+1}\left[\ \tilde{\lambda}_{t_n}^{h_{t_n}}\ \right] = \ddot{\mathbb{E}}_{i,t+1}\left[\ \tilde{\theta}_{t_n}\tilde{\lambda}_{t_{n-2}}^{h_{t_{n-2}}} + \tilde{\varepsilon}_{t_{n-1}}^{h_{t_{n-1}}}\ \right] \tag{78}$$

But, as above, we want to aggregate to a monthly version to match our data and better understand what

[48] There are exceptions. The Bank of Canada specified the monetary policy rule from its main DSGE model used to generate its economic projections (§1.3) along with corresponding parameter estimates (§2.2.2) here. To provide information about the historical pattern of its monetary policy reaction function, The Norges Bank included estimates of a simple interest rate rule in its December 2021 Monetary Policy Report here.

we will ultimately estimate:[49]

$$\tilde{\lambda}^{h}_{t_n} = \tilde{\theta}_{t_n}\tilde{\lambda}^{h}_{t_{n-2}} + \tilde{\varepsilon}^{h}_{t_{n-1}} \tag{79}$$

$$= \tilde{\theta}_{t_n}\left(\tilde{\theta}_{t_{n-2}}\tilde{\lambda}^{h}_{t_{n-4}} + \tilde{\varepsilon}^{h}_{t_{n-3}}\right) + \tilde{\varepsilon}^{h}_{t_{n-1}} \tag{80}$$

$$= \tilde{\theta}_{t_n}\tilde{\theta}_{t_{n-2}}\tilde{\lambda}^{h}_{t_{n-4}} + \tilde{\theta}_{t_n}\tilde{\varepsilon}^{h}_{t_{n-3}} + \tilde{\varepsilon}^{h}_{t_{n-1}} \tag{81}$$

$$\vdots$$

$$= \bar{\theta}_{t_n}\tilde{\lambda}^{h}_{t_{n-2R}} + \sum_{r=1}^{R}\tilde{\theta}^{r-1}\tilde{\varepsilon}^{h}_{t_{n-(2r-1)}} \quad \text{where} \quad \bar{\theta}_{t_n} := \tilde{\theta}_{t_n}\cdot\tilde{\theta}_{t_{n-2}}\cdot\ldots \tag{82}$$

where R denotes the total number of sub-periods within a given month. This yields the monthly-level law of motion:

$$\tilde{\lambda}^{h_t}_{t} = \bar{\theta}_t\tilde{\lambda}^{h_{t-1}}_{t-1} + \bar{\varepsilon}^{h_t}_{t}, \quad \text{where} \quad \bar{\theta}_t := \tilde{\theta}_{t_n}\cdot\tilde{\theta}_{t_{n-2}}\cdot\ldots, \quad \bar{\varepsilon}^{h_t}_{t} := \sum_{r=1}^{R}\tilde{\theta}^{r-1}\tilde{\varepsilon}^{h}_{t_{n-(2r-1)}} \tag{83}$$

Now we take expectations with respect to $\ddot{\Omega}_{it}$ such that t and t_n (from Eq. 83) coincide:

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\lambda}^{h_t}_{t}\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \bar{\theta}_t\tilde{\lambda}^{h_{t-1}}_{t-1} + \bar{\varepsilon}^{h_t}_{t}\ \right] \Rightarrow \tag{84}$$

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\lambda}^{h_t}_{t}\ \right] = \theta_{it}\ddot{\mathbb{E}}_{it}\left[\ \tilde{\lambda}^{h_{t-1}}_{t-1}\ \right] + \ddot{\mathbb{E}}_{it}\left[\ \bar{\varepsilon}^{h_t}_{t}\ \right] \tag{85}$$

And, finally, employing shorthand notation once again $\left(\ddot{\lambda}^{h_t}_{it} := \ddot{\mathbb{E}}_{it}\left[\ \tilde{\lambda}^{h_t}_{t}\ \right]\right)$:

$$\ddot{\lambda}^{h_t}_{it} = \theta_{it}\ddot{\lambda}^{h_{t-1}}_{i,t-1} + \varepsilon^{h_t}_{it} \tag{86}$$

The interpretation of which is that forecaster-expected FOMC judgment is the sum of the perceived judgment persistence times past forecaster-expected FOMC judgment, and the cumulated forecaster-expected FOMC judgment shocks that occurred in the sub-periods ($(t-1)_n$) between month $t-1$ and month t. $\varepsilon^{h_t}_{it}$ **a cumulated variant of the monetary policy shocks that are extracted from high-frequency changes in market prices**, as in Kuttner (2001) and the many improvements since then. Indeed, it is $\varepsilon^{h_t}_{it}$ that most interests us in this paper. **The perceived reaction function-implied monetary policy shocks that we will extract from the data, test, and use are**:

$$\varepsilon^{h_t}_{it} := \ddot{\mathbb{E}}_{it}\left[\ \bar{\varepsilon}^{h_t}_{t}\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \sum_{r=1}^{R}\tilde{\theta}^{r-1}\tilde{\varepsilon}^{h}_{t_{n-(2r-1)}} \ \middle|\ \tilde{\Omega}_{it} = \tilde{\omega}_{it}\ \right] \tag{87}$$

Interpretation of the Instrument. The estimated forecaster-perceived monetary policy shocks, $\varepsilon^{h_t}_{it}$, are a measurable function of the underlying, true FOMC judgment shocks, $\tilde{\varepsilon}^{h}_{t_{n-(2j-1)}}$, filtered through the forecaster's (realized) information set, $\tilde{\Omega}_{it} = \tilde{\omega}_{it}$. Specifically, the perceived shocks are the conditional expectation of a weighted sum of the true shocks, where the weights reflect the perceived persistence of judgment.

Econometrically, $\varepsilon^{h_t}_{it}$ can be interpreted as an instrument that satisfies the relevance condition by construction, as it is a non-degenerate function of the underlying FOMC FFR projection shocks. Under the maintained assumption that forecaster-perceived FOMC judgment shocks are conditionally exogenous

[49] For simplicity, we use h in this derivation instead of h_{t_n} as we will ultimately aggregate to the monthly level, t.

with respect to other determinants of private interest rate expectations, the instrument satisfies the exclusion restriction. That is, its only effect on the private interest rate forecast operates through the perceived FOMC projection shocks. As such, it recovers a projected version of the structural monetary policy shocks that are visible and influential to private forecasters.

More formally, the instrument is an *informational projection* of the structural FOMC judgment shocks onto the forecaster's sigma-algebra, $\ddot{\sigma}(\tilde{\Omega}_{it})$. This projection maps the full-information shock process —measurable with respect to a richer, higher-dimensional sigma-algebra due to its inclusion of more sources of information relevant to the structural shocks and finer temporal resolution (t_n rather than t), such as internal deliberations, confidential briefings, model outputs, and real-time data updates—onto a coarser, lower-dimensional sub-σ-algebra that reflects the forecaster's limited access to those FOMC-relevant dimensions. The projection filters out the unobservable components of the structural shocks and retains only those features that are measurable with respect to the forecaster's information set, that is, the components forecasters can potentially observe or infer and that influence FOMC participants' FFR projections.

A.3 Mapping The Expected Reaction Function to the Forecast Data

Assumption 11. ***Factor Loadings***. *The factor loadings* $\mathbf{\Gamma_i}$ *are independently and identically distributed across i, and of the common factors,* $\mathbf{f_t}$*, for all i and t, with mean* $\mathbf{\Gamma}$.

Assumption 12. ***Weakly Exogenous Regressors***. $\mathbf{x_{it}^H}$ *is a vector of weakly exogenous regressors.*

A section in the online appendix will cover why the regressors are weakly exogenous. In the meantime, we refer the reader to Carvalho, Nechio, and Tristao (2021), Sutherland (2023), and Bauer, Pflueger, and Sunderam (2024), which we elaborate on below..

First, Carvalho, Nechio, and Tristao (2021) argue in favor of OLS estimation of policy rules. The authors show analytically in a three-equation New Keynesian model that the asymptotic OLS estimation bias is proportional to the fraction of the variance of regressors due to monetary policy shocks. Yet, they argue that since there is ample evidence that such shocks explain only a small fraction of the variance of the regressors typically included in estimation of monetary policy rules (and those included in this paper), their analytical finding suggests that any endogeneity bias is likely to be small.

Second, Sutherland (2023) addresses the potential endogeneity concerns due to the simultaneity between interest rate forecasts and macroeconomic forecasts, such as inflation or growth. Yet, because monetary policy is assumed to operate with a lag, and especially when the interest rate forecast horizon actually exceeds the inflation and growth forecast horizons (as is the case with the interest rate forecasts at the 12-month forecast horizon), the paper argues that there should be very little risk of endogeneity due to this simultaneity, which applies here for the same reasons.[50]

Third, Bauer, Pflueger, and Sunderam (2024) build on Carvalho, Nechio, and Tristao (2021) by using a model that captures the potential endogeneity bias described by Carvalho, Nechio, and Tristao (2021) and show that such bias is unlikely to affect the their main results.

[50]The corollary is that there may be omitted variable bias due to omitted long-term inflation expectations, but we observe that, in the data, the (lower frequency, quarterly) longer-term inflation expectations do not vary much.

Assumption 13. ***Individual-Specific Errors***. *The individual specific errors, ε_{it}, which are forecaster-perceived monetary policy rate path shocks are independently distributed of $\mathbf{v}^{\mathbf{h}}_{\mathbf{jt}'}$ $\forall i, j, t, t'$. The vector of errors, $\boldsymbol{\varepsilon}_{\mathbf{t}} = \{\varepsilon_{1t}, \varepsilon_{2t}, \ldots, \varepsilon_{Nt}\}$ is spatially correlated.*

Lemma A.1.

First, consider the term of interest, $\Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right]$, in levels, $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right]$, expand using Eq. 8:[51]

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] = \rho_{it}^h\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-2+h}\right]\right] + \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\left(\mathbf{x_{t-1+H}} - \mathbf{x^*_{t-1+H}}\right)'\right]\right]\boldsymbol{\ddot{\beta}^h_{it}} + \phi_{it}^{r^*}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[r^*_{t-1+h}\right]\right] + \phi_{it}^{\pi^*}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\pi^*_{t-1+h}\right]\right] + \lambda_{it}^h \quad (88)$$

Applying Eq. 13, 14, and 15:

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] = \rho_{it}^h\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-2+h}\right]\right] + \left(\mathbf{A^H_{it}x^H_{it}} + \boldsymbol{\upsilon}^{\mathbf{H}}_{\mathbf{it}}\right)'\boldsymbol{\ddot{\beta}^h_{it}} + \phi_{it}^{r^*}r_t^* + \phi_{it}^{\pi^*}\pi_t^* + \lambda_{it}^h \quad (89)$$

Now, to match the equation of interest, take first differences:

$$\begin{aligned}\Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] &= \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] - \ddot{\mathbb{E}}_{i,t-1}\left[\tilde{\mathbb{E}}_{t-1}\left[\tilde{r}_{t-1+h}\right]\right] \\ &= \left(\rho_{it}^h\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-2+h}\right]\right] + \left(\mathbf{A^H_{it}x^H_{it}} + \boldsymbol{\upsilon}^{\mathbf{H}}_{\mathbf{it}}\right)'\boldsymbol{\ddot{\beta}^h_{it}} + \phi_{it}^{r^*}r_t^* + \phi_{it}^{\pi^*}\pi_t^* + \lambda_{it}^h\right) \\ &- \left(\rho_{i,t-1}^h\ddot{\mathbb{E}}_{i,t-1}\left[\tilde{\mathbb{E}}_{t-1}\left[\tilde{r}_{t-2+h}\right]\right] + \left(\mathbf{A^H_{i,t-1}x^H_{i,t-1}} + \boldsymbol{\upsilon}^{\mathbf{H}}_{\mathbf{i,t-1}}\right)'\boldsymbol{\ddot{\beta}^h_{i,t-1}} + \phi_{i,t-1}^{r^*}r_{t-1}^* + \phi_{i,t-1}^{\pi^*}\pi_{t-1}^* + \lambda_{i,t-1}^h\right)\end{aligned} \quad (90)$$

To simplify the derivation, define:

$$R_{it}^{h,0} := \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] \quad (91)$$

And, more generally,

$$R_{it}^{h,\ell} := \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h-\ell}\right]\right] \qquad \ell = 0, 1, 2, \ldots \quad (92)$$

So, $R_{it}^{h,0}$ is the policy rate expectation entering the equation of interest, while $R_{it}^{h,1}$ is the corresponding one-period earlier lagged policy rate expectation. Next, define the forecast-revision operator as

$$\Delta R_{it}^{h,\ell} := R_{it}^{h,\ell} - R_{i,t-1}^{h,\ell} \quad (93)$$

Now, with this notation, the differenced equation above becomes:

$$\Delta R_{it}^{h,0} = \rho_{it}^h\Delta R_{it}^{h,1} + \Delta\rho_{it}^h R_{i,t-1}^{h,1} + \mathcal{S}_{it}^h \quad (94)$$

$$\Delta\rho_{it}^h := \rho_{it}^h - \rho_{i,t-1}^h \quad (95)$$

$$\mathcal{S}_{it}^h := \left[\left(\mathbf{A^H_{it}x^H_{it}} + \boldsymbol{\upsilon}^{\mathbf{H}}_{\mathbf{it}}\right)'\boldsymbol{\ddot{\beta}^h_{it}} - \left(\mathbf{A^H_{i,t-1}x^H_{i,t-1}} + \boldsymbol{\upsilon}^{\mathbf{H}}_{\mathbf{i,t-1}}\right)'\boldsymbol{\ddot{\beta}^h_{i,t-1}}\right] + \left(\phi_{it}^{r^*}r_t^* - \phi_{i,t-1}^{r^*}r_{t-1}^*\right) + \left(\phi_{it}^{\pi^*}\pi_t^* - \phi_{i,t-1}^{\pi^*}\pi_{t-1}^*\right) + \Delta\lambda_{it}^h \quad (96)$$

If we apply the same structure one lag back, we obtain:

$$\Delta R_{it}^{h,1} = \rho_{it}^h\Delta R_{it}^{h,2} + \Delta\rho_{it}^h R_{i,t-1}^{h,2} + \mathcal{S}_{it}^{h,1} \quad (97)$$

where $\mathcal{S}_{it}^{h,1}$ is defined analogously for the forecast target date $t-2+h$. Substituting Eq. 97 into Eq. 94 gives:

$$\Delta R_{it}^{h,0} = \rho_{it}^h\left[\rho_{it}^h\Delta R_{it}^{h,2} + \Delta\rho_{it}^h R_{i,t-1}^{h,2} + \mathcal{S}_{it}^{h,1}\right] + \Delta\rho_{it}^h R_{i,t-1}^{h,1} + \mathcal{S}_{it}^h \quad (98)$$

Proceeding recursively J steps yields:[52]

$$\Delta R_{it}^{h,0} = \sum_{j=0}^{J}\left(\prod_{k=0}^{j-1}\rho_{it}^h\right)\mathcal{S}_{it}^{h,j} + \sum_{j=0}^{J}\left(\prod_{k=0}^{j-1}\rho_{it}^h\right)\Delta\rho_{it}^h R_{i,t-1}^{h,j+1} + \left(\prod_{k=0}^{J}\rho_{it}^h\right)\Delta R_{it}^{h,J+1} \quad (99)$$

[51] Throughout, the h superscript indexes the forecast horizon associated with the expectation formed at the indexed date (t, $t-1$, etc.). So, an expectation formed at date t uses parameters indexed by (i, t, h), while an expectation formed at date $t-1$ uses parameters indexed by $(i, t-1, h)$.

[52] When j=0, the product is empty and equals one.

Where $\mathcal{S}_{it}^{h,0} \equiv \mathcal{S}_{it}^{h}$, and $\mathcal{S}_{it}^{h,j}$ denotes the corresponding non-lagged block for the forecast target date $t-1+h-j$.

Next, to connect this recursive representation to the common-factor structure, substitute Eq. 19 into the non-lagged block. For the leading target date, this gives:

$$\mathcal{S}_{it}^{h} = \left[\left(\mathbf{A}_{\mathbf{it}}^{\mathbf{H}}\left[\mathbf{c}_{\mathbf{xi}} + \mathbf{\Gamma}_{\mathbf{it}}^{\mathbf{H}}\mathbf{f}_{\mathbf{t}} + \mathbf{v}_{\mathbf{it}}^{\mathbf{H}}\right] + \upsilon_{\mathbf{it}}^{\mathbf{H}}\right)' \ddot{\beta}_{\mathbf{it}}^{\mathbf{h}} - \left(\mathbf{A}_{\mathbf{i,t-1}}^{\mathbf{H}}\left[\mathbf{c}_{\mathbf{xi}} + \mathbf{\Gamma}_{\mathbf{i,t-1}}^{\mathbf{H}}\mathbf{f}_{\mathbf{t-1}} + \mathbf{v}_{\mathbf{i,t-1}}^{\mathbf{H}}\right] + \upsilon_{\mathbf{i,t-1}}^{\mathbf{H}}\right)' \ddot{\beta}_{\mathbf{i,t-1}}^{\mathbf{h}}\right] + \left(\phi_{it}^{r^*} r_t^* - \phi_{i,t-1}^{r^*} r_{t-1}^*\right) + \left(\phi_{it}^{\pi^*} \pi_t^* - \phi_{i,t-1}^{\pi^*} \pi_{t-1}^*\right) + \Delta\lambda_{it}^{h} \quad (100)$$

Analogous expressions apply for $\mathcal{S}_{it}^{h,j}$ at the recursively lagged forecast dates. These terms collect the effects of changes in the perceived macro state, changes in the mapping from macro factors to forecast-relevant variables, changes in coefficients, and forecaster-specific residual components. So, the recursive expression above can be written compactly as a factor-based representation.

We can now write the expression more compactly in order to substitute Eq. 101 below back into Eq. 23:

$$\Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{r}_{t-1+h}\right]\right] = \Delta R_{it}^{h,0} = \varphi_{it}^{r^*} r_t^* + \varphi_{it}^{\pi^*} \pi_t^* + \underbrace{\mathbf{\Psi}_{it}^{h}\mathbf{f}_t + \check{\mathbf{\Psi}}_{it}^{h}\Delta\mathbf{f}_t}_{\text{Macro shocks}} + \underbrace{\Xi_{it}^{h}}_{\substack{\text{Composite} \\ \text{error}}} \quad (101)$$

Interpretation.

- $\mathbf{\Psi}_{it}^{h}$ collects the persistence-weighted loadings on the common macro factors $\mathbf{f}_t$ arising from the factor representation of $\mathbf{x}_{it}^{H}$ and from the mapping between the forecast-relevant macro state and the expected policy-rate path.
- $\check{\mathbf{\Psi}}_{it}^{h}$ collects the persistence-weighted loadings on common macro factor innovations $\Delta\mathbf{f}_t$. These terms arise from changes in the common macro factors and from local changes in the mapping from factors to forecast-relevant variables.
- $\varphi_{it}^{r^*}$ and $\varphi_{it}^{\pi^*}$ collect the persistence-weighted loadings on r_t^* and π_t^*, respectively.
- Ξ_{it}^{h} is a composite disturbance collecting the remaining components of the expansion. These include forecaster-specific deviations, judgment adjustments, changes in persistence parameters, coefficient-drift terms, recursively lagged terminal terms, and other residual components not captured by the leading macro-factor terms.

The components entering Ξ_{it}^{h} are either idiosyncratic or involve changes in quantities that are assumed to evolve smoothly over time:

1. Persistence parameters $\rho_{it}^{h} \in (0, 1)$ are bounded.
2. Changes in coefficients, including $\Delta\rho_{it}^{h}$, $\Delta\beta_{it}^{h}$, and $\Delta\phi_{it}$, are assumed to be small under smooth parameter evolution.
3. Forecaster-specific deviations, judgment adjustments, and terminal-recursion components are not systematically related to the common macro factors.

Consequently, under bounded persistence and smooth parameter evolution, the common macro-factor terms and anchor terms are expected to account for the dominant first-order variation in forecast revisions. The remaining higher-order, idiosyncratic, and terminal-recursion components are collected in Ξ_{it}^{h}. This simplification allows the expression to be substituted back into the main equations without introducing unnecessary algebraic clutter, while preserving the first-order macro drivers of forecast revisions.

A.4 Why can the within-window residuals be serially uncorrelated while the final shock estimate is serially correlated?

Let $w_t = [t - W + 1, t]$ denote the rolling window ending at date t. Suppose that, after quasi-differencing and the DCCE transformation, the ideal estimating equation is:

$$q_{it} = \mathbf{X}_{it}^{'}\boldsymbol{\beta}_i + \varepsilon_{it}, \qquad \mathbb{E}_{t-1}\left[\varepsilon_{it}\right] = 0 \tag{102}$$

Within a given rolling window, the transformed residual is:

$$\widehat{\varepsilon}_{it}^{[w_t]} = q_{it} - \mathbf{X}_{it}^{'}\widehat{\boldsymbol{\beta}}_i^{[w_t]} \tag{103}$$

Adding and subtracting the true coefficient vector gives:

$$\widehat{\varepsilon}_{it}^{[w_t]} = \varepsilon_{it} - \mathbf{X}_{it}^{'}\left(\widehat{\boldsymbol{\beta}}_i^{[w_t]} - \boldsymbol{\beta}_i\right) \tag{104}$$

Define the generated estimation component:

$$m_{it}^{[w_t]} := -\mathbf{X}_{it}^{'}\left(\widehat{\boldsymbol{\beta}}_i^{[w_t]} - \boldsymbol{\beta}_i\right) \tag{105}$$

So that:

$$\widehat{\varepsilon}_{it}^{[w_t]} = \varepsilon_{it} + m_{it}^{[w_t]} \tag{106}$$

Even if the primitive innovations satisfy:

$$\mathrm{Cov}\left(\varepsilon_{it}, \varepsilon_{i,t-1}\right) = 0 \tag{107}$$

The rolling-window residuals need not satisfy:

$$\mathrm{Cov}\left(\widehat{\varepsilon}_{it}^{[w_t]}, \widehat{\varepsilon}_{i,t-1}^{[w_{t-1}]}\right) = 0 \tag{108}$$

This is because adjacent rolling windows overlap heavily. For example, with an 84-month rolling window,

$$w_t = \{t - 83, \ldots, t\}, \qquad w_{t-1} = \{t - 84, \ldots, t - 1\} \tag{109}$$

The two estimation samples therefore share 83 of 84 observations. And so:

$$\widehat{\boldsymbol{\beta}}_i^{[w_t]} \approx \widehat{\boldsymbol{\beta}}_i^{[w_{t-1}]} \tag{110}$$

Which implies that the corresponding estimation-error components $m_{it}^{[w_t]}$ and $m_{i,t-1}^{[w_{t-1}]}$ are highly correlated.

So, a sequence of terminal rolling-window residuals can exhibit substantial serial correlation even when the underlying structural innovations, ε_{it}, are serially uncorrelated. The source of persistence can arise mechanically from the overlapping-window estimation procedure rather than from persistence in the underlying economic shock itself.