Two Sides of the Same Coin: Expected Reaction Functions and Perceived Monetary Policy 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 effect 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 perceived monetary policy expectations shock measures, which we show to be both relevant and exogenous instruments. Finally, using our two instruments, we demonstrate that the impulse response function estimates are consistent with New Keynesian theory.

January 2, 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, csutherland@fas.harvard.edu.
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, 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, Luba Petersen, Carolin Pflueger, Giovanni Ricco, Neil Shephard, James Stock, Adi Sunderam, Jon Witmer, and Xu Zhang. Thank you also to seminar participants at The BSE Advances in Structural Shocks Identification Workshop, 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. Perhaps the current state-of-the-art measures were introduced by Miranda-Agrippino and Ricco (2021) and Bauer and Swanson (2023).

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 Bianchi, Ludvigson, and Ma (2025) for related work). 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. Further, to the extent that narrative measures focus on the federal funds rate, they may not be an ideal summary statistic for the stance of monetary policy, which is better captured by a term structure of interest rate expectations, not just one point.

So, this paper proposes a complementary approach (ideally we would utilize all three approaches) to estimating monetary policy shock measures that attempts to balance relevance and exogeneity. We develop a framework designed for the structural estimation of the perceived FOMC reaction functions implied by monetary policy expectations. The framework allows us to isolate separate information, macroeconomic response, forward guidance, smoothing/gradualism, and persistent judgment channels (where smoothing/gradualism and persistent judgment each contribute to monetary policy *inertia*). Private forecasts are likely to have been constructed using simple rules and heuristics (unlike the projections formed by central banks), which makes disentangling these channels more straightforward. We then extract the residuals from the expected reaction function 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, perceived FOMC reaction functions implied by monetary policy expectations. 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*." This paper proposes strategies to manage these measurement and specification errors. First, our proposed central bank reaction function framework (§A) considers all forms of heterogeneity: the reaction function parameters vary by forecaster, by country, by forecast horizon, and by time. As in Bauer, Pflueger, and Sunderam (2024), we allow for time-varying parameters

using rolling windows. Then, we go further by estimating forecaster-expected, forward-looking reaction functions forecaster-by-forecaster and horizon-by-horizon. We also define gradualism/smoothing and central bank judgment channels, which allows us to both consider concepts such as forward guidance, and to account for the persistent nature of the residuals.

Based on this framework, we propose an econometric model (§2) that is designed to eliminate the many potential sources of estimation bias outlined in the reaction function framework: common latent factors, ensuing cross-sectional (forecaster) dependence, and failure to account for parameter heterogeneity. For instance, using Pesaran (2015) test statistics, we detect strong dependence between forecasters, precluding independent and identically distributed error terms. Failure to account for cross-sectional dependence causes the error term to be autocorrelated and the expected reaction function coefficient estimates to be biased. So, we utilize the Chudik and Pesaran (2015) dynamic panel estimator, which accounts for cross-sectional dependence. We also parse the data by forecaster, country, time, and forecast horizon to account for extensive parameter heterogeneity. We argue that our econometric 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. 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 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, unlike the high-frequency or narrative monetary shock measures available to date, at least, our shock measures are not weak instruments.

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, we argue that macroeconometricians should use Lewis and Mertens (2026) / Montiel Olea and Pflueger (2013) weak instruments tests when using monetary policy instruments. 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. We also do not reject the hypothesis that our instruments are exogenous. Interestingly, while the residuals from the Bauer, Pflueger, and Sunderam (2024) regressions cannot be treated as shocks a priori, we find evidence that they are nonetheless a valid monetary policy instrument.

Finally, we use our monetary policy instruments to estimate impulse response functions based on the approach of Gertler and Karadi (2015). Our estimates are broadly consistent with New Keynesian theory, as are those formed using the Aruoba and Drechsel (2025) shocks, and interestingly, as are those formed using Bauer, Pflueger, and Sunderam (2024) residuals.

The paper is structured as follows. Section (§) 2 presents the econometric model. §3 presents expected reaction function and perceived shock estimates. §4 discusses instrument validity, and compares impulse response functions to those from the literature. §5 concludes. §A presents 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 very general specification of a *Federal Open Market Committee* (FOMC) reaction function:

$$\underbrace{\tilde{p}_t}_{\substack{\text{Fed Funds}\\ \text{Rate}}} = \underbrace{\tilde{\varphi}_t \tilde{p}_{t-\ell}}_{\substack{\text{Monetary Policy}\\ \text{Gradualism}}} + \underbrace{\tilde{\phi}_t^{\pi^*} \pi_t^*}_{\substack{\text{Inflation}\\ \text{Target}}} + \underbrace{\tilde{\phi}_t^{r^*} r_t^*}_{\text{r-star}} + \underbrace{\left(\tilde{\mathbb{E}}_{t^-}\left[\mathbf{x_{t+h_x}} - \mathbf{x^*_{t+h_x}} \;\middle|\; \tilde{\Omega}_{t^-}\right]\right)'}_{\substack{\text{Macroeconomic}\\ \text{Projections}}} \underbrace{\tilde{\boldsymbol{\beta}}_{\mathbf{t}}}_{\substack{\text{Macro}\\ \text{Response}}} + \underbrace{\tilde{\varepsilon}_t}_{\substack{\text{Monetary Policy}\\ \text{Shock}}} \tag{1}$$

Notation. $\sim$ denotes objects attributable to The FOMC. t denotes a time period of arbitrary length. ℓ denotes lag length. $*$ denotes long-run equilibrium variables. t^- is some time before the monetary policy shock. t is when the shock is realized. t^+ is some time 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 $\tilde{\varphi}_t \tilde{p}_{t-\ell}$, s.t. $|\tilde{\varphi}_t| < 1$. As such, $\tilde{\phi}_t^{\pi^*}, \tilde{\phi}_t^{r^*}, \tilde{\boldsymbol{\beta}}_{\mathbf{t}} \propto (1 - \tilde{\varphi})$, so estimates of $\tilde{\phi}_t^{\pi^*}$, $\tilde{\phi}_t^{r^*}$, and $\tilde{\boldsymbol{\beta}}_{\mathbf{t}}$ are implicitly scaled by $(1 - \tilde{\varphi}_t)$.

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. The usual assumption of constant parameters is too restrictive, so unlike in a standard specification, we remove this constraint, allowing for a more general reaction function. $\tilde{\varepsilon}_t$ 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

[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.

periods, and taking the conditional expectation at time t:

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

where $\tilde{\mathbb{E}}_t\left[\tilde{\varepsilon}_{t+h} \;\middle|\; \tilde{\Omega}_t\right] = 0$, and so drops out, and where $h' = h_x + h$. Finally, observe that FOMC members are free to add judgment to their FFR projections:

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

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 takes the following form:

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

$$\tilde{\lambda}^h_t = \tilde{\rho}_t\tilde{\lambda}^h_{t-\ell} + \tilde{\varepsilon}^h_t \qquad |\tilde{\rho}_t| < 1 \qquad \tilde{\varepsilon}^h_t \sim i.i.d.\left(0, \sigma^2_{\tilde{\varepsilon}_t}\right) \tag{4}$$

where $\tilde{\varepsilon}^h_t$ is a shock to judgment, and can also be interpreted as a **monetary policy expectations shock**. For our purposes, ℓ will be set relative to t such that Eq. 4 is an AR(1) process. 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:

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

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. **Risk management** due to elevated projection uncertainty (e.g. anticipated insurance rate cuts due to asymmetric risks to the outlook).
2. Past and/or new **forward guidance**.[4]
3. **Model uncertainty** (e.g. about the monetary policy rule itself, or the r-star , etc.).

[3] Campbell et al. (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).

4. **Rounding effects** (the FFR tends to change in (sometimes multiple) 25-basis-point increments, but the output of a reaction function will not tend to come in 25-basis point increments).

5. **Dissent** to make a point (to the market and/or to other FOMC participants).

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). $\tilde{\lambda}^h_t$ 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, to derive the *expected* reaction function, we take expectations of Eq. 3 with respect to the information set of forecaster i at time t.

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

where $\ddot{\mathbb{E}}_{it}\left[. \mid \ddot{\Omega}_{it}\right]$ is forecaster i's ($i = 1, ..., N$) expectation conditional on their information set, $\ddot{\Omega}_{it}$. Double dots denote objects attributable to forecasters. We assume that both the FOMC and forecasters know their own expected coefficients ($\tilde{\Omega}_{t^+}$–measurable and $\ddot{\Omega}_{it^+}$–measurable, respectively—Key Assumption 6 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, $\ddot{\varphi}^h_{it}$. The analogous definitions apply for all other coefficients in Eq. 7.

We denote $\ddot{\mathbb{E}}_{it^+}\left[\tilde{\lambda}^h_{t^+}\right] := \ddot{\lambda}^h_{it^+} = \ddot{\rho}_{it^+}\ddot{\lambda}^h_{i,t^+-\ell} + \ddot{\varepsilon}^h_{it^+}$. 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 63 to 73). 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]

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h}\right]\right] = \ddot{\varphi}^h_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-\ell}\right]\right] + \ddot{\phi}^{\pi^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\pi^*_{t+h}\right]\right] + \ddot{\phi}^{r^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[r^*_{t+h}\right]\right] + \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\left(\mathbf{x}_{\mathbf{t+h}'} - \mathbf{x}^*_{\mathbf{t+h}'}\right)'\right]\right]\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}} + \ddot{\lambda}^h_{i,t} \quad (7)$$

2.3 Mapping The Expected Reaction Function to the Data

Third, we take the forecaster-expected reaction function (Eq. 7), 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 $\ell = 1$. For now, we use the *Consensus Economics* data (in this draft), but the *Blue Chip Financial Forecasts* data have similar attributes, and may also be suited to our framework. We begin by defining the relationship between an observable-yet-noisy measure of the dependent variable, $\ddot{y}^h_{it}$, and the

[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, The FOMC knows its own macro responsiveness), then it follows that $\ddot{\mathbb{E}}_{it^+}\left[\tilde{\mathbb{E}}_t\left[\left(\tilde{\mathbb{E}}_{t^-+h}\left[\mathbf{x}_{\mathbf{t+h}'} - \mathbf{x}^*_{\mathbf{t+h}'}\right]\right)'\tilde{\beta}_{t+h}\right]\right] = \ddot{\mathbb{E}}_{it^+}\left[\left(\tilde{\mathbb{E}}_t\left[\mathbf{x}_{\mathbf{t+h}'} - \mathbf{x}^*_{\mathbf{t+h}'}\right]\right)'\tilde{\beta}_{t+h}\right]$.

unobservable dependent variable from above, $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h}\right]\right]$. $\ddot{y}^h_{it}$ 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.

$$\underbrace{\ddot{y}^h_{it}}_{\substack{\text{Private Interest}\\ \text{Rate Forecast}}} = \underbrace{\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h}\right]\right]}_{\substack{\text{Expected FOMC}\\ \text{FFR Projection}}} + \underbrace{\ddot{\mathbb{E}}_{it}\left[f^e_{t+h}\right]}_{\substack{\text{Expected FOMC}\\ \text{FFR Projection Error}}} \tag{8}$$

$\ddot{\mathbb{E}}_{it}\left[f^e_{t+h}\right]$ is an expected FFR projection error term, which could also include a small term premium expectation, although evidence suggests that this component is small. If $\ddot{\mathbb{E}}_{it}\left[f^e_{t+h}\right] \neq 0$, the forecaster thinks the FOMC is wrong about future policy rates (see Caballero and Simsek (2022)). Next, we substitute the right-hand side of Eq. 7 for $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h}\right]\right]$ into Eq. 8. To simplify notation, we assume one lagged dependent variable (for now, but later two), [and simplify the time-dependent forecast horizon notation from h_t to just h].

$$\ddot{y}^h_{it} = \underbrace{\ddot{\varphi}^h_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right] + \ddot{\phi}^{r^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[r^*_{t+h}\right]\right] + \ddot{\phi}^{\pi^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\pi^*_{t+h}\right]\right] + \ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t^-+h}\left[\left(\mathbf{x_{t+h}} - \mathbf{x^*_t}\right)'\right]\right]\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}} + \ddot{\lambda}^h_{i,t}}_{\text{Expected FOMC FFR Projection}} + \ddot{\mathbb{E}}_{it}\left[f^e_{t+h}\right] \tag{9}$$

Now we continue to map the expected reaction function to the data. We focus on four main subcomponents, and explain each in turn. We also relabel some terms in an ongoing effort to simplify notation where possible as we progress.

$$\ddot{y}^h_{it} = \underbrace{\ddot{\varphi}^h_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right]}_{\ddot{\varphi}^h_{it}\ddot{y}^h_{i,t-1} + \ddot{\varphi}^h_{it}\left(\ddot{\gamma}^{\mathbf{\Lambda}}_{\mathbf{it}}\right)'\mathbf{f_t} - \ddot{\varphi}^h_{it}\ddot{\gamma}^e_{i,t-1}f^e_{t-1}} + \underbrace{\ddot{\phi}^{r^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[r^*_{t+h}\right]\right] + \ddot{\phi}^{\pi^*}_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\pi^*_{t+h}\right]\right]}_{\ddot{\phi}^{\pi^*}_{it}\pi^*_t + \ddot{\phi}^{r^*}_{it}r^*_t} + \underbrace{\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_{t^-+h}\left[\left(\mathbf{x_{t+h}} - \mathbf{x^*_t}\right)'\right]\right]\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}}}_{\left(\ddot{\mathbf{x}}^{\mathbf{h}}_{\mathbf{it}}\right)'\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}}} + \ddot{\lambda}^h_{it} + \underbrace{\ddot{\mathbb{E}}_{it}\left[f^e_{t+h}\right]}_{\ddot{\gamma}^e_{it}f^e_t} \tag{10}$$

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, a m x 1 vector of latent common factors, which contemporaneously orthogonal and possibly serially correlated. $\mathbf{f_t}$ 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.[6] Accordingly, we postulate that we can approximate a number of the elements of the expected reaction functions using time-varying dynamic factor models.

First, we approximate $\ddot{\mathbb{E}}_{it}\left[f^e_{t+h}\right]$ with the product of a common, time-varying factor, f^e_t, and a forecaster-specific loading, $\ddot{\gamma}^e_{it}$.

Second, consider the expected (forward-looking) monetary policy gradualism term, $\ddot{\varphi}^h_{it}\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right]$. Observe that we can shift Eq. 8 back one period:

$$\ddot{y}^h_{i,t-1} = \ddot{\mathbb{E}}_{i,t-1}\left[\tilde{\mathbb{E}}_{t-1}\left[\tilde{p}_{t-1+h}\right]\right] + \ddot{\mathbb{E}}_{i,t-1}\left[f^e_{t+h-1}\right] \Rightarrow \tag{11}$$

Disaggregating the first term on the right-hand-side:

$$\ddot{y}^h_{i,t-1} = \underbrace{\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right] - \Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right]}_{\ddot{\mathbb{E}}_{i,t-1}\left[\tilde{\mathbb{E}}_{t-1}\left[\tilde{p}_{t+h-1}\right]\right]} + \ddot{\mathbb{E}}_{i,t-1}\left[f^e_{t+h-1}\right] \Rightarrow \tag{12}$$

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

Rearranging and approximating two of the terms as forecaster-weighted common factors:

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right] = \ddot{y}^h_{i,t-1} + \underbrace{\Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right]}_{\approx \left(\ddot{\gamma}^{\mathbf{\Lambda}}_{\mathbf{it}}\right)' \mathbf{f_t}} - \underbrace{\ddot{\mathbb{E}}_{i,t-1}\left[f^e_{t+h-1}\right]}_{\approx \ddot{\gamma}^e_{i,t-1} f^e_{t-1}} \tag{13}$$

In the first case, again following Giacomini, Lu, and Smetanina (2024), we postulate that $\Delta\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t+h-1}\right]\right]$ can be approximated by $\left(\ddot{\gamma}^{\mathbf{\Lambda}}_{\mathbf{it}}\right)' \mathbf{f_t}$, The vector $\mathbf{f_t}$ is composed of three structural shocks: $\mathbf{f_t}' = \left[f^s_t, f^d_t, f^e_t\right]$ where f^s_t is a supply factor, f^d_t is a demand factor, and f^e_t is the anticipated FOMC projection error factor. $\left(\ddot{\gamma}^{\mathbf{\Lambda}}_{\mathbf{it}}\right)'$ is a vector of time-varying, forecaster-specific factor loadings. The second case is simply $\ddot{\gamma}^e_{it} f^e_t$ from Eq. 10 lagged one period.

Third, we postulate that the vector of macroeconomic forecasts, $\ddot{\mathbf{x}}^{\mathbf{h}}_{\mathbf{it}}$ is generated by the following linear covariance stationary dynamic heterogeneous panel data model:

$$\ddot{\mathbf{x}}^{\mathbf{h}}_{\mathbf{it}} = \mathbf{c}_{\omega i} + \alpha_{\mathbf{i}} \ddot{y}^h_{i,t-1} + \ddot{\mathbf{\Gamma}}'_{\mathbf{it}} \mathbf{f_t} + \ddot{\mathbf{v}}^{\mathbf{h}}_{\mathbf{it}} \tag{14}$$

where $\ddot{\mathbf{v}}^{\mathbf{h}}_{\mathbf{it}}$ follows a general linear covariance stationary process distributed independently of the idiosyncratic errors, $\ddot{\varepsilon}_{it}$. $\mathbf{c}_{\omega i}$ is an individual fixed effect. $\ddot{\mathbf{\Gamma}}_{\mathbf{it}}$ is an $m\ x\ k$ matrix of factor loadings.

Now, the elements of $\ddot{\mathbf{x}}^{\mathbf{h}}_{\mathbf{it}}$—the private inflation, growth, and unemployment forecasts—can also be rewritten in the same manner as Eq. 8, that is, by including a forecaster-expected FOMC projection error component. This means that they may also be influenced by f^e_t. And so, all the macroeconomic forecasts can include the expected FOMC projection error factor, f^e_t, as well. That is, the FOMC is perceived as being wrong for one principal reason (e.g. inflation is not transitory, but persistent), which transmits through all forecasts linearly. Note that a fourth structural shock, a forecaster-expected judgment shock, $\ddot{\varepsilon}_{it}$, which is a function of the FOMC-expected judgment shock, $\tilde{\varepsilon}_{it}$, also enters the system through $\ddot{y}^h_{it}$ and $\ddot{y}^h_{i,t-1}$.

Fourth, given that the central bank regularly publishes its r^*_t estimates, let us assume that:

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

where $\tilde{\mathbb{E}}_t\left[r^*_{t+h}\right]$ could be the last-published estimates from The SEP, which is what use do for long-run equilibrium values of real growth and unemployment. Yet, our data frequency is monthly, and the SEP data are only available from 2012. So, instead, we use the updated Laubach and Williams (2003) r-star estimates from Holston, Laubach, and Williams (2023). This is of course observable, and so can be substituted directly into Eq. 17. Similarly, we can observe that:

$$\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\pi^*_{t+h}\right]\right] = \tilde{\mathbb{E}}_t\left[\pi^*_{t+h}\right] := \pi^*_t \tag{16}$$

the prevailing inflation target, which is observable, and can also be substituted directly into Eq. 17. Consistent with the literature on monetary policy rules, $\ddot{\phi}^{\pi^*}_{it} \propto \left(1 - \sum_{\ell=1}^{L} \ddot{\varphi}^{\ell}_{it}\right)$, $\ddot{\phi}^{r^*}_{it} \propto \left(1 - \sum_{\ell=1}^{L} \ddot{\varphi}^{\ell}_{it}\right)$.

Putting these four components together (we substitute the right-hand-side of Eq. 13 into $\ddot{\mathbb{E}}_{it}\left[\tilde{\mathbb{E}}_t\left[\tilde{p}_{t-1+h}\right]\right]$ in Eq. 10—consistent with its labels), and re-ordering terms:

$$\ddot{y}^h_{it} = \underbrace{\ddot{\varphi}^h_{it} \ddot{y}^h_{i,t-1} + \ddot{\phi}^{\pi^*}_{it} \pi^*_t + \ddot{\phi}^{r^*}_{it} r^*_t + \left(\ddot{\mathbf{x}}^{\mathbf{h}}_{\mathbf{it}}\right)' \ddot{\beta}^{\mathbf{h}}_{\mathbf{it}}}_{\text{Observable}} + \underbrace{\ddot{\varphi}^h_{it}\left(\ddot{\gamma}^{\mathbf{\Lambda}}_{\mathbf{it}}\right)' \mathbf{f_t} + \ddot{\gamma}^e_{it} f^e_t - \ddot{\varphi}^h_{it} \ddot{\gamma}^e_{i,t-1} f^e_{t-1}}_{\text{Latent common factors} := \sum_{\ell=0}^{L=1} \left(\ddot{\gamma}_{\mathbf{i,t-}\ell}\right)' \mathbf{f_{t-}\ell}} + \ddot{\lambda}^h_{it} \tag{17}$$

Finally, we can observe that the three terms, $\ddot{\varphi}_{it}^{h}\left(\ddot{\gamma}_{\mathbf{it}}^{\mathbf{\Lambda}}\right)^{'}\mathbf{f_t} + \ddot{\gamma}_{it}^{e}f_t^e - \ddot{\varphi}_{it}^{h}\ddot{\gamma}_{i,t-1}^{e}f_{t-1}^{e}$, are simply a linear combination of the three factors in the vectors $\mathbf{f_t}$, and its lag, $\mathbf{f_{t-1}}$. As such, they can be rewritten as $\sum_{\ell=0}^{L=1}\left(\ddot{\gamma}_{\mathbf{i,t-\ell}}\right)^{'}\mathbf{f_{t-\ell}}$, where $\ddot{\gamma}_{\mathbf{i,t-\ell}}$ is a vector of aggregated factor loadings. The inclusion of the lagged factor, $\mathbf{f_{t-1}}$, also has the added benefit of accounting for any unmodeled persistence in the factors.

2.4 Identification Strategy

We start from Eq. 17, and we can now abstract from the double-dot notation, which is now redundant as all variables and parameters with an i subscript pertain to individual forecasters.

$$y_{it} = \varphi_{it}y_{i,t-1} + \phi_{it}^{\pi^*}\pi_t^* + \phi_{it}^{r^*}r_t^* + \left(\mathbf{x_{it}}\right)^{'}\beta_{\mathbf{it}} + \sum_{\ell=0}^{L=1}\left(\gamma_{\mathbf{i,t-\ell}}\right)^{'}\mathbf{f_{t-\ell}} + \lambda_{it} \tag{18}$$

Assumption 2. *The parameters of interest follow a random coefficient model:*

$$\varphi_{\mathbf{it}} = \varphi + \eta_{\varphi\mathbf{i}} + \nu_{\varphi\mathbf{t}}, \quad \eta_{\varphi\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_{\eta_\varphi}\right), \quad \nu_{\varphi\mathbf{t}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_{\nu_\varphi}\right) \tag{19}$$

$$\beta_{\mathbf{it}} = \beta + \eta_{\beta\mathbf{i}} + \nu_{\beta\mathbf{t}}, \quad \eta_{\beta\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\eta_\beta}\right), \quad \nu_{\beta\mathbf{t}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\nu_\beta}\right) \tag{20}$$

$$\gamma_{\mathbf{it}} = \gamma + \eta_{\gamma\mathbf{i}} + \nu_{\gamma\mathbf{t}}, \quad \eta_{\gamma\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\eta_\gamma}\right), \quad \nu_{\gamma\mathbf{t}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_{\nu_\gamma}\right) \tag{21}$$

$$\rho_{\mathbf{it}} = \rho + \eta_{\rho\mathbf{i}} + \nu_{\rho\mathbf{t}}, \quad \eta_{\rho\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_{\eta_\rho}\right), \quad \nu_{\rho\mathbf{t}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_{\nu_\rho}\right) \tag{22}$$

where each parameter fluctuates around a global mean. This helps clarify that, in the model, there are **two main sources of perceived monetary policy shocks**: shocks to the reaction function **parameters**—the η_i and ν_t terms, and shocks to FOMC **judgment** (λ_{it})—the ε_{it} terms.

Identification Challenge 1. Overparameterization. This generalized forecaster-expected reaction function is plausibly an accurate representation of the true expected reaction function, but it is necessarily overparamaterized because parameters almost certainly vary by country (c), forecaster (i), time (t), and forecast horizon (h). 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. 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 average 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 1 and 2).

Identification Challenge 3. Forecast Horizon (h) Heterogeneity. We opt not to pool data across h. As such, we can abstract from the h superscript in Eq. 18 because we will estimating the model separately for each h. Arguably, imposing parameter homogeneity across forecast horizons is an 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 (§3.6), 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. 41), which allows us to test the exogeneity of the instruments, 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.[7] Hence, we can further simplify Eq. 18 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:

$$y_{it} = \varphi_i y_{i,t-1} + \phi^{\pi^*} \pi_t^* + \phi^{r^*} r_t^* + (\mathbf{x_{it}})' \beta_\mathbf{i} + \sum_{\ell=0}^{L=1} (\gamma_\mathbf{i})' \mathbf{f_{t-\ell}} + \lambda_{it} \tag{23}$$

Accordingly, the heterogeneous coefficients now follow a simpler random coefficients model within each window:[8]

$$\varphi_\mathbf{i} = \varphi + \eta_{\varphi\mathbf{i}}, \quad \eta_{\varphi\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_\varphi\right) \tag{24}$$

$$\beta_\mathbf{i} = \beta + \eta_{\beta\mathbf{i}}, \quad \eta_{\beta\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_\beta\right) \tag{25}$$

$$\gamma_\mathbf{i} = \gamma + \eta_{\gamma\mathbf{i}}, \quad \eta_{\gamma\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \mathbf{\Sigma}_\gamma\right) \tag{26}$$

$$\rho_\mathbf{i} = \rho + \eta_{\rho\mathbf{i}}, \quad \eta_{\rho\mathbf{i}} \sim \mathbf{iid}\left(\mathbf{0}, \sigma_\rho\right) \tag{27}$$

Identification Challenge 5. Cross-Sectional Dependence *A priori*, the inclusion of $\mathbf{f_t}$ in Eq. 23 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 detect cross-sectional dependence across forecasts using the test from Pesaran (2015).[9]

To account for cross-sectional dependence across forecasters, that is, exposure to common shocks (f_t), we employ 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 us to account for both common effects and parameter heterogeneity.

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

[8] 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, φ_i lies strictly inside the unit circle.

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

In the DDCE 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 to be certain. c_{yi} are forecaster fixed effects. The λ_{it}^{h} can be spatially correlated.

y_{it} and $\mathbf{x}_{it}$ are assumed to be generated by the following linear covariance stationary dynamic heterogeneous panel data model with weakly exogenous regressors:

$$y_{it} = c_{yi} + \varphi_i y_{i,t-1} + \phi^{\pi^*}\pi_t^* + \phi^{r^*} r_t^* + \beta_i^{'}\mathbf{x}_{it} + u_{it} \tag{28}$$

$$u_{it} = \sum_{\ell=0}^{L=1} (\gamma_{\mathbf{i}})^{'}\, \mathbf{f_{t-\ell}} + \lambda_{it} \tag{29}$$

$$\lambda_{it} = \rho_i \lambda_{i,t-1} + \varepsilon_{it} \tag{30}$$

$$\omega_{\mathbf{it}} = \begin{bmatrix} \mathbf{x_{it}} \\ \mathbf{g_{it}} \end{bmatrix} = \mathbf{c}_{\omega i} + \alpha_{\mathbf{i}} y_{i,t-1} + \mathbf{\Gamma_i^{'} f_t} + \mathbf{v_{it}} \tag{31}$$

This model can be consistently estimated using the *Dynamic Common Correlated Effects* (DCCE) (mean-group) estimator (Pesaran (2006), Chudik and Pesaran (2015)). The $m\ x\ 1$ 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.[10]

$\alpha_{\mathbf{i}}$ is a $k\ x\ 1$ vector of unknown coefficients i.i.d. across i, and independently distributed of each η_i, $\gamma_{\mathbf{j}}$, $\mathbf{\Gamma_j}$, ε_{jt}, and $\mathbf{f_t}$, $\forall i, j, t$. $\mathbf{g_{it}}$ is a vector of variables, such as consumption expenditure and business investment forecasts, that do not enter Eq. 28 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 i.i.d. 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 becomes 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. 28.[11] We denote the number of lags needed by $p_T = \sqrt[3]{T}$. 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.

$$y_{it} = c_{yi} + \varphi_i y_{i,t-1} + \phi^{\pi^*}\pi_t^* + \phi^{r^*} r_t^* + \beta_{\mathbf{i}}^{'}\mathbf{x_{it}} + \sum_{\ell=1}^{p_T} \kappa_{i\ell}\bar{y}_{t-\ell} + \sum_{\ell=0}^{p_T} \bar{\mathbf{x}}_{t-\ell}^{'}\beta_{\mathbf{i}\ell} + e_{it} \tag{32}$$

$$e_{it} = \lambda_{it} + \sum_{\ell=p_T+1}^{\infty} \kappa_{i\ell}\bar{y}_{t-\ell} + \sum_{\ell=p_T+1}^{\infty} \bar{\mathbf{x}}_{t-\ell}^{'}\beta_{\mathbf{i}\ell} + O_p\left(N^{-1/2}\right) \tag{33}$$

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 elim-

[10] 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 $km\ x\ 1$.

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

inated.[12] 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.[13] 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.

The key difference from Chudik and Pesaran (2015) is that our error term, λ_{it}^{h}, is serially correlated by definition (Eq. 5). Yet, if we quasi-difference Eq. 32 we get:[14]

$$y_{it}-\rho_i y_{i,t-1}=\left(c_{yi}+\varphi_i y_{i,t-1}+\phi^{\pi^*}\pi_t^*+\phi^{r^*}r_t^*+\beta_{\mathbf{i}}'\mathbf{x_{it}}+\sum_{\ell=1}^{p_T}\kappa_{i\ell}\bar{y}_{t-\ell}+\sum_{\ell=0}^{p_T}\bar{\mathbf{x}}_{t-\ell}'\beta_{\mathbf{i}\ell}+e_{it}\right)-$$
$$\rho_i\left(c_{yi}+\varphi_i y_{i,t-2}+\phi^{\pi^*}\pi_{t-1}^*+\phi^{r^*}r_{t-1}^*+\beta_{\mathbf{i}}'\mathbf{x_{i,t-1}}+\sum_{\ell=1}^{p_T}\kappa_{i\ell}\bar{y}_{t-1-\ell}+\sum_{\ell=0}^{p_T}\bar{\mathbf{x}}_{t-1-\ell}'\beta_{\mathbf{i}\ell}+e_{i,t-1}\right) \quad (34)$$

Rearranging gives the final form of our estimator:

$$y_{it}=c_{yi}(1-\rho_i)+(\rho_i+\varphi_i)y_{i,t-1}-\rho_i\varphi_i y_{i,t-2}+\left(\phi^{\pi^*}\pi_t^*-\rho_i\phi^{\pi^*}\pi_{t-1}^*\right)+\left(\phi^{r^*}r_t^*-\rho_i\phi^{r^*}r_{t-1}^*\right)+\beta_{\mathbf{i}}'\mathbf{x_{it}}-\rho_i\beta_{\mathbf{i}}'\mathbf{x_{i,t-1}}$$
$$+\left(\sum_{\ell=1}^{p_T}\kappa_{i\ell}\bar{y}_{t-\ell}-\rho_i\sum_{\ell=1}^{p_T}\kappa_{i\ell}\bar{y}_{t-1-\ell}\right)+\left(\sum_{\ell=0}^{p_T}\bar{\mathbf{x}}_{t-\ell}'\beta_{\mathbf{i}\ell}-\rho_i\sum_{\ell=0}^{p_T}\bar{\mathbf{x}}_{t-1-\ell}'\beta_{\mathbf{i}\ell}\right)+(e_{it}-\rho_i e_{i,t-1}) \quad (35)$$

Expanding the error term gives:

$$(e_{it}-\rho_i e_{i,t-1})=\lambda_{it}-\rho_i\lambda_{i,t-1}+$$
$$\left(\sum_{\ell=p_T+1}^{\infty}\kappa_{i\ell}\bar{y}_{t-\ell}-\rho_i\sum_{\ell=p_T+1}^{\infty}\kappa_{i,\ell+1}\bar{y}_{t-1-\ell}\right)+\left(\sum_{\ell=p_T+1}^{\infty}\bar{\mathbf{x}}_{t-\ell}'\beta_{\mathbf{i}\ell}-\rho_i\sum_{\ell=p_T+1}^{\infty}\bar{\mathbf{x}}_{t-1-\ell}'\beta_{\mathbf{i},\ell+1}\right)+O_p^t\left(N^{-1/2}\right)-\rho_i O_p^{t-1}\left(N^{-1/2}\right) \quad (36)$$

where we can assume that the terms in the second line (essentially, the effect of distant lagged regressors) above are vanishingly small,[15] then:

$$e_{it}-\rho_i e_{i,t-1}\approx\lambda_{it}-\rho_i\lambda_{i,t-1} \quad (37)$$

Substituting in Eq. 30 gives:

$$e_{it}-\rho_i e_{i,t-1}\approx(\rho_i\lambda_{i,t-1}+\varepsilon_{it})-\rho_i\lambda_{i,t-1}\Rightarrow \quad (38)$$

[12] Given the form our shock measure (Eq. 40), we cannot include $\bar{y}_{it}$ because we would purge the very shock measure we seek to extract. The consequence is that our shock measure (40) could still be somewhat contaminated by f_t^e, although because f_t^e appears in $\mathbf{f_t}$, which appears in $\mathbf{x_{it}}$, this should not be the case. As such, terms pertaining $\kappa_{i0}\bar{y}_{it}$ to do not appear in e_{it}.

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

[14] Quasi-differencing 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.

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

$$e_{it} - \rho_i e_{i,t-1} \approx \varepsilon_{it} \tag{39}$$

where ε_{it} becomes our new error term, which is not serially correlated. As such, we can estimate Eq. 35 with the DCCE estimator. Then, the forecast horizon-specific (h) perceived shock estimates at time t are:

$$\widehat{\varepsilon_t^h} = \frac{1}{N} \sum_i^N \widehat{\varepsilon_{it}^h} \tag{40}$$

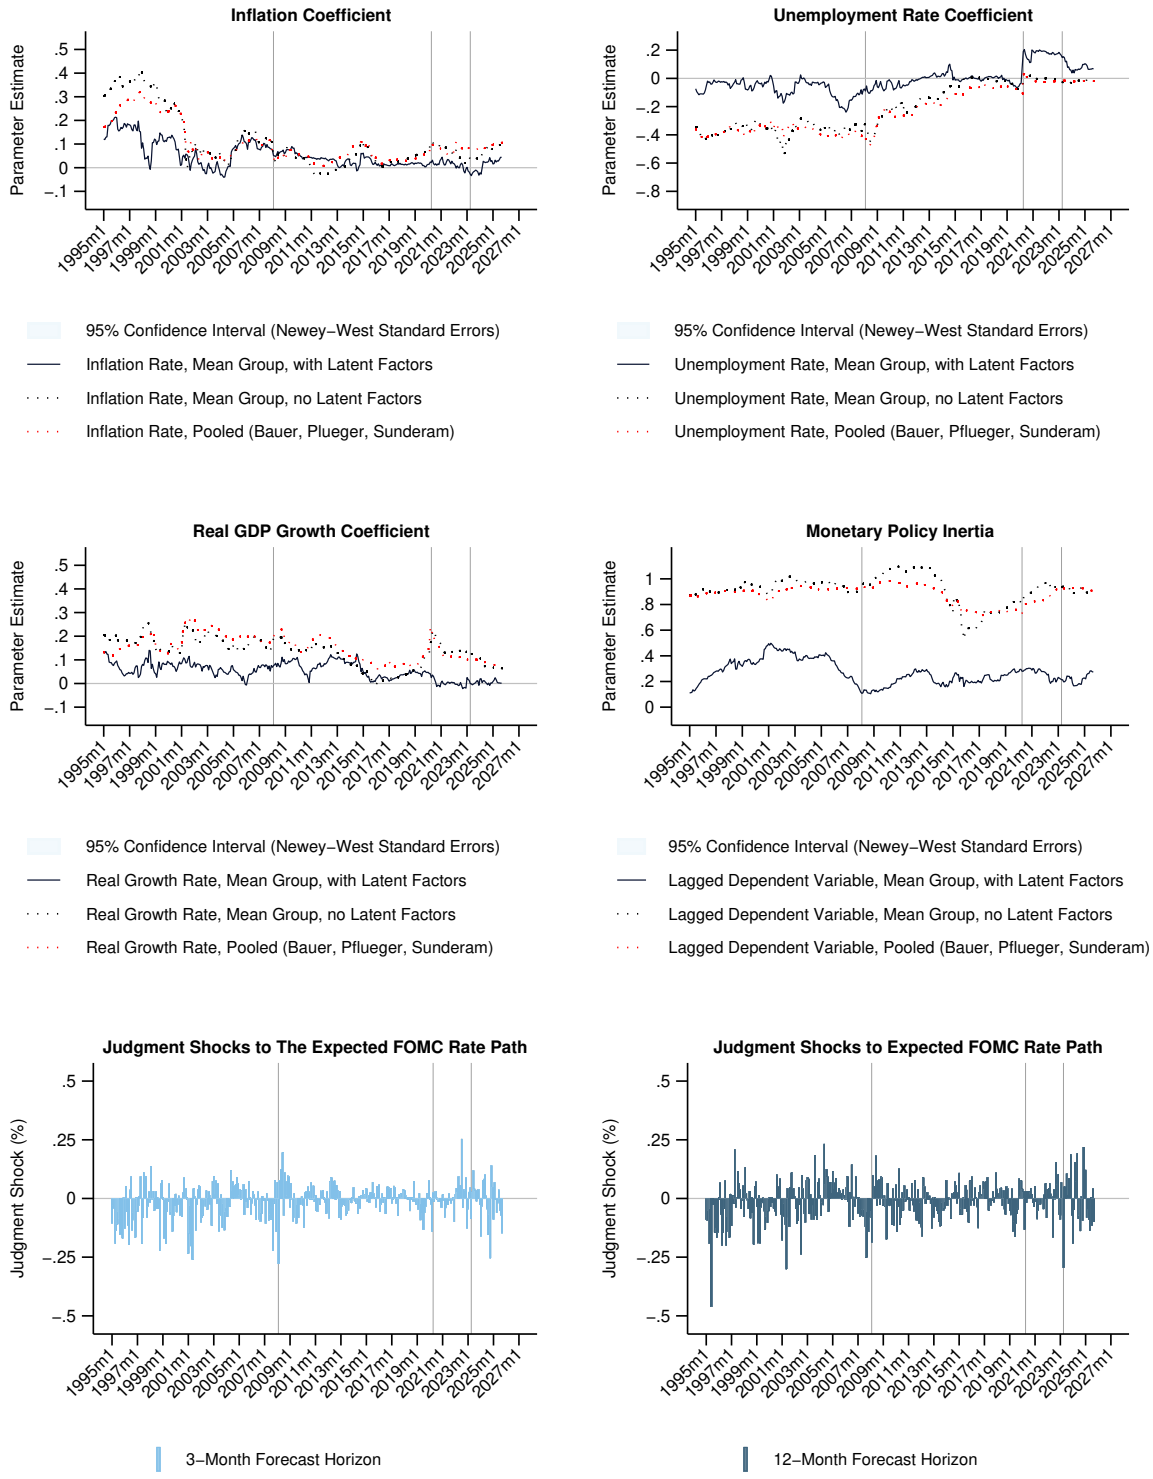
where the forecaster-specific parameter and ε_{it}^h estimates come from the rolling-window estimate whose last observation is t (and first observation is $t-W+1$), where W is window length in months. The mean-group coefficients and ε_{it}^h estimates are averaged in the same manner.

3 Estimation

In this section, we present the forecast data, the forecaster-expected reaction function parameter estimates, and the monetary 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, 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*.

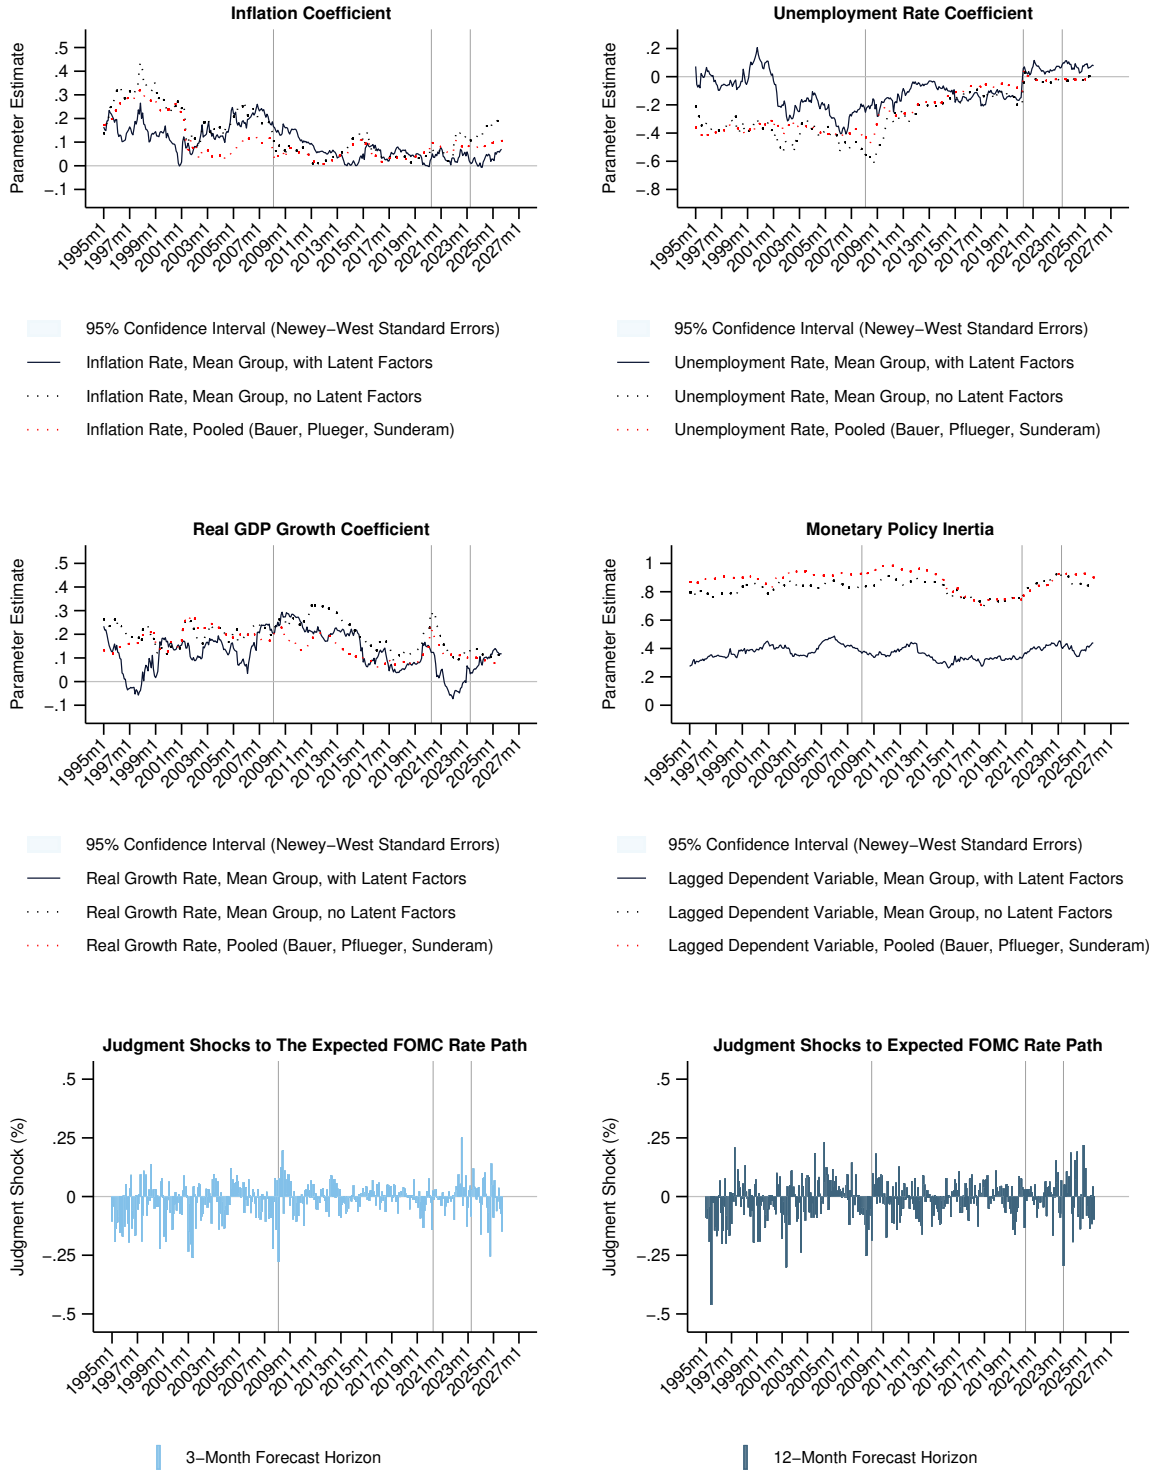
3.1 Expected FOMC Reaction Function Estimates

Figure 1: **Expected FOMC Reaction Function 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 2025m9. Forecasts made on second or third Monday of each month.

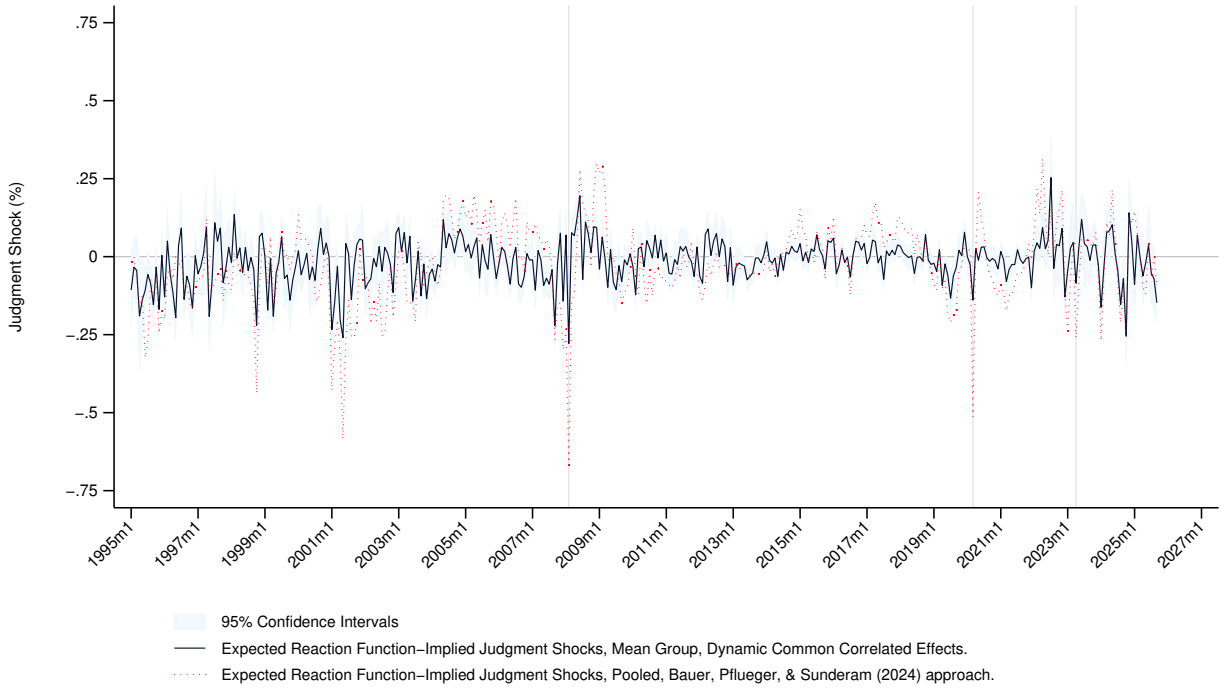
Figure 2: **Expected FOMC Reaction Function Estimates (12-month forecast horizon)**



Note 1: The thin, grey, vertical lines denote the case studies described below.
Note 2: results available for months 1995m1 to 2025m9. Forecasts made on second or third Monday of each month.

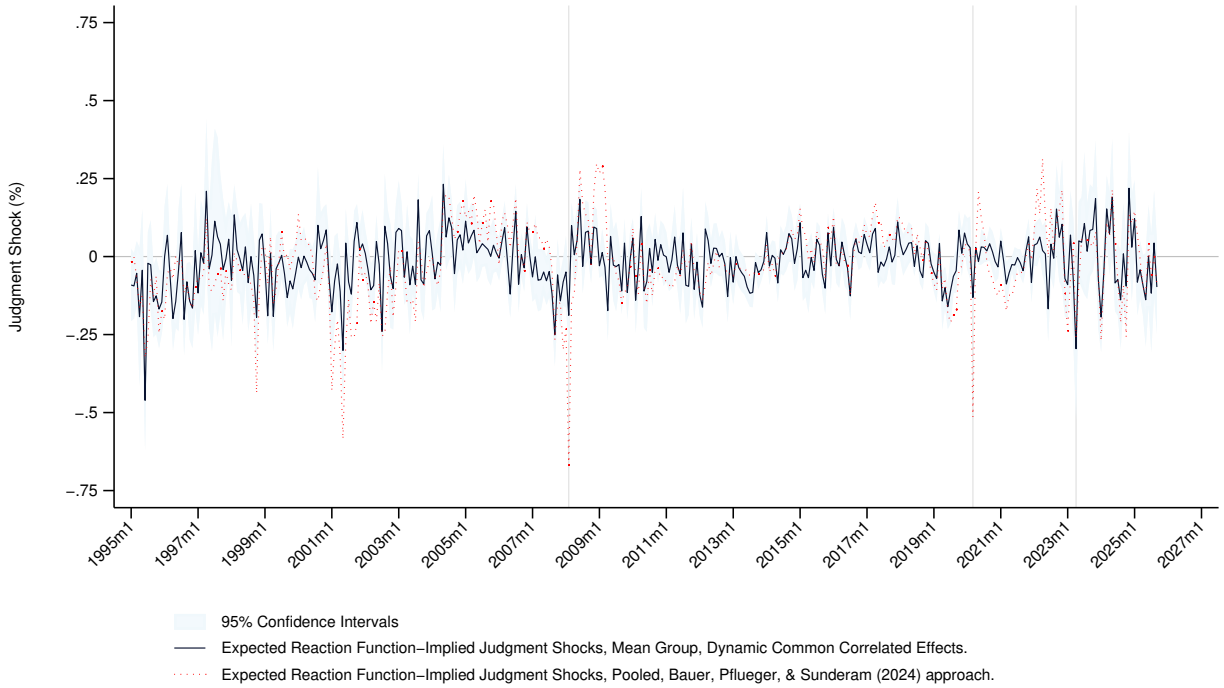
3.2 Expected FOMC Reaction Function-Implied Judgment Shocks

Figure 3: **FOMC Judgment Shock Estimates (3-month Forecast Horizon)**



Sample period: 1995m1 to 2025m9. The thin, grey, vertical lines denote the case studies described below.

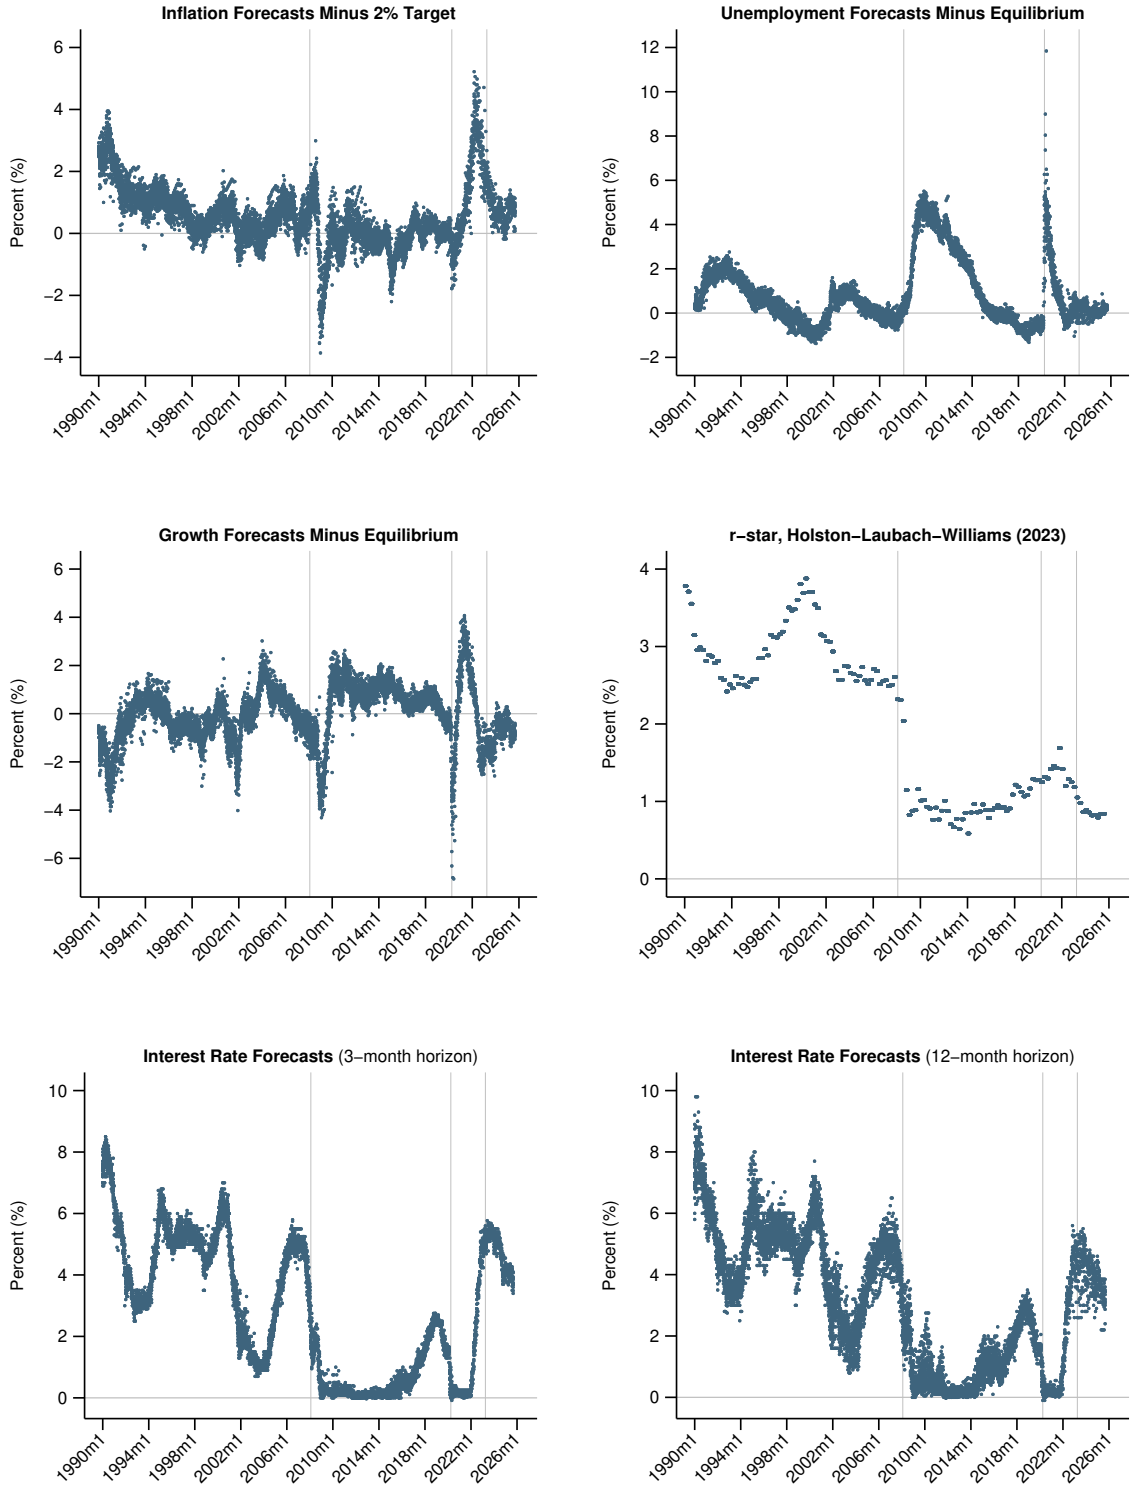
Figure 4: **FOMC Judgment Shock Estimates (12-month Forecast Horizon)**



Sample period: 1995m1 to 2025m9. The thin, grey, vertical lines denote the case studies described below.

3.3 Individual Forecasts at the 12-month Forecast Horizon

Figure 5: **Forecast Data**



Each dot represents an individual forecaster's forecast. The thin, grey, vertical lines denote the case studies described below.

3.4 Why are the Estimated Inflation Coefficients Less Than 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 1 and 2, 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) "present three reasons why a central bank with firmly anchored long-run inflation expectations optimally varies the nominal interest rate less than one-for-one with inflation in certain circumstances. 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 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.

Next, as Coibion and Gorodnichenko (2012), Bauer, Pflueger, and Sunderam (2024), and others have observed, the dynamic nature of the monetary policy rule could imply that the macro response coefficients are short-run in nature—that is, if we assume they are proportional to $(1 - \tilde{\varphi}_t)$. Now, since the estimates of expected $\tilde{\varphi}_t$ in Figures 1 and 2 oscillate around $\frac{1}{3}$ most of the time, that implies that the long-run coefficients are approximately $\frac{3}{2}$ times the magnitude of the short-run expected coefficient estimates. Yet, this is not nearly enough to increase the coefficients above one.

3.5 Why are the Inertia Coefficients So Low?

The coefficient on the lagged dependent variable, $(\varphi_i + \rho_i)$, captures the persistence of forecaster i's own projections (**idiosyncratic inertia**), whereas the coefficient on the cross-sectional average of the lagged dependent variable, κ_{i1}, reflects the influence of common shocks and aggregate conditions that drive all forecasters' projections (**common inertia**).

The estimates from Figures 1 and 2 therefore have a structural and behavioural interpretation: they show how forecaster i's current projection responds to their own past projection (idiosyncratic persistence) and to their own perceived fundamentals (idiosyncratic sensitivity).

By contrast, estimates in the literature that report persistence of around 0.8–0.9 (Clarida et al. (2000), Coibion and Gorodnichenko (2012)) are reduced-form measures of statistical persistence. Such estimates conflate idiosyncratic inertia with common-factor-driven persistence, because they do not control for shared aggregate shocks. Our DCCE approach, by purging common dynamics via the cross-sectional means, isolates the structural, forecaster-specific inertia and thus yields a more interpretable behavioural parameter.

3.6 Case Studies

February 2008. We estimate that there was a large, negative judgment shock at both the three-month (Figure 3) and twelve-month (Figure 4) forecast horizons. The expected reaction function coefficient estimates were fairly stable around this time (Figures 1 and 2). Amidst a weakening economic outlook, increasing downside risks to growth, strains in short-term funding markets, deteriorating financial market conditions, tightening credit conditions, deepening housing contraction, and a softening in labor markets, the FOMC cut the federal funds rate at an unscheduled meeting on January 22, 2008 by 75 basis points—the largest single-day reduction in over two decades. Immediately after the first cut, futures contracts implied a 72 percent chance of another 50 basis points rate cut at the scheduled January 30 meeting (source). And, indeed The FOMC cut another 50 basis points. Later, during testimony to Congress, Chairman Bernanke noted that "*The FOMC has responded aggressively to the weaker outlook for economic activity, having reduced its target for the federal funds rate by 225 basis points since last summer*." "*The FOMC will be carefully evaluating incoming information bearing on the economic outlook and will act in a timely manner as needed to support growth and to provide adequate insurance against downside risks*."

March-April 2020, Early-Pandemic Shocks. In March 2020 (2020m3), we observe negative judgment shocks following a 50-basis-point FFR decrease (Figures 3, 4). In early April (2020m4), we observe large, downward forecast adjustments (Figure 5), and sharp upward adjustment to the FOMC's perceived unemployment responsiveness (Figures 1, 2). Due to these negative judgment shocks, the expected FOMC rate path over the next year or so was lower than that implied by the systematic component of the forecaster-expected reaction function.

These patterns are partly attributable to FOMC monetary policy. First, at an unscheduled meeting on March 3, 2020 (2020m3), the FFR was unexpectedly lowered by 50 basis points to the effective lower bound (ELB), helping explain the negative judgment shocks observed in early March. Commensurately, forecasters lowered their three-month rate forecasts to the ELB, and so could not go lower (Figure 5); twelve-month rate forecasts also fell to the ELB, but dropped from a higher starting point. Second, at another unscheduled meeting on March 15, 2020 (2020m4), strong FOMC forward guidance was provided.[16] These downward interest rate forecast changes are also partly attributable to the differences in perceived FOMC responsiveness to unemployment forecasts across the two forecast horizons (3-month horizon: from -0.04 in February 2020 (2020m2), to 0.08 in March 2020 (2020m3), to 0.19 in April 2020 (2020m4); 12-month horizon: from -0.15 in February 2020, to -0.03 in March 2020, to 0.06 in April 2020). This coefficient heterogeneity demonstrates the utility in estimating the model separately across the two forecast horizons.[17]

April 2023. At its March 22, 2023 meeting, the FOMC raised rates for a ninth consecutive time. Yet, it signaled that banking-system stress might curtail its tightening cycle sooner than expected, or at a miminum, that there was greater uncertainty about future rate increases. Forward guidance shifted from "*The Committee anticipates that ongoing increases*" to "*The Committee anticipates that some additional policy firming may be appropriate*." We estimate that there was a large, negative monetary shock at the twelve-month forecast horizon (Figure 4), and a moderate shock at the three-month horizon (Figure 3). The macro coefficient estimates were stable.

[16] "*The Committee expects to maintain this target range until it is confident that the economy has weathered recent events and is on track to achieve its maximum employment and price stability goals*." (Source).

[17] In moving from the seven-year rolling window ending March 2020 to that ending April 2020, the data from April 2013 drop out. Fortunately, the data and coefficient estimates at this time are highly stable; the monetary shock estimates corresponding to this period were each near-zero, at about 0.01.

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).[18]

$$y_{n,t+h} = \mu_n^h + \theta_n^h r_t + \sum_{\ell=0}^{3} \mathbf{y}_{\mathbf{t}-\ell}^{'} \kappa_{\mathbf{l}}^{\mathbf{h}} + \xi_{nt}^h \tag{41}$$

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.[19] 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 $\mathbf{z}_\mathbf{t}^{'} = \left[\widehat{\varepsilon_t^{h=3}}, \widehat{\varepsilon_t^{h=12}}\right]$ be the vector of instrumental variables. The conditions for valid estimation of Eq. 41 are:[20]

1. $\mathbb{E}\left[\xi_{1t}^{h,\perp} \mathbf{z}_\mathbf{t}^{\perp}\right] = \alpha > \alpha^*$ (instrument **relevance** for monetary shock, $\xi_{1t}^{h,\perp}$).
2. $\mathbb{E}\left[\xi_{nt}^{h,\perp} \mathbf{z}_\mathbf{t}^{\perp}\right] = 0$ (**contemporaneous exogeneity** of instrument to other shocks, $\xi_{nt}^{h,\perp}, n \neq 1$).
3. $\mathbb{E}\left[\xi_{n,t+j}^{h,\perp} \mathbf{z}_\mathbf{t}^{\perp}\right] = 0 \quad \forall j{\neq}0$ (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. 41 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

[18] 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.

[19] 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.

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

and Mertens (2025) test[21]). We provide three main results here. First, we find evidence that the forecast-derived perceived monetary policy shocks are not weak instruments (Figure 6).[22] 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.

The weak instruments tests are based on our Local Projections Instrumental Variables estimator (Eq. 41). During estimation, we take steps to adjust for potential bias from weak instruments. Eq. 41 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 et al. (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).

[21] 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.

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

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,[23] but in this case, compare it to the Montiel Olea and Pflueger (2013) critical values (two dashed lines, Figure 6), which are LIML-specific.[24] 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.[25]

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 θ_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.

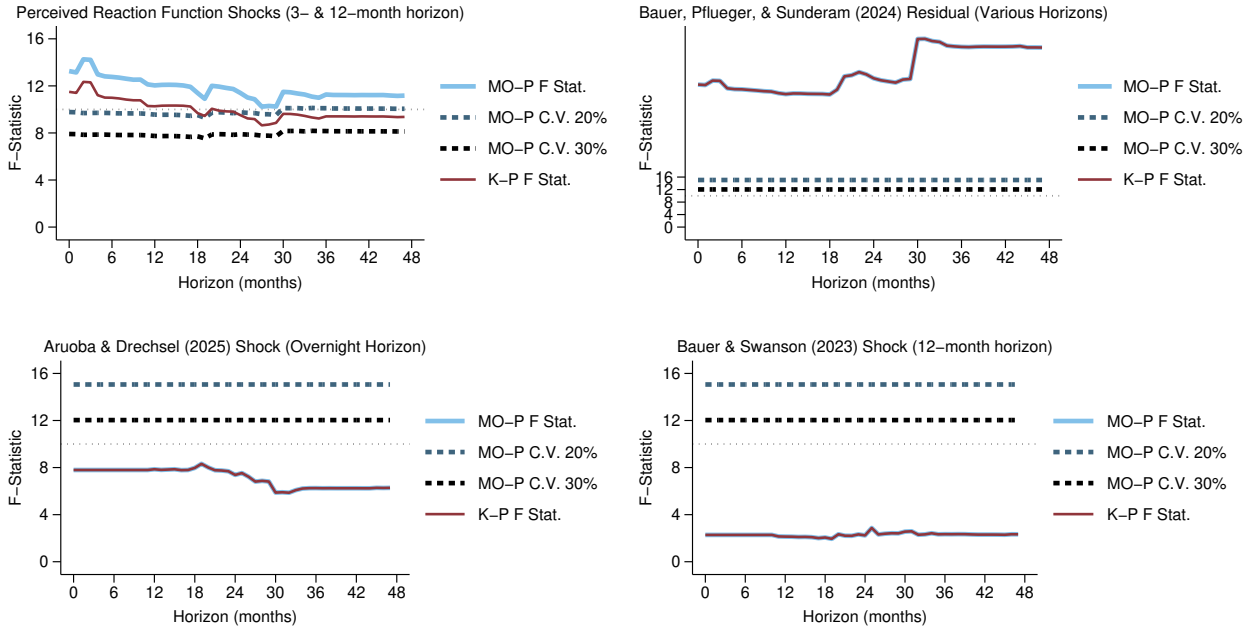
[23] If just identified, the Kleibergen-Paap Wald F statistic equals the effective F statistic.

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

[25] 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 et al. (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).

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 2025m3.

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.

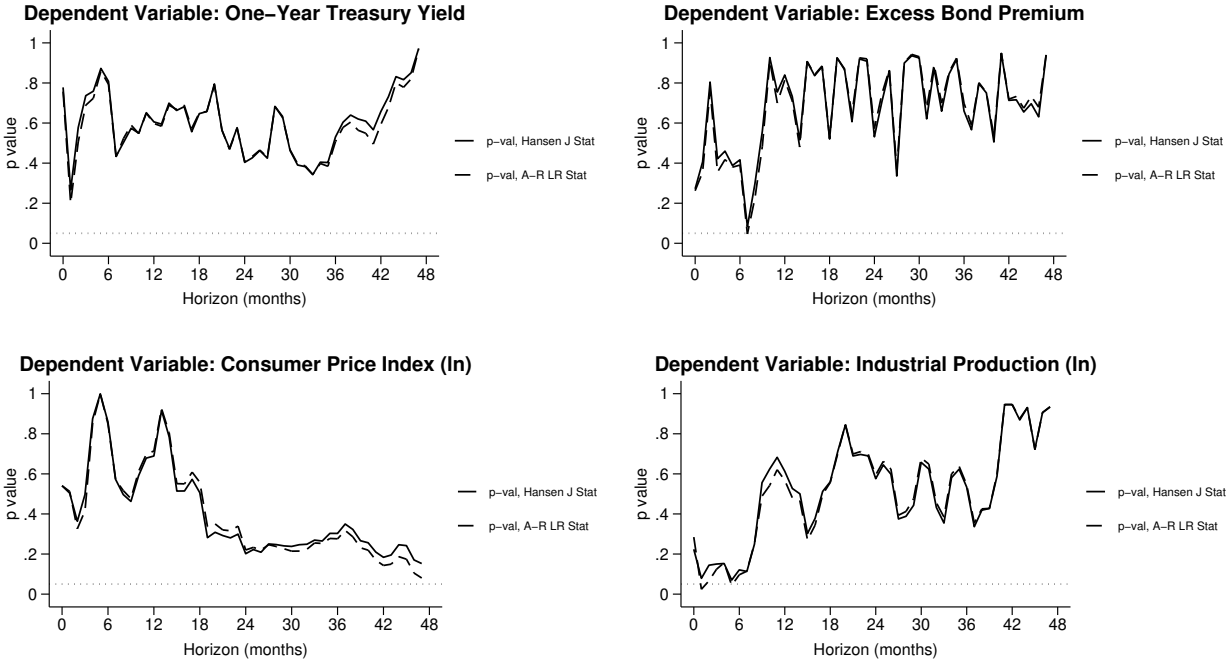
In Figures 7 and 8 below, we plot two sets of statistics.

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**

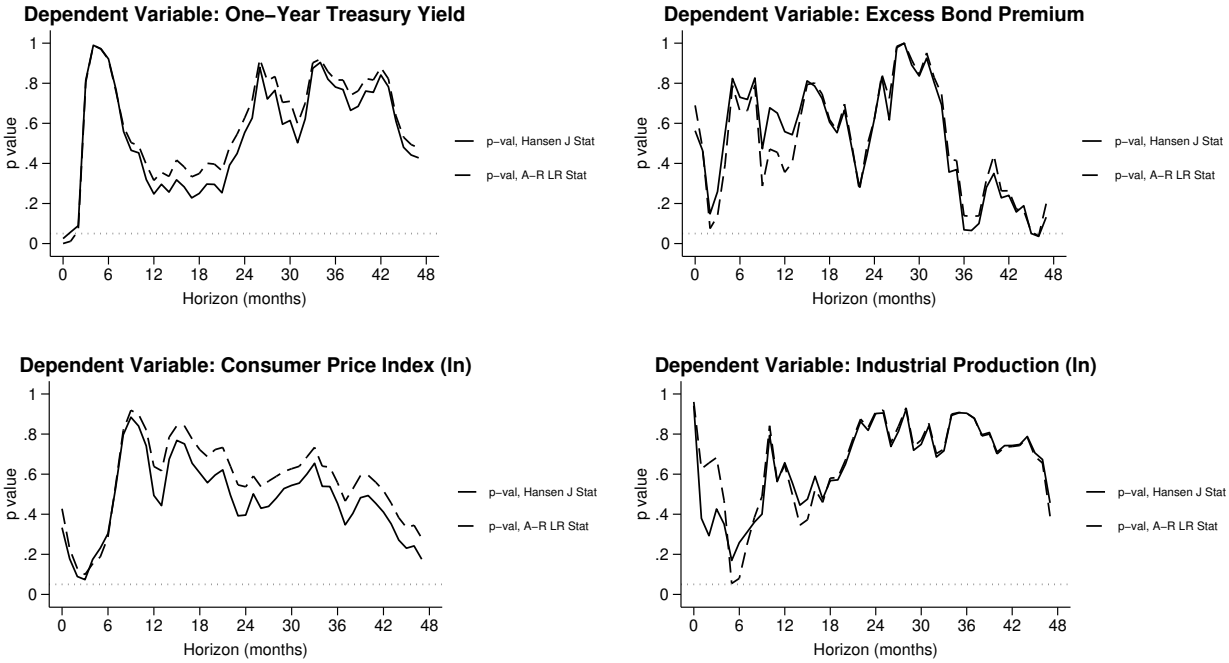
Test statistics based on forecaster-perceived monetary policy shocks instruments.



Dotted, straight line denotes p=0.05. Sample period (monthly data): 1990m1 to 2025m3.
Hansen J−test: robust to heteroskedasticity. Anderson−Rubin test: robust to weak instruments.

Figure 8: **Tests of Overidentifying Restrictions**

Test statistics using a set of Bauer and Swanson (2023), Bauer, Pflueger, and Sunderam (2024), Aruoba and Drechsel (2025) instruments.



Dotted, straight line denotes p=0.05. Sample period (monthly data): 1990m1 to 2025m3.
Hansen J−test: robust to heteroskedasticity. Anderson−Rubin test: robust to weak instruments.

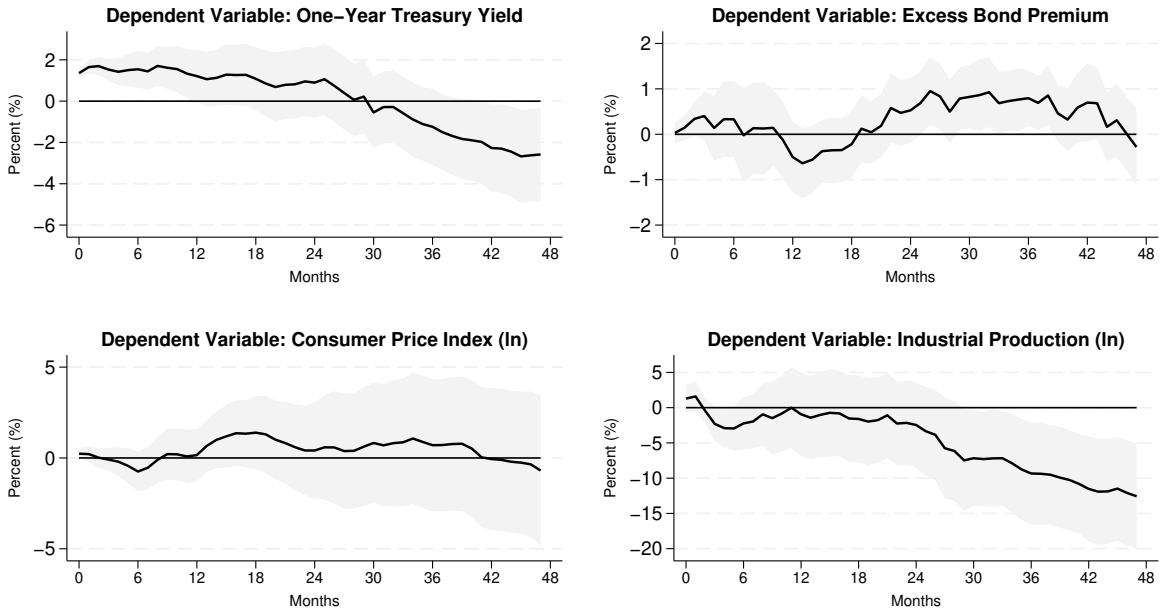
4.3 Impulse Response Functions

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 relationship with the Excess Bond Premium is less clear.

The estimates created using the Bauer, Pflueger, and Sunderam (2024) residuals (Figure 10) are broadly similar to those in Figure 9, although with narrower confidence intervals attributable to the stronger instrument. Conversely, the estimates created using the Bauer and Swanson (2023) shocks (Figure 11) have rather wide confidence intervals given the weakness of the instrument, so it is difficult to draw conclusions. The estimates created using the Aruoba and Drechsel (2025) shocks (Figure 12) are also broadly consistent with New Keynesian theory.

Figure 9: **Impulse Response Functions**

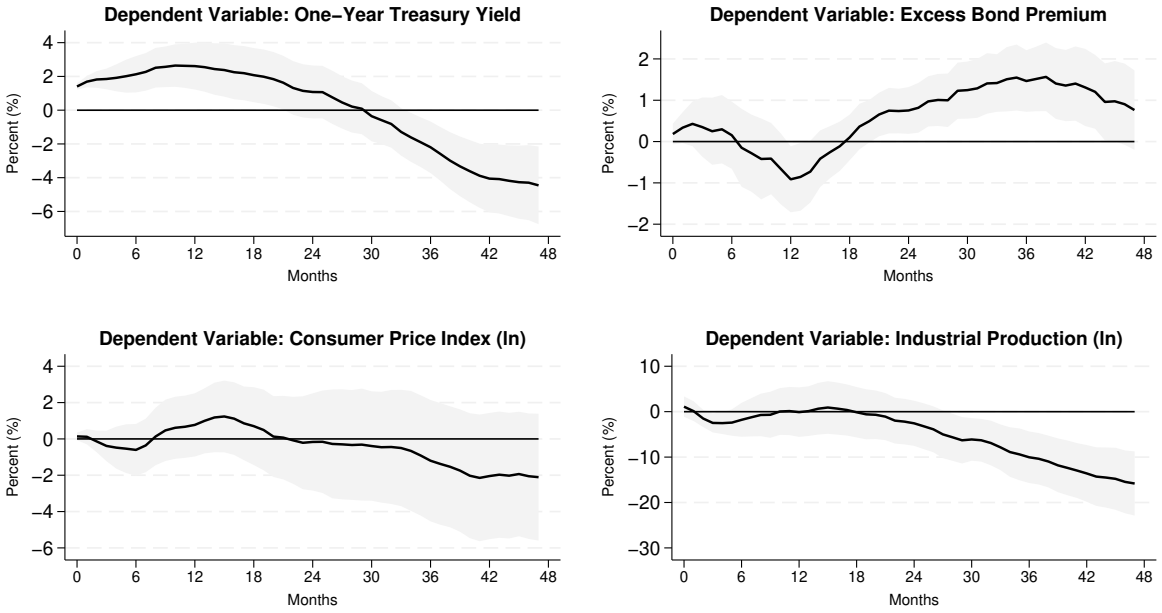
Instruments: perceived monetary policy shocks instruments (based on expected reaction function estimates using forecast data.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2025m3.

Figure 10: **Impulse Response Functions**

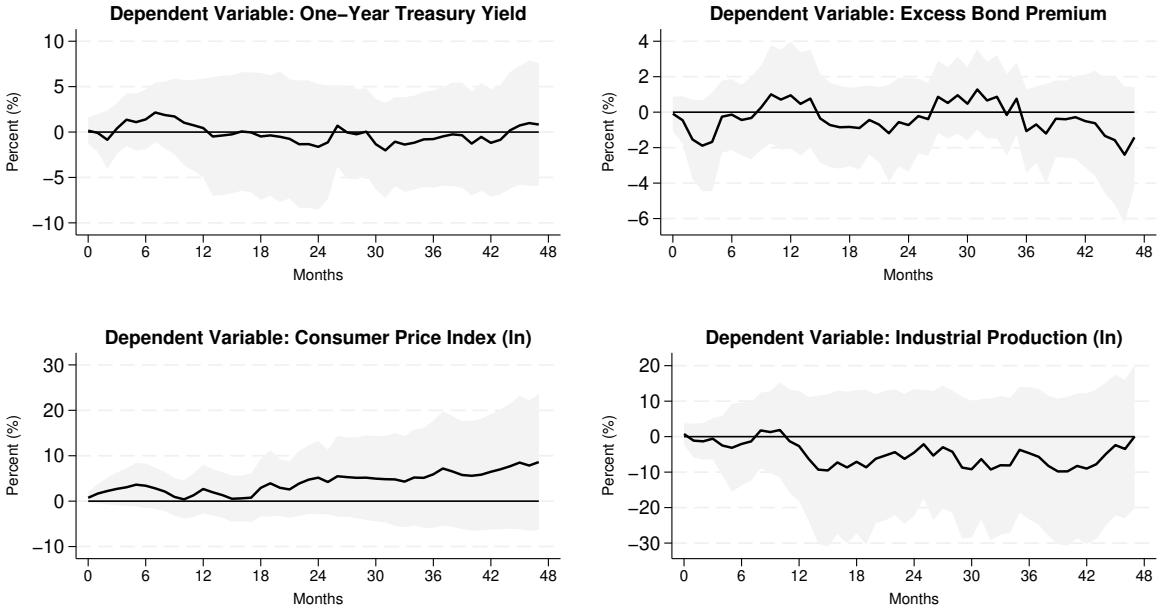
"Instrument": Bauer, Pflueger, and Sunderam (2024) residuals.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2025m3.

Figure 11: **Impulse Response Functions**

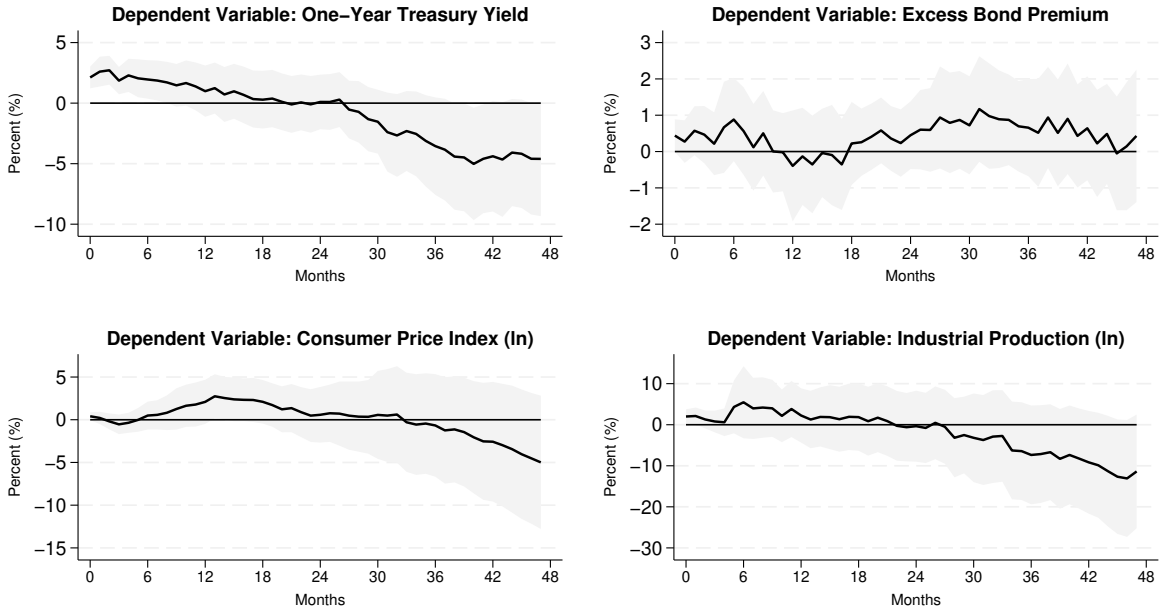
Instrument: Bauer and Swanson (2023) shock measure.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2025m3.

Figure 12: **Impulse Response Functions**

Instrument: Aruoba and Drechsel (2025) shock measure.



90% confidence intervals.
HAC−robust standard errors.
Sample period (monthly data): 1990m1 to 2025m3.

5 Conclusion

This paper has proposed a number of ideas. First, the reaction function framework incorporates central bank judgment and multidimensional parameter heterogeneity. 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, all results based on such instruments must be read through the lens of weak instruments (estimates bi-

ased towards OLS with understated standard errors). Fortunately, more recent efforts to incorporate more monetary policy events hold promise (for example the 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.

Testing monetary policy instruments for relevance and exogeneity. *Relevance*. To test this, either the Montiel Olea and Pflueger (2013) or the Lewis and Mertens (2025) robust tests for weak instruments for 2SLS should be used. *Exogeneity*. It is useful to use numerous monetary policy instruments, 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, what if parameters follow a random walk instead of a random coefficients model? Second, what if coefficients are not known at the time forecasts are made? Third, can we improve our Local Projections estimates by, for example, using Bayesian Local Projections? We also 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," .

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

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.

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.

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

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.

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

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.

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

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.

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.

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

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.

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*.

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*.

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., 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, $\tilde{\Omega}_{\tilde{i}t_n}$.[26] 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, $\tilde{\mathbb{E}}_{t_n}\left[\,.\;\middle|\;\tilde{\Omega}_{t_n}\right]$.[27] This conditional expectation represents an FOMC macroeconomic projection at time t_n.

$$\tilde{f} : \tilde{\Omega}_{t_n} \rightarrow \tilde{\mathbb{E}}_{t_n}\left[\;\mathbf{x}_{\mathbf{t_n}+\mathbf{h_{t_n}}}\;\middle|\;\tilde{\Omega}_{t_n}\;\right] \tag{42}$$

$\mathbf{x}_{\mathbf{t_n}+\mathbf{h_{t_n}}}$ 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.[28]

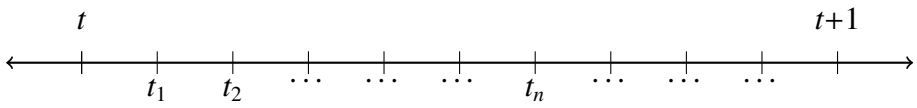
Time. To make sense of evolving information sets (filtrations), shocks, and what is known at the time a projection is made, we must be very careful with the sequence of events. Throughout the paper, t=$\{1, 2, \ldots, T\}$, $t \in \mathbb{N}^+$ is a standard indexing set, which denotes monthly periods, 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.

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

[27] 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.

[28] 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:

$$\tilde{f} : \tilde{\omega}_{t_1} \in \tilde{\Omega}_{t_1} \rightarrow \tilde{\mathbb{E}}_{t_1}\left[\mathbf{x}_{\mathbf{t_1}+\mathbf{h}^{\mathbf{0}}_{\mathbf{t_1}}} \;\middle|\; \tilde{\Omega}_{t_1} = \tilde{\omega}_{t_1} \right] \in \mathbb{R}^k \tag{43}$$

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 $\tilde{\Omega}_{t_1} \subset \tilde{\Omega}_{t_2}$.[29] 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**:

$$\tilde{g} : \mathbb{R}^k \rightarrow \mathbb{R} \tag{44}$$

$$\tilde{g} : \tilde{\mathbb{E}}_{t_1}\left[\mathbf{x}_{\mathbf{t_1}+\mathbf{h}^{\mathbf{0}}_{\mathbf{t_1}}} \;\middle|\; \tilde{\Omega}_{t_1} = \tilde{\omega}_{t_1} \right] \rightarrow \tilde{\mathbb{E}}_{t_2}\left[\tilde{p}_{t_2+h^0_{t_2}} \;\middle|\; \tilde{\Omega}_{t_2} = \tilde{\omega}_{t_2} \right] \in \mathbb{R} \tag{45}$$

Or, equivalently:

$$\tilde{g}\left(\tilde{f}\left(\tilde{\omega}_{t_1}\right)\right) = \tilde{g}\left(\tilde{\mathbb{E}}_{t_1}\left[\mathbf{x}_{\mathbf{t_1}+\mathbf{h}^{\mathbf{0}}_{\mathbf{t_1}}} \;\middle|\; \tilde{\Omega}_{t_1} = \tilde{\omega}_{t_1}\right]\right) = \tilde{\mathbb{E}}_{t_2}\left[\tilde{p}_{t_2+h^0_{t_2}} \;\middle|\; \tilde{\Omega}_{t_2} = \tilde{\omega}_{t_2}\right] \in \mathbb{R} \tag{46}$$

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.[30] That is, the final projection from Eq. 45 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:

$$\tilde{\mathbb{E}}_{t_n}\left[\tilde{\lambda}_{t_n+h_{t_n}} \;\middle|\; \tilde{\Omega}_{t_n} = \tilde{\omega}_{t_n} \right] = \tilde{\rho}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\left[\tilde{\lambda}_{t_{n-2}+h_{t_{n-2}}} \;\middle|\; \tilde{\Omega}_{t_{n-2}} = \tilde{\omega}_{t_{n-2}} \right] + \tilde{\varepsilon}^{h_{t_{n-1}}}_{t_{n-1}}, \qquad \tilde{\varepsilon}^{h_{t_n}}_{t_n} \overset{iid}{\sim} \left(0, \sigma^2_{\tilde{\varepsilon}}\right) \tag{47}$$

$\tilde{\varepsilon}^{h_{t_n}}_{t_n}$ is a shock to FOMC-expected judgment. When h_{t_n}=0, then $\tilde{\mathbb{E}}_{t_2}\left[\tilde{p}_{t_2+h^0_{t_2}} \;\middle|\; \tilde{\Omega}_{t_2} = \tilde{\omega}_{t_2}\right]$ 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 $\tilde{\varepsilon}^{h_{t_n}}_{t_n}$.

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

[30] 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. $\tilde{\Omega}_{t_2}\subset\tilde{\Omega}_{t_3}\subset\tilde{\Omega}_{t_4}$).

$$\tilde{\mathbb{E}}_{t_4}\left[\ \tilde{p}_{t_4+h^0_{t_4}}\ \middle|\ \tilde{\Omega}_{t_4}=\tilde{\omega}_{t_4}\ \right]=\tilde{\mathbb{E}}_{t_2}\left[\ \tilde{p}_{t_2+h^0_{t_2}}\ \middle|\ \tilde{\Omega}_{t_2}=\tilde{\omega}_{t_2}\ \right]+\tilde{\mathbb{E}}_{t_3}\left[\ \tilde{\lambda}_{t_3+h^0_{t_3}}\ \middle|\ \tilde{\Omega}_{t_3}=\tilde{\omega}_{t_3}\ \right] \tag{48}$$

Next, the projection produced using Eq. 48 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.

$$\tilde{f}:\tilde{\omega}_{t_5}\in\tilde{\Omega}_{t_5}\rightarrow\tilde{\mathbb{E}}_{t_5}\left[\ \mathbf{x}_{\mathbf{t_5}+\mathbf{h}^1_{\mathbf{t_5}}}\ \middle|\ \tilde{\Omega}_{t_5}=\tilde{\omega}_{t_5}\ \right]\in\mathbb{R}^k \tag{49}$$

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:

$$\tilde{g}:\tilde{\mathbb{E}}_{t_5}\left[\ \mathbf{x}_{\mathbf{t_5}+\mathbf{h}^1_{\mathbf{t_5}}}\ \middle|\ \tilde{\Omega}_{t_5}=\tilde{\omega}_{t_5}\ \right]\rightarrow\tilde{\mathbb{E}}_{t_6}\left[\ \tilde{p}_{t_6+h^1_{t_6}}\ \middle|\ \tilde{\Omega}_{t_6}=\tilde{\omega}_{t_6}\ \right]\in\mathbb{R} \tag{50}$$

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.

$$\tilde{\mathbb{E}}_{t_8}\left[\ \tilde{p}_{t_8+h^1_{t_8}}\ \middle|\ \tilde{\Omega}_{t_8}=\tilde{\omega}_{t_8}\ \right]=\tilde{\mathbb{E}}_{t_6}\left[\ \tilde{p}_{t_6+h^1_{t_6}}\ \middle|\ \tilde{\Omega}_{t_6}=\tilde{\omega}_{t_6}\ \right]+\tilde{\mathbb{E}}_{t_7}\left[\ \tilde{\lambda}_{t_7+h^1_{t_7}}\ \middle|\ \tilde{\Omega}_{t_7}=\tilde{\omega}_{t_7}\ \right] \tag{51}$$

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

$$\begin{aligned}&\tilde{\mathbb{E}}_{t_n}\left[\ \tilde{p}_{t_n+h_{t_n}}\ \right]=\tilde{\mathbb{E}}_{t_{n-1}}\left[\ \tilde{\lambda}_{t_{n-1}+h_{t_{n-1}}}\ \right]+\\&\tilde{\mathbb{E}}_{t_{n-2}}\left[\ \tilde{\varphi}^{h_{t_{n-2}}}_{t_{n-2}}(L)\tilde{p}_{t_{n-2}+h_{t_{n-2}}}+\tilde{\phi}^{\pi^*}_{t_{n-2}}\pi^*_{t_n}+\tilde{\phi}^{r^*}_{t_{n-2}}r^*_{t_{n-2}}+\left(\mathbf{x}_{\mathbf{t_{n-2}}+\mathbf{h}_{\mathbf{t_{n-2}}}}-\mathbf{x}^*_{\mathbf{t_{n-2}}}\right)'\tilde{\boldsymbol{\beta}}^{\mathbf{h}_{\mathbf{t_{n-2}}}}_{\mathbf{t_{n-2}}}\ \right]\end{aligned} \tag{52}$$

$\tilde{\varphi}^{h_{t_n}}_{t_n}(L)$ is a lag polynomial such that $\tilde{\varphi}^{h_{t_n}}_{t_n}(L)\tilde{p}_{t_n+h_{t_n}}=\sum_{\ell=1}^{L}\tilde{\varphi}^{\ell}_{t_n}\tilde{p}_{t_n+h_{t_n}-l}$. $\tilde{\varphi}^{\ell}_{t_n}$ is a monetary policy gradualism parameter.[31] Accordingly, $\tilde{\phi}^{\pi^*}_{t_n}\propto\left(1-\sum_{\ell=1}^{L}\tilde{\varphi}^{\ell}_{t_n}\right)$, $\tilde{\phi}^{r^*}_{t_n}\propto\left(1-\sum_{\ell=1}^{L}\tilde{\varphi}^{\ell}_{t_n}\right)$ (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.

[31] 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. 52 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:

$$\tilde{\mathbb{E}}_t\left[\ \tilde{p}_t\ \right] = \tilde{\mathbb{E}}_t\left[\ \tilde{\varphi}(L)\tilde{p}_t + \tilde{\phi}^{\pi^*}\pi^*_t + \tilde{\phi}^{r^*}r^*_t + \left(\mathbf{x_t} - \mathbf{x^*_t}\right)'\tilde{\beta}\ \right] \Rightarrow \tag{53}$$

$$\tilde{p}_t = \tilde{\varphi}(L)\tilde{p}_t + \tilde{\phi}^{\pi^*}\pi^*_t + \tilde{\phi}^{r^*}r^*_t + \left(\mathbf{x_t} - \mathbf{x^*_t}\right)'\tilde{\beta} \tag{54}$$

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

$$\tilde{p}_t = \tilde{\varphi}(L)\tilde{p}_t + \tilde{\phi}^{\pi^*}\pi^*_t + \tilde{\phi}^{r^*}r^*_t + \left(\mathbf{x_t} - \mathbf{x^*_t}\right)'\tilde{\beta} + \tilde{\varepsilon}_t \tag{55}$$

Assumption 4. *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 3, 4, and the linear property of the conditional expectation operator implies:

$$\begin{aligned}\tilde{\mathbb{E}}_{t_n}\left[\ \tilde{p}_{t_n+h_{t_n}}\ \right] = \tilde{\varphi}^{h_{t_n}}_{t_n}(L)\tilde{\mathbb{E}}_{t_{n-2}}\left[\tilde{p}_{t_{n-2}+h_{t_{n-2}}}\ \right] + \tilde{\phi}^{\pi^*}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\left[\ \pi^*\ \right] + \tilde{\phi}^{r^*}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\left[\ r^*_{t_{n-2}}\ \right] + \\ \tilde{\mathbb{E}}_{t_{n-2}}\left[\ \left(\mathbf{x_{t_{n-2}+h_{t_{n-2}}}} - \mathbf{x^*_{t_{n-2}}}\right)'\ \right]\tilde{\beta}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} + \tilde{\mathbb{E}}_{t_{n-1}}\left[\ \lambda^{h_{t_{n-1}}}_{t_{n-1}}\ \right]\end{aligned} \tag{56}$$

Assumption 5. *The coefficients in Equation 56 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:

$$\tilde{\beta}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} = \tilde{\beta} + \tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} \tag{57}$$

where $\tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} \in \mathbb{R}^k$ is a vector of macro-response coefficient shocks, where $\tilde{\upsilon}^{h_{t_n}}_{t_n} \sim i.i.d.\left(0, \sigma^{\mathbf{2}}_{\tilde{\upsilon}}\right)$.[32] 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**, $\tilde{\mathbb{E}}_{t_n}\left[\tilde{p}_{t_n+h_{t_n}}\right]$ (left-hand side of Eq. 52), which is time-contingent FG, i.e., the FOMC Fed Funds Rate (FFR) projections released quarterly in The Summary of Economic Projections (SEP).

[32] 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).[33]

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

$$\Delta_n \tilde{\beta}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} = \tilde{\beta}_{\mathbf{t_n}} - \tilde{\beta}_{\mathbf{t_{n-1}}} = \left(\tilde{\beta} + \tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}}\right) - \left(\tilde{\beta} + \tilde{\upsilon}^{\mathbf{h_{t_{n-1}}}}_{\mathbf{t_{n-1}}}\right) = \Delta_n \tilde{\upsilon}^{\mathbf{h_{t_n}}}_{\mathbf{t_n}} \tag{58}$$

(from Eq. 52 and 57), 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).[34]

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).[35]

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.[36] 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.[37] Although new macroeconomic data

[33] $\tilde{\lambda}^{h_{t_n}}_{t_n}$ 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).

[34] 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.

[35] 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.

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

[37] 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 $\Delta_n \tilde{\mathbb{E}}_{t_n}\left[\mathbf{x}_{\mathbf{t_n}+\mathbf{h}_{\mathbf{t_n}}}\right]$, and by how much. So, the public usually cannot disentangle which of the three channels drove $\Delta_n \tilde{\mathbb{E}}_{t_n}\left[\mathbf{x}_{\mathbf{t_n}+\mathbf{h}_{\mathbf{t_n}}}\right]$. 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. 52 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.[38] 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. 52 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, $\ddot{\Omega}_{it}$.[39] We use the double dots to refer to a forecaster expectation. So, the forecaster-expected central bank reaction function is:[40]

$$\begin{aligned}&\ddot{\mathbb{E}}_{i,t+1}\left[\;\tilde{\mathbb{E}}_{t_n}\left[\;\tilde{p}_{t_n+h_{t_n}}\;\right]\;\middle|\;\ddot{\Omega}_{i,t+1}\right] = \ddot{\mathbb{E}}_{i,t+1}\left[\;\tilde{\mathbb{E}}_{t_{n-1}}\left[\;\tilde{\lambda}_{t_{n-1}+h_{t_{n-1}}}\;\right]\;\middle|\;\ddot{\Omega}_{i,t+1}\;\right] + \\ &\ddot{\mathbb{E}}_{i,t+1}\left[\;\tilde{\mathbb{E}}_{t_{n-2}}\left[\;\tilde{\varphi}^{h_{t_n}}_{t_{n-2}}(L)\tilde{p}_{t_{n-2}+h_{t_{n-2}}} + \tilde{\phi}^{\pi^*}_{t_{n-2}}\pi^*_t + \tilde{\phi}^{r^*}_{t_{n-2}}r^*_{t_{n-2}} + \left(\mathbf{x}_{\mathbf{t_{n-2}}+\mathbf{h}_{\mathbf{t_{n-2}}}} - \mathbf{x}^*_{\mathbf{t_{n-2}}}\right)'\tilde{\beta}^{\mathbf{h}_{\mathbf{t_{n-2}}}}_{\mathbf{t_{n-2}}}\;\right]\;\middle|\;\ddot{\Omega}_{i,t+1}\;\right]\end{aligned} \tag{59}$$

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

[39] 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. 60 and write $\ddot{\mathbb{E}}_{i,t+1}\left[\tilde{p}_{t+h} \;\middle|\; \ddot{\Omega}_{i,t+1}\right]$ 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.

[40] 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.[41] And so we can simplify the notation in Eq. 59 as follows:

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \tilde{p}_{t+h_t}\ \right]\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \tilde{\varphi}_t^{h_t}(L)\tilde{p}_{t+h_t} + \tilde{\phi}_t^{\pi^*}\pi_t^* + \tilde{\phi}_t^{r^*}r_t^* + \left(\mathbf{x_{t+h_t}} - \mathbf{x_t^*}\right)'\tilde{\boldsymbol{\beta}}_{\mathbf{t}}^{\mathbf{h_t}}\ \right]\ \right] + \ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \tilde{\lambda}_{t+h_t}\ \right]\ \right] \tag{60}$$

Forecasters are assumed to form their projections and to update their information sets in a manner consistent with the FOMC (Equations 42 to 57), which will be important to bear in mind as we consider forecaster-expected FOMC judgment later on (Eq. 63).

Assumption 6. *The coefficients are known given* $\ddot{\Omega}_{it}$ *($\ddot{\Omega}_{it}$–measurable). This implies that:*

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \left(\mathbf{x_{t+h_t}} - \mathbf{x_t^*}\right)'\ \right]\ \tilde{\boldsymbol{\beta}}_{\mathbf{t}}^{\mathbf{h_t}}\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \tilde{\mathbb{E}}_t\left[\ \left(\mathbf{x_{t+h_t}} - \mathbf{x_t^*}\right)'\ \right]\ \right]\ \ddot{\boldsymbol{\beta}}_{\mathbf{it}}^{\mathbf{h_t}} \tag{61}$$

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:

$$\ddot{\boldsymbol{\beta}}_{\mathbf{it}}^{\mathbf{h_t}} := \ddot{\mathbb{E}}_{it}\left[\ \tilde{\boldsymbol{\beta}}_{\mathbf{t}}^{\mathbf{h_t}}\ \right] \tag{62}$$

By the linearity of the expectation operator, individuals can form a separate expectation for each component of Eq. 60. 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).

[41] 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): $\ddot{\varphi}^{\ell}_{it}$, $\ddot{\beta}^{\mathbf{h}}_{\mathbf{it}}$, and $\ddot{\varepsilon}^{h}_{it}$, "*or because it changes their forecast of economic conditions*" (the information channel): $\ddot{\mathbb{E}}_{it}\Big[\ \tilde{\mathbb{E}}_t\Big[\ \Big(\mathbf{x_{t+h}} - \mathbf{x^*_t}\Big)'\ \Big]\ \Big]$.

$\ddot{\mathbb{E}}_{it}\Big[\ \tilde{\varphi}^{h_t}_t(L)\ \tilde{\mathbb{E}}_t[\ \tilde{p}_{t+h_t}\ ]\ \Big]$ 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: $\ddot{\varphi}^{h_t}_{it}(L) := \ddot{\mathbb{E}}_{it}\Big[\ \tilde{\varphi}^{h_t}_t(L)\ \Big]$.

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

$\ddot{\mathbb{E}}_{it}\Big[\ \tilde{\mathbb{E}}_t\Big[\ \Big(\mathbf{x_{t+h_t}} - \mathbf{x^*_t}\Big)'\ \Big]\ \Big]$ 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).

$\ddot{\mathbb{E}}_{it}\Big[\ \tilde{\beta}^{\mathbf{h_t}}_{\mathbf{t}}\ \Big] := \ddot{\beta}^{\mathbf{h_t}}_{\mathbf{it}}$ 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.[42] 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. 47, is:

$$\ddot{\mathbb{E}}_{i,t+1}\Big[\ \tilde{\mathbb{E}}_{t_n}\Big[\ \tilde{\lambda}_{t_n+h_{t_n}}\ \Big]\ \Big] = \ddot{\mathbb{E}}_{i,t+1}\Big[\ \tilde{\mathbb{E}}_{t_n}\Big[\ \tilde{\rho}_{t_n}\tilde{\mathbb{E}}_{t_{n-2}}\Big[\ \tilde{\lambda}_{t_{n-2}+h_{t_{n-2}}}\ \Big] + \tilde{\varepsilon}^{h_{t_{n-1}}}_{t_{n-1}}\ \Big]\ \Big] \tag{63}$$

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

$$\ddot{\mathbb{E}}_{i,t+1}\Big[\ \tilde{\lambda}^{h_{t_n}}_{t_n}\ \Big] = \ddot{\mathbb{E}}_{i,t+1}\Big[\ \tilde{\rho}_{t_n}\tilde{\lambda}^{h_{t_{n-2}}}_{t_{n-2}} + \tilde{\varepsilon}^{h_{t_{n-1}}}_{t_{n-1}}\ \Big] \tag{64}$$

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

[42] 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:[43]

$$\tilde{\lambda}^{h}_{t_n} = \tilde{\rho}_{t_n}\tilde{\lambda}^{h}_{t_{n-2}} + \tilde{\varepsilon}^{h}_{t_{n-1}} \tag{65}$$

$$= \tilde{\rho}_{t_n}\left(\tilde{\rho}_{t_{n-2}}\tilde{\lambda}^{h}_{t_{n-4}} + \tilde{\varepsilon}^{h}_{t_{n-3}}\right) + \tilde{\varepsilon}^{h}_{t_{n-1}} \tag{66}$$

$$= \tilde{\rho}_{t_n}\tilde{\rho}_{t_{n-2}}\tilde{\lambda}^{h}_{t_{n-4}} + \tilde{\rho}_{t_n}\tilde{\varepsilon}^{h}_{t_{n-3}} + \tilde{\varepsilon}^{h}_{t_{n-1}} \tag{67}$$

$$\vdots$$

$$= \bar{\rho}_{t_n}\tilde{\lambda}^{h}_{t_{n-2R}} + \sum_{r=1}^{R}\tilde{\rho}^{r-1}\tilde{\varepsilon}^{h}_{t_{n-(2r-1)}} \quad \text{where} \quad \bar{\rho}_{t_n} := \tilde{\rho}_{t_n}\cdot\tilde{\rho}_{t_{n-2}}\cdot\ldots \tag{68}$$

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

$$\tilde{\lambda}^{h_t}_{t} = \bar{\rho}_t\tilde{\lambda}^{h_{t-1}}_{t-1} + \bar{\varepsilon}^{h_t}_{t}, \quad \text{where} \quad \bar{\rho}_t := \tilde{\rho}_{t_n}\cdot\tilde{\rho}_{t_{n-2}}\cdot\ldots, \quad \bar{\varepsilon}^{h_t}_{t} := \sum_{r=1}^{R}\tilde{\rho}^{r-1}\tilde{\varepsilon}^{h}_{t_{n-(2r-1)}} \tag{69}$$

Now we take expectations with respect to $\ddot{\Omega}_{it}$ such that t and t_n (from Eq. 69) coincide:

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\lambda}^{h_t}_{t}\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \bar{\rho}_t\tilde{\lambda}^{h_{t-1}}_{t-1} + \bar{\varepsilon}^{h_t}_{t}\ \right] \Rightarrow \tag{70}$$

$$\ddot{\mathbb{E}}_{it}\left[\ \tilde{\lambda}^{h_t}_{t}\ \right] = \ddot{\rho}_{it}\ddot{\mathbb{E}}_{it}\left[\ \tilde{\lambda}^{h_{t-1}}_{t-1}\ \right] + \ddot{\mathbb{E}}_{it}\left[\ \bar{\varepsilon}^{h_t}_{t}\ \right] \tag{71}$$

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

$$\ddot{\lambda}^{h_t}_{it} = \ddot{\rho}_{it}\ddot{\lambda}^{h_{t-1}}_{i,t-1} + \ddot{\varepsilon}^{h_t}_{it} \tag{72}$$

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. $\ddot{\varepsilon}^{h_t}_{it}$ **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 $\ddot{\varepsilon}^{h_t}_{it}$ 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**:

$$\ddot{\varepsilon}^{h_t}_{it} := \ddot{\mathbb{E}}_{it}\left[\ \bar{\varepsilon}^{h_t}_{t}\ \right] = \ddot{\mathbb{E}}_{it}\left[\ \sum_{r=1}^{R}\tilde{\rho}^{r-1}\tilde{\varepsilon}^{h}_{t_{n-(2r-1)}}\ \middle|\ \tilde{\Omega}_{it} = \tilde{\omega}_{it}\ \right] \tag{73}$$

Interpretation of the Instrument. The estimated forecaster-perceived monetary policy shocks, $\ddot{\varepsilon}^{h_t}_{it}$, 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, $\ddot{\varepsilon}^{h_t}_{it}$ 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

[43] 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, $\sigma(\tilde{\Omega}_{it})$. 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 7. ***Common Effects***. $\mathbf{f_t}$ *is an m x* 1 *vector of unobserved common factors, which can be interpreted as structural shocks (Giacomini, Lu, and Smetanina (2024)). Here,* $m = 3$*, and the three structural shocks are a supply shock, demand shock, and a perceived FOMC projection error shock,* f^e_{it}*.* $\mathbf{f_t}$ *is covariance stationary with absolute summable autocovariances, and is distributed independently of the individual-specific errors and* $\mathbf{v^h_{it}}$ $\forall i, t$ *and* t'*.*

Assumption 8. ***Factor Loadings***. *The factor loadings* $\gamma_\mathbf{i}$ *and* $\mathbf{\Gamma_i}$ *are independently and identically distributed across i, and of the common factors,* $\mathbf{f_t}$*, for all i and t, with mean* γ *and* $\mathbf{\Gamma}$*, respectively.*

Assumption 9. ***Weakly Exogenous Regressors***. $\mathbf{\ddot{x}^h_{it}}$ *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.[44]

[44]The corollary is that there may be omitted variable bias due to omitted long-term inflation expectations, but

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.

Assumption 10. ***Individual-Specific Errors***. *The individual specific errors, $\ddot{\varepsilon}_{it}$, 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, $\ddot{\boldsymbol{\varepsilon}}_{\mathbf{t}} = \{\ddot{\varepsilon}_{1t}, \ddot{\varepsilon}_{2t}, \ldots, \ddot{\varepsilon}_{Nt}\}$ is spatially correlated.*

To sum up, we are left with a structural equation representing forecasters' perceived reaction functions. At this point, we can also abstract from the double-dot notation, as all data and parameters with an i subscript pertain to individual forecasters.

$$y^h_{it} = \varphi^h_{it}\ddot{y}^h_{i,t-1} + \phi^{\pi^*}_{it}\pi^*_t + \phi^{r^*}_{it}r^*_t + \left(\ddot{\mathbf{x}}^{\mathbf{h}}_{\mathbf{it}}\right)' \beta^{\mathbf{h}}_{\mathbf{it}} + \lambda^h_{it} + \gamma^h_i \Delta f^e_t \tag{74}$$

where both λ^h_{it} and $\gamma^h_i \Delta f^e_t$ are unobservable and therefore is absorbed by the residual. We also allow for the possibility that other, outside latent factors are present in the data, which could be proxied by other forecasts by forecaster i at time t.

we observe that, in the data, the (lower frequency, quarterly) longer-term inflation expectations do not vary much.