Estimating Impulse Responses to Bundled Forward Guidance

Paul Bousquet*

August 28, 2026
[Click here for latest version]

Abstract

When central banks announce an interest rate decision, it's typically alongside forward guidance, communication informative about future interest rate decisions not already implicit in the rate set today. Standard specifications to estimate the effects of monetary policy only capture a slice of this multi-dimensional action. This paper formalizes how to modify impulse response estimation to jointly account for all effects. This approach reduces estimation uncertainty, allows macroeconomic responses to be conditioned on different shapes for the expected path of policy, and is straightforward to implement. An application to impulse response matching demonstrates its potential to strengthen identification of deep parameters.

*Department of Economics, University of Virginia, pbousquet@virginia.edu. I especially thank Christian Wolf for helpful conversations.

1 Introduction

A large literature studies the effects of monetary policy using high-frequency identification, instrumenting endogenous interest rates with measured changes in the nominal yield curve in a narrow window around central bank announcements. When leads of macro time series are regressed on these instruments, intuition for what the responses should resemble typically is derived from one-time impulses to the intercept of a monetary policy rule, as in Christiano et al. (2005); Smets and Wouters (2007). Adams and Barrett (2025) point out this is the wrong mental model— central bank announcements often contain forward guidance, providing news relevant for how future interest rates will be set not already implicit in the interest rate set today. Just as the issue of "multiple treatments" in microeconometrics, care must be taken to ensure the effects of multi-dimensional policy announcements are being fully accounted for.

This paper shows how to account for the sum total of forward guidance in announcements by combining high-frequency identification with the paradigm of McKay and Wolf (2023); Manuel and Wolf (2026), who show simultaneous treatments can always be summarized by the expected path of interest rates they collectively induce.[1] The McKay and Wolf (2023) approach is typically implemented by regressing leads of interest rates $\{i_{t+h}\}_{h=1}^{H}$ on a valid instrument z_t to get the policy path. The high-frequency literature uses interest rate futures markets to construct instruments from a measurement of $z_t^h = \Delta\mathbb{E}_t[i_{t+h}]$. The main conceptual cross-application is to recognize that under standard assumptions, the collection of instruments $\boldsymbol{p}_t = \{z_t^h\}_{h=1}^{H}$ is *itself* the policy impulse response. We can also condition directly on these paths because they are fit by a two-dimensional factor model with essentially no estimation error. So in a large class of linear models, these factors are sufficient to estimate the total effects on the macroeconomy. Because the strong factor fit was recognized early in the high-frequency literature by Gürkaynak et al. (2005), implementation requires minimal departures from current practices, mostly amounting to using factor loadings to represent the conditioned policy path and factor values to flexibly vary its shape.[2] An application to impulse response matching demonstrates how this approach tractably compresses a family of responses for the model to match, delivering stronger identification of deep parameters.

[1] I also focus on "conventional" forward guidance occurring away from the zero-lower bound; see Cieslak and Khorrami (2026) for a discussion of why the high-frequency approach may not be appropriate to study unconventional monetary policy.

[2] One recommended change is to not rotate the factors, as the rotation angle varies substantially over time. The unrotated factor model is stable, relatively interpretable, and can represent announcements containing no forward guidance without rotation.

2 Sketching the Problem (and Solution)

Suppose that after a central bank announcement, changes in futures market pricing imply the expectation of interest rates 12 months ahead increased 25 basis points. In the status quo, the central bank almost surely did not commit to unconditionally pegging interest rates 25 bp higher than the previous implied expectation. Instead, consider an inertial Taylor (1993) rule

$$i_t = \rho i_{t-1} + \phi \pi_t + \nu_t. \tag{2.1}$$

An impulse to ν_t affects expectations of i_{t+12} through the persistence of interest rates and the effect of ν_t on economic conditions, affecting π_{t+12}. Interest rate expectations can also shift from forward guidance, news about future interest rate decisions (e.g., ν_t may also contain previous period's news shocks). Clearly, these measured yield curve changes do not correspond to a primitive process in a structural model (Adams and Barrett, 2025). So when we observe $z_t^{12} \equiv \Delta\mathbb{E}_t[i_{t+12}] = .25$, how can we disentangle these forces to study the effects of policy on the economy? Put differently, impulse responses are typically estimated by regressing leads of K different low-frequency macro time series $\{\{y_{k,t+h}\}_{k=1}^K\}_{h=1}^H$ on a scalar instrument z_t, producing a collection of estimands $\{\{\alpha_h^k\}_{k=1}^K\}_{h=1}^H$. When we use z_t^{12} or $z_t^h = \Delta\mathbb{E}_t[i_{t+h}]$ as an instrument, what are we estimating?

To fix ideas, we follow McKay and Wolf (2023); Manuel and Wolf (2026) and begin with a standard sequence space economy, taking the following form under their assumptions

$$\begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{i} \end{pmatrix} = \underbrace{- \begin{pmatrix} H_x & H_i \\ A_x & A_i \end{pmatrix}^{-1} \begin{pmatrix} H_\eta & 0 \\ 0 & I \end{pmatrix}}_{\equiv \Theta} \begin{pmatrix} \boldsymbol{\eta} \\ \boldsymbol{\nu} \end{pmatrix}, \qquad \Theta = \begin{pmatrix} \Theta_{x,\eta} & \Theta_{x,\nu} \\ \Theta_{i,\eta} & \Theta_{i,\nu} \end{pmatrix}. \tag{2.2}$$

Here, $\boldsymbol{x}$ collects private-sector aggregates, $\boldsymbol{i}$ is the policy vector, $\boldsymbol{\eta}$ are non-policy shocks, and $\boldsymbol{\nu}$ are current and news shocks to the policy rule. Setting $\boldsymbol{\eta} = 0$ yields a policy causal effect basis

$$\Theta_{x,i} = \Theta_{x,\nu} \cdot \Theta_{i,\nu}^{-1}. \tag{2.3}$$

So over the span of $\boldsymbol{\nu}$, we have a universe of policy causal effects. The earlier $\{\{\alpha_h^k\}_{k=1}^K\}_{h=1}^H$ recover an element of this span. If we let the $k = 1$ time series be interest rates i_{t+h}, $\{\{\alpha_h^k\}_{k=2}^K\}_{h=1}^H$ are the macroeconomic responses conditioned on a policy path $\{\{\alpha_h^1\}\}_{h=1}^H$. For high-frequency instruments, the policy path selected, and therefore the macroeconomic responses, depend on which z_t^h was used. So while these exercises do summarize all the general equilibrium effects associated with a change in z_t^h, they do not recover the effects of the policy path induced by the announcement itself, which is the collective $\boldsymbol{p}_t = \{z_t^h\}_{h=1}^H$. Appropriately combining the information from the individual z_t^h estimates to construct the response to $\boldsymbol{p}_t$ requires the variance-covariance matrix across all estimated parameters, which will not be invertible in practice.[3]

The main idea is the following. We assume some primitive impulse occurred in a high-frequency window. Under standard identifying assumptions, the observed $\boldsymbol{p}_t$ is exactly the impulse response of interest rates to this primitive shock and, in many data generating processes like (2.2), tells us all we need to know to estimate the response of low-frequency macroeconomic aggregates.[4] To apply this to McKay and Wolf (2023), we can exploit that $\boldsymbol{p}_t$ empirically admits a precisely-estimated, low dimensional factor structure. This allows for the policy path being conditioned on to be visualized directly, with minimal estimation involved, and for the system of responses to be conditioned on policy paths that are not equivalent up to scale. Section 3 formalizes how to implement this result and the differences from other approaches.

3 Estimating Impulse Responses

3.1 Procedure

We typically model a monetary policy shock as an impulse to the intercept of a fixed interest rate rule. We can instead consider a primitive impulse with general state-dependent effects on the term structure of interest rates. Specifically, let $\boldsymbol{p}_t \equiv (\Delta\mathbb{E}_t[i_{t+1}], \ldots, \Delta\mathbb{E}_t[i_{t+H}])' \in \mathbb{R}^H$ be an induced interest rate path of length H to some primitive impulse ε_t

$$\boldsymbol{p}_t = g_t(\varepsilon_t), \tag{3.1}$$

where $g_t : \mathbb{R} \times \mathbb{R}^S \to \mathbb{R}^H$ depends on the S-dimensional state of the world at time t.[5]

Assumption 3.1 (Linearity and path sufficiency)**.** *For each horizon h, the causal effect of a monetary policy event on outcome y is linear in the event's expected policy path:*

$$y_h(\varepsilon_t) = \theta_h' \boldsymbol{p}_t. \tag{3.2}$$

Directly estimating θ_h is generally infeasible because $\boldsymbol{p}_t$ will be approximately low-dimensional,

[3] Section 3.3 discusses choosing a sparse weight matrix in the context of impulse response matching still poses challenges.
[4] The assumptions are (i) efficient markets (ii) no time-varying premia (iii) the event window only captures an impulse to ε_t.
[5] Note that ε_t need not correspond to the earlier Taylor (1993) rule shock ν_t.

so we will derive a result for the best feasible estimation. One can always decompose $\boldsymbol{p}_t$ by

$$\boldsymbol{p}_t = V f_t + r_t, \tag{3.3}$$

for some $f_t \in \mathbb{R}^F$ and $V \in \mathbb{R}^{H \times F}$ such that $\mathbb{E}[f_t f_t']$ is diagonal and $F < H$. For a scalar outcome, the horizon-h Jordà (2005) local projection with f_t is

$$y_{t+h} = b_h + \beta_h' f_t + u_{t+h}, \tag{3.4}$$

Proposition 3.1. *Under Assumption 3.1, estimation of* (3.4) *satisfies*

$$\hat{\beta}_h' f_t = \theta_h' \boldsymbol{p}_t + O_p(||r_t||) + o_p(1). \tag{3.5}$$

Proposition 3.1 formalizes a simple result: if $\boldsymbol{p}_t$ is exactly fit by the factor model, we can recover the response of y_{t+h} to observed policy paths despite θ_h not being identified.[6] In the data, the $F = 2$ factor model fits $\boldsymbol{p}_t$ well (i.e., $r_t \approx 0$). Measurements of $\boldsymbol{p}_t$ come from high-frequency changes in interest rate futures around Federal Reserve announcements from 1994-2025 (Acosta et al., 2025). The collection of contracts is FF2 and ED2-ED8, which are the futures contracts for the average federal funds rate one month later and 1 to 7 quarters later, respectively, and observations where ED2 moves less than one basis point are excluded. The first and second principal components explain 85.5 and 10.8% of contract variation, respectively, bringing the combined total to 96.3%. This suggests (3.4) is a tractable specification to estimate the effects of $\boldsymbol{p}_t$.

Next, incorporating Proposition 3.1 into the McKay and Wolf (2023) paradigm has two immediate benefits: a more precise sense of what policy path the macroeconomic responses are conditioned on (i.e., taking the policy path as essentially observed instead of estimating it using low-frequency interest rate data) and the path itself summarizes the entire policy content of announcements. Implementation is simple precisely because the differences from current practice are subtle.

Recall that the McKay and Wolf (2023) policy path is simply $\{\alpha_h^1\}_{h=1}^H$. Current practice normalizes the system of responses somewhat arbitrarily, e.g., scaling responses by c so that $c\alpha_1^1$ corresponds to an impact effect on interest rates of +100 bp. This takes the impulse response shapes as fixed and scales them up and down. Instead, (3.4) allows us to evaluate the response system at different policy path shapes that match those observed in real time. Define $\beta_h^k(f_t) = \beta_h^{k\prime} f_t$ from regressing $y_{k,t+h}$ on f_t. Then the policy path is $V f_t$ from (3.3) and the macroeconomic responses

[6] There is a related Lucas (1976) limitation on counterfactuals. Assumption 3.1 allows (3.2) to depend on $\boldsymbol{p}_t$'s data generating process (unlike (2.2)), in which case nothing can be said about the effects of a $\boldsymbol{p}$ outside the support of the fitted factor model.

are $\{\{\beta_h^k(f_t)\}_{k=2}^K\}_{h=1}^H$. By picking f_t to match something in the support of the factor model, we can observe the macroeconomic response to fundamentally different policy paths.

3.2 Existing Factor Model Approaches

Factor models are already commonly used, either by rotating f_t to facilitate economic interpretation or by using only the first factor for the simplicity of a scalar instrument. In these cases, the main novelty of combining high-frequency identification with the McKay and Wolf (2023) approach is simply through the precision and flexibility of using $\widehat{V} f_t$ for the policy path rather than $\{\widehat{\alpha}_h^1\}_{h=1}^H$. There are also practical reasons to follow Proposition 3.1 and use the raw set of all f_t.

Namely, Gürkaynak et al. (2005) find similar strong explanatory power for the yield curve up to 3 quarters out and propose rotating the factors so that the second component has no impact on the shortest end of the yield curve. Including both rotated factors on a common set of interest rate futures contracts will deliver identical paths $\{\{\beta_h^k(f_t)\}_{k=1}^K\}_{h=1}^H$, but typically both are not included to focus on different aspects of monetary policy in isolation. When including both, the rotation can still complicate inference, as Figure 1 shows it is unstable over time. In contrast, Figure 2 shows the unrotated factor model is stable. Gürkaynak et al. (2005) perform the rotation to enhance interpretability of the factors. Figure 2 shows in the unrotated case, the first component acts as roughly as a level shift while the second "tilts" the yield curve, so the second component captures situations where the central banks holds rates steady but provides forward guidance about future actions.[7] Appropriate choices of f_t can accommodate a desire to condition on different types of paths (e.g., that involve no forward guidance) without the imprecision of the rotation.

Using the first principal component of an unrotated factor model is also common, popularized by Nakamura and Steinsson (2018a). Newly available data from Acosta et al. (2025) allows for increasing we can increase the terminal horizon to 7 quarters ahead (from the standard of 3), and the second component can also be easily included in standard local projections to capture the timing effects of policy announcements (i.e., whether guidance is front or back loaded). By capturing more of the policy action, these changes in principle could help with the statistical power problem of high-frequency identification (Ramey, 2016; Nakamura and Steinsson, 2018a,b).[8]

[7]Further exercises and details are available in Online Appendix A. In particular, corrections for mismeasurement due to fixed event window lengths increase stability and explanatory power further (Tran, 2026).

[8]I show in Online Appendix A this unfortunately does not appear to be the case with currently available data.

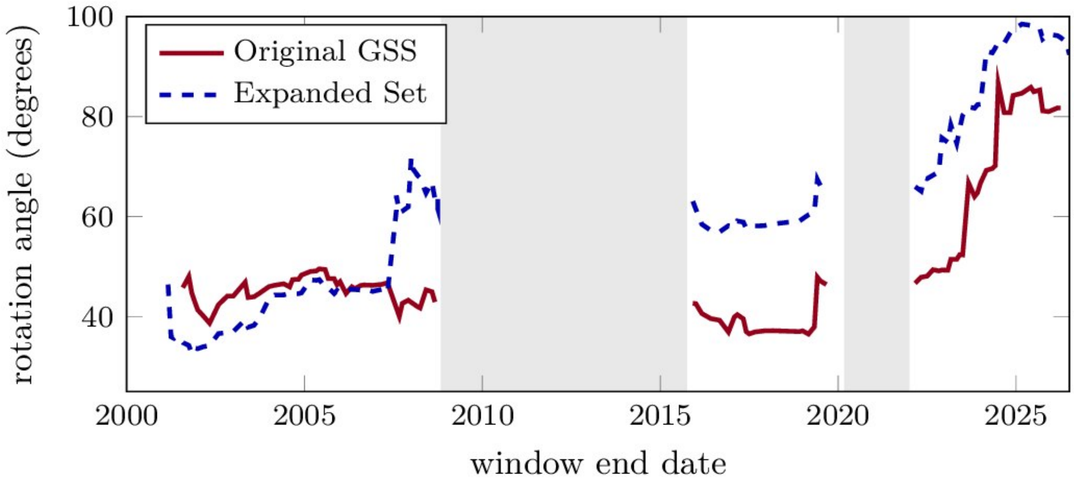

Figure 1

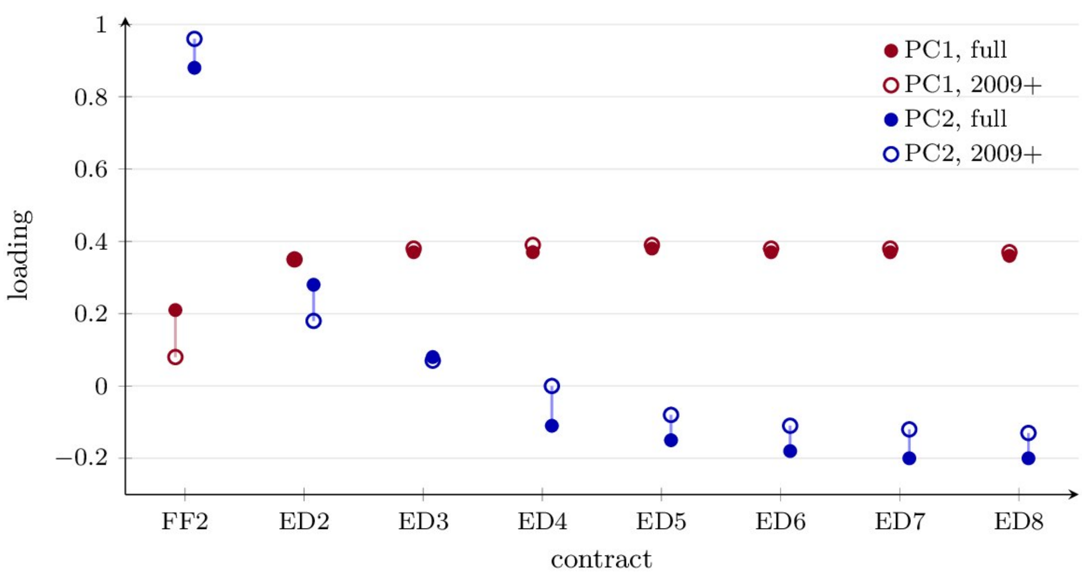

Figure 2

3.3 Application: Impulse Response Matching

Another useful application is impulse response matching; the factor model structure compresses a family of policy paths for the model to match, addressing weak-identification problems. To fix ideas, we again consider (2.2). Let ψ denote the deep parameters of the model which $\Theta = \Theta(\psi)$ depends on. For a candidate ψ, choose the matrix $N(\psi)$ of implementing shock sequences to satisfy consistency with the empirical factor loadings $\widehat{V}$ estimated from (3.3)

$$\Theta_{i,\nu}(\psi)N(\psi) = \widehat{V}. \tag{3.6}$$

Invertibility gives the following expression for model-implied macro responses

$$\Theta_{x,\nu}(\psi)N(\psi) = \Theta_{x,i}(\psi)\widehat{V}. \tag{3.7}$$

Thus, if $\widehat{B}$ stacks the estimated coefficients from (3.4), a minimum-distance estimator is

$$\widehat{\psi} = \arg\min_{\psi} \operatorname{vec}\left[\widehat{B} - \Theta_{x,i}(\psi)\widehat{V}\right]' W \operatorname{vec}\left[\widehat{B} - \Theta_{x,i}(\psi)\widehat{V}\right]. \tag{3.8}$$

Relative to matching a single policy impulse response, this extension has two benefits. As before, $\widehat{V}$ comes directly from the cross-section of announcement-window futures changes and will be more precisely estimated than the low-frequency policy impulse response to a scalar instrument. More importantly, the factor structure compresses a representation of a family of impulse responses the model should match. In population, this would be no different from stacking all the responses from individual instruments. In practice, the system is likely to be computationally unwieldy even for tractable choices for the weight matrix W (e.g., an identity matrix of dimension equal to the number of estimated parameters). In a class of linear models like (2.2), simply taking the responses to a single instrument may be sufficient to identify the parameters in population, but in practice the identification is likely to be far weaker. In Online Appendix B, I modify the replication codes of Caravello et al. (2026) to show how this procedure better identifies parameters in a standard New Keynesian model.

4 Conclusion

The proposed procedure proposed bridges the gap between the universe of possible policy paths and the policy paths we are interested in. Returning to the anecdotal illustration, suppose we observe $z_t^{12} = .25$ following an announcement. McKay and Wolf (2023) provides the tools to say "an impulse of 25 basis points to z_t^{12} on average induces this expected path of policy and these macroeconomic responses". But in real time, we observe the actual change in the expected path of policy the announcement caused, which necessarily will be different from this estimated path. This procedure allows the econometrician to use information observed in real time to make more precise statements, so long as the announcement's forward guidance, or term structure of future interest rate news, was within historical norms (Lucas, 1976). Simple modifications to common practice allow for this approach to be incorporated into common practice for impulse response estimation and matching.

References

Acosta, Miguel, Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino. 2025. "Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database." Working Paper 2025-30, Federal Reserve Bank of San Francisco.

Adams, Jonathan J., and Philip Barrett. 2025. "What Are Empirical Monetary Policy Shocks? Estimating the Term Structure of Policy News." June, Working Paper.

Caravello, Tomás E., Alisdair McKay, and Christian K. Wolf. 2026. "Evaluating Monetary Policy Counterfactuals: (When) Do We Need Structural Models?." Working paper.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Cieslak, Anna, and Paymon Khorrami. 2026. "Monetary Shocks and Long-Run Neutrality in Asset Pricing." Working paper.

Gürkaynak, Refet S., Brian Sack, and Eric 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(1): .

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95(1): 161–182.

Lucas, Jr., Robert E. 1976. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy*, 1(1): 19–46.

Manuel, Ed, and Christian K. Wolf. 2026. "Identifying Policy Causal Effects from Rule Changes." Working Paper.

McKay, Alisdair, and Christian K. Wolf. 2023. "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica*, 91(5): 1695–1725.

Nakamura, Emi, and Jón Steinsson. 2018a. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Nakamura, Emi, and Jón Steinsson. 2018b. "Identification in Macroeconomics." *Journal of Economic Perspectives*, 32(3): 59–86.

Ramey, Valerie. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics.* Chap. 2.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

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

Tran, Paul L. 2026. "How Long Do Markets Need to Fully React to Monetary Policy Announcements?", Working paper.