Online Appendix

Estimating Impulse Responses to Bundled Forward Guidance*

A. Factor Model and High-Frequency Analysis

A.1 Data

A large literature beginning with Kuttner (2001); Cochrane and Piazzesi (2002) uses changes in interest rate futures contracts around central bank announcements to measure changes in interest rate expectations. There are two types of contracts used: those tied to the federal funds rate, denoted FFm for the average rate that will prevail over $m-1$ months ahead, and those tied to the overnight lending rate offered in the international repurchase market for U.S. government debt, denoted EDq for the average rate that will prevail over $q-1$ quarters ahead. The high-frequency changes in the prices of these futures are now available since 1988 (Jarociński, 2024; Acosta et al., 2026). I use an extended set of contracts only available since 1994: FF2 and ED2-ED8.[1] I also exclude observations where the change in ED2 is below one basis point. Gürkaynak et al. (2005), a canonical paper in the high-frequency literature, uses a shorter set of contracts: MP1, MP2, ED2, ED3, and ED4. At the short end, the set uses MP1, the forecast error of the Fed's target announcement implied by interest rate futures contracts FF1 and FF2, and MP2, the futures-implied change in expectation of the next Fed target announcement; see Brennan et al. (2025) for more discussion.

A.2 Standard Factor Model

We assume there is a primitive, unpredictable impulse ε_t inside the announcement window that causes a shift in the near-term expected path of the policy rate $\boldsymbol{p}_t \equiv \big(\Delta\mathbb{E}_t[i_{t+1}], \ldots, \Delta\mathbb{E}_t[i_{t+H}]\big)'$

*The replication codes are available at `https://github.com/paulbousquet/BundledFG`

[1] Swanson and Jayawickrema (2024) also discuss that data before 1993 is less congruent than subsequent data because it was previously not immediately known the Fed had changed its federal funds target.

via the following data generating process

$$\boldsymbol{p}_t = g_t(\epsilon_t) \tag{A.1}$$

where $g_t : \mathbb{R} \times \mathbb{R}^S \to \mathbb{R}^H$ may depend on some S-dimensional state at t. Futures surprises record $\boldsymbol{p}_t$ at discrete maturities. The application for estimating impulse responses can be understood within the following context. Let $y_{t+h}(\epsilon_t)$ be the potential outcome for y at horizon h when the announcement-window impulse equals ϵ_t, and let $y_{t+h}(0)$ be the outcome under no impulse. In many linear macroeconomic models, this difference will be exactly equal to the impulse response:

$$y_{t+h}(\epsilon_t) - y_{t+h}(0) = \theta_h' \boldsymbol{p}_t. \tag{A.2}$$

In the data, we observe that realized paths concentrate in a low-dimensional subspace of $\mathbb{R}^H$. Specifically, we can estimate the following decomposition

$$\boldsymbol{p}_t = V f_t + r_t, \tag{A.3}$$

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 the identity matrix, $\mathbb{E}[f_t r_t'] = 0$, and $F < H$. Table A.1 reports the elements of V before and after 2009. The factor loadings are relatively stable, with FF2 showing the most signs of instability. Table A.2 reports the share of strip variance captured by PC1 (85.9%) and by PC1 and PC2 (96.6%). Another way to demonstrate the fact that $\boldsymbol{p}_t$ empirically admits a low-dimensional representation is to estimate a simple polynomial equation in maturity horizon. For FF2 and ED2-ED8, we can index these contracts by $h = 0, \ldots, 7$ and fit

$$\Delta \mathrm{i}_{t,h} = a_t + b_t\, h + c_t\, h^2. \tag{A.4}$$

Table A.2 places that quadratic beside the principal-component shares, showing the fit is near perfect.

What little lack of fit does remain is in part affected by measurement issues. Figure A.1 shows two particular dates with clear measurement outliers. Issues related to window length for measurement warrant further discussion (Cieslak and Khorrami, 2026; Tran, 2026); because the spanning hypothesis must be invoked with confidence, they are a first-order concern (Duffee, 2013). But for these exercises in particular, making some manual corrections didn't seem to change results materially. The other instance in which the factor model is not fit well is during the zero lower bound period, where there are a few dates that make it through the ED2 magnitude filtering.[2]

Even though the second component adds a non-trivial amount of explanatory power for the yield curve, it does not seem to add detectable explanatory power for asset prices. Table A.3 shows the explanatory power for the high-frequency S&P 500 changes is nearly non-existent. Naturally, there are similar conclusions when using low-frequency data.

A.3 Gürkaynak et al. (2005) Rotation

Gürkaynak et al. (2005) map the two principal components into a target factor τ_t and a path factor π_t. For $F = 2$, with scores $f_t = (f_{1t}, f_{2t})'$ and front-contract row $v_0' = (v_{01}, v_{02})$ of V,

$$R(\phi) = \begin{pmatrix} \cos\phi & -\sin\phi \\ \sin\phi & \cos\phi \end{pmatrix}, \qquad \widetilde{V} = VR(\phi), \qquad \begin{pmatrix} \tau_t \\ \pi_t \end{pmatrix} = R(\phi)' f_t. \tag{A.5}$$

Requiring π_t to put zero weight on the front contract pins

$$-v_{01}\sin\phi + v_{02}\cos\phi = 0, \qquad \tan\phi = \frac{v_{02}}{v_{01}}. \tag{A.6}$$

The column space of V is unchanged; only the basis inside that space rotates. Under the

[2]The replication codes have files to run the same figures and tables where the zero lower bound is excluded.

restriction, τ_t loads on the shortest contract in the strip. For the extended and Gürkaynak et al. (2005) contract sets, these are FF2 and MP1, respectively. Because ϕ is pinned to v_0, subsample movements in front-contract loadings redefine both rotated factors. Complementing Figure 3 in the main text, Table A.4 shows the rotations are unstable. Given the stability of the raw loadings, this suggests a need for sensitivity analysis to see which types of paths within the family generated by an empirical estimation of (A.3) are less precisely estimated.

Standard errors for the fit to a particular policy path, as in Figure 1 in the main text, can be computed via a simple bootstrap procedure. Each bootstrap draw rebuilds the principal component sample by sampling announcement dates with replacement: if the estimation sample has T meetings, T meetings are drawn from that list, allowing repeats. The path is then re-fit on this new sample, holding the original factor values fixed. Different sample stability tests can complement these sorts of procedures.

A.4 Tables and Figures

Table A.1: Raw PC1 and PC2 time-variation in factor loadings

Panel A: Original Gürkaynak et al. (2005) set

Component	Sample	MP1	MP2	ED2	ED3	ED4
PC1	Full	0.33	0.43	0.50	0.49	0.47
PC1	pre-2009	0.37	0.43	0.49	0.48	0.46
PC1	2009+	0.14	0.41	0.53	0.52	0.51
PC2	Full	0.76	0.38	−0.14	−0.32	−0.39
PC2	pre-2009	0.76	0.33	−0.14	−0.34	−0.42
PC2	2009+	0.82	0.44	−0.10	−0.22	−0.26

Panel B: Expanded set

Component	Sample	FF2	ED2	ED3	ED4	ED5	ED6	ED7	ED8
PC1	Full	0.21	0.35	0.37	0.37	0.38	0.37	0.37	0.37
PC1	pre-2009	0.26	0.35	0.37	0.36	0.37	0.37	0.37	0.36
PC1	2009+	0.07	0.35	0.38	0.39	0.39	0.38	0.38	0.37
PC2	Full	0.88	0.27	0.08	−0.12	−0.15	−0.18	−0.19	−0.19
PC2	pre-2009	0.84	0.30	0.08	−0.17	−0.18	−0.20	−0.22	−0.22
PC2	2009+	0.96	0.17	0.06	−0.00	−0.08	−0.10	−0.11	−0.11

Notes: Panel A: standardized PCA of MP1, MP2, and ED2–ED4 on scheduled FOMC meetings since 1994 with $|\Delta \text{ED2}| > 1$ bp. Panel B: standardized PCA of the untransformed FF2+ED2–ED8 event sample. Observations dated 2009–2014 are excluded.

Table A.2: Share of variance captured by the components and by an event-by-event quadratic

Representation	Original GSS	Expanded Set
PC1	76.0%	85.9%
PC1+PC2	93.0%	96.6%
Quadratic	—	99.2%

Notes: Full 1994–2025 sample excluding 2009–2014. PC rows are full-set standardized variance shares. Factor and quadratic fit are respectively (A.3) and (A.4).

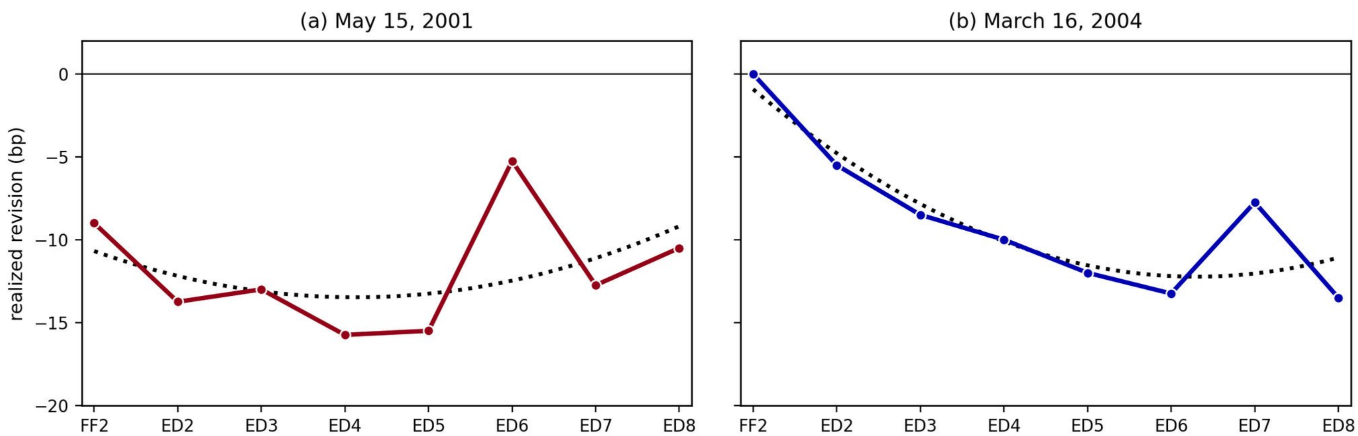

Figure A.1: Measured yield curve movements following Federal Reserved announcements (Acosta et al., 2026). The dotted line is the fit by (A.4).

Table A.3: S&P 500 returns and principal components

	(1)	(2)
PC1	-0.050^{***}	-0.051^{***}
	(0.006)	(0.006)
PC2		-0.015
		(0.018)
Observations	156	156
R^2	0.341	0.344

Notes: PC1 is scaled so that a unit increase corresponds to a one basis-point increase in ED4. PC2 is scaled by the same factor. The sample excludes observations dated 2009–2014 and trims the two largest and two smallest S&P 500 returns. All regressions include a constant. HC1 standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A.4: Rotation angle and target factor's far-end loading, before and after 2009

Set	Period	Rotation angle	Target far loading
Original GSS	pre	39.2°	0.51
Original GSS	post	74.0°	−0.02
Expanded Set	pre	45.1°	0.55
Expanded Set	post	79.1°	0.07

Notes: Original GSS is the Gürkaynak et al. (2005) contract set (MP1,MP2, ED2-ED5); Expanded set is FF2 and ED2-ED8. Scheduled Fed announcements from 1994-2025 excluding 2009–2014.

B. Impulse Response Matching

B.1 Model

The model is the behavioral representative-agent New Keynesian model of Caravello et al. (2026), which embeds Gabaix (2020) cognitive discounting into the model of Smets and Wouters (2007). Households discount the effect of news h quarters ahead by m_d^h, and price and wage setters discount their forward-looking terms by m_f. With $m_d = m_f = 1$ the model is exactly Smets and Wouters (2007) (i.e., expectations are rational). Specifically, the price and wage Phillips curves are

$$\pi_t - \iota_p \pi_{t-1} = \kappa_p \, mc_t + \beta_p \left(\pi_{t+1} - \iota_p \pi_t\right), \qquad \pi_t^w - \iota_w \pi_{t-1} = \kappa_w \, \chi_t + \beta_w \left(\pi_{t+1}^w - \iota_w \pi_t\right), \quad \text{(B.1)}$$

where π_t and π_t^w are price and wage inflation, mc_t is real marginal cost, χ_t is the gap between the household's marginal rate of substitution and the after-tax real wage, and $\beta_j = \beta \xi_j m_f \left[1 + \kappa_j / (1 - \beta \xi_j m_f)\right]$ for $j \in \{p, w\}$. The investment block is given by

$$q_t \;=\; m_d \Big[-\left(i_t^n - \pi_{t+1}\right) + \tfrac{1-\delta}{1+r} \, q_{t+1} + \left(1 - \tfrac{1-\delta}{1+r}\right) r_{t+1}^k \Big] \qquad q_t \;=\; \kappa \Big[\left(1 + \tfrac{1}{1+r}\right) i_t - \tfrac{m_d}{1+r} \, i_{t+1} - i_{t-1} \Big], \quad \text{(B.2)}$$

where q_t is the price of capital, i_t^n is the nominal interest rate, r_t^k is the return on capital, and i_t is investment.

In this exercise, nine parameters are estimated using impulse response matching. Table B.1 lists them with the values used to generate the data below; the remaining parameters and solution method are taken directly from the replication codes of Caravello et al. (2026).

Table B.1: Estimated parameters and the values used to generate the data

Parameter	Description	Value
h	habit in consumption	0.75
ξ_p	Calvo price parameter	0.85
ι_p	price indexation parameter	0.9
ξ_w	Calvo wage parameter	0.9
ι_w	wage indexation parameter	0.9
κ	investment adjustment cost	5.5
φ	Frisch elasticity	0.5
m_d	cognitive discounting, households	0.65
m_f	cognitive discounting, price and wage setters	0.65

Notes: Parameters rounded to values near replication codes of Caravello et al. (2026) with the exception of the indexation parameters (set away from corner of .99) and setting $m_d = m_f$.

B.2 Policy experiments

Monetary policy in the model follows

$$i_t^n = \rho\, i_{t-1}^n + (1-\rho)\big[\phi_\pi \pi_t + \phi_y y_t + \phi_{\Delta y}(y_t - y_{t-1}) + \overline{\nu}_t\big], \tag{B.3}$$

Following McKay and Wolf (2023); Manuel and Wolf (2026), the monetary policy shock term is equal to $\overline{\nu}_t = \sum_{h=0}^{H} \nu_{t-h,t}$, where $\nu_{t,t+h}$ denotes a shock to nominal interest rates at $t+h$ revealed at t. Monetary shock "announcements" at t are therefore a vector $\boldsymbol{\nu}_t = (\nu_{t,t}, \dots, \nu_{t,t+H})'$.

For clarity, we will repeat some derivations from the main text. Let θ index all structural parameters in the model. Stacking expected responses over horizons $0, \dots, H$, the linearized model admits the sequence-space representation

$$\begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{i} \end{pmatrix} = \underbrace{- \begin{pmatrix} G_x(\theta) & G_i(\theta) \\ A_x(\theta) & A_i(\theta) \end{pmatrix}^{-1} \begin{pmatrix} G_\eta(\theta) & 0 \\ 0 & I \end{pmatrix}}_{\equiv \Theta(\theta)} \begin{pmatrix} \boldsymbol{\eta} \\ \boldsymbol{\nu} \end{pmatrix}, \qquad \Theta(\theta) = \begin{pmatrix} \Theta_{x,\eta}(\theta) & \Theta_{x,\nu}(\theta) \\ \Theta_{i,\eta}(\theta) & \Theta_{i,\nu}(\theta) \end{pmatrix}. \tag{B.4}$$

Here, $\boldsymbol{x}$ stacks the responses of private-sector aggregates, including inflation and output, $\boldsymbol{i}$ is the nominal policy-rate response, and $\boldsymbol{\eta}$ collects non-policy shocks.

Setting $\boldsymbol{\eta} = 0$, the blocks $\Theta_{i,\nu}(\theta)$ and $\Theta_{x,\nu}(\theta)$ map announcements into policy-rate and private-sector responses, respectively. Their column $s+1$ contains the responses to a unit shock $\nu_{t,t+s}$, announced at t and entering the rule at $t+s$. Assuming invertability, the mapping from policy-rate paths to private-sector responses is

$$\Theta_{x,i}(\theta) \equiv \Theta_{x,\nu}(\theta)\Theta_{i,\nu}(\theta)^{-1}. \tag{B.5}$$

A policy experiment is defined by a fixed expected policy-rate path $\boldsymbol{p} = (p_0, \ldots, p_H)'$. At each candidate θ, choose the announcement vector that implements this path:

$$\boldsymbol{\nu}(\theta) = \Theta_{i,\nu}(\theta)^{-1}\boldsymbol{p}. \tag{B.6}$$

The corresponding private-sector responses are

$$\boldsymbol{x}(\theta) = \Theta_{x,\nu}(\theta)\boldsymbol{\nu}(\theta) = \Theta_{x,i}(\theta)\boldsymbol{p}. \tag{B.7}$$

Note that $\boldsymbol{p}$ is held fixed while varying candidate parameter values.

B.3 Impulse Response Targets

For a candidate θ and a fixed $V \in \mathbb{R}^{H \times F}$, choose the matrix $N(\theta)$ of implementing shock sequences to satisfy

$$\Theta_{i,\nu}(\psi) N(\psi) = V. \tag{B.8}$$

Invertibility gives the following expression for model-implied macro responses

$$\Theta_{x,\nu}(\theta) N(\theta) = \Theta_{x,i}(\theta) V. \tag{B.9}$$

Let $K = \dim(\boldsymbol{x})$. A minimum-distance estimator for a fixed $B \in \mathbb{R}^{K \times H \times F}$ is

$$\widehat{\theta} = \arg\min_{\theta} \operatorname{vec}[B - \Theta_{x,i}(\theta)V]' \, W \operatorname{vec}[B - \Theta_{x,i}(\theta)V] \, . \tag{B.10}$$

In the conventional scalar-instrument application, B contains coefficients from regressions of leads of macroeconomic variables on an instrument z_t, and V contains coefficients from regressions of leads of the policy rate on the same instrument. Allowing multiple columns generalizes this procedure to multiple policy experiments. The application in the main text uses $F = 2$: V is estimated from a two-factor model of announcement-window changes in interest rate futures, and B contains the coefficients from regressions of macroeconomic variables on both factors.

We compare two policy experiments through two choices of V. The factor specification uses V_F, fixed at the empirically estimated loadings described in Appendix A. The scalar specification uses $V_S = V_F c$, with $c'c = 1$. Here, we pick c so that V_2 corresponds to the scalar policy impulse path associated with a unit change in interest rates three quarters ahead. The empirical analog of this choice would be using the ED4 futures contract as an instrument. The macro responses are for inflation and output ($K = 2$).

B.4 Results

In a linear model, identification is exact in population. But in practice, identification strength will affect estimation. The main body of the paper claims using V_F instead of V_S will generally strengthen identification of structural parameters θ. The formal statement is the Jacobian of (B.9) is more sensitive to changes in parameters. This affects identification because it is more difficult for changes in parameters to reproduce paths induced by V_F than the scalar path induced by V_S. A parameter change may alter inflation and output responses substantially, yet leave the model weakly identified if changes in other parameters can reproduce those effects. Figure B.1 demonstrates this exercise by comparing how the fit of (B.9) deteriorates with a minor perturbation in each parameter. The fit is benchmarked by re-computing (B.10) while holding the given parameter at its perturbed value and seeing how well the paths still fit. For several of the parameters, the deterioration in fit is roughly an order of magnitude larger when using V_F compared to V_S. Figure B.2 shows this visualization while varying the household cognitive discounting parameter, showing loss is much more sensitive for V_F.

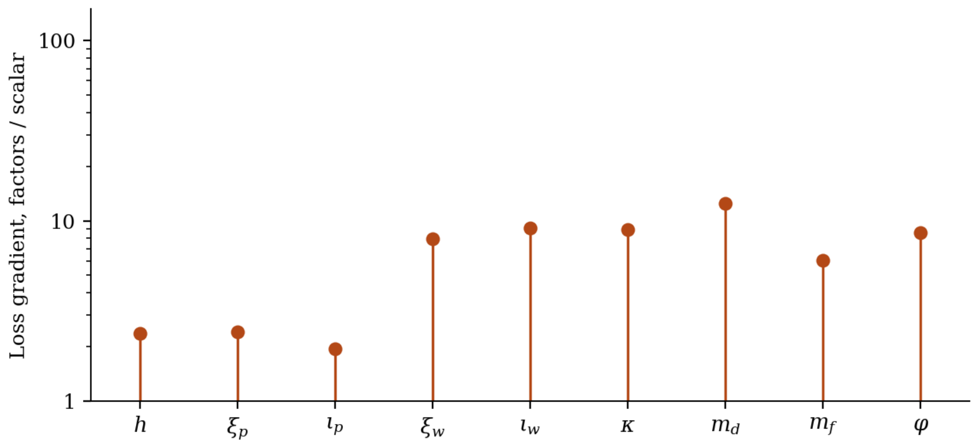

Figure B.1: Sensitivity of fit to each structural parameter, factor paths relative to the scalar path. Each bar is the curvature of the minimum-distance loss (B.10) with respect to the named parameter under V_F, divided by the same quantity under V_S, with the other parameters re-optimized at every perturbation.

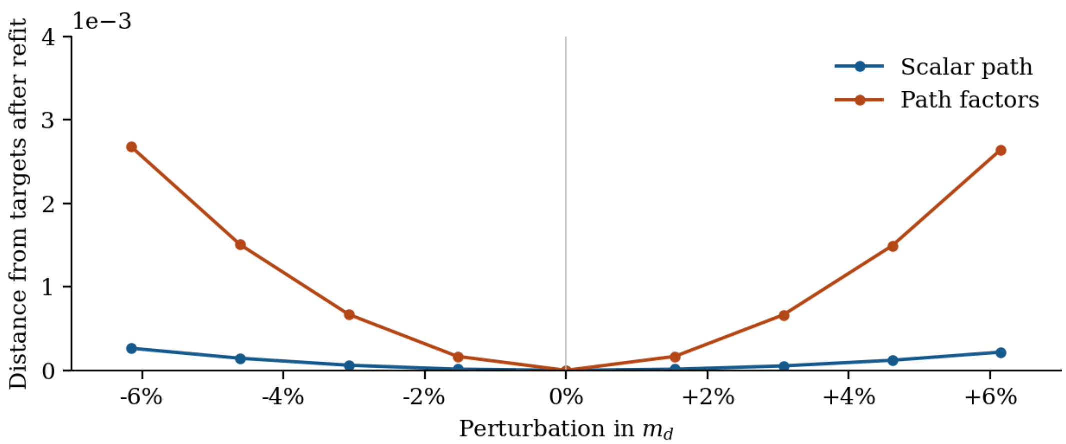

Figure B.2: Minimum-distance loss as the household cognitive discounting parameter m_d is moved away from its true value, under the factor paths V_F and the scalar path V_S. At each point the remaining parameters are re-estimated, so the curves show only the misfit that other parameters cannot absorb.

References

Acosta, Miguel, Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino. 2026. "Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database." working paper.

Brennan, Connor M., Margaret M. Jacobson, Christian Matthes, and Todd B. Walker. 2025. "Monetary Policy Shocks: Data or Methods?", 25(2): 595–659.

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

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

Cochrane, John H., and Monika Piazzesi. 2002. "The Fed and Interest Rates - A High-Frequency Identification." *American Economic Review*, 92(2): 90–95.

Duffee, Gregory. 2013. "Forecasting Interest Rates." In *Handbook of Economic Forecasting*. eds. by Graham Elliott, and Allan Timmermann, 2, Chap. 7 385–426.

Gabaix, Xavier. 2020. "A Behavioral New Keynesian Model." *American Economic Review*, 110(8): 2271–2327.

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): .

Jarociński, Marek. 2024. "Estimating the Fed's unconventional policy shocks." *Journal of Monetary Economics*, 144, p. 103548.

Kuttner, Kenneth N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics*, 47(3): 523–544.

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.

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

Swanson, Eric T., and Vishuddhi Jayawickrema. 2024. "Speeches by the Fed Chair Are More Important Than FOMC Announcements: An Improved High-Frequency Measure of U.S. Monetary Policy Shocks." Working paper.

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