Housing Finance and the Dynamic Transmission of Monetary Policy

Han Zhang*

University of California, Riverside

Job market paper — November 2026

Abstract

Monetary policy reaches consumption through housing finance along several margins, and refinancing activity is the one most often read as its barometer. Using a Bayesian structural VAR with stochastic volatility on monthly U.S. data for 1998–2023, I separate housing valuation, home-purchase financing, and rate-and-term refinancing and trace each to consumption with a first-entry path decomposition. Under a common persistent rise in mortgage rates, refinancing contracts within six months by an amount equivalent to 6% of its mean monthly intensity and then fades, while purchase financing declines gradually, with a three-year contraction equivalent to 4% of its mean monthly intensity. Yet refinancing-entry paths contribute least to consumption; valuation- and purchase-entry paths contribute most and keep accumulating, and the three together account for a −0.16% three-year consumption-level contribution, partly offset by paths that bypass housing finance. The margin that moves most is not the margin that carries consumption.

JEL codes: E52, E44, G21, R31

Keywords: monetary policy transmission, housing finance, refinancing, structural VAR, stochastic volatility

*I am grateful to Marcelle Chauvet for her guidance throughout this project, and to Jana Grittersová and Dongwon Lee for helpful comments. All errors are my own. Email: han.zhang@email.ucr.edu.

1. Introduction

Housing finance transmits monetary policy to household spending through more than one decision. A change in financing conditions revalues the housing stock, alters the terms on which a home is bought, and changes the payoff from replacing an existing mortgage. These margins do not operate on the same clock. Valuation moves without anyone transacting; a purchase involves search, matching, and closing; a refinance is a discrete choice that depends on the gap between new and existing contract rates. The 2022 tightening made the distinction visible: as mortgage rates rose, refinancing applications collapsed—the Mortgage Bankers Association's Refinance Index fell 86 percent year over year by early December 2022 (Mortgage Bankers Association, 2022). A single mortgage-origination series would conflate these two adjustments. The question this paper asks is how each margin responds to a common policy impulse, and how much of the eventual consumption response each one carries as transmission unfolds. The answer is that the margin that moves most is not the margin that carries consumption: refinancing activity, the usual barometer of mortgage-market transmission, responds fastest and most visibly but contributes least to consumption, while valuation and purchase financing, whose own responses are slower or less precisely estimated, carry most of it.

I answer it with an estimated monthly system for January 1998–December 2023. Purchase and non-cash-out refinance originations from FHFA's National Mortgage Database enter separately, each scaled by the lagged mortgage stock; real FHFA house-price growth measures valuation. These three series are modeled jointly with consumption, industrial production, core inflation, the one-year Treasury yield, and the thirty-year mortgage rate in a Bayesian structural VAR with twelve lags and a separate stochastic-volatility process for each structural innovation. Contemporaneous zero restrictions isolate a policy innovation, scaled to raise the one-year yield by 25 basis points on impact, while leaving mortgage pricing and financing quantities free to feed back on each other within the month. To compare the margins in common units, I apply the transmission-channel framework of Wegner et al. (2025) and allocate every propagation path to the housing-finance variable it enters first; paths that never enter these three variables form the complement.

The first finding is a timing asymmetry between the two financing margins. The policy innovation raises the mortgage rate by about 9 basis points on impact and 18 after a year, and the rate is still elevated after three years. Under this common pricing path, refinancing contracts quickly: six months after the shock its decline is equivalent to nearly 6% of its sample mean monthly intensity, the response is below zero at the 90% level from month 1 through month 17, and it then returns toward zero. Purchase financing contracts slowly: its decline is equivalent to about 1% of its sample mean monthly intensity at six months and about 4% at three years, and it is below zero

from month 3 through month 36. The pattern survives replacing the FHFA index with Case-Shiller, allowing the mortgage rate to enter the policy equation contemporaneously in either direction, and restricting the posterior to stable states. Real house-price growth turns negative within the first year and accumulates to a level gap of about one percent after three years.

The second finding is the paper's central one: in the posterior medians, refinancing-entry paths contribute the least to consumption over three years and valuation-entry paths the most, with purchase-entry paths building to second place—the reverse of what the financing responses alone would suggest. A fall in refinancing activity is not itself a consumption effect, and a modest valuation response can be carried a long way by subsequent collateral and financing feedback. The first-entry accounting makes the margins comparable on a consumption scale while retaining their interactions: a valuation–refinancing–consumption path belongs to valuation, with its refinancing leg kept. The resulting composition shifts with the horizon. Valuation-entry contributions turn negative early and keep accumulating; purchase-entry contributions start near zero and build over longer windows; refinancing-entry contributions are comparatively even through the middle of the horizon and moderate late. Over three years the posterior medians order the entries as valuation, purchase, refinancing, and this ordering holds across four prespecified accounting representations and the stable-state subset. The posterior does not sharply rank them—the three cumulative contributions have 90% intervals that include zero, and Section 5 reports the paired comparisons—so what the accounting establishes is a pattern rather than a precise ranking, and that pattern is already informative. The three entries together contribute about -0.16 percent to the three-year consumption level in posterior mean, partly offset by positive propagation along paths that never enter the three entry variables, so the modest net response of about -0.09 percent conceals a larger gross housing-finance contribution.

The third finding is that the three margins have distinct volatility histories. In the posterior-median paths of the observable equations' innovation standard deviations, refinancing peaks in July 2003, during the refinancing wave; house-price growth peaks in January 2009; and purchase financing peaks in October 2020. Several innovation scales are higher on average in NBER contractions, but the housing-finance peaks do not follow a single recession chronology. A common mortgage rate does not imply a common volatility history for the activities it prices, and a single housing-market indicator would hide the difference. Two further calculations place the policy innovation within these histories. Its time-averaged share of thirty-six-month forecast-error variance is about 29% for purchase financing and 13% for refinancing but only 2% for consumption growth, so policy is a material source of uncertainty about financing activity and a minor one for monthly spending growth. A historical decomposition ties the realized policy innovations to the 2003 wave—their contribution to refinancing intensity in July 2003 is about +0.25 percent-

age points, with a 90% interval above zero—and shows the timing asymmetry in the data: after the 2020–21 easing, the policy contribution to refinancing peaks in mid-2020 and has reversed by 2023, while the contribution to purchase financing peaks almost two years later.

The policy interpretation rests on timing and exclusion restrictions applied to a market yield. The obvious alternative, a proxy-SVAR with high-frequency announcement surprises as the instrument, is not available with ordinary inference on this sample: the first-stage HC1 F for the current-target surprise is about 11 on 1999–2023 and halves when either January 2008 or March 2020 is omitted. The weakness is a property of the post-1999 sample rather than of this system: in a four-variable auxiliary system of the kind used by Bauer and Swanson (2023), their own surprise series has an ordinary first-stage F of about 8 over 1988–2019 and about 0.1 over 1999–2019, and their orthogonalized surprises are likewise weak in two earlier policy-indicator specifications (Appendix C.5). I therefore use the surprises for what they can support. As a label check, the surprise correlates more closely with the restriction-identified policy innovation than with any other structural innovation in every one of the 12,000 posterior states (median 0.37 against 0.15). As an exploratory comparison, a proxy identification on an earlier Case-Shiller specification also shows a refinancing contraction, a gradual purchase-financing decline, a persistent mortgage-rate increase, and a hump in the one-year yield; its industrial-production impact changes sign when March 2020 is omitted, while the consumption impact remains positive under both single-month omissions (Appendix C.5). Re-estimating with contemporaneous mortgage-rate feedback in the policy equation changes pass-through magnitudes but not the financing pattern; the early rise in industrial production is taken up in Section 7.3.

The paper connects two literatures. Micro evidence shows that mortgage payment reductions raise spending (Di Maggio et al., 2017), that mortgagor households account for much of the consumption response to policy (Cloyne, Ferreira, and Surico, 2020), that regional housing equity shapes refinancing and its spending effects (Beraja et al., 2019), and that the strength of the refinancing channel depends on the distribution of refinancing incentives (Berger et al., 2021; Eichenbaum, Rebelo, and Wong, 2022). Ringo (2024) separates payment transmission to marginal and inframarginal borrowers; Berger et al. (2018) and Favara, Graham, and Li (2025) study house prices and home-purchase expenditure, respectively. Aggregate studies document the responses of house prices, mortgage credit, and consumption to policy (Musso, Neri, and Stracca, 2011; Calza, Monacelli, and Stracca, 2013), and Rella (2025) distinguishes purchase from refinancing responses. Structural household models quantify housing and refinancing choices in transmission (Kinnerud, 2025; Maxted, Laibson, and Moll, 2025). Relative to this work, I place the three margins in one aggregate system under one identified shock, with purpose-specific monthly financing measures, and measure how their contributions to consumption accumulate over time. The decomposition

framework is Wegner et al.'s; the contribution here is the empirical comparison it makes possible, not a new decomposition method.

Section 2 describes the data and the three margins. Section 3 presents the model, the policy experiment, and the external announcement evidence. Section 4 reports the housing-finance responses, Section 5 the consumption accounting, and Section 6 the volatility histories together with forecast-variance shares and historical contributions of the policy innovation. Section 7 reports robustness and discusses limitations, including state dependence of the refinancing response, and Section 8 concludes.

2. Housing finance and measurement

2.1 Three margins of adjustment

Housing connects monetary policy to consumption through decisions that need not occur together. A change in financing costs alters the value of housing, the financing of a purchase, and the terms on an existing mortgage, and the three margins differ in how they reach expenditure. Buying an existing home is an asset transaction; moving, furnishing, and housing-service decisions connect it to consumption, while down-payment funding and subsequent debt service reduce the resources available for other spending. The net consumption response associated with purchase financing is therefore an empirical question. A rate-and-term refinance replaces an existing mortgage and changes future payments without financing a purchase or extracting equity. House-price movements affect wealth and collateral. Asking how the margins' contributions to consumption accumulate after one policy shock requires measuring them separately.

The empirical system therefore includes a house-price measure and two purpose-specific origination flows. The thirty-year fixed mortgage rate is the common financing price through which purchase and refinancing opportunities respond, and house prices enter both financing equations, so valuations and financing interact without all originations being treated as one activity.

Separate financing measures complement the micro evidence that mortgage contracts and household balance sheets shape spending responses. Di Maggio et al. (2017) study payment reductions and household expenditure and deleveraging; Beraja et al. (2019) show how the distribution of housing equity affects refinancing and spending. The monthly aggregate system here measures the joint adjustment of valuations and financing activity and places their propagation to aggregate consumption on a common scale.

2.2 Monthly sample and variable definitions

The input sample contains 312 U.S. monthly observations from January 1998 through December 2023. Twelve initial observations supply lagged values, leaving 300 response observations, January 1999–December 2023, in the likelihood. Growth rates use actual monthly source observations, including December 1997 levels for the initial difference; no quarterly series is interpolated.

Write the observable vector as

$$y_t = (C_t, IP_t, \pi_t, r_t, M_t, H_t, P_t, R_t)'. \tag{1}$$

Here C_t and IP_t denote the model-input growth rates, including the fixed extreme-observation treatment of Section 2.4. Uppercase P_t denotes purchase-financing intensity throughout; the consumption-price deflator is written P_t^{cons}.

Table 1. Variables, measurement, and timing. Growth and inflation rates are monthly and nonannualized. Interest rates enter in percent per annum as quoted; financing intensities enter as percentages of the lagged mortgage stock. A response of 0.01 in a financing-intensity equation is 0.01 percentage points.

Symbol	Economic object and source	Model input and timing
C	Real PCE quantity index, BEA DPCERA3M086SBEA	100Δ log, with fixed growth caps
IP	Industrial production, Federal Reserve INDPRO	100Δ log, with fixed growth caps
π	Core PCE price index, BEA PCEPILFE	100Δ log
r	One-year Treasury yield, H.15 DGS1	Last finite quotation in each month, level
M	Thirty-year fixed mortgage rate, Freddie Mac PMMS	Mean of weekly quotations dated within each month, level
H	FHFA national purchase-only HPI, HPIPONM226S	$100\Delta \log(HPI/P^{\text{cons}})$, with headline PCEPI as deflator
P	NMDB purchase origination volume; Z.1 mortgage liabilities	$100\times$ seasonally adjusted monthly volume / previous-quarter-end stock
R	NMDB non-cash-out refinance volume; same Z.1 stock	$100\times$ unadjusted monthly volume / previous-quarter-end stock

The data are a fixed-vintage historical panel: their dates identify the periods measured, and the contemporaneous restrictions in Section 3 are economic assumptions about within-month timing rather than a record of when information arrived.

Real PCE is measured with the BEA's seasonally adjusted real quantity index, and industrial production supplies a monthly measure of real activity. Core PCE inflation enters the model's price-information set, while the headline consumption-price index deflates housing values, which keeps the inflation control separate from the relative price of housing. Source definitions and retrieval dates are in Appendix A.

The one-year Treasury yield is the policy indicator because it carries the expected near-term path of short rates and stays informative when the overnight rate is at its lower bound, as in Gertler and Karadi (2015). Section 3 defines the innovation to it that I interpret as a monetary-policy shock.

For mortgage rates, weekly PMMS observations receive equal weight within their calendar month. PMMS measures rates on conventional conforming home-purchase mortgages with a specified borrower profile, and I use it as a common market-pricing indicator for purchase and refinancing activity. The November 2022 transition from lender surveys to application-based information is retained and documented (Freddie Mac, 2022). The series prices new loans; the attractiveness of refinancing depends on the gap between this rate and the coupon on existing debt, which the model leaves implicit.

2.3 Housing valuations and purpose-specific financing flows

The FHFA purchase-only index uses repeat purchase transactions associated with mortgages acquired or securitized by Fannie Mae and Freddie Mac, so it excludes refinance appraisals. I use its official seasonal adjustment and deflate it by headline PCE prices:

$$H_t = 100\Delta \log\left(100\frac{HPI_t^{FHFA}}{P_t^{\text{cons}}}\right). \tag{2}$$

This measures housing appreciation relative to the consumption basket. Because a house-price response can transmit through wealth and collateral, I refer to the corresponding entry point as *housing valuation*. A Case-Shiller alternative examines sensitivity to housing-price measurement; its published monthly index uses a three-month transaction window, so very short timing differences between the two indexes should not be read as adjustment delays.

FHFA's National Mortgage Database distinguishes funded purchase originations from refinancing and separates refinancing purposes. Let V_t^P be purchase volume, V_t^{ref} total refinance volume, and $s_t^{RT,V}$ the published dollar-weighted non-cash-out share in percent. Rate-and-term volume is

$$V_t^{RT} = V_t^{ref}\frac{s_t^{RT,V}}{100}. \tag{3}$$

The dollar-weighted share matches the volume numerator; published rounding carries through to the constructed volume. Appendix A reconciles the constructed volumes against published all-origination and quarterly totals.

Both financing indicators use the preceding quarter-end stock of household and nonprofit one-to-four-family mortgage liabilities, $B_{q(t)-1}$, from the Financial Accounts:

$$P_t = 100\frac{SA_{X13}(V_t^P)}{B_{q(t)-1}}, \qquad R_t = 100\frac{V_t^{RT}}{B_{q(t)-1}}. \tag{4}$$

Numerator and denominator are in millions of nominal dollars, and all three months in a quarter use the same inherited stock. The denominator places the two flows on a common market-size basis and keeps contemporaneous quarter-end debt out of the monthly ratio. These are aggregate financing intensities on a common benchmark.

The series measure financing activity, not net credit creation. Refinancing can generate a large new origination while replacing almost the same outstanding balance, and a change in origination dollars may reflect loan counts, average size, or their composition. The refinancing measure isolates non-cash-out originations, separating replacement financing from borrowing that extracts housing equity.

Purchase originations exhibit seasonality, so I seasonally adjust the purchase-dollar numerator with X-13 before dividing by the stock; the adjusted-numerator QS test detects no significant residual seasonality. Rate-and-term refinancing intensity is retained without seasonal adjustment: its calendar-month test over January 1998–September 2025 does not detect significant seasonality. Appendix A provides the adjustment specification and test details.

2.4 Extreme observations, persistence, and the consumption estimand

Consumption and industrial-production growth receive fixed caps calibrated on pre-pandemic observations. For each treated variable x, the model uses

$$\widetilde{x}_t = \min\{\max(x_t, L_x), U_x\}, \tag{5}$$

with thresholds computed from January 1998–February 2020 using type-7 empirical quantiles. PCE uses the second and ninety-eighth percentiles, approximately -0.47705 and 0.87567; IP uses the first and ninety-ninth percentiles, approximately -2.15572 and 1.21818. The rule applies to the full sample and preserves every observation date. It changes 22 PCE observations and 15 IP observations out of 312; the other six variables, including both mortgage intensities, are unchanged.

The caps limit the influence of extreme monthly growth observations while allowing innovation-specific volatility to vary. They are a measurement choice as well as an estimation choice, since bounding growth compresses some economically important variation. Lenza and Primiceri (2022) and Carriero et al. (2024) treat pandemic outliers through different volatility models; Section 6 reports how the volatility estimates change without the caps.

The consumption IRF is therefore the response of the treated growth series supplied to the model, and summing it gives the log-level gap implied by that measure. It is in consumption units, and the same definition applies to every contribution in Section 5.

Interest rates and financing intensities remain in levels to retain their economic interpretation. The data are persistent: the 312-month purchase and refinancing means are 0.86 and 0.63 percent of lagged mortgage stock, and their lag-one correlations are 0.99 and 0.95. ADF and KPSS diagnostics are reported with their actual windows in Appendix A. The model accommodates persistence through twelve lags and shrinkage, and Section 7.2 examines the role of nonstable posterior states in the finite-horizon analysis.

3. Empirical framework

3.1 A common dynamic system with changing innovation variances

The model combines unrestricted lagged feedback with a structured contemporaneous relationship:

$$y_t = c + \sum_{\ell=1}^{12} \Phi_\ell y_{t-\ell} + u_t, \qquad A_0 u_t = \varepsilon_t = D_t^{1/2} z_t, \quad z_t \sim N(0, I_8), \tag{6}$$

where A_0 is invertible and has a unit diagonal. All eight lagged variables enter each equation alongside an intercept, so consumption can respond to lagged housing valuations, mortgage rates, purchase financing, and refinancing. The specification has 776 conditional-mean coefficients; twelve lags allow a full year of monthly dynamics, and shrinkage regularizes the parameter set. A twelve-lag Bayesian specification also appears in the consumption application of Miranda-Agrippino and Ricco (2021), whose external-instrument identification differs from the identification used here. A matched six-lag comparison appears in Section 7.2.

The diagonal structural variance matrix is $D_t = \text{diag}\{\exp(h_{1t}), \ldots, \exp(h_{8t})\}$. Each log variance follows

$$h_{jt} = \mu_j + \varphi_j(h_{j,t-1} - \mu_j) + \sigma_j \eta_{jt}, \quad \eta_{jt} \sim N(0,1), \quad |\varphi_j| < 1, \tag{7}$$

with the stationary initial distribution $h_{j0} \sim N(\mu_j, \sigma_j^2/(1-\varphi_j^2))$. Observation and volatility innovations are independent. This is the stationary SV building block of Kastner and Frühwirth-Schnatter (2014). Separate volatility processes let the scale of a refinancing innovation evolve differently from that of a purchase or housing-price innovation.

A_0 and the lag coefficients are fixed over time. The model therefore distinguishes changes in innovation variances from changes in transmission coefficients, a distinction emphasized in Sims and Zha (2006), and it describes distinct volatility histories without estimating a different dynamic system in every period. Section 6 uses those histories as a complementary result.

3.2 Contemporaneous restrictions and their economic content

The restrictions say which variables can respond to one another within a month. Their economic content is that quoted rates move first and completed transactions follow, while mortgage pricing and financing quantities remain free to feed back on each other. The maintained matrix is

	C	IP	π	r	M	H	P	R
C	1	0	0	0	0	0	0	0
IP	$\star$	1	0	0	0	0	0	0
π	$\star$	$\star$	1	0	0	0	0	0
r	$\star$	$\star$	$\star$	1	0	$\star$	0	0
M	$\star$	$\star$	$\star$	$\star$	1	$\star$	$\star$	$\star$
H	$\star$	$\star$	$\star$	0	0	1	0	0
P	$\star$	$\star$	$\star$	0	$\star$	$\star$	1	0
R	$\star$	$\star$	$\star$	0	$\star$	$\star$	$\star$	1

. (8)

Stars denote unrestricted coefficients with no imposed signs. The 28 zeros concern direct contemporaneous relationships after conditioning on the equation's other variables and all lags; they do not eliminate indirect or subsequent transmission. Excluding the one-year yield from a financing equation, for example, still permits an impact response through the mortgage rate.

The restrictions have four roles. First, the slow block $S = (C, IP, \pi, H)$ excludes the current policy indicator and financing variables: monthly spending, production, price adjustment, and completed housing transactions respond more slowly than quoted financial rates. Second, the policy equation allows contemporaneous macroeconomic and housing information but excludes direct feedback from M, P, R. Together these 19 zeros isolate a policy block following Christiano, Eichenbaum, and Evans (1999, Section 4.1). Third, six zeros complete the slow block recursively

as C, IP, π, H. Fourth, the financing equations exclude additional direct effects of r conditional on mortgage pricing, and purchase excludes direct current refinancing feedback.

Letting consumption enter the policy equation uses current activity to separate systematic yield movements from the policy disturbance, as in Ludvigson, Steindel, and Lettau (2002).

The retained interactions are the economically important ones. The mortgage-rate equation includes all other contemporaneous variables, so mortgage prices and financing volumes adjust jointly. Housing valuation enters both purchase and refinancing, retaining a collateral connection even for non-cash-out loans. Purchase activity enters the refinancing equation, and refinancing still affects purchase through mortgage pricing and all lagged relationships. The complete system is nonrecursive because mortgage pricing and financing quantities feed back on each other within the month.

Excluding the direct yield term from the purchase and refinancing equations gives the mortgage rate a specific role. Because the slow variables have zero impact responses to the policy innovation, its initial transmission to the two financing quantities operates through mortgage pricing and the retained quantity–price feedback; Appendix C derives the mapping. This initial rate-entry structure is maintained, while the signs, magnitudes, and subsequent dynamics of the financing responses are estimated, with all lagged relationships available for later propagation.

Two of the restrictions carry economic weight and are tested. Excluding mortgage-market feedback from the policy equation is an assumption about a market yield that can reprice with financial news; Section 7.1 re-estimates the model with that feedback in either direction. The purchase–refinancing asymmetry is an approximation to origination timing in which purchase closings respond to refinancing activity only through mortgage pricing; processing-capacity effects (Sharpe and Sherlund, 2015) are the alternative, and Appendix C.1 records each restriction and the mechanism it retains. No restriction fixes the sign of any estimated response.

3.3 Identification of the policy experiment

The policy experiment is a single structural innovation scaled to raise the one-year yield by 25 basis points on impact; subsequent rates and quantities follow the estimated dynamics. The scale does not impose a permanent rate increase or identify a literal 25-basis-point change in the federal funds target.

Identification of the policy direction can be shown directly. For a positive-definite innovation covariance Σ compatible with the model, orthogonality of the policy innovation to the slow block implies

$$a_{rS} = -\Sigma_{rS}\Sigma_{SS}^{-1}, \qquad d_r = \Sigma_{rr} - \Sigma_{rS}\Sigma_{SS}^{-1}\Sigma_{Sr} > 0, \tag{9}$$

where a_r is the policy row of A_0. Writing $B = A_0^{-1}$ gives

$$B_{\cdot r} = \frac{\Sigma a_r'}{d_r}, \qquad B_{S,r} = 0, \qquad B_{rr} = 1. \tag{10}$$

The 19 block zeros therefore pin down the policy row and column from the covariance alone, before the seven other shocks are named. The remaining nine restrictions complete the system used in estimation and channel accounting; Appendix C.2 gives a constructive identification argument for the complete matrix when $B_{M,r} \neq 0$, in the sense of Rubio-Ramírez, Waggoner, and Zha (2010).

The substantive assumption is that what remains of the yield innovation after the permitted current and lagged information is an exogenous policy disturbance. Three pieces of evidence support it: the timing and exclusion arguments above, the feedback sensitivity in Section 7.1, and announcement information that plays no role in estimation (Section 3.5).

At each joint posterior state s, responses are computed as

$$b^{25,(s)} = 0.25\frac{B_{\cdot r}^{(s)}}{B_{rr}^{(s)}}, \qquad g_h^{(s)} = \Psi_h^{(s)} b^{25,(s)}, \qquad \Psi_h = \sum_{\ell=1}^{\min(12,h)} \Phi_\ell \Psi_{h-\ell}, \quad \Psi_0 = I. \tag{11}$$

Normalization is performed within each state before any posterior summary, and responses run through month 36 with no filtering on sign. Pointwise medians and central 68% and 90% credible intervals summarize the state-level responses; cumulative quantities are summed within each state before taking quantiles.

Two alternatives deserve a word. A proxy-SVAR would use the announcement surprises as the identifying instrument. Appendix C.5 reports the first-stage strength of MP1, FF4, and the public USMPD statement and monetary-event candidates across samples, policy indicators, and control systems: on the post-1999 months the strength is concentrated in two months, varies with the ELB/non-ELB split in opposite directions for different indicators, and is already absent in the literature's own four-variable system once the sample is restricted to 1999 onward. Ratio-based proxy inference is unreliable under those conditions, and weak-instrument-robust sets that also carry the BVAR-SV parameter uncertainty remain to be constructed for this sample. The appendix also reports an exploratory proxy identification on an earlier Case-Shiller specification whose responses share broad financing and mortgage-rate patterns with those reported in Section 4. Identification through the model's own stochastic volatility (Bertsche and Braun, 2020) is a second alternative; it would deliver statistical identification of the innovations, with the policy label still to be established from outside information, and is left for future work. The present design therefore uses the

economic restrictions for identification and the surprises for labeling: Section 3.5 asks whether announcement information corresponds to the innovation the restrictions call policy, not whether that innovation is exogenous.

3.4 Priors, estimation, and numerical support

I estimate the conditional mean, contemporaneous matrix, and latent volatility jointly, so the structural matrix used for responses is the one the time-varying covariance was estimated with.

The lag prior is Minnesota-type (Bańbura, Giannone, and Reichlin, 2010): distant lags and cross-variable coefficients are shrunk harder, with cross-variable shrinkage scaled by equation-specific residual scales s_j^2 from AR(6) regressions on the 312 input months. Table 2 gives the exact priors.

Table 2. Prior specification. Normal distributions use variance as the second argument. The priors are conditional on the calibrated scales. Their components are independent except for latent volatility paths conditional on SV parameters.

Parameter	Prior or restriction
Own lag $(\Phi_\ell)_{jj}$	$N(0, 0.16/\ell^2)$
Cross lag $(\Phi_\ell)_{jk}$	$N(0, 0.004 s_j^2/(\ell^2 s_k^2))$
Intercept c_j	$N(0, 100)$
Free contemporaneous coefficient $(A_0)_{jk}$	$N(0, (0.20 s_j/s_k)^2)$
SV level μ_j	$N(\log s_j^2, 1.5^2)$
SV persistence φ_j	$(\varphi_j + 1)/2 \sim \text{Beta}(20, 1.5)$
SV innovation SD σ_j	$\lvert N(0, 0.20^2)\rvert$
Initial log variance	Stationary distribution conditional on $(\mu_j, \varphi_j, \sigma_j)$

With $\varepsilon_t = A_0(y_t - c - \sum_\ell \Phi_\ell y_{t-\ell})$, the conditional Gaussian log likelihood is

$$\log L = T \log |\det A_0| - \frac{1}{2} \sum_{t=1}^{T} \sum_{j=1}^{8} \{h_{jt} + \varepsilon_{jt}^2 e^{-h_{jt}}\} + \text{constant}. \tag{12}$$

The determinant term matters because the system is nonrecursive. Posterior simulation updates the 776 mean coefficients jointly, the free rows of A_0 with slice updates that keep the determinant term (Neal, 2003), and each SV block with interweaving (Kastner and Frühwirth-Schnatter, 2014); Appendix B gives the conditionals and implementation checks.

Four chains of 5,000 warmup and 30,000 post-warmup sweeps, thinned by ten, give 12,000 joint posterior states. Rank-normalized $\widehat{R}$ and bulk and tail effective sample sizes (Vehtari et al., 2021) meet the convergence criteria for all 3,236 state components and for the impulse responses at horizons 1, 6, 12, 24, and 36 months (Appendix B, which also reports simulation-based checks of the code). About 6% of states (726) imply a nonstable VAR; they are kept, and Section 7.2 shows that the stable subset gives the same financing and late-consumption patterns.

3.5 External announcement information and the policy innovation

Announcement surprises provide an external check on which fitted innovation carries policy information. I use the statement-window MP1 (current-target) and FF4 (fourth federal-funds futures contract) surprises of Acosta et al. (2026), summed within each calendar month; 298 months are common to both series.

For each posterior state I reconstruct the structural innovations $\varepsilon_t^{(s)} = A_0^{(s)} u_t^{(s)}$ and correlate each of the eight with the surprises. Announcement information plays no role in estimation, so the comparison asks whether it lines up with the innovation the restrictions already call policy, in the spirit of Bertsche and Braun (2020) and Chițu et al. (2023).

Table 3. Announcement information and the predetermined policy innovation. Entries are medians across the 12,000 retained posterior states. F is the single-slope auxiliary-regression statistic, reported with HC1 and calendar-lag HAC12 covariance estimates. January 2008 and March 2020 are omitted separately from the diagnostic sample, holding fitted parameter states fixed; each omission removes the entire monthly observation, and means and auxiliary regressions are recomputed on the remaining months.

Surprise	Monthly support	Correlation	HC1 F	HAC12 F
MP1	Common, $N = 298$	0.373	11.13	26.13
MP1	Exclude January 2008, $N = 297$	0.258	5.39	12.34
MP1	Exclude March 2020, $N = 297$	0.318	6.48	16.07
FF4	Common, $N = 298$	0.277	4.09	6.62
FF4	Exclude January 2008, $N = 297$	0.199	3.23	4.88
FF4	Exclude March 2020, $N = 297$	0.234	3.13	5.84

Both surprises line up with the policy innovation. In all 12,000 states its absolute correlation exceeds that of every other innovation; the median largest nonpolicy correlation is 0.152 for MP1 and 0.125 for FF4, against 0.373 and 0.277 for policy.

January 2008 and March 2020 are influential: omitting either lowers the statistics in Table 3, yet the policy covariance stays positive in every state and the policy innovation stays the most closely associated one in almost all states (Appendix C.3). MP1 also correlates with the consumption innovation (median 0.151 against 0.373 for policy), so announcement information touches several disturbances while matching the policy innovation most closely. This supports the policy label; Appendix C.5 explains why the surprises are too weak on this sample to serve as instruments.

4. Monetary policy and housing-finance adjustment

4.1 A common pricing environment, different financing responses

The policy experiment raises mortgage rates persistently, and purchase financing and refinancing contract along different time profiles: refinancing declines early and moves back toward zero, while purchase financing contracts progressively and remains depressed after three years. These dynamics set the financing environment for the consumption accounting in Section 5.

The one-year yield rises by 25 basis points initially and reaches a median of about 41 basis points after a year; a single innovation generates this hump through endogenous propagation, with no additional policy shocks injected. Mortgage rates rise by 8.5 basis points on impact, 17.8 at month 12, and 14.2 at month 36, with pointwise 90% intervals above zero throughout. The two financing activities therefore adjust under persistently higher new-mortgage rates rather than a brief change in borrowing costs.

Refinancing intensity falls by a median 0.036 percentage points at month 6 and 0.033 at month 12, with 90% intervals below zero from months 1 through 17. The median then moves toward zero, reaching -0.005 at month 36 with a 90% interval of $[-0.046, 0.027]$. Higher new-loan rates reduce the gain from replacing an existing mortgage; past contract rates and the distribution of borrower balance sheets shape the available opportunities, consistent with Berger et al. (2021).

Purchase intensity falls more gradually: its median responses are -0.009 at month 6, -0.016 at month 12, and -0.033 percentage points at month 36, with 90% intervals below zero from months 3 through 36. Purchase financing accompanies transactions involving search, matching, and completion, which distinguishes its adjustment from refinancing an existing contract.

The magnitudes are clearer once the denominator is made explicit. Both intensities scale new monthly originations by the entire outstanding mortgage stock, so a small percentage-point response is not a negligible change in monthly activity. The fixed mean intensities in the 300-month response sample are 0.84664% for purchase financing and 0.62293% for rate-and-term refinanc-

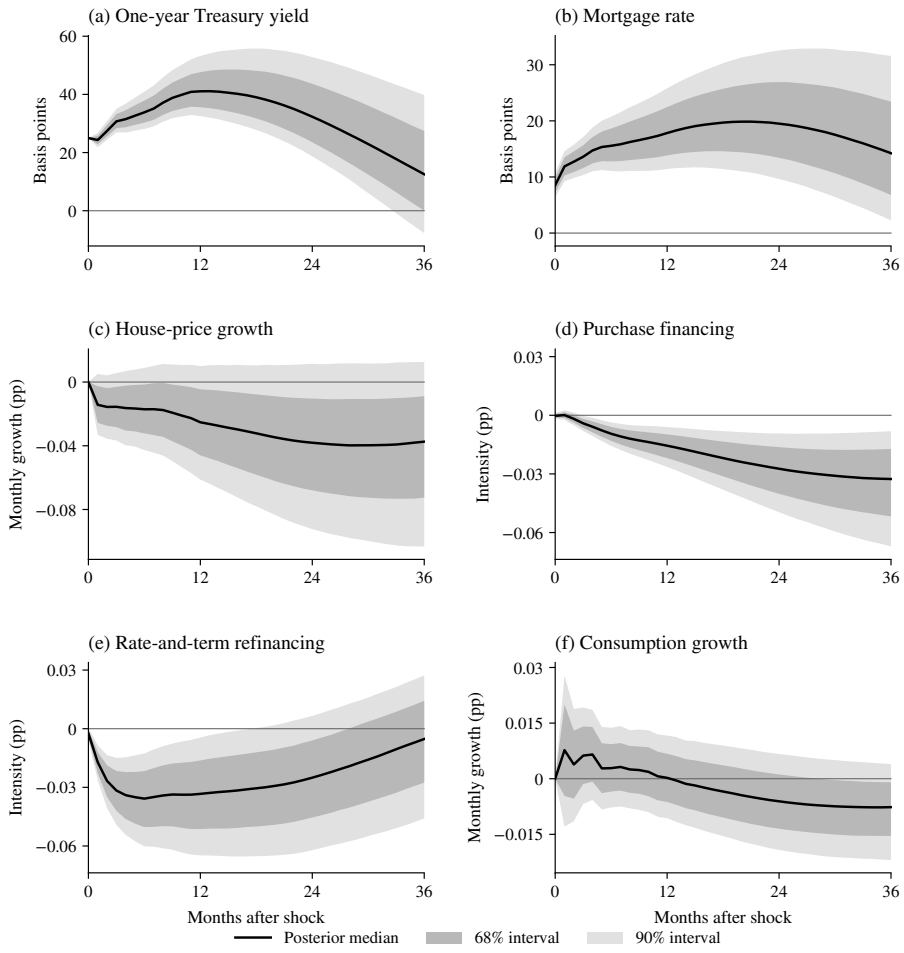


Figure 1. Responses to a policy innovation raising the one-year yield by 25 basis points on impact. The solid line is the posterior median; dark and light bands are central 68% and 90% pointwise credible intervals. Rates are shown in basis points, house-price and consumption growth in monthly percentage points, and financing intensities in percentage points of lagged mortgage stock. Purchase and refinancing share a vertical scale. The model uses the FHFA purchase-only index, January 1998–December 2023 inputs, twelve lags, and all 12,000 posterior states. The consumption series incorporates the treatment in Section 2.4. The zero impact responses of house-price and consumption growth follow from the restrictions.

ing. Dividing each response by its own mean gives a six-month refinancing decline equivalent to 5.74% of usual monthly intensity and a three-year purchase-financing decline equivalent to 3.85%, with 90% intervals of [−9.66%, −2.04%] and [−7.93%, −0.96%] on the same fixed benchmarks. These are adjustments of several percent of typical monthly financing intensity under a common initial 25-basis-point yield innovation; Appendix C retains the full rescaled paths with both interval levels.

The relative size of the median purchase and refinancing responses therefore changes with the horizon. Whether that timing carries over to consumption is the question of Section 5.

4.2 Valuation and aggregate consumption

Real house-price growth has a negative median response after impact: −0.025 percentage points at month 12 and −0.037 at month 36. The 68% intervals lie below zero over months 1–36, while the 90% intervals include zero. Persistence gives these small monthly growth differences an economically interpretable level implication: the three-year cumulative median is approximately −1.07%, with a 68% interval of [−1.89%, −0.33%] and a 90% interval of [−2.55%, 0.23%], approximating the model's $100 \times$ log level gap. The model points to a sustained valuation adjustment that accumulates to about a one-percent level difference; its contribution to consumption is traced in Section 5.

The consumption-growth response is initially small and imprecise: at month 12 its median is 0.0002 percentage points with a 90% interval of [−0.011, 0.012]. It subsequently becomes negative, reaching −0.008 at month 36 with a 90% interval of [−0.022, 0.004], and its 68% intervals lie below zero over months 29–36. This later negative adjustment survives both mortgage-rate-feedback alternatives in Section 7.1. Accumulating within posterior states gives a three-year median consumption-level gap of approximately −0.089%; both its 68% and 90% intervals contain zero.

Monthly and cumulative responses answer different questions. Earlier positive responses partly offset later negative growth responses, leaving a moderate net level gap even where later growth has negative posterior evidence, and because accumulation precedes posterior summarization, the cumulative median is not the sum of monthly medians. A second source of offset concerns propagation through different paths: Section 5 shows that the posterior-mean negative contribution of the three housing-finance entries is partly offset by positive propagation outside those entries, so the net consumption response alone does not reveal the magnitude or composition of the model's housing-finance transmission.

The other macroeconomic responses provide context for the policy interpretation. Industrial-production growth rises at month 1, with a median of 0.041 percentage points and a 90% interval of [0.006, 0.077]. Core inflation has a small positive early median whose 90% interval includes zero, followed by negative 90% intervals at months 32–36. The early production response persists among stable states (Appendix C retains the complete paths); it is the one feature of Figure 1 that does not fit a contractionary policy shock, and Section 7.3 takes it up.

Related studies place these responses in context. Musso, Neri, and Stracca (2011, Section 3.3.1) report a short-run US price puzzle while studying housing, mortgage debt, and consumption. Miranda-Agrippino and Ricco (2021, Section 4, Figures 3–4) find early output increases under some narrative and high-frequency identifications, which their information-adjusted instrument removes. Christiano, Eichenbaum, and Evans (1999, Section 4.4.2) relate price puzzles to omitted forward-looking information. The same early directions appear here in core inflation and IP growth; they are a recurring identification and information-set problem in monthly monetary VARs, not one specific to housing finance.

5. Dynamic contributions to consumption

5.1 From intermediate responses to a common consumption scale

This section asks how much of the consumption response each housing-finance margin carries. A financing variable's own response is a different quantity from its consumption contribution: a large refinancing response can carry little consumption if its downstream links are weak, and a modest valuation response can carry a lot through subsequent financing and macroeconomic feedback. The accounting below measures, for each margin, the consumption that propagates along paths entering that margin first, within the same fitted system and under the same policy innovation.

Every contribution ends at consumption and is expressed in consumption units, which is what makes the margins comparable despite their different measurements. Channel exercises in the tradition of Ludvigson, Steindel, and Lettau (2002) and Bachmann and Sims (2012) evaluate propagation in a common outcome in the same way; Appendix D.5 relates their counterfactual operations to the path allocation used here.

The calculation uses the transmission-channel framework of Wegner et al. (2025). The economic question is where transmission first enters housing valuation, purchase financing, or rate-and-term refinancing, and how it subsequently reaches consumption. A first-entry rule matches

that question: it assigns ownership once, retains later interactions, and counts no path more than once.

For each posterior state, represent propagation on a time-expanded graph whose nodes identify a variable and a month. A path ending at consumption in month h receives a signed weight determined by its source loading and propagation coefficients. Its label is the *first* variable it visits in $\{H, P, R\}$; a path that never visits any of these belongs to Other. The contribution of category j is the sum of its path weights:

$$c_j^{(s,o)}(h) = \sum_{p:\, f(p)=j,\, p \to C_h} w^{(s,o)}(p), \qquad j \in \{H, P, R, O\}. \tag{13}$$

The index o denotes the contemporaneous accounting ordering. A label is fixed at first entry while subsequent feedback is unrestricted: policy $\to H_1 \to R_2 \to C_3$ belongs to H, with refinancing part of its subsequent propagation, and policy $\to P_1 \to H_2 \to C_3$ belongs to P. Paths may revisit core variables, pass through industrial production, inflation, or interest rates, and visit consumption before their final endpoint; first entry need not occur on impact.

FRM30 remains an endogenous mortgage-pricing node that can precede or follow a core entry; it is not a fourth category, and visiting it first is not required for inclusion in P or R. Other contains all consumption paths that never enter H, P, or R, including propagation through the remaining macroeconomic and rate variables and consumption's own dynamics. Within each state, ordering, and horizon the four categories satisfy

$$c_H^{(s,o)}(h) + c_P^{(s,o)}(h) + c_R^{(s,o)}(h) + c_O^{(s,o)}(h) = g_C^{(s)}(h). \tag{14}$$

Each category therefore measures the consumption that reaches the household sector through paths that enter housing finance at that margin, with everything that happens afterward—including visits to the other margins—kept. H includes later purchase and refinancing links when valuation is visited first; R measures refinancing-first paths, which is a narrower object than the refinancing channel of the micro literature.

5.2 Representation and accumulation

Because the structural system is nonrecursive, contemporaneous propagation needs an ordered representation; I use a QL transformation that preserves the total response, with primary order $IP, \pi, r, M, H, P, R, C$ and three prespecified alternatives (Appendix D.1 and D.3).

Contributions to monthly growth are in percentage points. Cumulative contributions are

$$K_j^{(s,o)}(h) = \sum_{\tau=0}^{h} c_j^{(s,o)}(\tau), \qquad W_j^{(s,o)}(a,b) = \sum_{\tau=a}^{b} c_j^{(s,o)}(\tau). \tag{15}$$

$K_j(h)$ is a contribution to the $100 \times$ log consumption-level gap, approximately a percentage level gap. Figure 2 accumulates from $h = 0$ and shows posterior means, for which equation (14) holds exactly; the window comparisons in Section 5.4 use months 1–12 and 1–36 and report medians and intervals of state-level sums (Appendix Table D6). All accounting and pairwise comparisons are done within a joint state before summarizing, and Appendix D verifies the add-up in every state and ordering.

5.3 Different accumulation patterns

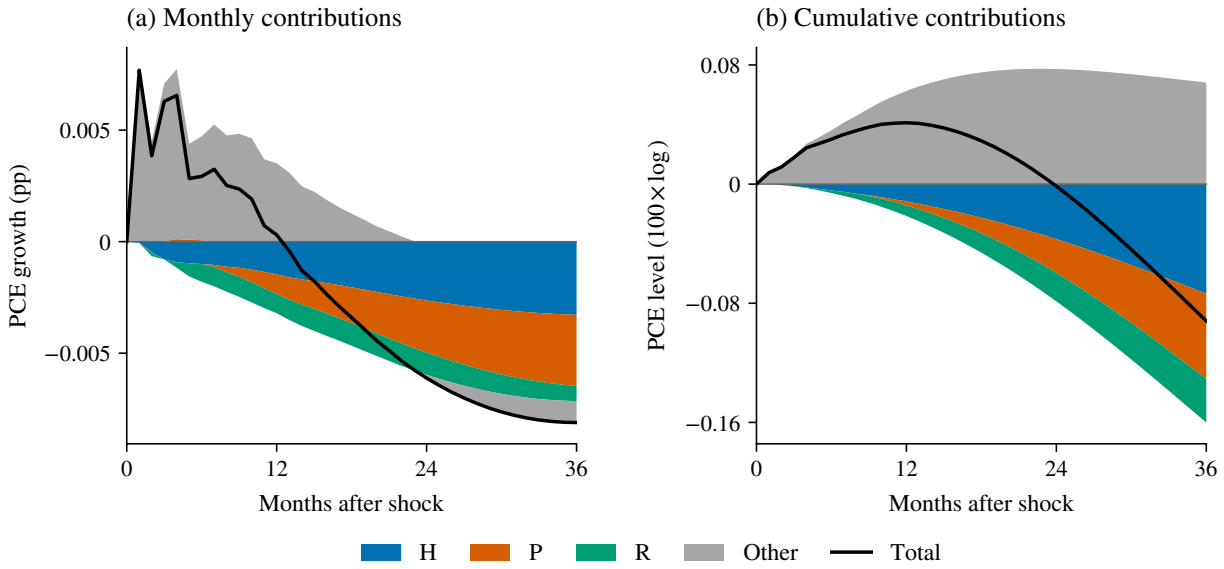


Figure 2. Monthly and cumulative first-entry contributions to consumption. H is housing valuation (blue), P purchase financing (orange), R rate-and-term refinancing (green), and Other the complement (gray). The black line is Total. Positive and negative posterior-mean contributions are stacked separately; the thickness and side of zero identify each contribution. Monthly units are percentage points of model-input PCE growth; cumulative units are contributions to its $100 \times$ log level gap over months 0 through h. The figure uses the primary order $IP, \pi, r, M, H, P, R, C$, all 12,000 states, and the same +25-basis-point policy innovation as Figure 1. The shaded areas are signed contributions, not uncertainty bands. Appendix Tables D3, D4, and D6 report the complete totals, paired comparisons, and window intervals.

The three entries accumulate differently. Housing valuation becomes negative relatively early and continues to accumulate; purchase financing begins near zero and becomes increasingly negative later; refinancing contributes a comparatively even drag over the middle of the horizon and

moderates toward its end. The consumption response therefore has a changing composition even under a fixed-coefficient model and one common shock.

For housing valuation, the monthly posterior-mean contribution is approximately -0.0010 percentage points at month 6 and -0.0033 at month 36, and continued negative growth contributions make the cumulative gap grow. Existing housing assets can be revalued without a household buying a home or replacing a loan that month, and subsequent collateral and financing feedback extend the propagation; Berger et al. (2018) provide a household framework connecting house prices to consumption. Because the H category retains mixed paths such as H→R→C, its estimate describes valuation-entry propagation rather than a pure wealth effect.

Purchase financing has a monthly mean contribution of approximately $+0.0001$ at month 6 and -0.0032 percentage points at month 36: an early value close to zero and later negative contributions that build a larger cumulative gap. This pattern is consistent with transaction-related expenditure occurring over time; Favara, Graham, and Li (2025) document increases in home-related expenditure around purchases.

The refinancing mean contribution is approximately -0.0008 percentage points at month 6, -0.0009 at month 12, -0.0010 at month 24, and -0.0007 at month 36. Its late moderation concerns the current contribution to growth; as long as that contribution remains negative, the cumulative level gap continues to deepen, more slowly. This separates the recovery of refinancing activity in Figure 1 from the persistence of its accumulated consumption contribution in Figure 2.

Rate-and-term origination activity covers one part of mortgage transmission. It omits cash-out borrowing as a separate entry and does not directly measure payment changes on loans that would have been originated anyway, the inframarginal borrowers emphasized by Ringo (2024). Paths reaching valuation before refinancing are assigned to H. A small R contribution is therefore a statement about refinancing-first paths, not about mortgage payment transmission as a whole.

5.4 Why the observation horizon matters

Window comparisons summarize different parts of the propagation process. Over months 1–12, the median contributions are approximately -0.009 for valuation, -0.002 for purchase financing, and -0.006 for refinancing. Over months 1–36 they are -0.064, -0.047, and -0.023, in contributions to the model-implied $100 \times$ log consumption-level gap. These windows exclude impact and differ from the figure's cumulative means over months 0 through h; Appendix Table D6 reports the full window distributions.

Persistence changes the comparison. Refinancing's negative median contribution exceeds purchase financing's over the first year, whereas purchase financing's exceeds refinancing's over three years; valuation has the largest negative median in both windows. The size of an initial financing adjustment does not summarize its subsequent consumption propagation: contributions that begin near zero can build through continuing feedback, while a comparatively even monthly contribution accumulates at a different rate.

These are comparisons of estimated profiles, not a precisely established ranking. The 90% intervals for the three cumulative contributions include zero, and the saved paired differences, computed within each posterior state, favor the median ordering with posterior probabilities of 0.56 for H below P, 0.66 for H below R, and 0.60 for P below R (Appendix Table D4).

The complete accounting reveals transmission that the net response conceals. At month 36, cumulative posterior-mean contributions over months 0–36 are −0.074 for H, −0.057 for P, and −0.029 for R, summing to approximately −0.160 in $100 \times$ log consumption-level units, or about −0.16%. Other contributes +0.068, leaving a total of approximately −0.092. The modest net decline thus combines a larger negative housing-finance contribution with offsetting positive propagation along paths that never enter the three entry variables, and reading only the total would miss this composition. All quantities here are posterior means from the same states, representation, and window, so they add up.

Other's monthly contribution is negative by month 36 even though its cumulative contribution is positive, and its breadth leaves its internal composition an open diagnostic: the early positive production response points to IP-related paths, which the four-category accounting does not separate.

Taken together, the financing responses and the consumption accounting describe different stages of the same policy experiment. Refinancing activity contracts relatively early, purchase-entry consumption contributions build over longer horizons, and valuation-entry contributions retain subsequent financing feedback. Assessing housing-finance transmission requires specifying the entry rule, propagation dynamics, and observation horizon jointly.

6. Volatility histories and the policy innovation's role

6.1 Distinct volatility histories

Housing valuation, purchase financing, and rate-and-term refinancing have distinct conditional-volatility histories. For each posterior state and month I compute the conditional standard devia-

tion of each observable equation's innovation, $\sqrt{[BD_tB']_{ii}}$ with $B = A_0^{-1}$, a full-sample smoothed measure rather than a real-time one (Appendix E.1). The posterior-median paths peak for refinancing in July 2003, at 0.465 percentage points of lagged mortgage stock (90% pointwise interval 0.303–0.765), during the refinancing wave documented by Krainer and Marquis (2003); for house-price growth in January 2009, at 0.560 percentage points of monthly growth (0.401–0.832); and for purchase financing in October 2020, at 0.097 percentage points (0.067–0.155). The macroeconomic equations follow the business cycle more closely: contraction-month averages of the median paths exceed expansion-month averages by factors of 1.3 to 2.3 for industrial production, core inflation, the one-year yield, and housing valuation, against 1.21 for refinancing and 0.97 for purchase financing (Appendix Table E2). The refinancing peak survives replacing FHFA with Case-Shiller, the one-year with the two-year yield, capped with uncapped growth, and the baseline with halved lag-prior variances; the house-price and purchase peak months move under Case-Shiller, to March 2010 and July 2020, and the pandemic macroeconomic scales depend on the caps (Appendix E.4). A common mortgage rate therefore does not impose a common volatility history on the activities it prices. These histories are narrower objects than broad macroeconomic-uncertainty measures such as Jurado, Ludvigson, and Ng (2015): they say which parts of the system experienced unusually large conditional variability, not that high volatility means stronger transmission.

6.2 Forecast-variance shares and historical contributions

How much does the policy innovation matter relative to the other seven shocks, and when did it matter? Forecast-error variance decomposition answers the first question and historical decomposition the second (Brignone and Piffer, 2025). Both use 2,000 equally spaced joint posterior states (500 per chain, 119 nonstable, all kept). Because innovation variances change over time, the variance share depends on the forecast origin: for each state and month I compute twelve- and thirty-six-step shares, integrating future variances under the fitted log-volatility law, and average over the 300 origins within each state before taking quantiles. Table 4 reports the result; Appendix E.6 gives the formulas, the stable-state subset, and the alternative that freezes origin variances.

The policy innovation is a minor source of uncertainty about monthly consumption growth—half a percent of forecast-error variance at twelve months, 2% at thirty-six—and a much larger one for financing activity: the purchase median rises from 6% to 29% with the horizon, the refinancing median from 8% to 13%. The widening purchase share goes with the slow, persistent purchase response of Section 4, and the small consumption share is the variance counterpart of Section 5: financing activity responds clearly to policy, and the uncertainty sits in the propagation from that

Table 4. Time-averaged forecast-error variance share of the policy innovation. Percent; posterior median [90% interval] across 2,000 states. Each state's share is averaged over 300 monthly forecast origins. Consumption is monthly growth; financing variables are intensity levels. Future variances follow the estimated SV law. Appendix Tables E7 and E8 report selected dated shares, stable-state results, and the frozen-origin-variance comparison.

Response	H = 12 months	H = 36 months
Consumption growth	0.52	2.06
	[0.13, 2.07]	[0.39, 9.75]
Purchase intensity	6.36	28.75
	[1.27, 14.42]	[6.62, 54.89]
Rate-term refinance intensity	8.08	13.31
	[1.70, 18.91]	[1.56, 44.57]

activity to aggregate spending. The intervals are wide, and these are shares of forecast variance, a different object from the channel contributions of Section 5.

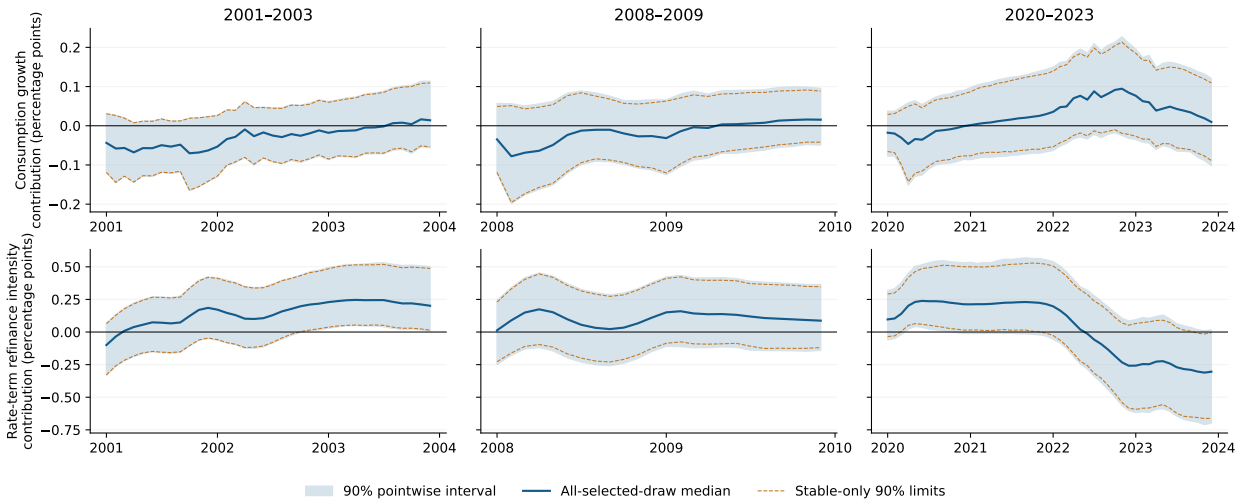


Figure 3. Historical contribution of the policy innovation to consumption growth and refinancing intensity. Solid blue lines are posterior medians and blue bands are pointwise 90% intervals across the 2,000 selected states; dashed orange lines mark the stable-state interval limits. Units are percentage points of monthly model-input real consumption growth and of refinancing intensity. All histories begin in January 1999 conditional on twelve observed initial lags and the estimated intercept, using structural innovations in their native scale. Earlier innovations carry into each displayed window. The three windows, 2001–03, 2008–09, and 2020–23, were fixed before the calculation.

Historical decomposition locates the realized policy innovations in calendar time (Figure 3). The 2003 refinancing wave is the one episode where they leave a clear mark: in July 2003 their contribution to refinancing intensity has a posterior median of +0.245 percentage points and a 90% interval of [0.041, 0.531], against an observed intensity of 3.165%; removing the policy innova-

tions while holding everything else fixed lowers the median path to 2.920% (Appendix Figure E4). The same mark appears in June 2020, at +0.239 [0.036, 0.517]. Averaged over whole windows the contributions are smaller and their intervals include zero—+0.15, +0.10, and +0.02 percentage points for 2001–03, 2008–09, and 2020–23 (Appendix Table E9)—so policy contributed to refinancing at the peaks of these episodes without accounting for the episodes as a whole.

The full-history paths also show the timing asymmetry of Section 4 operating in the data (Appendix Figure E3). After the 2020–21 easing, the policy contribution to refinancing peaks first, in June 2020, and has reversed to −0.304 [−0.698, 0.017] by December 2023 as the 2022 tightening propagates; the contribution to purchase financing peaks almost two years later, in April 2022, at +0.184 [−0.023, 0.512]. By December 2023 refinancing intensity would have stood near 0.344% rather than 0.040% without the policy innovations (accounting-path 90% interval [0.023%, 0.738%]). These are features of fitted contribution paths that carry all earlier innovations, which is why they complement rather than test the impulse responses.

7. Robustness and economic interpretation

7.1 Housing measurement and policy feedback

The financing pattern of Section 4 survives three changes a reader is likely to ask about: the house-price series, the exclusion of mortgage-rate feedback from the policy equation, and the nonstable part of the posterior.

With Case-Shiller replacing FHFA, the mortgage-rate 90% intervals remain positive throughout months 0–36, purchase intervals remain negative over months 3–36, and refinancing has its median trough near month 6, with negative 90% intervals through month 16 rather than 17. Late consumption-growth evidence is similar: negative 68% intervals at months 29–36, with 90% intervals spanning zero. Measurement affects housing magnitudes and the onset of its narrower negative band while preserving the principal financing profiles.

For policy feedback, I replace $A_{0,rM} = 0$ with

$$A_{0,rM} = \gamma \frac{s_r}{s_M}, \qquad \gamma \in \{-0.25, 0, +0.25\}, \qquad s_r/s_M = 1.06028. \tag{16}$$

A 10-basis-point mortgage-rate movement then enters the yield equation directly with about ∓ 2.65 basis points. Everything else is held fixed and the model is re-estimated jointly for each γ.

Table 5. Sensitivity to contemporaneous mortgage-rate feedback. Entries are posterior medians [90% pointwise intervals] from separate fits, each with 12,000 posterior states. Rates are basis points; financing and consumption-growth responses are percentage points. States are not paired across fits.

Response	$\gamma = -0.25$	$\gamma = 0$	$\gamma = +0.25$
Mortgage rate, month 12	14.12	17.83	20.91
	[7.08, 22.29]	[11.42, 25.81]	[14.60, 29.14]
Purchase intensity, month 36	−0.02465	−0.03261	−0.03230
	[−0.05401, −0.00383]	[−0.06714, −0.00815]	[−0.06792, −0.00906]
Refinancing intensity, month 12	−0.03156	−0.03334	−0.03325
	[−0.06388, −0.00158]	[−0.06482, −0.00454]	[−0.06383, −0.00689]
Consumption growth, month 36	−0.00693	−0.00766	−0.00744
	[−0.01825, 0.00173]	[−0.02190, 0.00396]	[−0.02150, 0.00386]

Both feedback alternatives preserve the financing contraction and the later negative consumption-growth adjustment while changing pass-through magnitudes. Purchase 90% intervals lie below zero from month 5 onward under negative gamma and from month 3 under the baseline and positive-gamma specifications; refinancing intervals remain negative through months 15, 17, and 18. All three specifications have negative 68% intervals for consumption growth over months 29–36, with 90% intervals that include zero. Month-12 mortgage-rate medians range from 14.12 to 20.91 basis points, and the median refinancing trough occurs at months 13, 6, and 5 across negative, zero, and positive gamma. The pattern of financing contraction and later consumption adjustment recurs across the three fits; the trough date and the pass-through magnitude do not. Appendix C retains the complete alternative IRFs.

7.2 Stable states and alternative transmission representations

The 726 nonstable states do not drive the results. Among the 11,274 stable states, the month-36 purchase median moves from -0.033 to -0.032 and consumption growth from -0.0077 to -0.0075, with the same interval patterns and slightly narrower long-horizon tails.

The accounting representation moves the allocation but not the ordering. Across the four pre-specified orderings the three-year medians stay ordered H, P, R; placing both financing entries before valuation gives -0.058, -0.051, and -0.023, and exchanging P and R moves R from -0.02251 to -0.01735, a 23% change in its absolute value (Appendix Table D2). Stable states give -0.062, -0.046, and -0.023.

Shortening the lag order or tightening the prior changes little that matters. A matched six-lag fit raises the nonstable share from 6.05% to 24.73% and leaves the refinancing residual's annual autocorrelation near its previous value (0.155 rather than 0.153), so twelve lags remain the reference. Halving the prior variances of all lag coefficients preserves every volatility peak month, with refinancing's marginal-SD change having a time median of +3.24% (Appendix E.4).

7.3 Interpretation and limitations

Three features of the framework shape the interpretation and its extensions.

First, the policy interpretation rests on timing and exclusion restrictions applied to a market yield. The re-estimated feedback alternatives preserve the central financing pattern and later consumption evidence, and announcement information unused in estimation supports the policy label. Announcement surprises may nevertheless reflect revisions to expected policy and to assessments of economic conditions, and the monthly VAR information set need not coincide with the information available to market participants before an announcement. The early rise in industrial production keeps this distinction economically relevant, although its sign alone does not identify an information component. Using the surprises to identify responses rather than to assess the label would require a separate exclusion argument and joint treatment of instrument-moment and model uncertainty, with weak-instrument-robust inference where needed. The exploratory proxy responses in Appendix C.5 display similar broad financing and mortgage-rate patterns under the earlier Case-Shiller specification. Their industrial-production impact changes sign when March 2020 is omitted, while the consumption impact remains positive under both single-month omissions. This comparison highlights similarities in subsequent financing adjustment alongside sensitivity in particular impact responses. Two kinds of further information could distinguish the alternatives: Jarociński and Karadi (2020) use announcement-window interest-rate and equity-price comovement, and Antolín-Díaz and Rubio-Ramírez (2018) use narrative information to restrict shocks or historical decompositions around selected events.

Second, the model uses one set of propagation coefficients across episodes, including the 2003 refinancing wave and the 2022–23 period in which most borrowers held mortgages far below market rates. It lets innovation variances change while holding the conditional-mean coefficients fixed, so the responses are those of a single coefficient set estimated on the full sample. The refinancing literature gives a specific reason to expect state dependence: the response of refinancing to a rate change depends on the share of borrowers for whom refinancing is in the money (Berger et al., 2021; Eichenbaum, Rebelo, and Wong, 2022). The natural extension is to condition transmission on a measure of refinancing incentives—for example the gap between the average coupon on out-

standing mortgages and the current mortgage rate, or a direct measure of the share of borrowers with an in-the-money refinancing option—either as the threshold variable in a regime-switching version of this system or as the state in state-dependent local projections driven by the shock identified here. The supplementary switching exercise below shows why such an extension must treat intercept and slope changes jointly.

A supplementary switching-model comparison, estimated on an earlier Case-Shiller, federal-funds/Wu–Xia, Cholesky specification, illustrates why changing volatility and changing mortgage-rate sensitivity must be distinguished. Allowing the twelve mortgage-rate lag coefficients in the refinancing equation to switch with the refinancing component's variance state gives a sum of coefficient differences whose 90% interval lies above zero when the intercept is held fixed, but spans zero once the intercept is also allowed to switch, with a posterior correlation of -0.929 between the intercept difference and the slope sum. Evidence for changing mortgage-rate slopes is therefore sensitive to how intercept shifts are modeled, which supports treating innovation volatility, conditional-mean shifts, and slope changes as distinct empirical questions. Appendix E.5 reports the specification and results.

Third, the decomposition is within-equilibrium accounting, in the sense of Wegner et al. (2025): it allocates the fitted response to the paths that carry it, and a path's weight is a feature of the estimated system rather than a policy-invariant household parameter. Connecting the entry margins to particular household decisions requires loan-level evidence, which is the natural next step.

8. Conclusion

Housing-finance margins respond to monetary tightening along distinct time profiles. Under the same persistently higher mortgage rates, refinancing contracts earlier in the median path and subsequently moves toward zero, while purchase financing contracts more gradually and remains depressed after three years. The financing responses are economically meaningful relative to usual monthly activity, and the contraction pattern survives the measurement, stability, and policy-feedback alternatives examined. Those alternatives retain negative 68% intervals for later consumption-growth responses, while 90% intervals include zero.

First-entry accounting connects the financing responses to consumption propagation: valuation-related contributions accumulate persistently, purchase-related contributions build over longer windows, and the monthly refinancing-related contribution moderates later. Negative housing-finance contributions are partly offset by other paths, so the net consumption response alone does not summarize the transmission operating inside the model. A financing variable's

own response and its consumption contribution are not interchangeable; entry points, feedback, and accumulation horizons jointly organize the transmission of monetary policy through housing finance.

The model also reveals distinct conditional-volatility histories for housing valuation, purchase financing, and rate-and-term refinancing, with posterior-median marginal innovation-SD paths peaking in 2009, 2020, and 2003. Several innovation scales are higher on average during NBER contractions, yet the financing peaks do not follow one recession chronology. At the posterior medians, the policy innovation accounts for a larger share of financing forecast-error variance than of consumption-growth forecast-error variance. Historical accounting finds positive refinancing contributions at selected dates in 2003 and 2020, alongside substantial uncertainty about average contributions over entire episodes. These results add calendar-time and relative-importance evidence to the comparison of transmission after a common policy innovation.

Three extensions follow directly. The first conditions transmission on the refinancing-incentive distribution, to ask whether the composition measured here changes when few borrowers can profitably refinance, as in 2022–23. The second explores identification through the model's own heteroskedasticity; the stochastic-volatility structure estimated here supplies the statistical ingredients, while attaching the policy label would still require external information or restrictions. The third connects the aggregate entry margins to loan-level refinancing and purchase decisions, so that the first-entry contributions can be read against household behavior rather than only against each other.

References

Acosta, Miguel, Andrea Ajello, Michael D. Bauer, Francesca Loria, and Silvia Miranda-Agrippino. 2026. "Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database." Federal Reserve Bank of San Francisco Working Paper 2025-30, August 2026 revision. Working paper.

Antolín-Díaz, Juan, and Juan F. Rubio-Ramírez. 2018. "Narrative Sign Restrictions for SVARs." *American Economic Review* 108(10): 2802–2829. Published article.

Bachmann, Rüdiger, and Eric R. Sims. 2012. "Confidence and the Transmission of Government Spending Shocks." *Journal of Monetary Economics* 59(3): 235–249. Published article. The inspected author manuscript is dated November 15, 2011. Author manuscript.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin. 2010. "Large Bayesian Vector Auto Regressions." *Journal of Applied Econometrics* 25(1): 71–92. Published article. The inspected working-paper version is *Large Bayesian VARs*, ECB Working Paper 966 (2008). ECB paper in institutional repository.

Bauer, Michael D., and Eric T. Swanson. 2023. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual* 37: 87–155. Published article.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2019. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *Quarterly Journal of Economics* 134(1): 109–183. Author-hosted paper.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. "House Prices and Consumer Spending." *Review of Economic Studies*. Published article.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra. 2021. "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy." *American Economic Review* 111(9): 2829–2878. Published article.

Bertsche, Dominik, and Robin Braun. 2020. "Identification of Structural Vector Autoregressions by Stochastic Volatility." Bank of England Staff Working Paper 869. Inspected working-paper version.

Bollerslev, Tim, Robert F. Engle, and Daniel B. Nelson. 1994. "ARCH Models." In *Handbook of Econometrics*, Volume IV, Chapter 49, edited by Robert F. Engle and Daniel L. McFadden. Chapter.

Brignone, Davide, and Michele Piffer. 2025. "A Structural VAR Model for the UK Economy." Bank of England Macro Technical Paper No. 3, July. Primary paper. Sections 3–4 discuss forecast-error variance and historical decompositions.

Bureau of Economic Analysis. 2019. "Personal Consumption Expenditures." Chapter 5 of the *National Income and Product Accounts Handbook*. Official chapter. Series identities and research retrieval dates appear in Appendix A.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca. 2013. "Housing Finance and Monetary Policy." *Journal of the European Economic Association* 11(Supplement 1): 101–122. Published article.

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens. 2024. "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility." *Review of Economics and Statistics* 106(5): 1403–1417. Federal Reserve working-paper record.

Chițu, Livia, Magdalena Grothe, Tatjana Schulze, and Ine Van Robays. 2023. "Financial Shock Transmission to Heterogeneous Firms: The Earnings-Based Borrowing Constraint Channel." European Central Bank Working Paper 2860. Working paper.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*, Volume 1A. Author manuscript.

Clark, Todd E. 2010. "Real-Time Density Forecasts from VARs with Stochastic Volatility." Federal Reserve Bank of Kansas City Research Working Paper 09-08, July 2010 revision. Inspected working paper.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87(1): 102–129. Published article.

Di Maggio, Matteo, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107(11): 3550–3588. Published article.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong. 2022. "State-Dependent Effects of Monetary Policy: The Refinancing Channel." *American Economic Review* 112(3): 721–761. Published article.

Favara, Giovanni, James Graham, and Geng Li. 2025. "Of House and Home-Related Goods: The Home Purchase Channel of Expenditure." Finance and Economics Discussion Series 2025-045, Board of Governors of the Federal Reserve System. Working paper.

Federal Housing Finance Agency. 2026. *NMDB Aggregate Statistics: Data Dictionary and Technical Notes*, 2026Q1 release. Official documentation.

Freddie Mac. 2022. "Freddie Mac's Newly Enhanced Mortgage Rate Survey." November 3. Official methodology announcement.

Fuster, Andreas, Stephanie H. Lo, and Paul S. Willen. 2016. "The Time-Varying Price of Financial Intermediation in the Mortgage Market." July 15 working-paper version. NBER conference manuscript.

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7(1): 44–76. Published article.

Geweke, John. 2004. "Getting It Right: Joint Distribution Tests of Posterior Simulators." *Journal of the American Statistical Association* 99(467): 799–804. Published article.

Hosszejni, Darjus, and Gregor Kastner. 2021. "Modeling Univariate and Multivariate Stochastic Volatility in R with stochvol and factorstochvol." *Journal of Statistical Software* 100(12): 1–34. Published article.

Jarociński, Marek, and Peter Karadi. 2020. "Deconstructing Monetary Policy Surprises—The Role of Information Shocks." *American Economic Journal: Macroeconomics* 12(2): 1–43. Published article.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng. 2015. "Measuring Uncertainty." *American Economic Review* 105(3): 1177–1216. Author-hosted article.

Kastner, Gregor, and Sylvia Frühwirth-Schnatter. 2014. "Ancillarity-Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Estimation of Stochastic Volatility Models." *Computational Statistics & Data Analysis* 76: 408–423. Author manuscript; published article.

Kinnerud, Karin. 2025. "The Effects of Monetary Policy through Housing and Mortgage Choices on Aggregate Demand." *Quantitative Economics* 16(2): 659–703. Published article.

Krainer, John, and Milton Marquis. 2003. "Mortgage Refinancing." *FRBSF Economic Letter* 2003-29, October 3. Federal Reserve article.

Lenza, Michele, and Giorgio E. Primiceri. 2022. "How to Estimate a Vector Autoregression after March 2020." *Journal of Applied Econometrics* 37(4): 688–699. Published article.

Ludvigson, Sydney, Charles Steindel, and Martin Lettau. 2002. "Monetary Policy Transmission through the Consumption-Wealth Channel." *Federal Reserve Bank of New York Economic Policy Review* 8(1): 117–133. Published article.

Maxted, Peter, David Laibson, and Benjamin Moll. 2025. "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy." *Quarterly Journal of Economics* 140(1): 691–743. Published article; author-hosted advance-access article. The issue year is 2025; advance access was in 2024.

McCracken, Michael W., and Serena Ng. 2016. "FRED-MD: A Monthly Database for Macroeconomic Research." *Journal of Business & Economic Statistics* 34(4): 574–589. Federal Reserve working paper.

Miranda-Agrippino, Silvia, and Giovanni Ricco. 2021. "The Transmission of Monetary Policy Shocks." *American Economic Journal: Macroeconomics* 13(3): 74–107. Inspected December 18, 2019 author manuscript.

Montiel Olea, José L., James H. Stock, and Mark W. Watson. 2021. "Inference in Structural Vector Autoregressions Identified with an External Instrument." *Journal of Econometrics* 225: 74–87. Author-hosted published article.

Mortgage Bankers Association. 2022. "Mortgage Applications Decrease in Latest MBA Weekly Survey." December 7. *Weekly Applications Survey*, week ending December 2, 2022. MBA release; contemporaneous reproduction.

Musso, Alberto, Stefano Neri, and Livio Stracca. 2011. "Housing, Consumption and Monetary Policy: How Different Are the US and the Euro Area?" *Journal of Banking & Finance* 35(11): 3019–3041. Bank of Italy working-paper version.

Neal, Radford M. 2003. "Slice Sampling." *Annals of Statistics* 31(3): 705–741. Author preprint; published article.

Rella, Giacomo. 2025. "Time-Varying Interactions between Monetary and Housing Credit Policy." *Journal of Macroeconomics* 86: 103723. Author-hosted article.

Ringo, Daniel. 2024. "Inframarginal Borrowers and the Mortgage Payment Channel of Monetary Policy." Finance and Economics Discussion Series 2024-069, Board of Governors of the Federal Reserve System. Working paper.

Rubio-Ramírez, Juan F., Daniel F. Waggoner, and Tao Zha. 2010. "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference." *Review of Economic Studies*. Author-hosted working-paper version.

Sharpe, Steve A., and Shane M. Sherlund. 2015. "Crowding Out Effects of Refinancing on New Purchase Mortgages." Finance and Economics Discussion Series 2015-017, Board of Governors of the Federal Reserve System. Working paper.

Sims, Christopher A., and Tao Zha. 2006. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review* 96(1): 54–81. Published article.

Talts, Sean, Michael Betancourt, Daniel Simpson, Aki Vehtari, and Andrew Gelman. 2018. "Validating Bayesian Inference Algorithms with Simulation-Based Calibration." arXiv:1804.06788. Inspected preprint.

Vehtari, Aki, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner. 2021. "Rank-Normalization, Folding, and Localization: An Improved R-hat for Assessing Convergence of MCMC (with Discussion)." *Bayesian Analysis* 16(2): 667–718. Author manuscript; published article.

Wegner, Enrico, Lenard Lieb, Stephan Smeekes, and Ines Wilms. 2025. "Transmission Channel Analysis in Dynamic Models." arXiv:2405.18987, version 3, May 2. Inspected version.

Technical and Empirical Appendices

Companion to *Housing Finance and the Dynamic Transmission of Monetary Policy*, November 2026.

Appendix A. Data construction and measurement

A.1 Sources, vintages, and sample alignment

The dataset records source identity, original units, frequency, seasonal status, retrieval date, and the stored research vintage. This documentation follows the reproducibility principle emphasized by McCracken and Ng (2016), while retaining transformations tailored to the economic objects in this paper. Current provider pages describe the series; they do not replace the archived observations used in estimation.

Table A1. Frozen source inventory. Dates denote retrieval unless otherwise stated. SA and NSA denote seasonally adjusted and not seasonally adjusted.

Source	Frozen date	Original measurement and use
BEA DPCERA3M086SBEA	2026-08-04	Monthly SA real-PCE quantity index, 2017=100; consumption
Federal Reserve INDPRO	2026-08-03	Monthly SA industrial-production index, 2017=100
BEA PCEPILFE	2026-07-23	Monthly SA core PCE price index, 2017=100
BEA PCEPI	2026-08-03	Monthly SA headline PCE price index; housing deflator
H.15 DGS1	2026-10-09	Daily one-year constant-maturity Treasury yield, percent
Freddie Mac MORTGAGE30US	2026-07-22	Weekly thirty-year fixed mortgage quote, percent
FHFA HPIPONM226S	2026-08-05	Monthly SA national purchase-only HPI
NMDB new-mortgage archive	2026-07-22	2026Q1 release vintage; national monthly volumes and subtype shares
Z.1 HMLBSHNO	2026-07-22	Quarterly end-of-period mortgage liabilities, millions of dollars
BEA PCEC96	2026-08-03	Real-PCE overlap validation only; never spliced into the model series
S&P Case-Shiller CSUSHPISA	2026-07-22	Officially SA national HPI for the measurement alternative

The construction archive contains 333 consecutive months, January 1998–September 2025. The paper input selects 312 months through December 2023, with no missing endogenous cells. December 1997 levels are retained where needed to construct initial growth. Conditioning on January–December 1998 lags produces the common 300-month response calendar, January 1999–December 2023. The six-lag comparison retains this same response calendar.

The real-PCE quantity index and the real-PCE level series are distinct provider objects. Their overlap validation uses 234 common levels and 233 growth pairs; the second series is a check rather than a

replacement. PCE includes housing services, including imputed owner-occupied services, while purchases of existing houses are asset transactions. The purchase-financing indicator can therefore be related to consumption through associated expenditure and subsequent propagation without being a component of PCE itself. Official measurement follows the BEA NIPA Handbook, Chapter 5.

A.2 Interest-rate and mortgage-flow timing

For DGS1, the last finite daily quotation within each calendar month is selected and its date retained. Missing quotations do not cause interpolation from adjacent months. The archived daily file contains 9,593 dates, of which 9,199 have finite quotations; selection yields every required monthly observation. Interest rates remain in quoted percent units, so a response of 0.25 is 25 basis points.

For FRM30, all finite weekly PMMS quotations dated within the month receive equal weight. Months with four and five observations are treated by the same arithmetic-mean rule. PMMS changed from lender-survey to application-based information on November 17, 2022, as documented by Freddie Mac (2022). The November mean spans that transition.

NMDB reports funded closed-end first-lien originations, covering Enterprise and non-Enterprise lending. Purchase and total-refinance volumes use the published total-dollar field. The non-cash-out volume multiplies refinance dollars by the dollar-weighted other-refinance percentage. The publication's rounding precision is retained. Purchase plus refinance totals and monthly sums are checked against separately published totals; the recorded discrepancies are at most one published unit within the specified tolerances. Source definitions are in the NMDB 2026Q1 technical notes.

Both financing numerators and the Z.1 denominator are in millions of nominal dollars. The preceding quarter's end-of-period stock is held fixed for the three months of the current quarter. This is a reference-period lag, not a reconstruction of when each revised release became publicly available. The flow and stock populations differ, which motivates the aggregate market-size interpretation.

For compatible unadjusted counts and volumes, $V_t = N_t \bar{L}_t$, so

$$\Delta \log I_t = \Delta \log N_t + \Delta \log \bar{L}_t - \Delta \log B_{q(t)-1}. \tag{A1}$$

For the actual adjusted purchase input, the exact identity is instead the change in the adjusted numerator less the change in the denominator. The raw count-size identity is not silently assigned to the separately adjusted numerator. The published financing indicators therefore combine transaction counts, average loan sizes, and stock normalization.

A.3 Seasonal adjustment and treatment of extreme growth

Purchase adjustment uses the X-11 option in X-13, a log transformation, and the frozen ARIMA $(0,1,1)(0,1,1)_{12}$ specification with recorded calendar and intervention terms. The adjustment window ends in September 2025, while the estimation input ends in December 2023. The software versions are seasonal 1.10.0 and x13binary 1.1.61.2. The exact call, full diagnostics, and endpoint comparisons remain in replication materials. The adjusted purchase numerator has a QS p-value of 1.00. For unadjusted rate-term intensity, the historical joint test of eleven calendar-month indicators, with January as the reference, gives $F(11, 321) = 0.179$ and $p = 0.9985$ over January 1998–September 2025. This test concerns month-of-year

differences in the observed level series. Officially seasonally adjusted macroeconomic and housing series retain the provider adjustment.

Growth caps use type-7 quantiles of January 1998–February 2020 and apply once to all 312 input months. Exact thresholds are $[-0.477050005690201, 0.875665244842141]$ for PCE and $[-2.1557201812255, 1.21818074987979]$ for IP. PCE changes comprise 12 observations before March 2020, six in March–December 2020, and four in 2021–2023. IP changes comprise six, six, and three observations in the same periods. These are variable-month counts; no dates are deleted.

The native-unit input exhibit, PCE treatment exhibit, and IP treatment exhibit show the values supplied to the model and their untreated companions. They are Appendix Figure A1–A3 source exhibits for typesetting. Full seasonal endpoint plots remain replication material.

A.4 Descriptive statistics and persistence

Table A2. Model-input descriptive statistics, 312 months. Units follow main-text Table 1. SD is the sample standard deviation; lag-one correlation uses 311 adjacent pairs. These are descriptive moments, not prior scales.

Variable	Mean	SD	Minimum	Maximum	Lag-one correlation
Treated real PCE growth	0.220328	0.324885	-0.477050	0.875665	-0.062496
Treated IP growth	0.082439	0.622594	-2.155720	1.218181	0.171286
Core inflation	0.164859	0.127298	-0.569854	0.701702	0.382405
One-year Treasury yield	2.152308	1.977917	0.050000	6.370000	0.994068
Mortgage rate	5.246373	1.446088	2.684000	8.515000	0.990949
Real FHFA growth	0.219584	0.525695	-1.585063	1.463226	0.753409
Purchase intensity	0.861488	0.289294	0.329000	1.345721	0.989671
Rate-term intensity	0.632645	0.517432	0.030892	3.164727	0.945552

The response-sample means used to scale financing IRFs are instead approximately 0.846641 and 0.622930. The distinction reflects conditioning on the initial twelve months, not a change in the underlying intensity definition.

FHFA growth rejects the ADF null at 5% in both reported samples, while its level-KPSS statistics are below the 5% critical value 0.463. Interest-rate levels instead show strong persistence. The level specification preserves the financing-cost interpretation; finite-horizon stability sensitivity addresses the resulting posterior dynamics. These diagnostics do not select transformations by the signs of subsequent policy responses.

Table A3. Selected persistence diagnostics and their actual samples. ADF tests a unit-root null; KPSS tests the indicated stationarity null. Results on the historical unadjusted purchase series are not evidence about the final adjusted purchase input.

Series and sample	ADF specification and statistic	KPSS specification and statistic
FHFA real growth, 312 months	Drift; BIC among up to 12 lags, selects 2; -3.60501	Level; bandwidth 12; 0.30970
FHFA real growth, 266 pre-COVID months	Same lag selection; -3.08167	Level; bandwidth 12; 0.29728
DGS1 level, 312 months	Drift / trend; selects 2; -1.91775 / -1.58040	Level / trend; bandwidth 12; 0.78467 / 0.28497
Change in DGS1, 311 months	Drift / trend; selects 1; -7.30853 / -7.44188	Level / trend; bandwidth 12; 0.22445 / 0.05116
FRM30 level, 312 months	Drift / trend; selects 2; -1.43238 / -0.60482	Level / trend; bandwidth 12; 1.54535 / 0.26934
Rate-term intensity, 333 months through 2025-09	Drift / trend; BIC selects 1; -4.24837 / -4.36241	Level / trend; Bartlett long bandwidth 16; 0.23082 / 0.04628
Historical NSA purchase intensity, same 333 months	Drift / trend; selects 12; -1.99794 / -1.85213	Level / trend; bandwidth 16; 0.55736 / 0.33928

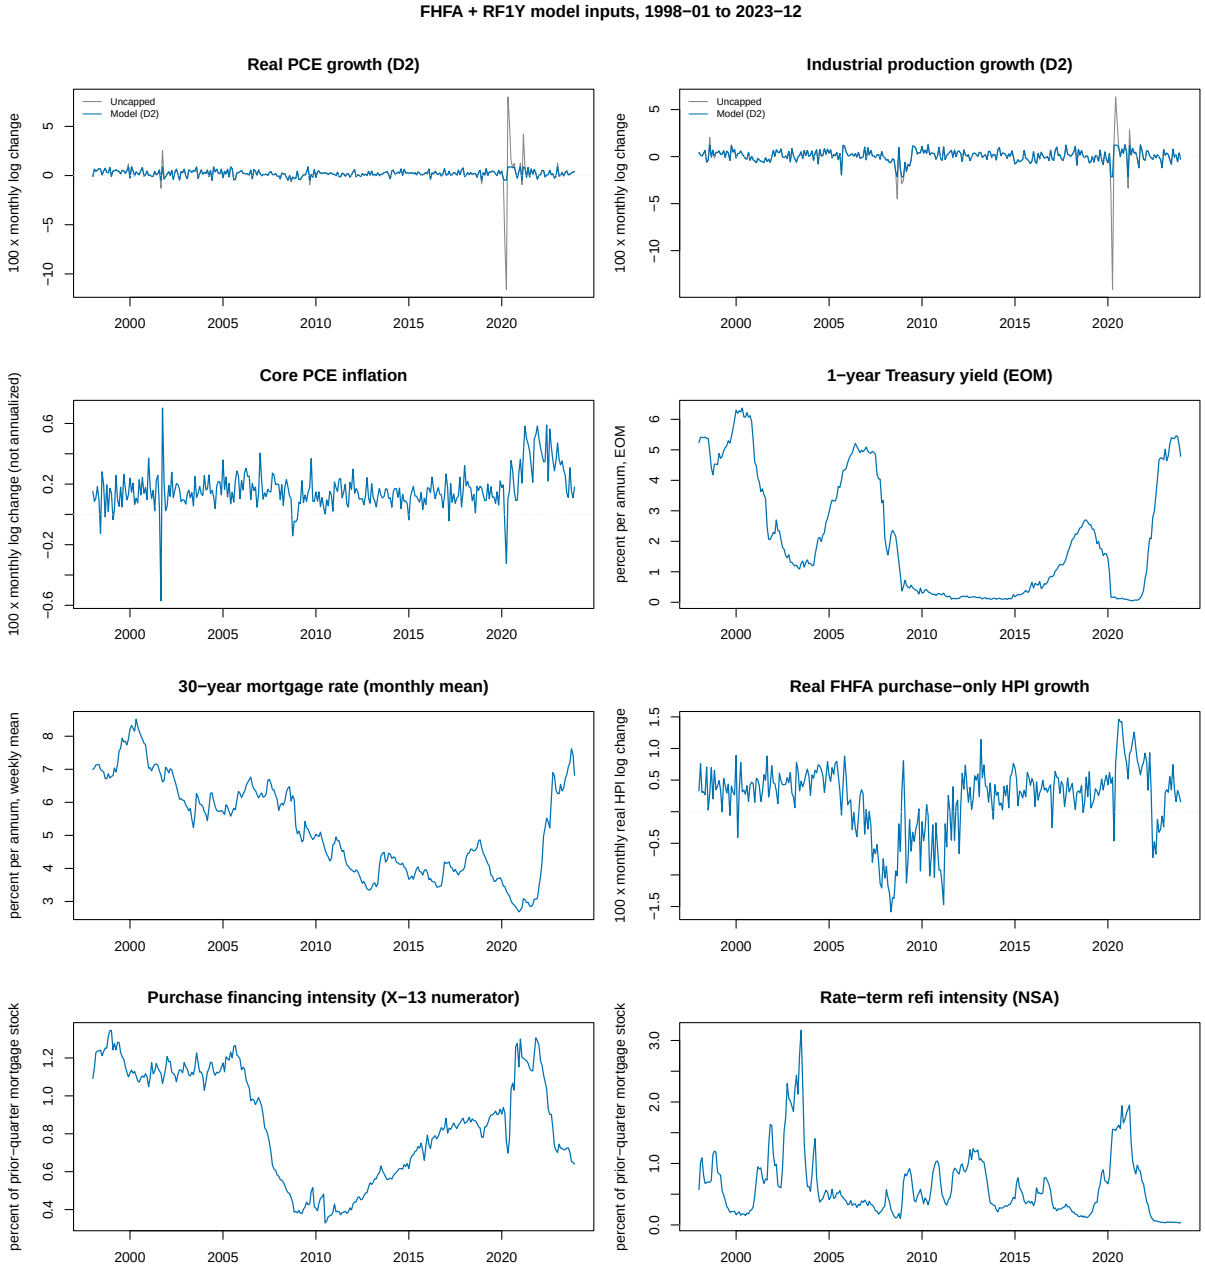


Figure A1. Model inputs in native units. Monthly observations from January 1998 through December 2023: growth and inflation are measured as 100 times monthly log differences, interest rates in percent, and financing intensities as percentages of lagged mortgage debt. The consumption and industrial-production inputs incorporate the fixed caps described in Appendix A; these are observed series, with no posterior draws or uncertainty intervals.

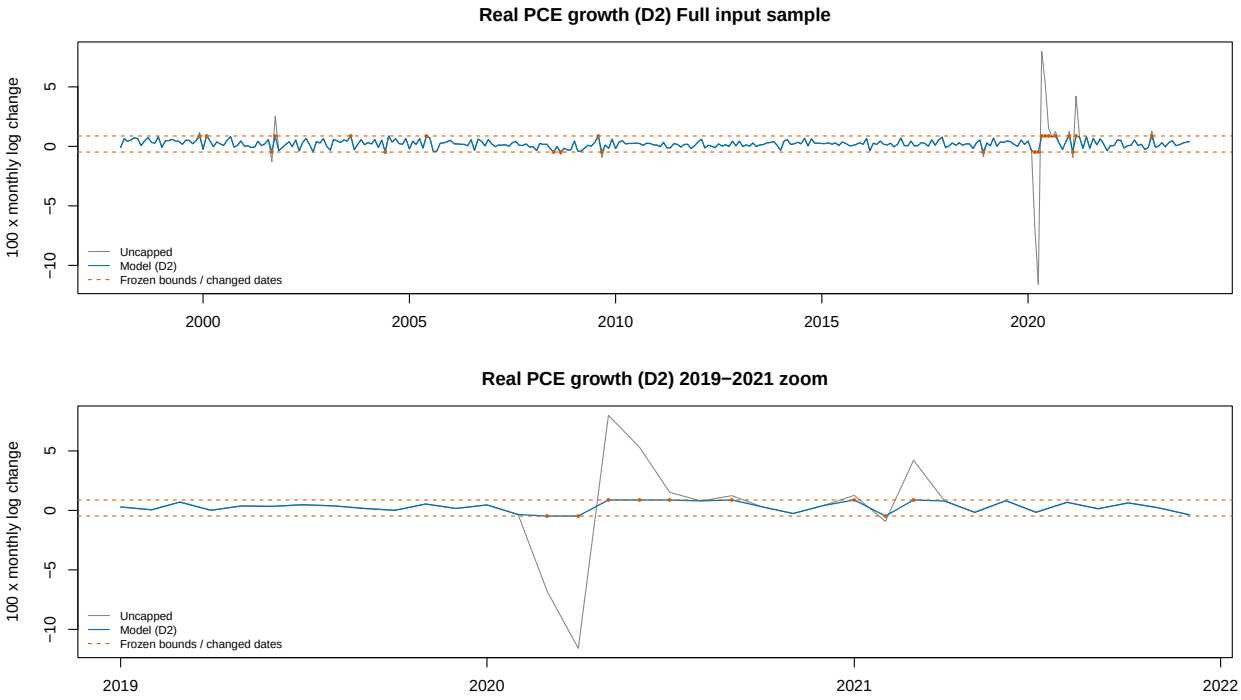


Figure A2. Treatment of consumption-growth extremes. Untreated and capped monthly real consumption growth, in percentage points of monthly log growth, over the January 1998–December 2023 input sample. The fixed thresholds use pre-March 2020 quantiles as described in Appendix A; the figure compares observed series and contains no posterior uncertainty intervals.



Figure A3. Treatment of industrial-production growth extremes. Untreated and capped monthly industrial-production growth, in percentage points of monthly log growth, over the January 1998–December 2023 input sample. The fixed thresholds use pre-March 2020 quantiles as described in Appendix A; the figure compares observed series and contains no posterior uncertainty intervals.

Appendix B. Priors, posterior simulation, and validation

B.1 Scale calibration and posterior target

Let $x_t = (y'_{t-1}, \ldots, y'_{t-12}, 1)'$ and let Φ be the 97×8 coefficient matrix, with lag-major rows and the intercept in row 97. Thus $u_t = y_t - \Phi' x_t$. With $b = \mathrm{vec}(\Phi)$ stacking equation columns, the prior covariance V_0 is diagonal and uses the variances in main-text Table 2. Prior means are zero, including own lags. The prior does not impose stable companion roots.

Table B1. Fixed calibration SDs. Each comes from a no-intercept AR(6) on the 312 input months, with $s_j^2 = RSS_j/306$. Estimation uses the full-precision saved values.

Variable	s_j
Consumption growth	0.327240557
IP growth	0.595515359
Core inflation	0.109225012
One-year Treasury yield	0.194110127
Mortgage rate	0.183075112
Real FHFA growth	0.321025240
Purchase intensity	0.0404422984
Rate-term intensity	0.155675108

The posterior is conditional on these data-calibrated scales; it does not integrate their estimation uncertainty. The half-normal SV innovation-scale prior is equivalently $\sigma_j^2 \sim \mathrm{Gamma}(1/2, 12.5)$ under shape-rate notation. Its support approaches zero but contains no point mass at constant volatility. SV stationarity is distinct from stability of the VAR mean dynamics.

For $\theta = (\Phi, A_0, \mu, \varphi, \sigma)$, the target is

$$p(\theta, h_{0:T} \mid Y, X; s) \propto L(Y \mid X, \Phi, A_0, h_{1:T}) p(\Phi \mid s) p(A_{0,free} \mid s) \prod_j \left[p(\mu_j \mid s_j) p(\varphi_j) p(\sigma_j) p(h_{j0} \mid \mu_j, \varphi_j, \sigma_j) \prod_{t=1}^{T} p(h_{jt} \mid h_{j,t-1}, \mu_j, \varphi_j, \sigma_j) \right]. \quad \text{(B1)}$$

The likelihood includes $T \log |\det A_0|$ and the stationary initial-state density enters the target through the path prior. Singular structural matrices have zero likelihood density; no positive-determinant restriction is added.

B.2 Coupled Gaussian update of mean coefficients

With $\Omega_t = A_0' D_t^{-1} A_0$, the Gaussian conditional is

$$\Lambda = V_0^{-1} + \sum_t \Omega_t \otimes (x_t x_t'), \qquad q = \sum_t (\Omega_t y_t) \otimes x_t, \qquad b \mid \cdot \sim N(\Lambda^{-1} q, \Lambda^{-1}). \quad \text{(B2)}$$

Off-diagonal elements of Ω_t couple equations. The implementation draws the full 776-dimensional vector using a Cholesky factor and triangular solves, rather than treating the equations as unrelated regressions. The time-varying-precision Gaussian structure has a VAR precedent in Clark (2010, Section 3.3); the coefficient layout and direct intercept prior differ from his steady-state parameterization. Positive prior precision establishes positive definiteness in exact arithmetic, not a guarantee of numerical conditioning for arbitrarily explosive synthetic data.

B.3 Nonrecursive contemporaneous-row updates

After the mean update, native innovations are recomputed and contemporaneous rows are visited in random order. For row n, let a_n collect its free coefficients, w_t the corresponding native innovations, and P_n its diagonal prior covariance. In the baseline, the structural innovation is $u_{nt} + a_n' w_t$. The Gaussian part has

$$V_n^{-1} = P_n^{-1} + \sum_t e^{-h_{nt}} w_t w_t', \qquad m_n = -V_n \sum_t e^{-h_{nt}} w_t u_{nt}. \tag{B3}$$

The row conditional is proportional to $|\det A_0(a_n)|^T N(a_n; m_n, V_n)$. Holding other rows fixed, the determinant is affine: $d_n + g_n' a_n$. Thus its non-Gaussian part depends on the scalar $\xi_n = g_n' a_n$:

$$p(\xi_n \mid \cdot) \propto |d_n + \xi_n|^T N(\xi_n; g_n' m_n, g_n' V_n g_n). \tag{B4}$$

Three slice transitions update this scalar using stepping out and shrinkage, following Neal (2003, Sections 4.1–4.3). With $v_n = g_n' V_n g_n > 0$, the row is reconstructed by

$$a_n = m_n + \frac{V_n g_n}{v_n}(\xi_n - g_n' m_n) + z_n - \frac{V_n g_n}{v_n} g_n' z_n, \qquad z_n \sim N(0, V_n). \tag{B5}$$

The slice width is $\sqrt{v_n}$ and stepping-out control is 50, with randomized left/right allowances. A determinant-independent row receives a Gaussian update; a row with no free coordinates remains fixed. A near-zero numerical branch applies at $v_n \le 10^{-300} \max(1, \|g_n\|^2)$. The determinant projection is an algebraic implementation device, not a new identification restriction. Finite slice transitions are Markov updates, not independent exact conditional draws.

B.4 Volatility update and chain settings

After updating the coefficients, the sampler forms $E = A_0 U$. Conditional on these structural innovations, the eight volatility blocks separate. Each updates its path, stationary initial state, and three SV parameters. Centered/noncentered interweaving follows Kastner and Frühwirth-Schnatter (2014, Section 3). The implementation uses a source-bound one-step stochvol adapter, with the auxiliary-mixture correction described by Hosszejni and Kastner (2021, Sections 2.3 and 4.6).

One outer sweep comprises the coupled mean update, randomized contemporaneous-row updates, and five consecutive one-step SV calls per equation. Path, parameter, initial-state, and adaptation state carry forward between calls. Adaptation remains enabled after warmup. The routine's internal `multi_asis=5` setting is distinct from the five outer calls. This describes the executed adaptive simulation, without claiming a fixed post-warmup kernel.

Table B2. Empirical simulation settings.

Setting	Value
Chains	Four
Seeds	2026102805–2026102808
Warmup per chain	5,000 sweeps
Post-warmup per chain	30,000 sweeps
Scalar monitors	102 saved each post-warmup sweep; 100 in the original screen
Complete-state storage	Every tenth post-warmup sweep
Total complete states	12,000
Complete-state contents	28 contemporaneous coefficients, 776 mean coefficients, 24 SV parameters, 8 initial log variances, 2,400 dated log variances
Reported sweep errors	Zero in each empirical chain

Mean coefficients initialize at zero; free contemporaneous coefficients use prior draws. Initial log-variance levels equal $\log s_j^2 + \delta_j$, $\delta_j \sim N(0, 0.5^2)$, and paths start flat at those levels. Initial persistence draws are clipped to $[-0.95, 0.95]$, and initial volatility innovation SDs are bounded below at 0.05. These rules initialize the chains rather than constrain their posterior support.

B.5 Empirical mixing, reconstruction, and stability

Table B3. Empirical posterior diagnostics. Full-state statistics cover all 3,236 components. IRF statistics cover eight variables at horizons 1, 6, 12, 24, and 36. Monitoring iterations, complete states, and effective sample sizes are distinct quantities.

Object	Maximum $\widehat{R}$	Minimum bulk ESS	Minimum tail ESS
FHFA unthinned monitors	1.001527	5,671.55	11,159.08
FHFA complete-state components	1.001294	5,071.73	7,674.78
FHFA selected IRF functions	1.000975	10,345.29	10,976.64
Case-Shiller unthinned monitors	1.001542	4,664.83	9,953.79
Case-Shiller complete-state components	1.001494	4,418.73	7,692.23

The statistics use rank-normalized split and folded diagnostics and bulk/tail ESS following Vehtari et al. (2021). The original empirical screen requires $\widehat{R} \leq 1.01$ and both ESS measures at least 1,000. The unthinned tail implementation includes equality at both tail thresholds; the separate complete-state review uses CDF indicators at both quantiles. The latter upper-tail ESS is equivalent to its strict upper-tail complement, not necessarily to an inclusive indicator when there are ties. The table preserves the verified convention for each set.

All complete states were matched to their original chain and iteration identifiers. Independent reconstruction of normalized impact vectors has maximum absolute error 5.00×10^{-16}, and main/stable IRF

summaries have maximum discrepancy 2.56×10^{-15}. These checks support the numerical mapping from the fitted states to the reported responses.

Stability is defined by the spectral radius of the 96×96 VAR companion matrix, excluding the intercept. The main FHFA population contains 11,274 stable and 726 nonstable states; the Case-Shiller population contains 10,987 and 1,013. Main summaries retain all states. Stable-only summaries condition and reweight within the stable subset. A finite-horizon response remains calculable for a nonstable state with finite parameters, while stationary infinite-horizon formulas require additional conditions.

B.6 Finite simulator checks and incomplete full-prior calibration

Joint-distribution checks compare direct prior simulation with successive conditional simulation, using functions for which the relevant variance exists. This follows Geweke (2004). The small test system has four variables, one lag, six response observations, and a nonrecursive cycle. Its prior scales differ from the empirical model. Each positive-control arm uses eight successive-conditional chains, 1,000 warmup and 50,000 retained iterations per chain, together with 400,000 independent prior draws.

The adaptive and frozen-adaptation arms have maximum absolute core comparison statistics of 2.838938 and 2.075464. Both meet the specified finite screen. An omitted-determinant negative control fails, with maximum 5.050174 and 13,907 recorded errors. This demonstrates detection of that defect, not guaranteed detection of all implementation errors.

The screen covers 102 of 114 recorded functions. Twelve raw squared-log-volatility functions are excluded because the specified persistence prior gives a divergent fourth moment for stationary log volatility, invalidating the ordinary finite-variance Monte Carlo comparison for those squares. Truncated stationary-deviation functions remain. The tail exclusion is part of the validation scope, rather than a result that the full prior tail has been checked.

Full-size simulation-based calibration generates eight-variable, twelve-lag datasets with 300 observations and zero initial lags, conditional on the empirical calibration scales. It uses 100 planned datasets for each of FHFA and Case-Shiller scales, two chains per replicate, 2,000 warmup and 6,000 retained sweeps per chain, and 1,200 stored states to evaluate 76 functions. The rank logic follows Talts et al. (2018), including the qualification that MCMC autocorrelation affects its ideal uniform-rank argument.

Table B4. Full-size calibration completion. Coverage averages concern the completed cases only.

Calibration panel	Completed / planned	Stable generators completed	Failed	Mean nominal-90% coverage among completed cases
FHFA scales	74 / 100	66 / 66	26	0.899360
Case-Shiller scales	75 / 100	67 / 67	25	0.890351

All 51 failures arise from nonstable generating systems and supply no posterior ranks. Severe floating-point conditioning in the Gaussian mean block is the principal localized failure. Every failed case remains recorded; parameters and data were not redrawn to replace it. Successful-case ranks and coverage therefore do not certify calibration over the full prior. These artificial-data failures are distinct from discarding empirical posterior states: no such empirical-state deletion is used in the main results.

Appendix C. Identification and supplementary response evidence

C.1 Restrictions and retained relationships

Writing $\beta_{ij} = -(A_0)_{ij}$ gives the right-side representation $u_i = \sum_{j \neq i} \beta_{ij} u_j + \varepsilon_i$. A zero excludes a conditional direct relationship, rather than every possible path between two variables.

Table C1. Nonoverlapping groups of contemporaneous zeros.

Group	Exclusions	Count
Slow-block timing	Each of C, IP, inflation, H excludes r, M, P, R	16
Policy financing feedback	r excludes M, P, R	3
Slow-block completion	C excludes IP, inflation, H; IP excludes inflation, H; inflation excludes H	6
Financing completion	P and R exclude r; P excludes R	3

The slow-block and policy exclusions provide the block argument for the policy column in main-text equations (9)–(10). The other nine restrictions complete the structural system. At a fixed compatible covariance, policy-column identification need not depend on naming all other shocks; re-estimating a different structural completion can nevertheless change the posterior covariance paths and policy responses.

Allowing housing to enter both financing equations retains valuation and collateral relationships. Refinancing without equity extraction can still depend on collateral eligibility. Allowing M to respond to both quantities retains a pricing-capacity link; applications and intermediary pricing evidence in Fuster, Lo, and Willen (2016) motivates that flexibility, without directly estimating a coefficient for completed national originations. Excluding direct R from the P equation is a completed-transaction timing approximation, while indirect R→M→P and lagged R effects remain. Capacity crowding is a specific alternative to that restriction.

For a policy disturbance with zero slow-block impact, financing responses satisfy

$$\kappa = 1 - \beta_{MP}\beta_{PM} - \beta_{MR}(\beta_{RM} + \beta_{RP}\beta_{PM}),$$
$$b_M^{25} = 0.25\beta_{Mr}/\kappa, \quad b_P^{25} = \beta_{PM} b_M^{25}, \quad b_R^{25} = (\beta_{RM} + \beta_{RP}\beta_{PM}) b_M^{25}. \tag{C1}$$

The invertible financing subsystem requires $\kappa \neq 0$. This expression shows why excluding direct yield effects from P and R does not exclude their impact responses. It also shows that mortgage pricing is the maintained initial rate-entry point for financing; later propagation uses all estimated lags. None of these coefficients has an imposed sign.

C.2 Constructive identification of the full matrix

Assume an invertible A_0 satisfying the stated mask and unit diagonal, positive diagonal structural variances, and $B_{M,r} \neq 0$. The first four slow-block rows, recovered in order C, IP, inflation, H, are uniquely deter-

mined by their triangular structure and covariance orthogonality. The policy row follows from its Schur-complement construction. These steps identify the policy column without requiring the added condition on $B_{M,r}$.

For purchase, let $q = (A_0)_{P,M}$. Its remaining financing coordinates in (r, M, P, R) order are $(0, q, 1, 0)$. The inverse identities give

$$B_{P,r} = -qB_{M,r}, \quad B_{P,P} = 1 - qB_{M,P}, \qquad \det \begin{pmatrix} B_{M,r} & B_{M,P} \\ B_{P,r} & B_{P,P} \end{pmatrix} = B_{M,r}. \tag{C2}$$

The nonzero condition supplies the required rank for the remaining purchase coefficient and for the policy/purchase constraints on the refinancing row. Orthogonality to the already determined rows, the mask, and the unit diagonal then determine P and R. Invertibility makes the remaining M row unique. This is a sufficient argument for this matrix, not a general theorem for arbitrary zero restrictions or a proved converse at $B_{M,r} = 0$.

Near-zero impact mortgage-rate transmission can weaken finite-sample information even away from exact degeneracy. Numerical invertibility and response-normalization checks do not measure distance from this exceptional set. The argument distinguishes conditional identification from the empirical precision and economic validity of the restrictions.

C.3 External diagnostic implementation

For each external series z_t and fitted state, the auxiliary regression is $q_t = \alpha + \beta z_t + v_t$, where q_t is either the native yield innovation or a structural innovation. The policy innovation is oriented by $\operatorname{sgn}(B_{44})$; other innovations retain their equation normalization. Let the row of X be $x_t' = (1, z_t)$ and $g_t = x_t \widehat{v}_t$. HC1 uses

$$\widehat{V}_{HC1} = \frac{n}{n-2}(X'X)^{-1} \left(\sum_t g_t g_t' \right) (X'X)^{-1}. \tag{C3}$$

HAC12 uses Bartlett weights with calendar-month rather than compressed-row lags:

$$\widehat{V}_{HAC12} = (X'X)^{-1} \left[\sum_t g_t g_t' + \sum_{\ell=1}^{12} \left(1 - \frac{\ell}{13} \right) \sum_{t-s=\ell} (g_t g_s' + g_s g_t') \right] (X'X)^{-1}. \tag{C4}$$

Means and regressions are recomputed on each omission support. The single-slope statistic is $F = \widehat{\beta}^2 / \widehat{V}_{\beta\beta}$. HC1 and HAC12 summarize the same association using different regression-error covariance estimates; a larger HAC12 statistic can arise from the serial covariance of the regression scores. The fitted parameter states remain unchanged when a month is omitted. These statistics describe conditional association rather than a calibrated weak-instrument test with a universal cutoff. Inference for a separately instrument-identified response must additionally account for instrument-moment uncertainty and its interaction with estimated model dynamics; see Montiel Olea, Stock, and Watson (2021).

The MP1/FF4 source is Acosta et al. (2026), using the September 17, 2026 research vintage. May 2010 and October 2019 lack complete eligible common records. No-event months otherwise remain zero. The two series use the same statement-event set, rather than differing event selections.

Table C2. Correlations of announcement surprises with all structural innovations. Brackets are fifth–ninety-fifth percentiles across fitted states, not independently estimated sampling confidence intervals.

Equation innovation	MP1 median [5%, 95%]	FF4 median [5%, 95%]
Consumption	0.151 [0.125, 0.175]	0.125 [0.096, 0.153]
Industrial production	0.022 [-0.005, 0.049]	0.021 [-0.007, 0.047]
Inflation	0.090 [0.064, 0.116]	0.025 [-0.002, 0.051]
Policy	0.373 [0.343, 0.399]	0.277 [0.241, 0.310]
Mortgage rate	-0.103 [-0.152, -0.057]	-0.033 [-0.076, 0.009]
Housing price	0.009 [-0.018, 0.036]	-0.042 [-0.067, -0.016]
Purchase financing	0.096 [0.071, 0.120]	0.083 [0.055, 0.110]
Rate-term refinancing	-0.027 [-0.049, -0.002]	-0.054 [-0.078, -0.028]

The native yield-innovation layer gives common-sample median correlations 0.378 and 0.282 for MP1 and FF4, with HC1/HAC12 statistics 11.42/26.54 and 4.20/6.69. For MP1 the absolute centered-cross-product shares attributable to January 2008 and March 2020 are 28.06% and 16.39%; for FF4 they are 18.60% and 11.43%. These shares concern native-yield association and are not automatically structural-policy contribution shares.

The policy innovation exceeds every other innovation in absolute correlation in all 12,000 common-sample states for both series. Omitting January 2008 gives 11,999 states for MP1 and 11,956 for FF4; omitting March 2020 gives all 12,000 for both. This preserves the relative correspondence while acknowledging the concentration of association strength in influential months.

C.4 Response magnitudes, cumulative units, and alternative fits

The complete macro-response exhibit supplies consumption, IP, and inflation; the Case-Shiller comparison shows all variables under the alternative housing measure. These are Appendix Figure C1–C2 source exhibits.

Cumulative growth responses are summed within each state. At month 36, cumulative H has median -1.074222 in $100 \times$ log level units. Cumulative consumption has median -0.088665 and a 90% interval $[-0.464598, 0.277697]$. The latter describes the level implied by the treated consumption-growth series. It is distinct from cumulative dollars of forgone expenditure. The cumulative-level exhibit preserves both 68% and 90% intervals.

Table C3. Selected baseline responses. Entries are posterior medians [90% pointwise intervals]. Rate responses are in basis points; other responses are percentage points. H denotes monthly house-price growth.

Variable	Month 0	Month 6	Month 12	Month 36
One-year yield	25.00	33.83	41.06	12.54
	[25.00,25.00]	[28.12,40.96]	[32.54,53.22]	[-7.67,39.76]
Mortgage rate	8.54	15.56	17.83	14.23
	[6.44,10.70]	[11.07,20.75]	[11.42,25.81]	[2.25,31.57]
H	0	-0.01710	-0.02527	-0.03743
	[0,0]	[-0.04271,0.00887]	[-0.06127,0.00998]	[-0.10322,0.01260]
P	-0.00028	-0.00946	-0.01557	-0.03261
	[-0.00172,0.00122]	[-0.01565,-0.00367]	[-0.02640,-0.00604]	[-0.06714,-0.00815]
R	-0.00241	-0.03578	-0.03334	-0.00527
	[-0.00603,0.00050]	[-0.06018,-0.01269]	[-0.06482,-0.00454]	[-0.04593,0.02727]

The financing-scale exhibit divides P/R responses by their fixed response-sample means and multiplies by 100. Positive fixed rescaling leaves the interval's relation to zero unchanged. It supplies a benchmark relative to usual monthly financing intensity rather than a new response estimate.

For stable-state sensitivity, the month-36 mortgage-rate median changes from 14.23 to 13.92 basis points and its 90% upper endpoint from 31.57 to 29.48. H changes from -0.03743 to -0.03607 percentage points; P from -0.03261 to -0.03179. The principal reported interval patterns are retained. The full gamma comparison and its stable-state counterpart retain the eight-variable feedback alternatives. Changing gamma changes the median refinancing trough from month 13 to month 6 to month 5 across negative, zero, and positive gamma; this variability cautions against interpreting one median trough month as an identified institutional delay.

The main figure and tables are pointwise posterior summaries, not simultaneous bands or predictive distributions. They integrate uncertainty within the fitted model and fixed calibration, rather than uncertainty across all possible identification schemes or model selections.

C.5 High-frequency surprises as instruments: first-stage strength and an exploratory proxy identification

This appendix documents why the paper does not use high-frequency monetary-policy surprises as the identifying instrument, and what the proxy route delivers where it can be run. All first-stage statistics are single-slope auxiliary regressions of the policy-indicator innovation on the surprise, as in Appendix C.3; "state median" summarizes the statistic across fitted posterior states. F values near or below 10 are a descriptive screen, not a calibrated test.

C.5.1 First-stage strength across instruments, samples, indicators, and systems

Table C5a. First-stage strength of high-frequency surprises. HC1 and HAC12 are the two regression-error covariance estimators of Appendix C.3; "conv." is the conventional homoskedastic F. Rows 1–4 are reference regressions with no VAR posterior; rows 5–12 and 15 are state medians; rows 13–14 are reference regressions using the native policy-indicator innovation averaged across saved states. The literature's four-variable system is log IP, log CPI, the excess bond premium, and the two-year zero-coupon yield; MPS is the Bauer–Swanson (2023) surprise series; STMT and ME are the public USMPD statement-window and statement-plus-press-conference candidates. In rows 13–15, P0 sums event-level orthogonalized surprises within each month, whereas M0 is the separately published monthly orthogonalized series. These rows use earlier eight-variable, p12, Case-Shiller Cholesky-SV models.

	Instrument	Policy indicator / system	Sample	N	conv. F	HC1 F	HAC12 F
1	MPS	Four-variable literature system	1988-02–2019-12	383	8.03	5.99	4.62
2	MPS	Four-variable literature system	1999-01–2019-12	252	0.14	0.10	0.09
3	MPS	Current eight variables, GSW two-year yield	1999-01–2019-12	252	0.06	0.05	0.06
4	STMT	Current eight variables, GSW two-year yield	1999-01–2019-12	252	0.12	0.08	0.14
5	STMT	Earlier model, federal-funds/Wu–Xia indicator	all 300 months	300	9.82	2.85	2.75
6	STMT	same	non-ELB months	192	10.12	7.53	4.31
7	STMT	same	ELB months	108	2.62	0.38	0.38
8	STMT	Earlier model, two-year yield indicator	ELB months	108	30.53	11.03	13.01
9	MP1	**Current model, one-year yield**	common 1999–2023	298	49.2	11.42	26.54
10	MP1	Current model	omit 2008-01	297	—	5.42	11.67
11	MP1	Current model	omit 2020-03	297	—	6.58	16.77
12	FF4	Current model	common 1999–2023	298	—	4.20	6.69
13	Orthogonalized MPS (P0)	Earlier model, federal-funds/Wu–Xia indicator	1999-01–2023-12, eligible months	289	0.83	0.37	0.21
14	Orthogonalized MPS (M0)	Earlier model, federal-funds/Wu–Xia indicator	same	289	0.34	0.22	0.12
15	Orthogonalized MPS (P0)	Earlier model, one-year yield indicator	same	289	3.57	1.51	1.76

Sources: rows 1–4, `results/runs/CC-20261013-02_U/tables/main_table.csv`; rows 5–8, `results/runs/CC-20261012-02_U/tables/table_01_stmt_03a.csv` and `table_03_stmt_rf2y.csv` (ME variants give the same picture); rows 9–12, `results/runs/CC-20261028-33_U/summary_type7.csv` (native yield innovation; Table 3 in the main text reports the structural-innovation version from 36_U, which is within 0.3 of these). Rows 13–14 use `results/runs/CC-20261008-03/02P/packet/relevance_reference.csv`, `mask=full`, `row_type=reference`, proxies `P0` and `M0`; row 15 reports medians over the 2,000 `proxy=mps_orth_p0`, `cell=native_all`

states in `results/runs/CC-20261018-03_I/screen/rf1y_sv/baseline_relevance.csv`. For rows 13–15, the 289 eligible months exclude September 2001 and March–December 2020 from the 1999–2023 span. P0 and M0 use the frozen July 23, 2026 SF Fed workbook; their distinct construction is documented in `results/runs/CC-20261008-03/02P/packet/source_qualification.md`. The 2,000- and 12,000-state medians are posterior dispersion summaries, not confidence intervals.

Three features of Table C5a matter for the design. First, the post-1999 weakness appears in the literature's own four-variable system with its own instrument (row 2 against row 1), so it is not created by moving to eight variables, by stochastic volatility, or by the choice of surprise. Second, the strength that remains on the current model is concentrated: omitting either of two months halves the HC1 F (rows 10–11), and in the earlier model the two months account for 58% of the absolute centered covariance between MP1 and the yield innovation (`CC-20261024-01_U`). Third, the ELB/non-ELB split does not produce a uniform pattern: with the federal-funds indicator the surprises are informative only outside the ELB (rows 6–7), with the two-year indicator only inside it (row 8), so no single subsample offers a strong instrument for the full-sample question. The gap between the conventional F of 49 and the HC1 F of 11 in row 9 is itself a symptom of this concentration: a few large-surprise months dominate the regression.

The orthogonalized-surprise diagnostics extend the comparison to two earlier policy-indicator specifications (rows 13–15). Both constructions yield HC1 and HAC12 F values below one with the federal-funds/Wu–Xia indicator. With the one-year yield, the event-orthogonalized construction yields state-median HC1 and HAC12 F values of 1.51 and 1.76. Weak relevance is therefore also present with orthogonalized MPS in these earlier specifications; these are separate diagnostics from the four-variable raw-MPS comparison in rows 1–2.

Weak-instrument-robust inference is the standard response (Montiel Olea, Stock, and Watson, 2021). For this model it would require Anderson–Rubin-type sets that carry the BVAR-SV parameter uncertainty along with the instrument-moment uncertainty. A prespecified synthetic study of that construction (`CC-20261025-04_U`) was inconclusive on two of three designs, and the construction has not been applied to the empirical sample. The paper therefore does not report proxy-identified credible sets.

C.5.2 An exploratory proxy identification

What the proxy route delivers can still be seen on an earlier specification where the computation was carried out: the Case-Shiller, one-year-yield, recursive (Cholesky) SV model, with 2,000 saved states, MP1 as the instrument, and impact vectors formed as $b = 0.25\,\mathrm{Cov}(u, z)/\mathrm{Cov}(u_r, z)$ within each state and propagated with that state's dynamics (`CC-20261022-04_U`). Only pointwise medians are available; no credible or confidence bands were constructed, for the reasons above. The same calculation was repeated with January 2008 and with March 2020 omitted from the instrument moments, holding the fitted states fixed.

Table C5b. Exploratory proxy-identified median responses, earlier Case-Shiller specification. Pointwise medians across 2,000 states of the earlier one-year-yield Cholesky-SV model; no uncertainty intervals; +25-basis-point yield normalization. Rates in basis points, other variables in the units of Table 1. Entries in brackets are the medians when January 2008 and March 2020, respectively, are omitted from the instrument moments.

Variable	Month 0	Month 6	Month 12	Month 36
One-year yield	25.0	33.9	42.1	19.1
		[33.9, 34.6]	[42.7, 43.1]	[21.5, 20.3]
Mortgage rate	2.4	10.7	15.3	18.0
	[−1.2, 3.2]	[8.1, 11.7]	[13.7, 16.3]	[18.6, 19.0]
Rate-term refinancing intensity	−0.020	−0.036	−0.038	−0.017
	[−0.036, +0.015]	[−0.041, −0.018]	[−0.043, −0.028]	[−0.020, −0.016]
Purchase intensity	+0.015	+0.005	−0.001	−0.023
	[+0.024, +0.017]	[+0.012, +0.008]	[+0.006, +0.001]	[−0.018, −0.022]
Real house-price growth (Case-Shiller)	+0.031 [−0.009, +0.061]	+0.015	−0.002	−0.046
Industrial-production growth	+0.113 [+0.184, −0.051]	+0.033	+0.004	−0.012
Consumption growth	+0.163 [+0.256, +0.131]	+0.014	+0.006	−0.006
Core inflation	+0.031 [+0.059, +0.015]	+0.005	+0.003	−0.007

Source: `results/runs/CC-20261022-04_U/comparison/outputs/median_lines.csv`, cells `native_full`, `native_without_2008_01`, `native_without_2020_03`, all states. The stable-state subset (1,860 states) gives the same picture.

Two things stand out. The earlier proxy calculation shares qualitative features with the restriction-based responses: refinancing contracts by about 0.04 percentage points within the first year and the contraction narrows by month 36; purchase financing is near zero early and about −0.02 at month 36; the mortgage rate rises by 15 basis points at month 12 and stays elevated; and the one-year yield climbs from 25 to about 42 basis points at month 12. A yield hump is thus also present in this earlier proxy calculation. The macroeconomic impact medians differ from the zeros imposed by the restriction-based identification: consumption, industrial production, and inflation are positive in the full-moment calculation. The industrial-production impact changes from +0.113 to −0.051 when March 2020 alone is omitted, whereas the consumption impact remains positive in the full sample (+0.163) and when January 2008 (+0.256) or March 2020 (+0.131) is omitted. Table C5b documents these broad similarities in subsequent financing adjustment and the sensitivity of particular impact responses.

C.5.3 What the surprises are used for

Given Table C5a, the paper does not identify responses with the surprises. It uses them, after estimation, to ask whether announcement information corresponds to the innovation the restrictions label as policy (Section 3.5 and Appendix C.3): the policy innovation is the most closely associated of the eight in every state, for both MP1 and FF4, and remains so in all but one state (MP1) or 44 states (FF4) of 12,000 when January 2008 is omitted. The current-model evidence supports the relative policy label among the eight innovations, while the earlier proxy exercise supplies an exploratory comparison of financing and mortgage-rate dynamics.

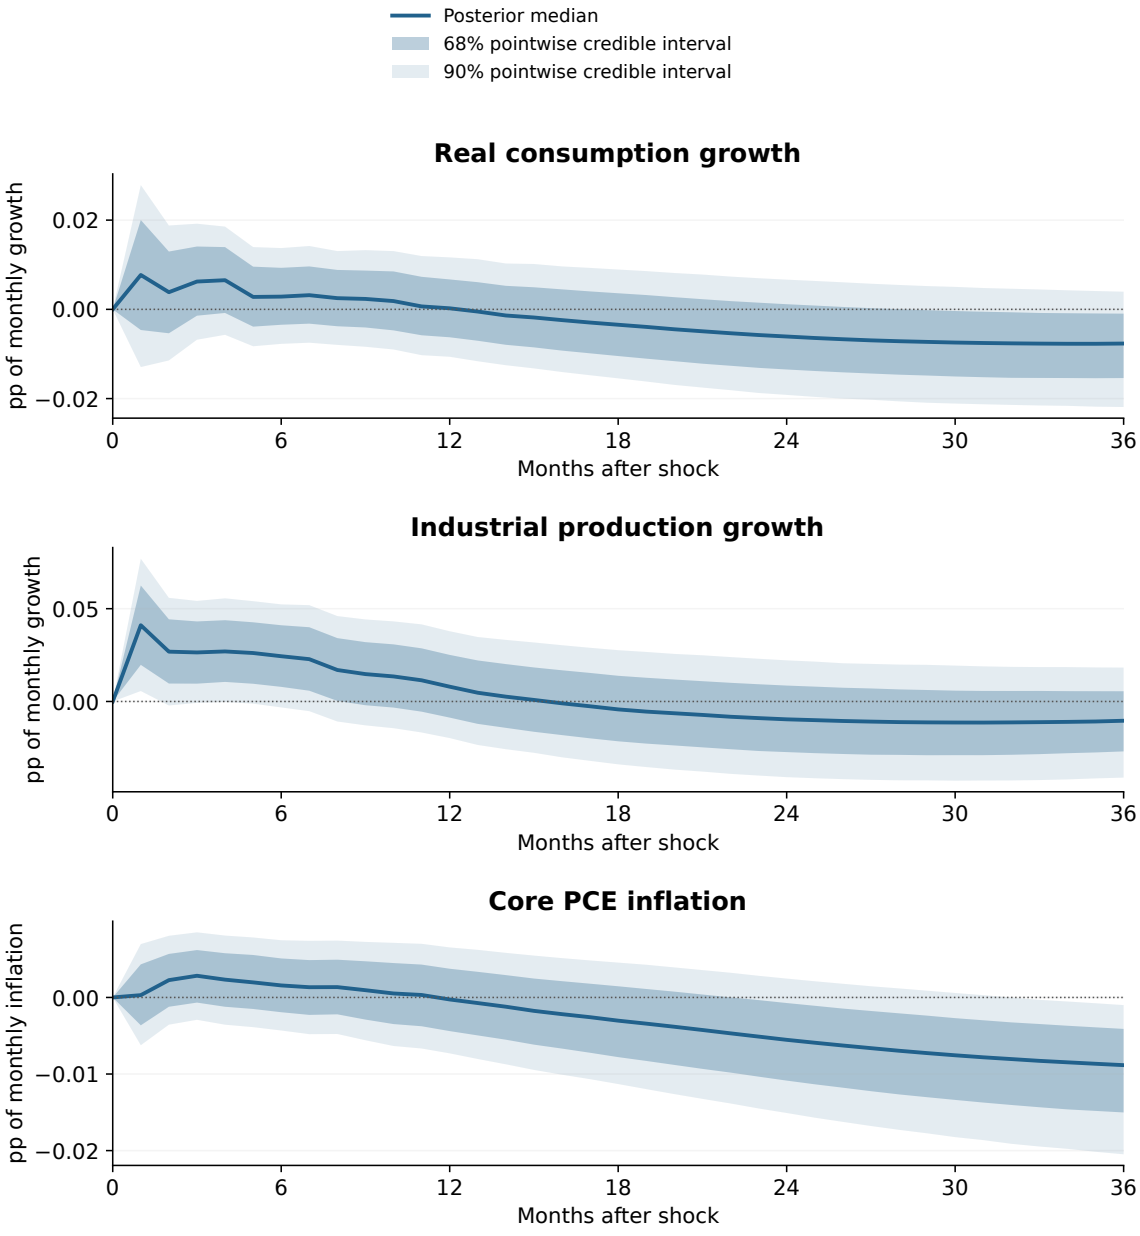


Figure C1. Aggregate macroeconomic responses. Responses of monthly real consumption growth, industrial-production growth, and core PCE inflation to a policy innovation that raises the one-year Treasury yield by 25 basis points on impact, in percentage points of monthly growth or inflation. Lines are medians and dark and light shading are central 68% and 90% pointwise credible intervals across all 12,000 joint posterior states of the FHFA specification, including nonstable states.

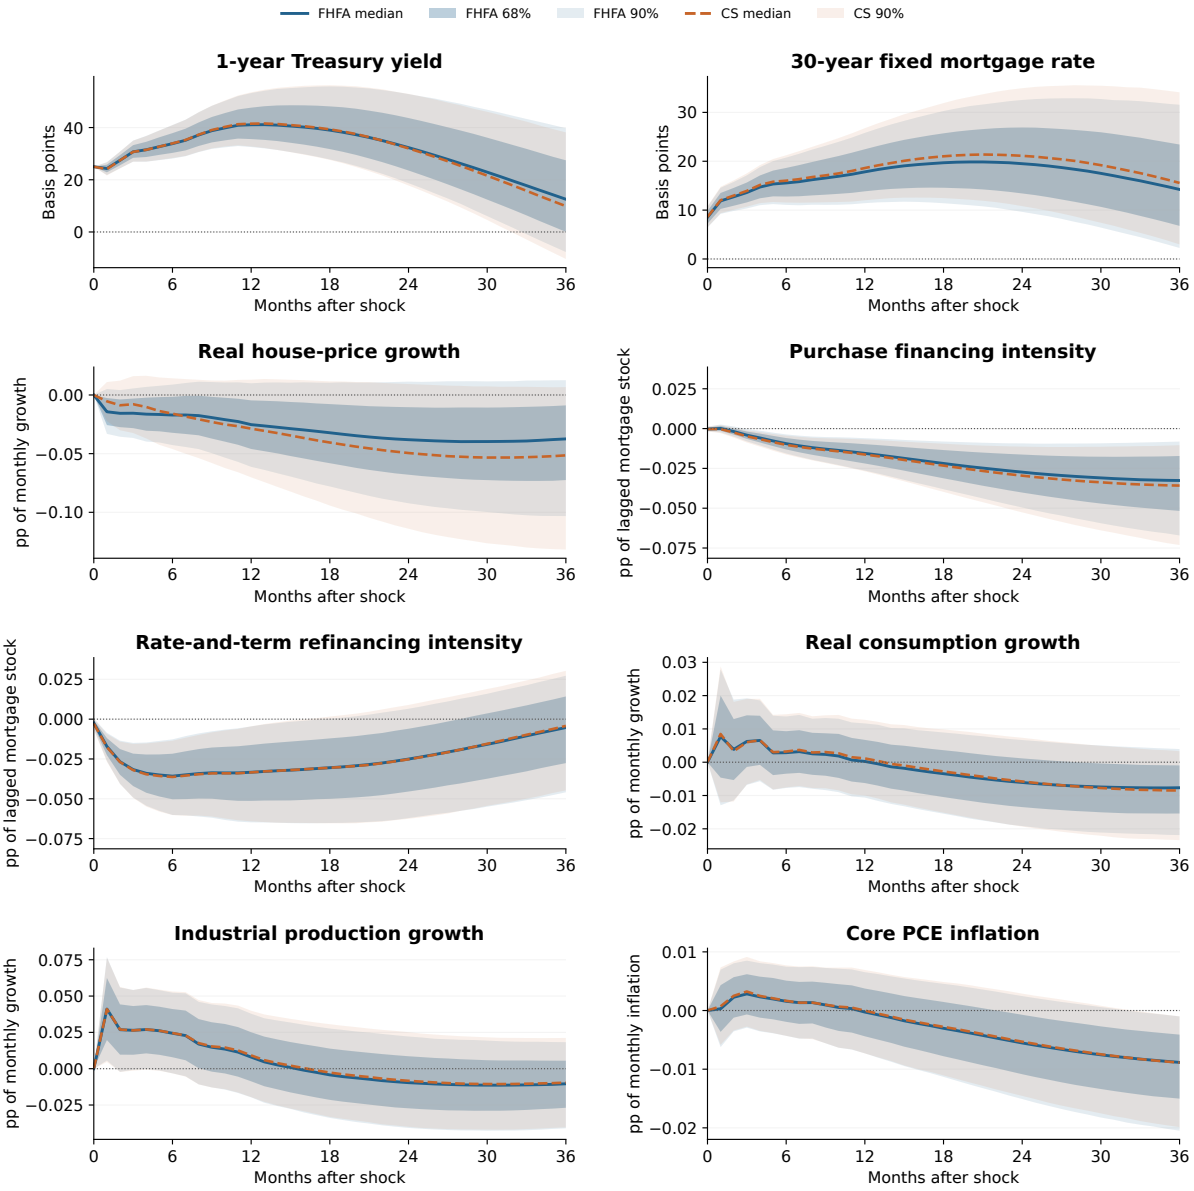


Figure C2. FHFA and Case-Shiller specifications. Responses of the eight model variables to the same 25-basis-point impact increase in the one-year Treasury yield, with interest rates in basis points, monthly growth and inflation in percentage points, and financing intensities in percentage points of lagged mortgage debt. Each specification uses all 12,000 joint posterior states: blue lines and dark/light shading show FHFA medians and central 68%/90% pointwise credible intervals, while orange dashed lines and shading show Case-Shiller medians and 90% intervals from its separately estimated posterior.

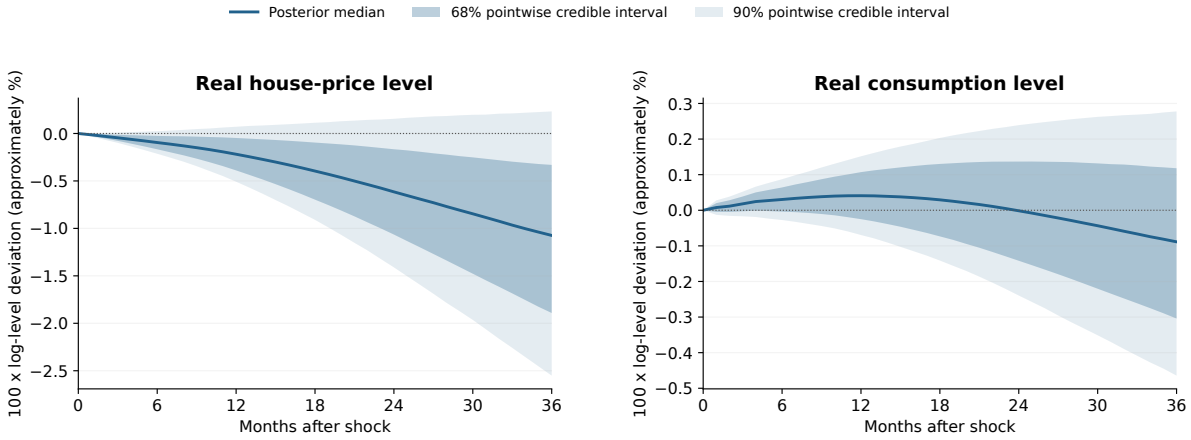


Figure C3. Cumulative consumption and housing-valuation responses. Monthly growth responses are summed from impact through each horizon within each posterior state, giving 100 times the consumption and real house-price log-level gaps following a 25-basis-point policy innovation. Lines show medians and dark/light shading shows central 68%/90% pointwise credible intervals across all 12,000 FHFA states; these are level gaps at each horizon, rather than spending losses summed over months.

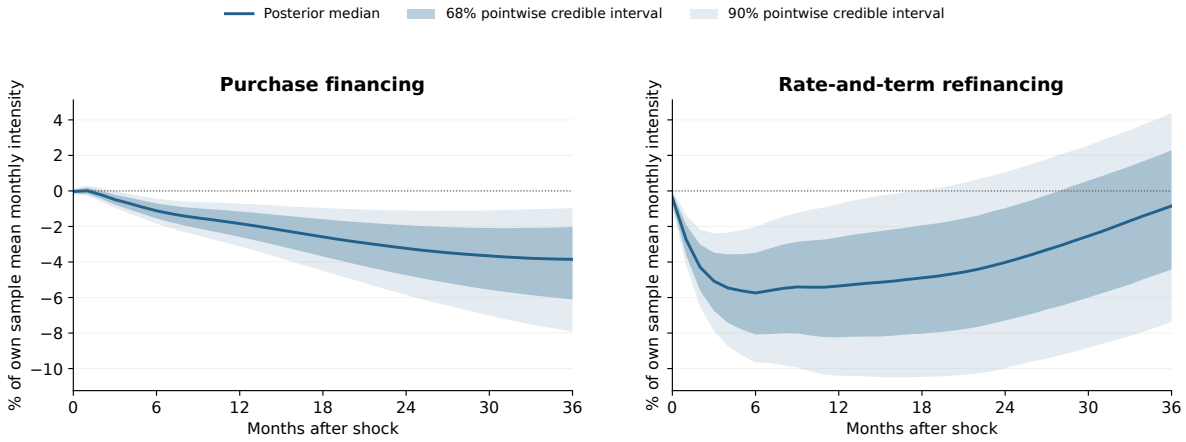


Figure C4. Financing responses relative to average monthly intensity. Purchase and rate-and-term refinancing responses to a 25-basis-point policy innovation are expressed as percentages of their respective mean monthly financing intensities over January 1999–December 2023, using fixed benchmarks of 0.846641 and 0.622930 percent of lagged mortgage debt. Lines show medians and dark/light shading shows central 68%/90% pointwise credible intervals across all 12,000 FHFA states; the rescaling describes aggregate flow-to-stock intensities, rather than household transaction probabilities.

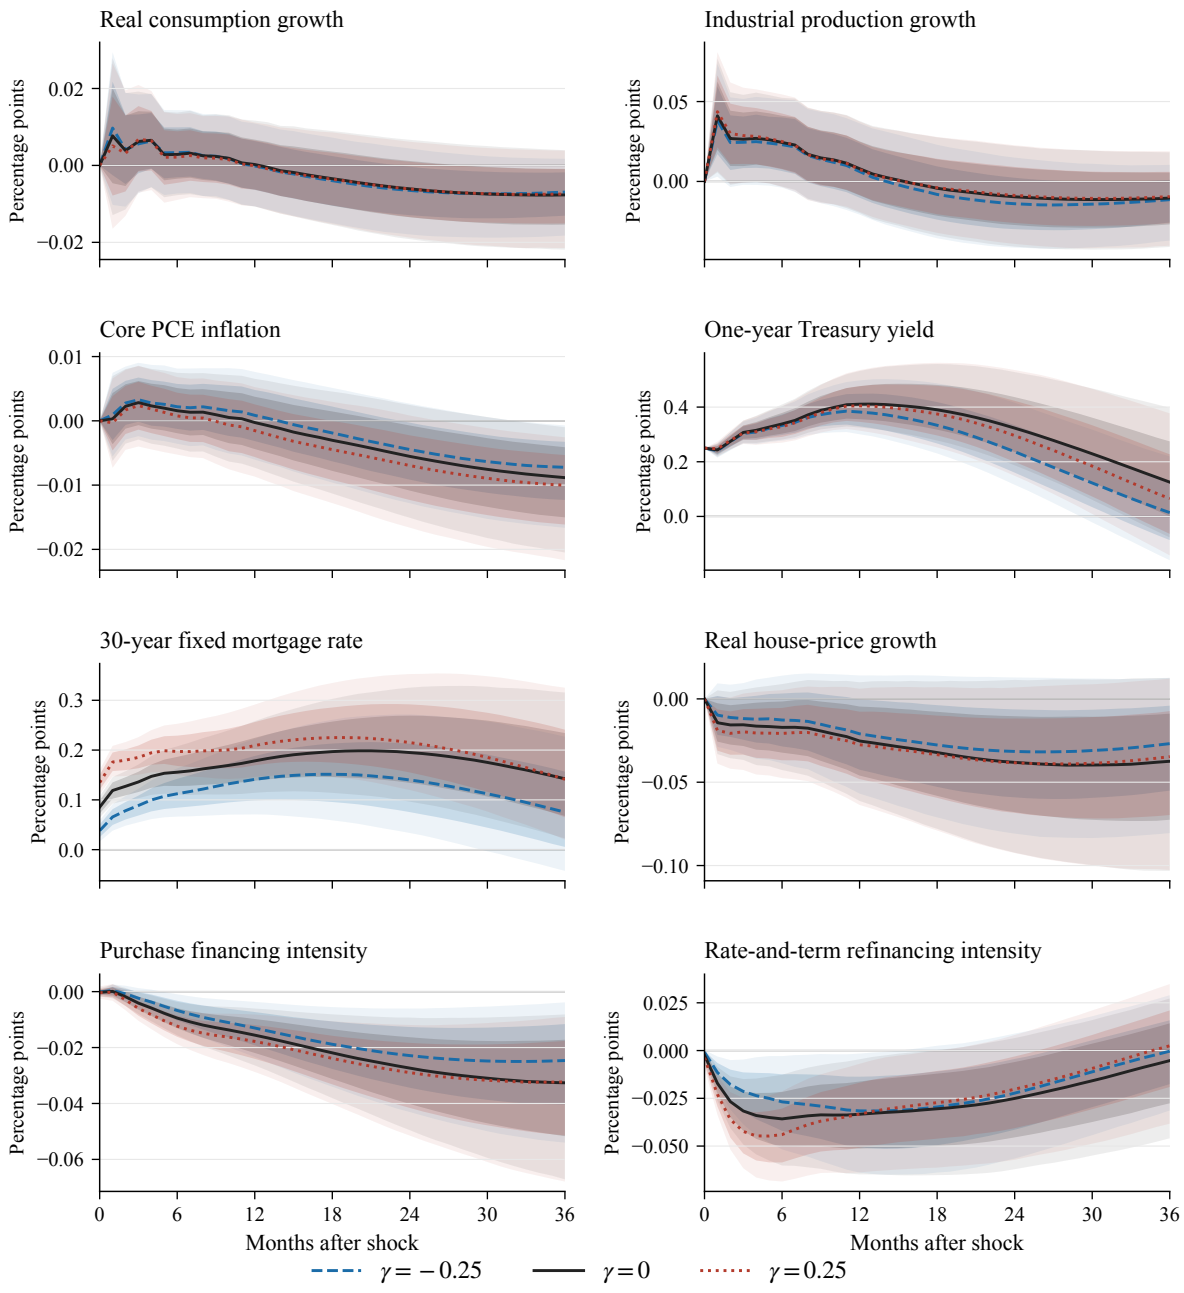


Figure C5. Contemporaneous mortgage-rate feedback. Responses under separately estimated fixed-feedback specifications with $\gamma = -0.25, 0, 0.25$, each normalized to a 25-basis-point impact increase in the one-year Treasury yield; all panels use percentage points of monthly growth or inflation, interest-rate levels, or financing intensity relative to lagged mortgage debt. Lines show posterior medians and dark/light shading shows central 68%/90% pointwise credible intervals across all 12,000 states per specification, with each interval summarizing its own posterior.

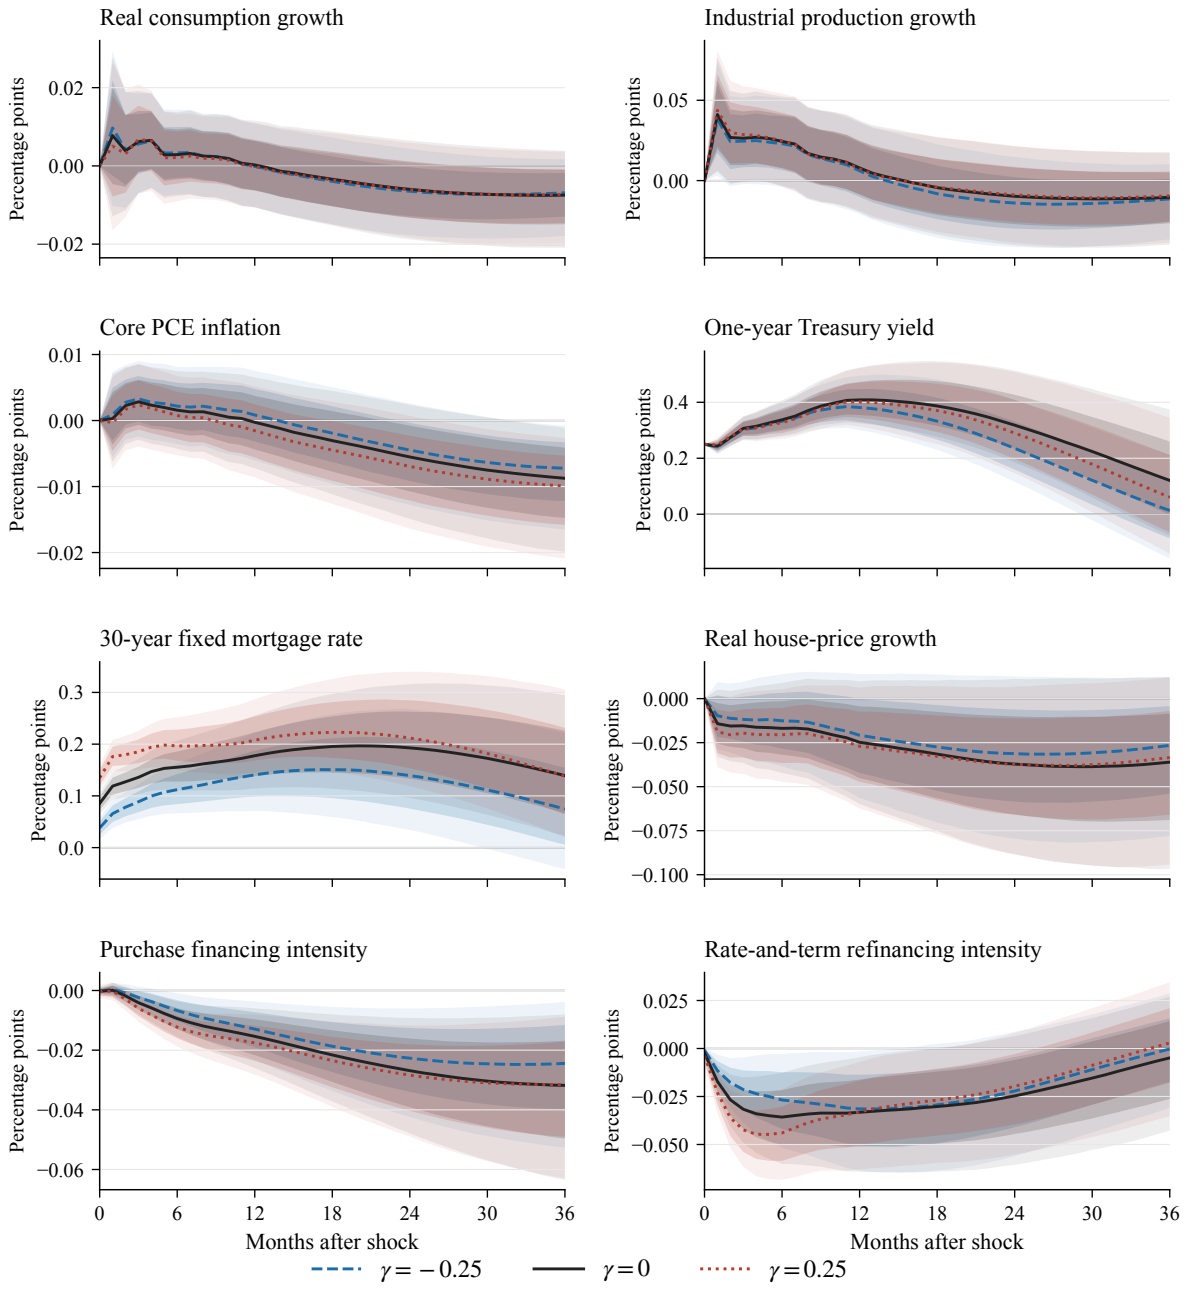


Figure C6. Mortgage-rate feedback: stable-state sensitivity. The responses in Figure C5 are restricted to stable states: 11,662 for $\gamma = -0.25$, 11,274 for $\gamma = 0$, and 11,216 for $\gamma = 0.25$, with the same 25-basis-point normalization and percentage-point units. Lines show medians and dark/light shading shows central 68%/90% pointwise credible intervals within each specification's stable-state posterior.

Appendix D. Transmission accounting and complete consumption results

D.1 Ordered representation of the fitted system

The accounting uses the policy response of each saved joint state, without re-estimating the model. Suppress the state index and let Π be the permutation matrix for one of the four representations below. If $v(h)$ is the original eight-variable response, write $x_h = \Pi v(h)$. Factor the unit-diagonal structural matrix as

$$A_0\Pi' = QL, \qquad Q'Q = I_8, \qquad L \text{ lower triangular}, \quad L_{ii} > 0,$$
$$W = \text{diag}(L_{11}^{-1}, \ldots, L_{88}^{-1}), \qquad \mathcal{M} = WL. \tag{D1}$$

The unit lower-triangular matrix $\mathcal{M}$ is an accounting operator; it is distinct from the observed mortgage rate M_t. Define

$$E_0 = I_8 - \mathcal{M}, \qquad E_\ell = \mathcal{M}\Pi\Phi_\ell\Pi', \qquad q = \mathcal{M}\Pi b_{25}. \tag{D2}$$

The original policy response then satisfies

$$x_h = E_0 x_h + \sum_{\ell=1}^{\min(12,h)} E_\ell x_{h-\ell} + q\,\mathbf{1}\{h = 0\}. \tag{D3}$$

Because E_0 is strictly lower triangular, the response can be computed by month and then by within-month variable order. This follows the QL representation in Wegner et al. (2025, Section 3.2). The original nonrecursive contemporaneous system is preserved in its total responses. The representation supplies a way to allocate those responses to paths.

Three identities clarify the mapping and provide independent checks:

$$\mathcal{M}^{-1}q = \Pi b_{25}, \qquad E_\ell = WQ'A_0\Phi_\ell\Pi', \qquad q = WQ'e_4\frac{0.25}{B_{44}}. \tag{D4}$$

In particular, the structural lag matrices are $A_0\Phi_\ell$, not the reduced-form matrices Φ_ℓ. Keeping this distinction explicit avoids omitting or applying the contemporaneous transformation twice.

For a finite horizon h, stacking $x_0, \ldots, x_h$ gives $X_h = \mathcal{E}_h X_h + \Pi_h$. The matrix $\mathcal{E}_h$ is strictly lower triangular, with dimension $n_h = 8(h+1)$, and hence

$$X_h = (I - \mathcal{E}_h)^{-1}\Pi_h = \sum_{m=0}^{n_h - 1} \mathcal{E}_h^m \Pi_h. \tag{D5}$$

This finite-graph identity explains why the accounting can be evaluated for all saved states through month 36. It does not require $\rho(I - A_0) < 1$ or an infinite expansion of the original contemporaneous feedback matrix. Nonstable mean dynamics can still have finite-horizon responses; their infinite-horizon limits are a separate issue.

D.2 First-entry labels and propagation

The classification tracks four labels, $\mathcal{A} = \{N, H, P, R\}$, where N means that no core entry has yet been visited. Let $\kappa(i)$ denote the identity of reordered node i: H, P, or R for a core variable and N otherwise. Upon arrival at a node, update the label according to

$$\lambda(a, i) = \begin{cases} \kappa(i), & a = N \text{ and } \kappa(i) \neq N, \\ a, & \text{otherwise.} \end{cases} \tag{D6}$$

Let $U^a_{h,i}$ be the contribution arriving at node (h, i) with label a. Chronological propagation gives

$$\begin{aligned} U^k_{h,i} = {} & \mathbf{1}\{h = 0\}\mathbf{1}\{k = \lambda(N, i)\}q_i \\ & + \sum_{j<i} \sum_{a\in\mathcal{A}} \mathbf{1}\{k = \lambda(a, i)\}(E_0)_{ij} U^a_{h,j} \\ & + \sum_{\ell=1}^{\min(12,h)} \sum_{j=1}^{8} \sum_{a\in\mathcal{A}} \mathbf{1}\{k = \lambda(a, i)\}(E_\ell)_{ij} U^a_{h-\ell,j}. \end{aligned} \tag{D7}$$

At the consumption node i_C, the contributions are $c_H(h) = U^H_{h,i_C}$, $c_P(h) = U^P_{h,i_C}$, $c_R(h) = U^R_{h,i_C}$, and $c_O(h) = U^N_{h,i_C}$. A source term entering directly at a core node receives that node's label immediately. A subsequent visit to a different core variable preserves the existing label. The recursion includes later feedback, repeated consumption visits, and paths that reach financing without first visiting the mortgage-rate node.

These definitions make the categories exhaustive and mutually exclusive within the chosen representation. A valuation–refinancing–consumption path belongs to H; a refinancing–valuation–consumption path belongs to R. No separate overlap category is needed because the first entry resolves ownership. Other collects consumption paths with no H/P/R visit, rather than an independently identified income or credit mechanism.

D.3 Representation and stable-state sensitivity

The four orderings vary relationships among the housing-finance nodes while retaining consumption as the last contemporaneous destination.

Table D1. Prespecified accounting representations. The yield r is the policy indicator and M is the mortgage rate. These are accounting representations of the fitted structural system, not four newly estimated models.

Representation	Within-month order
O1	IP, inflation, yield, mortgage rate, valuation, purchase, refinancing, consumption
O2	IP, inflation, yield, mortgage rate, valuation, refinancing, purchase, consumption
O3	IP, inflation, yield, mortgage rate, purchase, refinancing, valuation, consumption
O4	IP, inflation, yield, valuation, mortgage rate, purchase, refinancing, consumption

The three-year median ordering is preserved under these targeted alternatives, but the allocations are not invariant. Moving valuation after both financing entries reallocates some contribution away from H. Exchanging purchase and refinancing changes the R median from -0.02251 to -0.01735, a decline of about 23%

Table D2. Three-year cumulative median contributions across representations and states. Accumulation is over months 1–36, in 100 × log consumption-level units. All-state rows contain 12,000 states; the stable row conditions on the 11,274 states with companion roots inside the unit circle. Medians need not add.

Representation and states	Valuation H	Purchase P	Refinancing R
O1, all	-0.06364	-0.04708	-0.02251
O2, all	-0.06364	-0.05134	-0.01735
O3, all	-0.05830	-0.05140	-0.02267
O4, all	-0.06347	-0.04736	-0.02246
O1, stable	-0.06199	-0.04578	-0.02286

in its absolute value, while P moves from -0.04708 to -0.05134. These reallocations matter because propagation can involve more than one financing entry. The comparison supports describing the common median ordering across the four representations while treating the particular allocation as part of the estimand. It does not establish invariance across all permutations or structural-equation normalizations.

D.4 Complete totals, paired comparisons, and window definitions

Figure 2 reports posterior means of monthly contributions and their cumulative sums over months 0–h. Means preserve the accounting identity. Table D6 reports medians and intervals of state-level sums over months 1–12 and 1–36. The impact total is zero by restriction, but individual accounting contributions at impact need not be zero; the different window definitions are therefore kept explicit for the components.

Table D3. Total consumption response and complete mean accounting. Cumulative units are 100 × log consumption-level gaps. Panel A uses medians [90% intervals] of state-level sums. Panel B uses posterior means; its categories add apart from rounding.

Panel A: total, O1, all states	Median	90% interval
Months 1–12	0.04088	[-0.06958, 0.15168]
Months 1–36	-0.08867	[-0.46460, 0.27770]

Table D3 (continued, Panel B)

Panel B: means over months 0–36	H	P	R	Other	Total
Cumulative contribution	-0.073583	-0.057398	-0.029121	0.068195	-0.091906

The positive cumulative Other contribution offsets approximately 43% of the combined negative mean contributions of H, P, and R over months 0–36. It is not an offset defined against H alone. It also does not imply a positive Other contribution in every month: by month 36 its monthly mean is approximately -0.000963. Other collects paths that never enter H, P, or R, rather than a numerical fitting residual. The four-category results do not separately quantify its IP-related and non-IP-related paths; the early positive IP

response motivates that diagnostic but does not establish its outcome. Similarly, a late monthly consumption response whose 68% interval is below zero need not imply a three-year cumulative interval excluding zero. Monthly responses, accumulated gaps, and the composition of those gaps answer different questions.

Pairwise differences are calculated within the same state before taking quantiles. This preserves dependence across channels. Subtracting the entries in a median table would not produce the median of a paired difference.

Table D4. Paired differences in cumulative contributions over months 1–36. O1, all 12,000 states. Units match Table D2. Probabilities are posterior proportions of negative signed differences, not probabilities of uniformly larger absolute effects at every horizon.

Difference	Median	90% interval	Probability below zero	Batch Monte Carlo SE
H minus P	-0.01456	[-0.18484, 0.15184]	0.562	0.0042
H minus R	-0.03929	[-0.23919, 0.13440]	0.656	0.0049
P minus R	-0.02340	[-0.21919, 0.14958]	0.603	0.0045

These distributions retain uncertainty around the central cumulative comparisons. Minimum-response dates within a finite window provide a different statistic from onset. Substantial mass at the month-36 boundary for H and P also prevents identifying a global trough or declaring one entry the fastest from such dates.

D.5 Numerical verification and interpretation

Independent verification checks total responses, category addition, operators, and the tracking algorithm. For the tracker check, the reviewer separately propagates paths that have not yet entered a core node, computes the first-entry inflows, and then propagates those inflows through the full system. This supplies an algorithmic cross-check distinct from running the production tracker twice.

Table D5. Numerical checks of finite-horizon accounting. Maximum scaled error divides an absolute difference by $\max(1, |a|, |b|)$ and takes the largest value. The tolerance is 10^{-10}.

Check	Coverage	Maximum scaled error
Total versus independent reduced-form recursion	All states, eight variables, months 0–36	3.164136×10^{-15}
Categories sum to consumption	All states, four representations, months 0–36	5.655199×10^{-16}
Independent first-hit algorithm	Eight prespecified states, four representations, months 0–36	2.142383×10^{-16}
Independent reverse-Cholesky construction	Saved accounting-operator comparison	5.551115×10^{-16}
Window and paired summaries	128 specified rows, all/stable states, four representations	3.885781×10^{-16}

The QL factorization uses the original unit-diagonal A_0. It does not first whiten equations with a selected date's $D_t^{-1/2}$. The transformed innovations $WQ'\varepsilon_t$ are consequently not asserted to be a new set of

independent standardized shocks. The accounting follows the original fourth-equation innovation through its transformed source vector q.

The distinction matters for scope. Total-response equivalence follows from the finite algebra above, while category allocations depend on the specified representation and equation normalization. Arbitrary unequal row rescaling can change the QL representation even when the total response is unchanged. The stationary, standardized setting in the general transmission-channel framework does not by itself establish every stronger intervention interpretation for this stochastic-volatility model and its nonstable states. The results use the stated structural policy experiment and conditional first-entry accounting, with posterior uncertainty propagated through the saved joint states.

D.6 Window-specific channel contributions

Table D6. Cumulative contributions over one and three years. Entries are posterior medians [central 90% intervals] of within-state sums over the named months. Units are contributions to the model-implied $100 \times$ log consumption-level gap. The primary order and all 12,000 states are used. Impact is excluded; the statistic and window differ from Figure 2.

First entry	Months 1–12	Months 1–36
Housing valuation, H	-0.009	-0.064
	$[-0.037, 0.006]$	$[-0.223, 0.044]$
Purchase financing, P	-0.002	-0.047
	$[-0.021, 0.013]$	$[-0.187, 0.037]$
Rate-and-term refinancing, R	-0.006	-0.023
	$[-0.036, 0.019]$	$[-0.146, 0.070]$

All three channels' 90% intervals include zero in both windows. These distributions describe uncertainty around the window-specific contributions; Table D4 separately reports paired channel differences, computed within the same posterior state. A change in the ordering of marginal medians across windows does not by itself estimate the probability of a joint ordering reversal.

The cumulative quantity is the level gap implied by the model-input consumption-growth response, not the sum of foregone dollar expenditures. Its construction retains every month in the stated window, including months whose pointwise intervals include zero. Rescaling percentage units into basis points changes neither relative magnitudes nor whether an interval includes zero.

Appendix E. Volatility histories and model diagnostics

E.1 Conditional innovation scales and calendar-time summaries

For state s, the marginal innovation SD for observed equation j and the structural-innovation SD are, respectively,

$$\sigma_{j,t}^{u,(s)} = \sqrt{[B^{(s)} D_t^{(s)} B^{(s)\prime}]_{jj}}, \qquad \sigma_{j,t}^{\varepsilon,(s)} = \exp\{h_{j,t}^{(s)}/2\}. \tag{E1}$$

The first incorporates covariance transmitted through the contemporaneous system; the second is the equation-normalized structural scale. Both are conditional on the fitted mean and estimated with the full sample. They are smoothed historical objects, not real-time forecasts. The covariance architecture has a precedent in Clark (2010, Section 3.2); the stationary AR(1) log variances and nonrecursive matrix are distinct specification choices.

For each date, the exhibit reports the median and 90% interval across saved states. A reported peak is the date maximizing the pointwise median path. It is not the posterior distribution of the argmax. Across all eight native-SD paths, restricting attention to stable states changes the pointwise median by at most 0.7001% relative to the all-state median; the maximum occurs for the one-year yield in January 2021. The corresponding maximum for structural SDs is 0.6665%.

Table E1. Peaks in median marginal innovation-SD paths. Dates refer to January 1999–December 2023. Intervals describe the SD at the selected median-path peak, conditional on that date. Rate SDs are in percentage points; growth and financing-intensity SDs are in their respective percentage-point units. Cross-equation levels are not directly comparable economic volatility measures.

Equation	Full-sample peak	Median SD [90% interval]	Peak within 2001–2003
Consumption	2020-07	0.521 [0.384, 0.743]	2003-02
Industrial production	2008-12	0.963 [0.702, 1.403]	2003-01
Core inflation	2001-09	0.281 [0.206, 0.412]	2001-09
One-year yield	2008-01	0.505 [0.347, 0.802]	2001-09
Mortgage rate	2022-09	0.347 [0.258, 0.505]	2003-08
Housing valuation	2009-01	0.560 [0.401, 0.832]	2001-09
Purchase financing	2020-10	0.097 [0.067, 0.155]	2001-03
Rate-and-term refinancing	2003-07	0.465 [0.303, 0.765]	2003-07

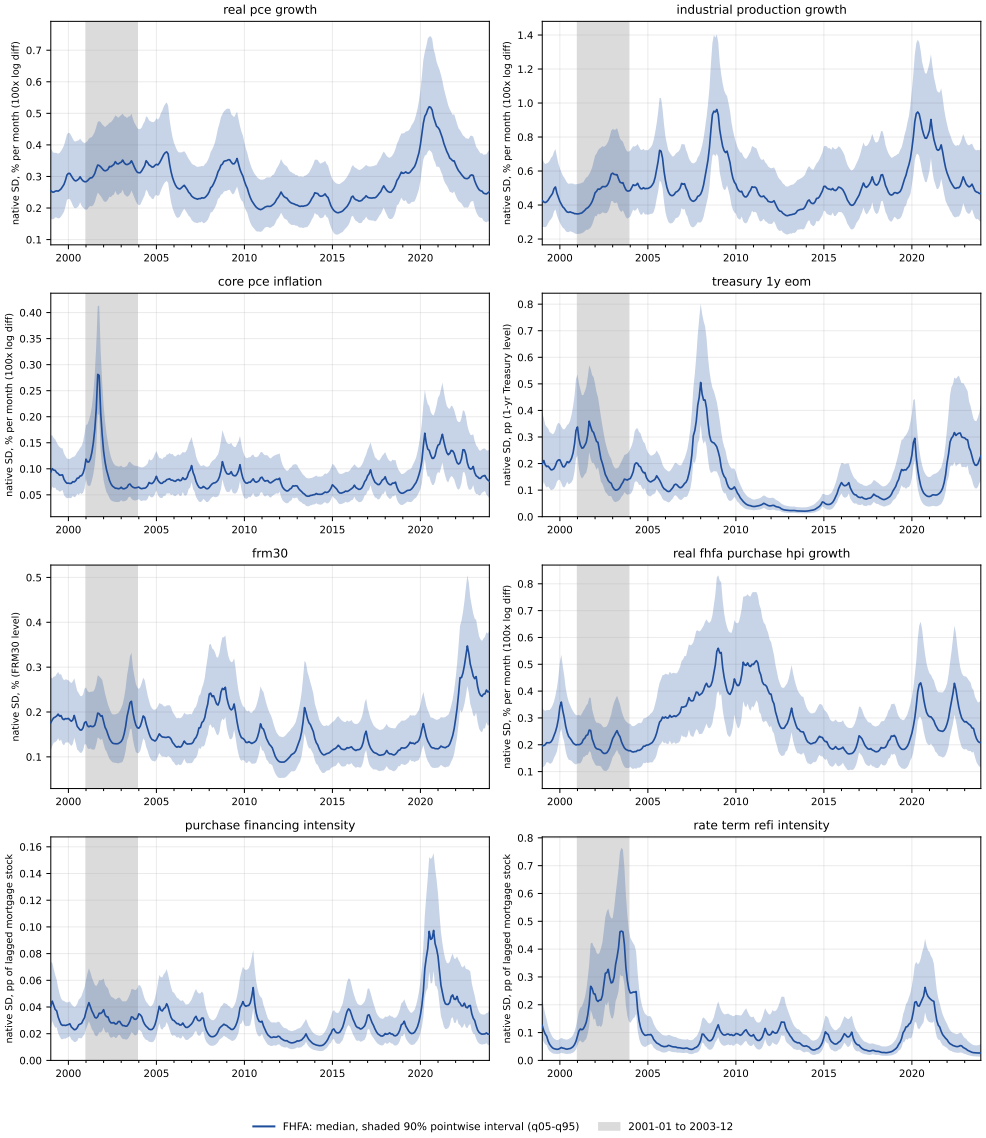


Figure E1. Marginal innovation standard deviations. Each panel shows the posterior median and central 90% pointwise credible interval of $\sqrt{[BD_tB']_{ii}}$ for January 1999–December 2023, using the FHFA specification and all 12,000 joint posterior states. Units follow the corresponding model inputs; consumption and production growth incorporate the caps described in Appendix A. Gray vertical shading highlights January 2001–December 2003. The paths are smoothed full-sample estimates.

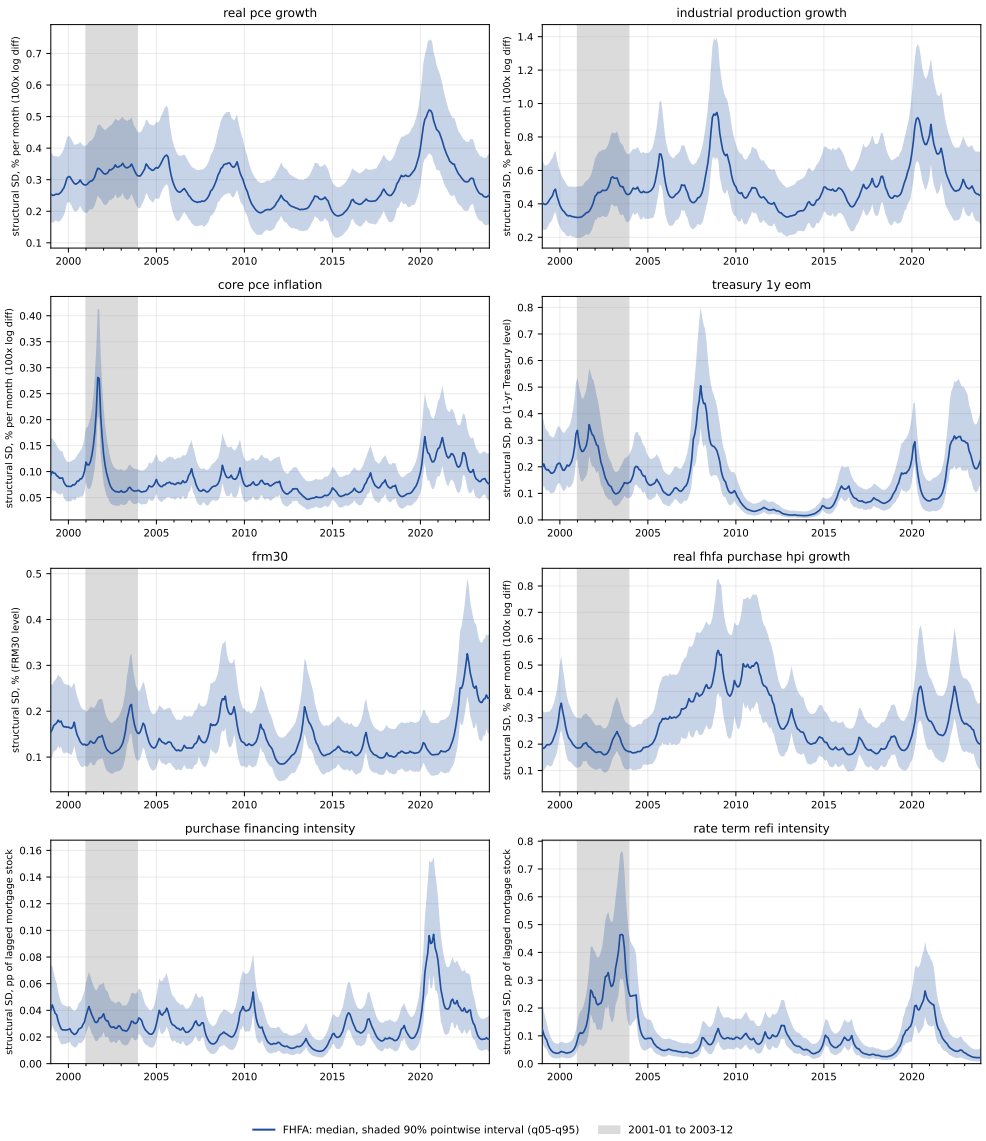


Figure E2. Structural innovation standard deviations. Each panel shows the posterior median and central 90% pointwise credible interval of $\exp(h_{it}/2)$ under the unit-diagonal normalization of A_0, for the same sample and all 12,000 joint posterior states. Gray vertical shading highlights January 2001–December 2003. Equation units differ, so the numerical scales should not be compared directly across panels.

The full eight-variable marginal-SD exhibit, structural-SD exhibit, and house-price-measure comparison retain the complete paths and intervals. The principal economic comparison concerns different histories for H, P, and R, rather than a comparison of their numerical units.

NBER chronology supplies a common calendar reference. Contraction months are April–November 2001, January 2008–June 2009, and March–April 2020, including the trough and excluding the peak month. This gives 28 contraction and 272 expansion months. The summary below divides the contraction-month average of the median SD path by its expansion-month average. It is a descriptive ratio of two path summaries, not the median of state-level ratios or a hypothesis test.

Table E2. Conditional innovation scale in contraction and expansion months.

Equation	Ratio: contraction average / expansion average
Consumption	1.166
Industrial production	1.297
Core inflation	1.520
One-year yield	2.274
Mortgage rate	1.392
Housing valuation	1.426
Purchase financing	0.966
Rate-and-term refinancing	1.208

Refinancing has a higher average scale in contractions but its median-path maximum occurs in 2003, outside the contraction months. Purchase financing peaks after the April 2020 trough and has a ratio slightly below one. The finding concerns an activity-specific chronology, not a uniform rise in all housing-finance volatility during recessions.

E.2 Signed and squared residual dependence

Residuals are reconstructed within each joint state. Three definitions separate original equation units from marginal scale adjustment and structural orthogonalization:

$$u_t = y_t - c - \sum_{\ell=1}^{12} \Phi_\ell y_{t-\ell}, \qquad v_{j,t} = u_{j,t}/\sqrt{\Sigma_{jj,t}}, \qquad z_t = D_t^{-1/2} A_0 u_t. \tag{E2}$$

For a residual sequence a_t of length $T = 300$, the lag-k ACF is

$$\widehat{\rho}_k(a) = \frac{\sum_{t=k+1}^{T}(a_t - \bar{a})(a_{t-k} - \bar{a})}{\sum_{t=1}^{T}(a_t - \bar{a})^2}, \qquad k = 1, \ldots, 24. \tag{E3}$$

The mean and denominator use the full sequence; there is no $T - k$ adjustment. Squared-residual dependence applies this formula to a_t^2, rather than squaring the signed ACF. Statistics are computed for every state and then summarized. The approximate iid reference $1.96/\sqrt{300} \simeq 0.113$ is a descriptive magnitude benchmark, not a calibrated simultaneous threshold for fitted residuals.

Table E3. Dependence across eight equations and 24 lags. Each row covers 192 equation–lag cells. The interval count records pointwise 90% intervals excluding zero and does not apply a multiple-testing correction.

Residual definition	Signed: absolute median ACF above 0.113	Signed: interval excludes zero	Squared: absolute median ACF above 0.113	Squared: interval excludes zero
Native u	17	81	51	129
Marginally standardized v	10	51	2	22
Structurally standardized z	7	44	1	7

Purchase financing illustrates the effect on squared dependence. At lag 3, the native squared-ACF median is 0.601 [0.529, 0.650], compared with 0.113 [0.025, 0.233] after marginal standardization and 0.063 [-0.016, 0.167] after structural standardization. This is within-fit evidence that changing scales account for much of the raw clustering.

Signed refinancing dependence remains at an annual horizon. Marginal-standardized medians at lags 12 and 13 are 0.149 [0.091, 0.204] and -0.178 [-0.228, -0.124]. Structural-standardized counterparts are 0.153 [0.093, 0.211] and -0.184 [-0.240, -0.124]. These patterns identify a feature that survives contemporaneous orthogonalization and volatility adjustment.

The complete exhibits show signed marginal ACFs, native squared ACFs, marginally standardized squared ACFs, and structurally standardized squared ACFs.

E.3 ARCH-type diagnostics

For each state and equation, regress the squared standardized residual on an intercept and its first q lagged squares, with $q \in \{6, 12, 24\}$. The statistic is $LM = (300 - q)R^2$. A nominal chi-square reference with q degrees of freedom gives a descriptive p-value, following the auxiliary-regression approach discussed by Bollerslev, Engle, and Nelson (1994). All 12,000 state-level calculations are valid in each reported cell; at $q = 12$, each regression has 288 observations.

The ARCH summary exhibit complements the ACF plots. The results show reduced variance dependence after standardization with meaningful differences across equations. They do not establish an out-of-sample forecast improvement over a separately fitted constant-variance model.

E.4 What alternative fits preserve and change

The comparisons retain the same response sample when that is part of their design. Their scope is specified because a completed volatility comparison does not automatically provide policy IRFs under the alternative.

The two-year fit has a 3.81% nonstable share. Its time-median ratios of marginal-SD medians relative to the one-year fit are 1.000 for consumption, 1.001 for H, 1.002 for P, 0.991 for R, and 0.955 for mortgage rates. The chronology comparison is therefore accompanied by a particularly noticeable change in mortgage-rate scale.

Table E4. Remaining ARCH-type dependence at twelve lags. Brackets give fifth and ninety-fifth percentiles of LM across states. The last column for each residual definition is the fraction of state-level nominal p-values below 0.05; it is not the posterior probability that a model is misspecified.

Equation	Marginal LM [90% range]	Share below 0.05	Structural LM [90% range]	Share below 0.05
Consumption	11.912 [5.949, 21.139]	5.17%	11.912 [5.949, 21.139]	5.17%
Industrial production	18.983 [10.398, 33.180]	38.22%	13.607 [6.971, 24.086]	10.93%
Core inflation	14.572 [7.730, 24.429]	12.62%	14.480 [7.731, 24.481]	12.64%
One-year yield	13.649 [7.135, 23.161]	8.91%	12.183 [6.151, 21.304]	5.46%
Mortgage rate	10.148 [5.093, 18.050]	1.86%	10.249 [5.029, 18.629]	2.26%
Housing valuation	14.384 [7.433, 23.767]	11.62%	15.607 [8.297, 25.369]	16.85%
Purchase financing	11.934 [5.787, 24.008]	8.87%	10.936 [5.487, 19.267]	2.58%
Rate-and-term refinancing	13.573 [6.774, 23.630]	10.38%	13.071 [6.334, 23.013]	8.54%

Table E5. Median-path peak dates under measurement and data-treatment alternatives. Each cell is the full-sample native-SD peak date. The uncapped comparison also changes empirically calibrated prior scales.

Specification	Consumption	Valuation H	Purchase P	Refinancing R
FHFA, one-year yield, capped inputs	2020-07	2009-01	2020-10	2003-07
Case-Shiller instead of FHFA	2020-08	2010-03	2020-07	2003-07
Two-year instead of one-year yield	2020-07	2009-01	2020-10	2003-07
Uncapped consumption and IP	2020-04	2009-01	2020-10	2003-07

The six-lag fit has 392 mean coefficients rather than 776 and uses the same 300 response months. Its nonstable share is 24.73%, compared with 6.05% at twelve lags. Refinancing's structural lag-12 ACF median is 0.155, compared with 0.153; the largest equation-specific median twelve-lag ARCH statistic is 16.5929, compared with 15.6075. These summaries do not provide a uniform diagnostic case for shortening the lag history. Policy IRFs have not been used to select between these two fits.

The uncapped fit has a 4.38% nonstable share. Between March and July 2020, the ratios of uncapped to capped native-SD medians range from 6.355 to 11.644 for consumption and from 2.977 to 4.074 for IP. Both peak in April 2020. H/P/R retain their peak dates; their peak medians are 0.546, 0.0903, and 0.4651, compared with 0.560, 0.0974, and 0.4647 in the baseline. Thus, the pandemic macroeconomic scale is strongly treatment-dependent, while the three housing-finance peak dates survive this comparison.

The uncapped exercise follows the same calibration rule rather than holding all numerical prior scales fixed. Univariate residual SDs increase by factors of 3.129 for consumption and 1.881 for IP. This changes 312 of 776 conditional-mean prior variances, 13 of 28 contemporaneous prior SDs, and two SV prior centers. Its interpretation is a combined data-and-calibration comparison. The large pandemic macroeconomic change is not presented as a pure likelihood-only effect.

A stronger-shrinkage fit multiplies the prior variances of all 768 lag coefficients by 0.5, preserving the intercept variance of 100, the frozen data and calibration scales, and the contemporaneous and SV priors. The multiplier applies to variances, so the corresponding prior SD multiplier is $\sqrt{0.5}$. All sixteen native and structural median-path peak dates remain the same in both the full sample and the 2001–2003 window. Time-median native-SD changes for the non-refinancing equations range from approximately 0.49% to 1.49%; refinancing increases by 3.24%. Its late path nevertheless decreases in places: in February 2023 the native median falls from 0.04068 to 0.02910 and the structural median from 0.03645 to 0.02436. The retained chronology is compatible with sensitivity in path levels.

The residual diagnostics do not show a uniform improvement. For structurally standardized residuals, the median ARCH(12) statistic rises from 11.91 to 13.33 for consumption, from 15.61 to 16.56 for housing valuation, and from 13.07 to 13.93 for refinancing. The corresponding fractions of draws with nominal diagnostic $p < 0.05$ rise from 5.17% to 9.14%, from 16.85% to 21.25%, and from 8.54% to 12.02%. These are descriptive summaries of fitted-residual diagnostics, not posterior probabilities that the model is misspecified. The 90% intervals intersect in all 1,152 ACF comparisons, which is compatible with broadly similar residual patterns but is not an equivalence test.

This stronger-prior fit has 683 nonstable states out of 12,000 (5.692%), all retained. The 100 unthinned monitors and 3,236 saved-state components meet the recorded convergence criteria; across saved-state components the maximum rank-normalized $\widehat{R}$ is 1.00105 and minimum bulk and tail ESS are 5,479 and 8,394. The completed comparison supports the main volatility chronology under stronger shrinkage while identifying sensitivity in some path levels and residual diagnostics. It does not supply a uniform diagnostic reason to replace the reference prior. The completed evidence here covers the variance multiplier of 0.5; the multiplier of two has not supplied results for this draft.

E.5 Supplementary switching-coefficient exercise

The supplementary comparison in Table E6 uses an earlier Case-Shiller, federal-funds/Wu–Xia, Cholesky specification with switching component variances. In both versions, the twelve FRM30 lag coefficients in the refinancing equation vary with the eighth component's variance state. Version A holds the intercept fixed; version B allows the intercept to vary with the same state. State labels are ordered by component

variance. The summary S_{12} is the sum of the twelve high-minus-low-state slope differences, measured in percentage points of refinancing intensity per percentage point of the mortgage rate. It is a conditional coefficient summary, not a cumulative policy response.

Table E6. Mortgage-rate coefficient differences and the intercept specification in supplementary switching models.

Refinancing conditional mean	**Median S_{12}**	**90% posterior interval**	$\Pr(S_{12} > 0)$
Fixed intercept; twelve switching mortgage-rate slopes	0.02068	[0.00909, 0.03269]	0.9978
Switching intercept and twelve mortgage-rate slopes	-0.00754	[-0.03776, 0.02229]	0.3333

Notes: Each fit retains 40,000 draws. States are ordered by the variance of the refinancing component, not by recessions or policy effectiveness. S_{12} sums differences in lag coefficients relating refinancing intensity in percentage points to the mortgage rate in percentage points; it is not a cumulative policy IRF or a long-run multiplier. Intervals are the fifth and ninety-fifth posterior percentiles within each fit. This is a supplementary Case-Shiller/federal-funds–Wu–Xia comparison, not a matched switching version of the main FHFA/one-year-yield model.

The 90% interval for S_{12} lies above zero in A but spans zero in B. In B, the posterior correlation between the intercept difference at the origin and S_{12} is -0.929. At the prespecified common mortgage-rate-history reference point, the conditional-mean difference between states has median 0.124 percentage points and a 90% interval [0.065, 0.186]. A conditional-mean difference can therefore remain clearly present at that reference point even when the slope-sum interval spans zero. The evidence concerns the allocation of conditional-mean variation between intercepts and slopes, not the disappearance of all differences between states.

This comparison explains why state-dependent mortgage-rate coefficients should be interpreted jointly with the rest of the conditional mean. It neither establishes constant slopes nor constitutes a matched switching-coefficient validation of the present FHFA/one-year-yield nonrecursive model. The current model's conditional-volatility findings remain grounded in its own posterior paths and matched sensitivity exercises.

E.6 Forecast-error variance and historical decomposition

E.6.1 State selection and forecast-variance definition

These calculations use the same FHFA, one-year-yield, twelve-lag structural-SV model as the main analysis. From each of the four chains, I select 500 equally spaced retained states with within-chain ordinals 3, 9, ..., 2,997. The selection is fixed independently of stability, response signs, or the decomposition results. It yields 2,000 equally weighted states, including 119 nonstable states; the separately reported stable subset has 1,881 states. The numerical summaries approximate the full 12,000-state posterior without re-estimating the model.

Write $\theta^{(s)}$ for the fixed coefficients and SV parameters in state s, $B = A_0^{-1}$ for the native impact matrix, and Ψ_ℓ for the moving-average matrices in equation (11). Conditional on the smoothed log variance h_{jt} and parameters at a forecast origin t, the fitted AR(1) law implies

$$m_{j,k} = \mu_j + \varphi_j^k(h_{jt} - \mu_j), \qquad v_{j,k} = \sigma_j^2 \sum_{r=0}^{k-1} \varphi_j^{2r}, \qquad \mathbb{E}[e^{h_{j,t+k}} \mid h_{jt}, \theta] = \exp(m_{j,k} + v_{j,k}/2), \quad k \geq 1. \tag{E4}$$

The variance term retains the lognormal Jensen correction. For response i, shock j, and forecast horizon H, the innovation-variance contribution is

$$V_{ij}(t,H) = \sum_{k=1}^{H} \left[(\Psi_{H-k}B)_{ij}\right]^2 \exp(m_{j,k} + v_{j,k}/2), \qquad \omega_{ij}(t,H) = \frac{V_{ij}(t,H)}{\sum_{a=1}^{8} V_{ia}(t,H)}. \tag{E5}$$

Thus $H = 12$ and $H = 36$ use moving-average coefficients 0 through $H - 1$. These are forecast horizons for monthly consumption growth and financing-intensity levels, not a decomposition of cumulative consumption-level uncertainty. The native impact matrix and innovation variances are used together; the +25-basis-point IRF scaling is not inserted into this accounting.

The main table reports posterior quantiles of the within-state equal-origin average

$$\overline{\omega}_{ir}^{(s)}(H) = \frac{1}{300} \sum_{t=\text{Jan1999}}^{\text{Dec2023}} \omega_{ir}^{(s)}(t,H). \tag{E6}$$

An average of dated ratios differs from a ratio of pooled variances. The volatility information at each origin is smoothed using the full sample, making the exercise retrospective. Each share decomposes innovation uncertainty conditional on a parameter state; posterior quantiles express uncertainty about that share. They do not add the cross-state variation in forecast conditional means to its denominator. The future-variance formula follows from this model's log-SV law. Brignone and Piffer (2025, Sections 3–4) provide a reference for the complementary economic uses of FEVD and historical decomposition under a different model specification.

The origins were selected before this calculation from previously reported volatility-peak dates: July 2003 for refinancing, January 2009 for housing valuation, and September 2022 for the mortgage rate. They were not selected for large policy shares. The main-text time-average table uses the same 300 origins for every response.

The central FEVD summaries are similar in the stable subset. Holding origin variances fixed changes the question: it removes their expected evolution over the forecast horizon. The thirty-six-month purchase median is 22.93% under that convention and 28.75% with expected future SV variances. Neither is substituted for the other. Appendix Figure E5 displays the dated shares under the main convention.

E.6.2 Historical contributions and accounting paths

Historical decomposition uses the twelve observed monthly vectors in 1998 as initial conditions and reconstructs native structural innovations from January 1999 onward. For shock j, its contribution at date t is

$$\mathcal{H}_{j,t} = \sum_{\tau=1}^{t} \Psi_{t-\tau} B_{\cdot j} \varepsilon_{j,\tau}, \qquad y_t = \mathcal{I}_t + \mathcal{C}_t + \sum_{j=1}^{8} \mathcal{H}_{j,t}. \tag{E7}$$

Table E7. Policy innovation's forecast-error variance shares at selected origins. Percent; posterior median [90% interval] across all 2,000 states. Future innovation variances follow equation (E4).

Origin	Response	H = 12	H = 36
2003-07	Consumption growth	0.25 [0.05, 1.24]	1.47 [0.28, 7.83]
2003-07	Purchase intensity	2.78 [0.42, 11.85]	23.25 [3.57, 59.10]
2003-07	Rate-term refinance intensity	0.53 [0.05, 3.27]	5.98 [0.27, 35.68]
2009-01	Consumption growth	0.70 [0.14, 3.44]	2.97 [0.49, 16.12]
2009-01	Purchase intensity	8.78 [1.59, 27.90]	40.95 [7.58, 77.47]
2009-01	Rate-term refinance intensity	7.44 [0.86, 29.82]	17.08 [1.42, 63.80]
2022-09	Consumption growth	1.30 [0.25, 5.88]	4.16 [0.66, 22.54]
2022-09	Purchase intensity	11.96 [1.72, 37.70]	51.84 [11.71, 84.23]
2022-09	Rate-term refinance intensity	23.50 [4.39, 60.72]	26.49 [3.10, 74.55]

Table E8. Time-average policy FEVD shares: state selection and future-variance convention. Percent; posterior median [90% interval]. The stable-only column uses 1,881 states and expected future SV variances. The frozen-origin column uses all 2,000 states and holds D_t fixed throughout the forecast horizon.

Response	H	All states	Stable-only states	Frozen-origin variance
Consumption growth	12	0.52 [0.13, 2.07]	0.51 [0.13, 2.08]	0.45 [0.12, 1.57]
Consumption growth	36	2.06 [0.39, 9.75]	2.01 [0.38, 9.75]	1.37 [0.31, 4.81]
Purchase intensity	12	6.36 [1.27, 14.42]	6.28 [1.28, 14.16]	6.74 [1.56, 13.91]
Purchase intensity	36	28.75 [6.62, 54.89]	28.20 [6.64, 53.32]	22.93 [5.64, 40.32]
Rate-term refinance intensity	12	8.08 [1.70, 18.91]	8.05 [1.76, 18.79]	10.65 [2.69, 20.89]
Rate-term refinance intensity	36	13.31 [1.56, 44.57]	13.18 [1.54, 43.83]	15.72 [4.01, 31.50]

Here $\mathcal{I}_t$ is the propagation of observed initial conditions with zero intercept and innovations, and $\mathcal{C}_t$ is the intercept path with zero initial conditions and innovations. Both follow the estimated autoregression. Their sum with the eight shock components reconstructs the observed model inputs within each state. The history is never restarted at the displayed window boundaries. Native $\varepsilon_{j,t}$ already contains its innovation scale, so the decomposition neither multiplies it by that scale again nor treats it as a repeated normalized 25-basis-point shock.

Table E9. Average monthly historical policy contributions in prespecified windows. Percentage points; posterior median [90% interval] across all 2,000 states. Each window average is calculated within a state before posterior summarization.

Window	Response	Average monthly policy contribution (pp)
2001-01–2003-12	Consumption growth	-0.027 [-0.077, 0.028]
2001-01–2003-12	Rate-term refinance intensity	0.148 [-0.022, 0.351]
2008-01–2009-12	Consumption growth	-0.016 [-0.080, 0.059]
2008-01–2009-12	Rate-term refinance intensity	0.101 [-0.139, 0.335]
2020-01–2023-12	Consumption growth	0.025 [-0.028, 0.100]
2020-01–2023-12	Rate-term refinance intensity	0.021 [-0.151, 0.274]

Every interval in this table includes zero. These are average monthly contributions, not cumulative consumption-level gaps or fractions of observed fluctuations. A positive contribution with a pointwise interval above zero at a particular date can coexist with an imprecise average over an entire window.

July 2003 is the prespecified refinancing-volatility origin. The additional dates describe features of the saved full-history paths: the purchase median reaches its maximum within 2001–04 in October 2004; the refinancing and purchase medians reach their maxima within 2020–23 in June 2020 and April 2022, respectively. These are maxima of pointwise median contribution paths, not posterior distributions of peak dates or tests of differences in peak timing. December 2023 records the sample endpoint. Each contribution includes all earlier policy innovations propagated to that date, rather than isolating the innovations of a selected easing or tightening episode.

The zero-policy-innovation accounting path is

$$y_t^{(-r,s)} = y_t - \mathcal{H}_{r,t}^{(s)}. \tag{E8}$$

It holds other reconstructed innovations, parameters, and observed initial conditions fixed. Since the observed input is fixed across states, its median and interval are obtained by subtracting the corresponding policy-contribution quantiles, reversing their order for interval endpoints. In July 2003, observed refinancing intensity is 3.164727% of lagged mortgage stock and the accounting-path median is 2.919743%, with a 90% interval of [2.633236%, 3.123709%]. In December 2023, the observed value is 0.040365% and the accounting-path median is 0.344489%, with interval [0.023434%, 0.738432%]. The latter interval includes

Table E10. Selected dated historical policy contributions. Percentage points; posterior median [90% pointwise interval] across all 2,000 states.

Date	Response	Policy contribution (pp)
2003-07	Consumption growth	-0.002 [-0.072, 0.088]
2003-07	Rate-term refinance intensity	0.245 [0.041, 0.531]
2004-10	Purchase intensity	0.100 [-0.128, 0.442]
2020-06	Rate-term refinance intensity	0.239 [0.036, 0.517]
2022-04	Purchase intensity	0.184 [-0.023, 0.512]
2023-12	Rate-term refinance intensity	-0.304 [-0.698, 0.017]

the observed input. These paths illustrate within-model attribution; they do not solve for behavior under an alternative policy regime.

E.6.3 Numerical support and sensitivity of long histories

All 2,000 selected states remain in the primary calculation. The maximum share-sum error is 4.44×10^{-16}, and the maximum scaled reconstruction error across dates, variables, and states is 2.40×10^{-12}. A separate implementation using the 96-dimensional companion system, explicit shock convolution, and direct forecast covariance construction cross-checks nine fixed sentinel states, including the state with the largest selected companion radius. Readback reproduces the saved FEVD table quantiles. These are numerical cross-checks of the accounting identities and implementation.

Across 36 monitored derived functions, the maximum rank/folded split $\widehat{R}$ is 1.00321; minimum bulk and tail effective sample sizes are 1,652 and 1,644. These diagnostics support the precision of the selected functions in the 2,000-state approximation, rather than every historical tail or the complete 12,000-state calculation.

Long histories amplify some nonstable states. In 29 selected states, the absolute policy contribution exceeds ten percentage points somewhere in the eight-variable history; large shock, intercept, and initial-condition components can offset one another while preserving exact reconstruction. The main plots therefore show stable-only interval limits alongside the all-state summaries. Over the full history, the largest all-versus-stable median differences are approximately 0.00115 percentage points for consumption growth and 0.00363 for refinancing intensity; their 90% interval endpoints are more sensitive. Add-up verifies the accounting, while this comparison records the economic sensitivity of long-history components. No state is silently removed or truncated.

E.6.4 Supplementary figures

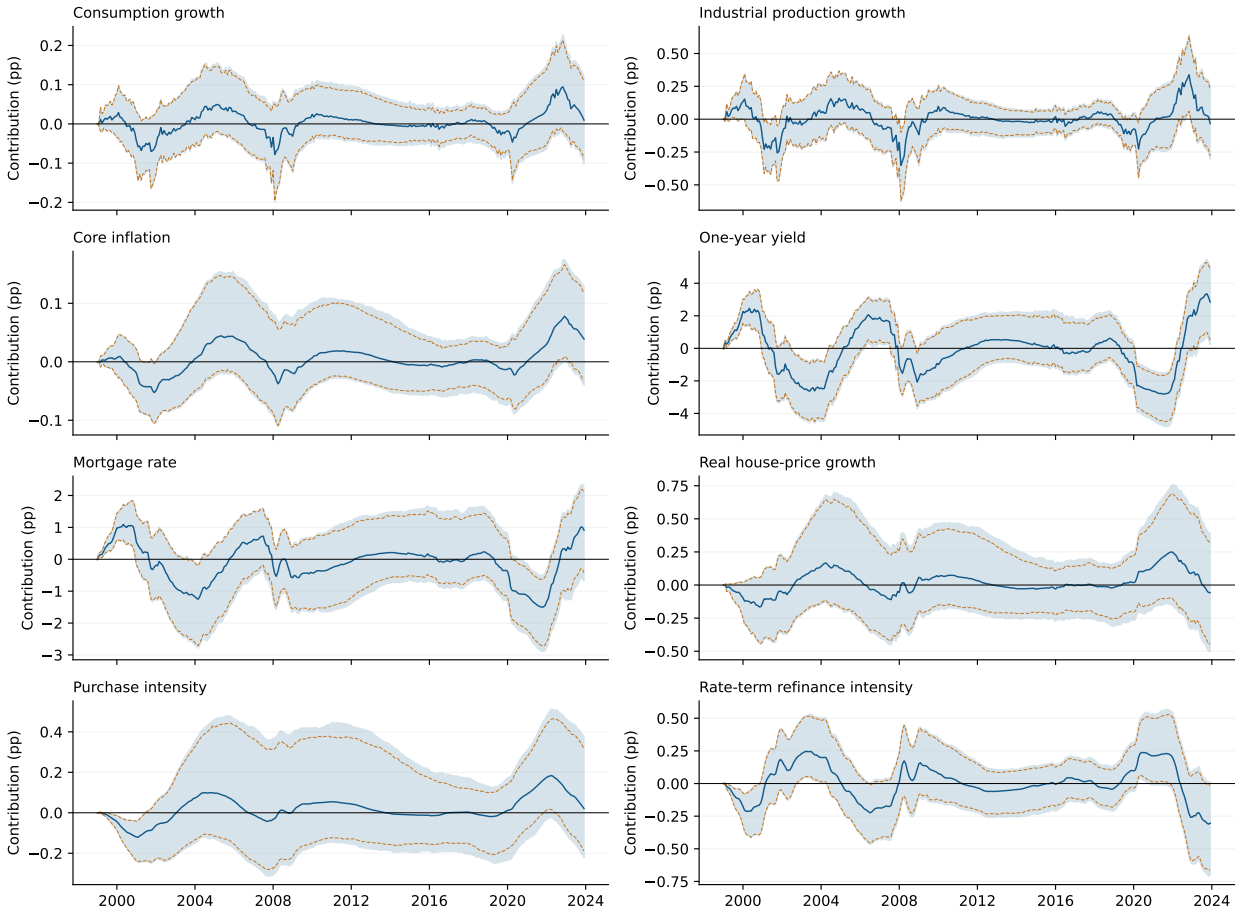


Figure E3. Historical contribution of the policy innovation to all eight model variables. Solid blue lines are medians and blue bands are pointwise 90% intervals across all 2,000 selected states; dashed orange lines are the stable-only interval limits. Histories start in January 1999 conditional on all twelve observed initial lags. Units follow the model inputs: C, IP, inflation, and housing appreciation are monthly growth or inflation percentage points; rates and financing intensities are percentage-point levels. Component medians need not add to observed data; equation (E7) holds within each state.

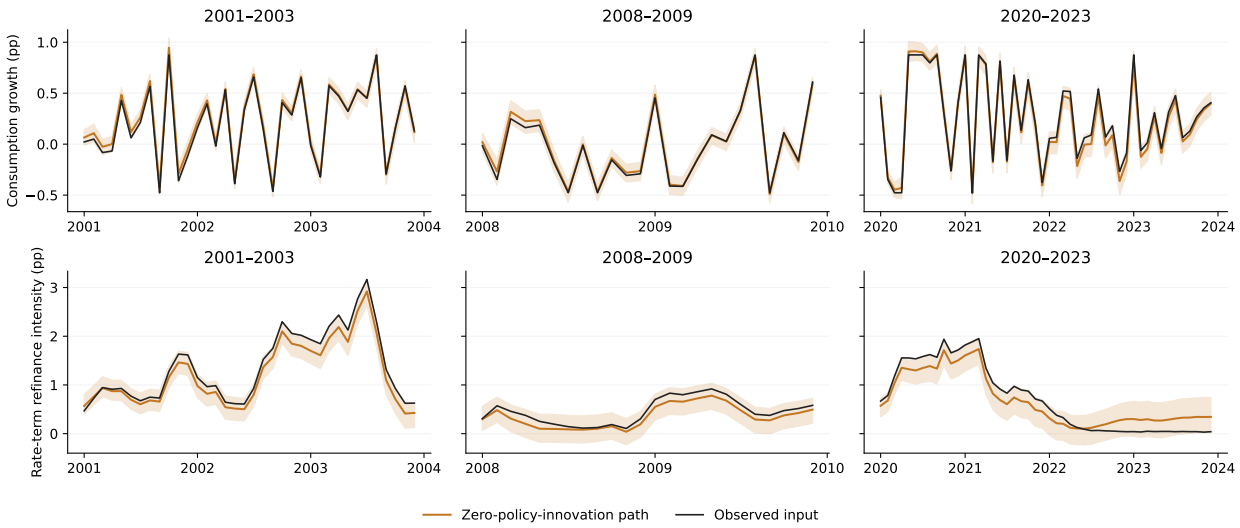


Figure E4. Observed inputs and zero-policy-innovation accounting paths. Black lines are the observed model inputs. Orange lines and bands are the posterior medians and pointwise 90% intervals of equation (E8), holding all other reconstructed innovations, coefficients, and initial conditions fixed. Rows show monthly real consumption growth and refinancing intensity in percentage points. Each displayed window retains the effects of innovations from earlier dates.

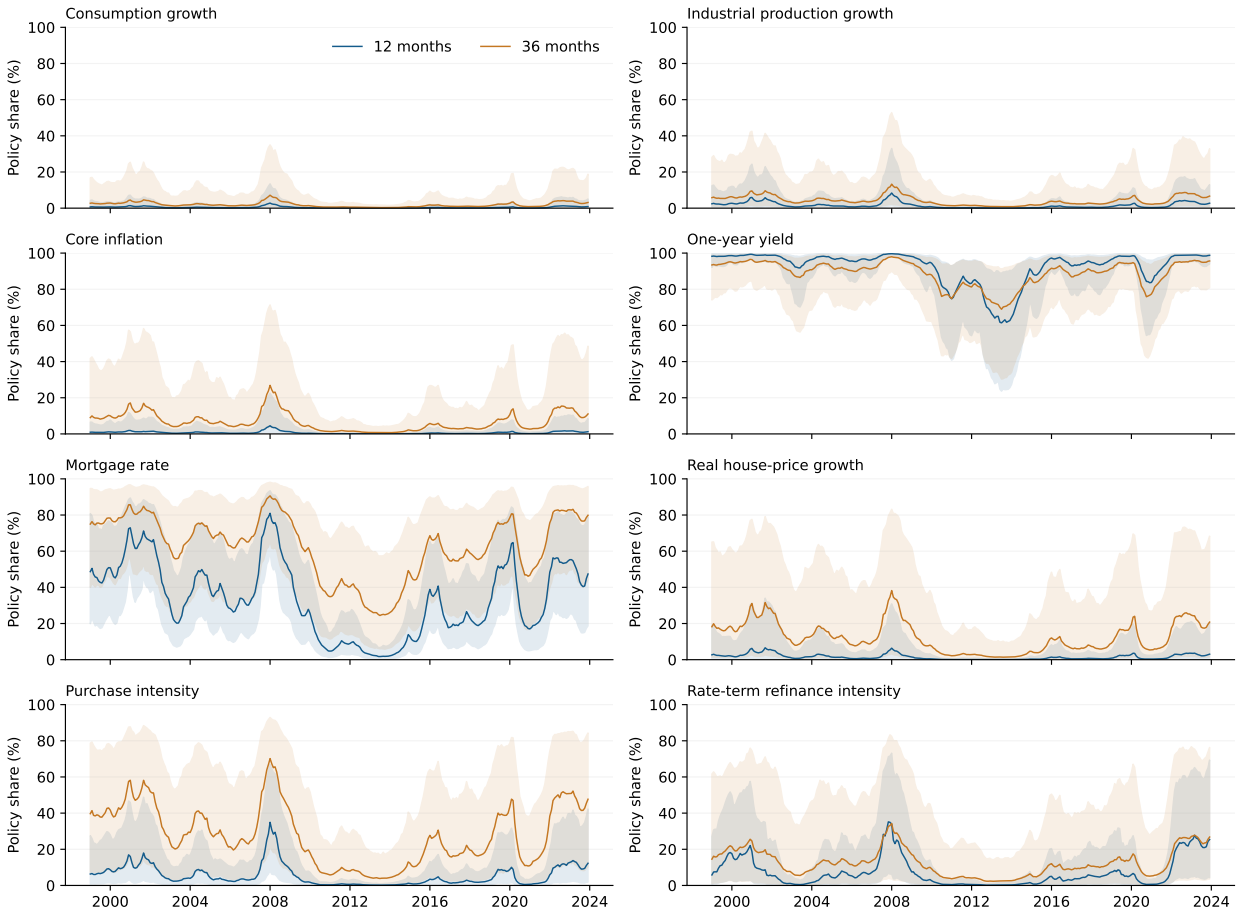


Figure E5. Policy innovation's dated forecast-error variance shares. Medians and pointwise 90% intervals across all 2,000 selected states for all eight model variables. Blue denotes the twelve-month forecast horizon and orange the thirty-six-month horizon. Shares use expected future SV variances and smoothed volatility at each historical origin. They are parameter-conditional innovation shares, expressed in percent.

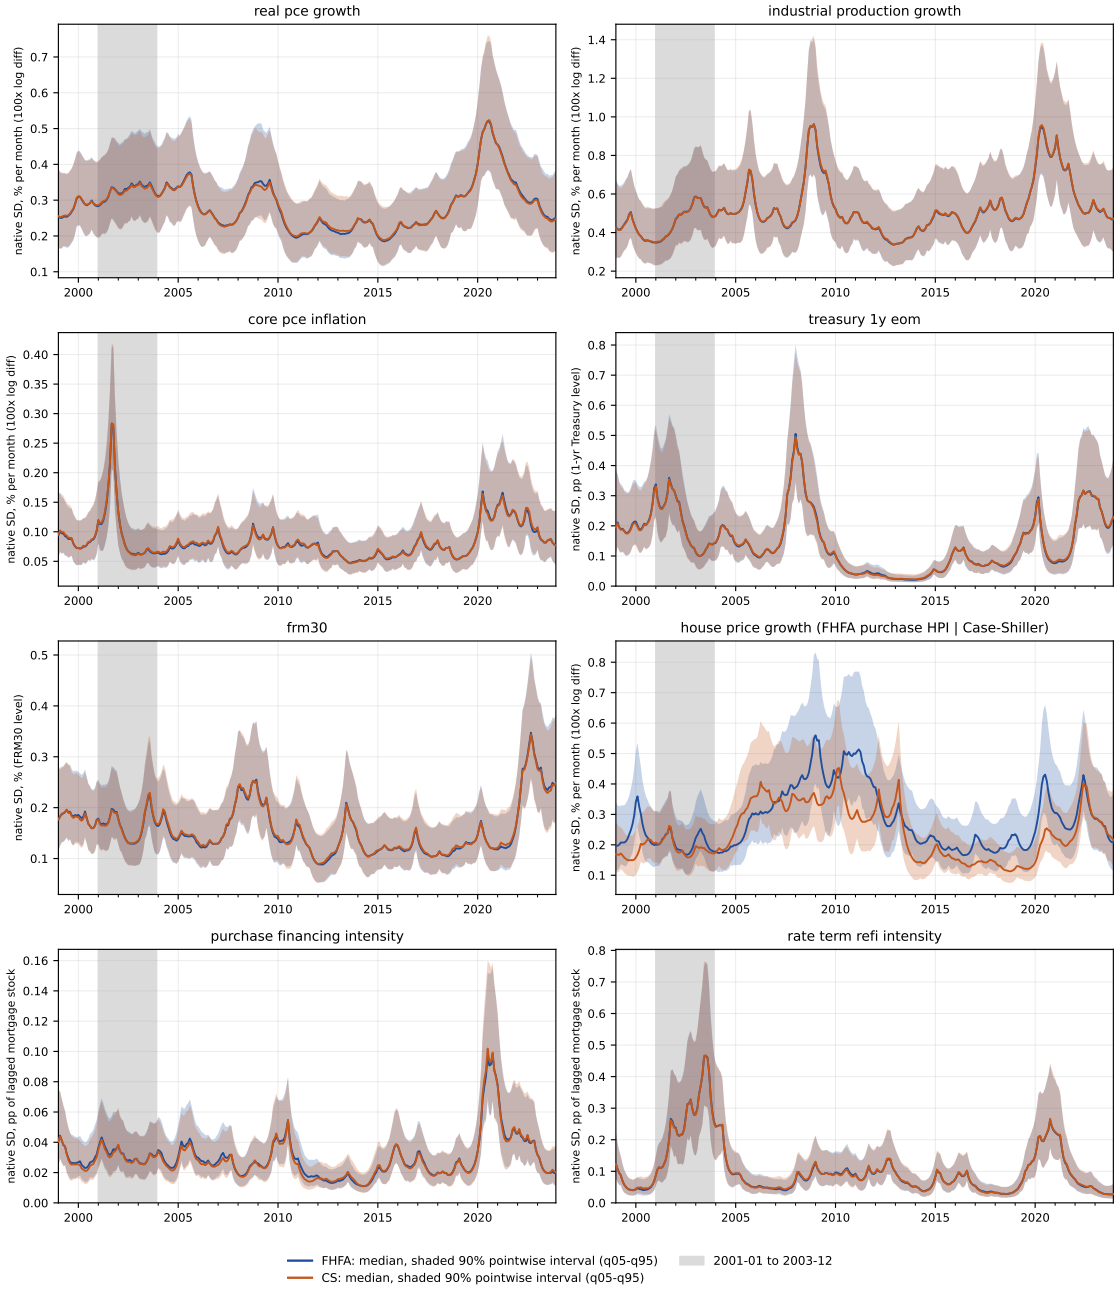


Figure E6. House-price measurement and marginal innovation volatility. Full-sample smoothed marginal innovation standard deviations for the FHFA and Case-Shiller specifications, January 1999–December 2023, in the units of each model input: monthly growth or inflation percentage points, interest-rate percentage points, and financing-intensity percentage points of lagged mortgage debt. Blue and orange lines are medians and shading gives central 90% pointwise credible intervals across all 12,000 joint states per specification; gray vertical shading marks January 2001–December 2003.

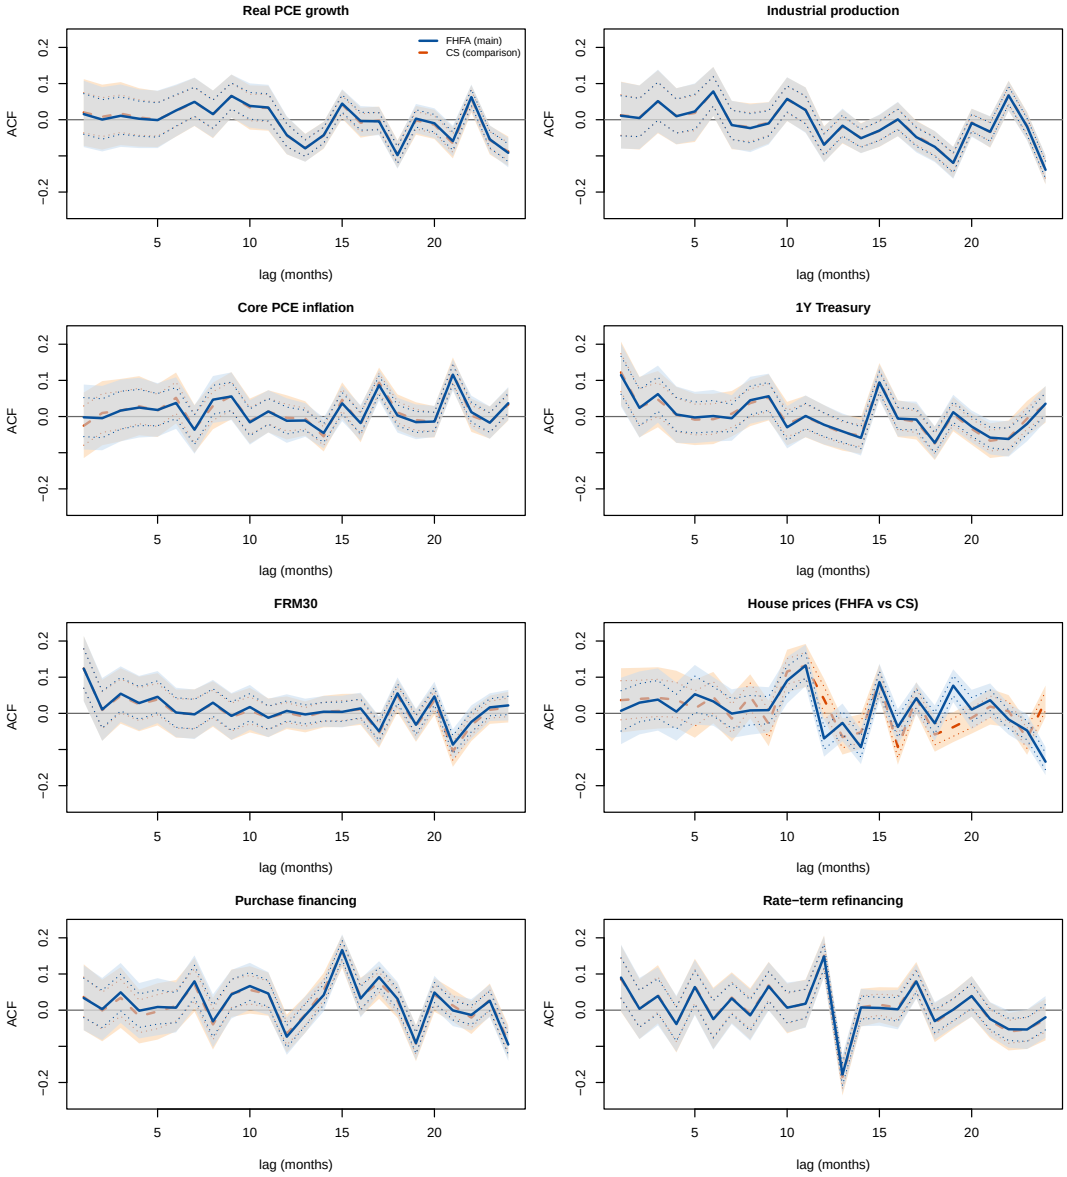


Figure E7. Signed autocorrelations of marginally standardized residuals. Dimensionless residual autocorrelations at lags 1–24 months after scaling each observable innovation by its time-varying marginal standard deviation, for the FHFA (blue) and Case-Shiller (orange) specifications. Lines show medians, shading shows central 90% pointwise intervals, and dotted limits show central 68% intervals across all 12,000 joint states per specification; these intervals describe posterior variation in fitted-residual diagnostics.

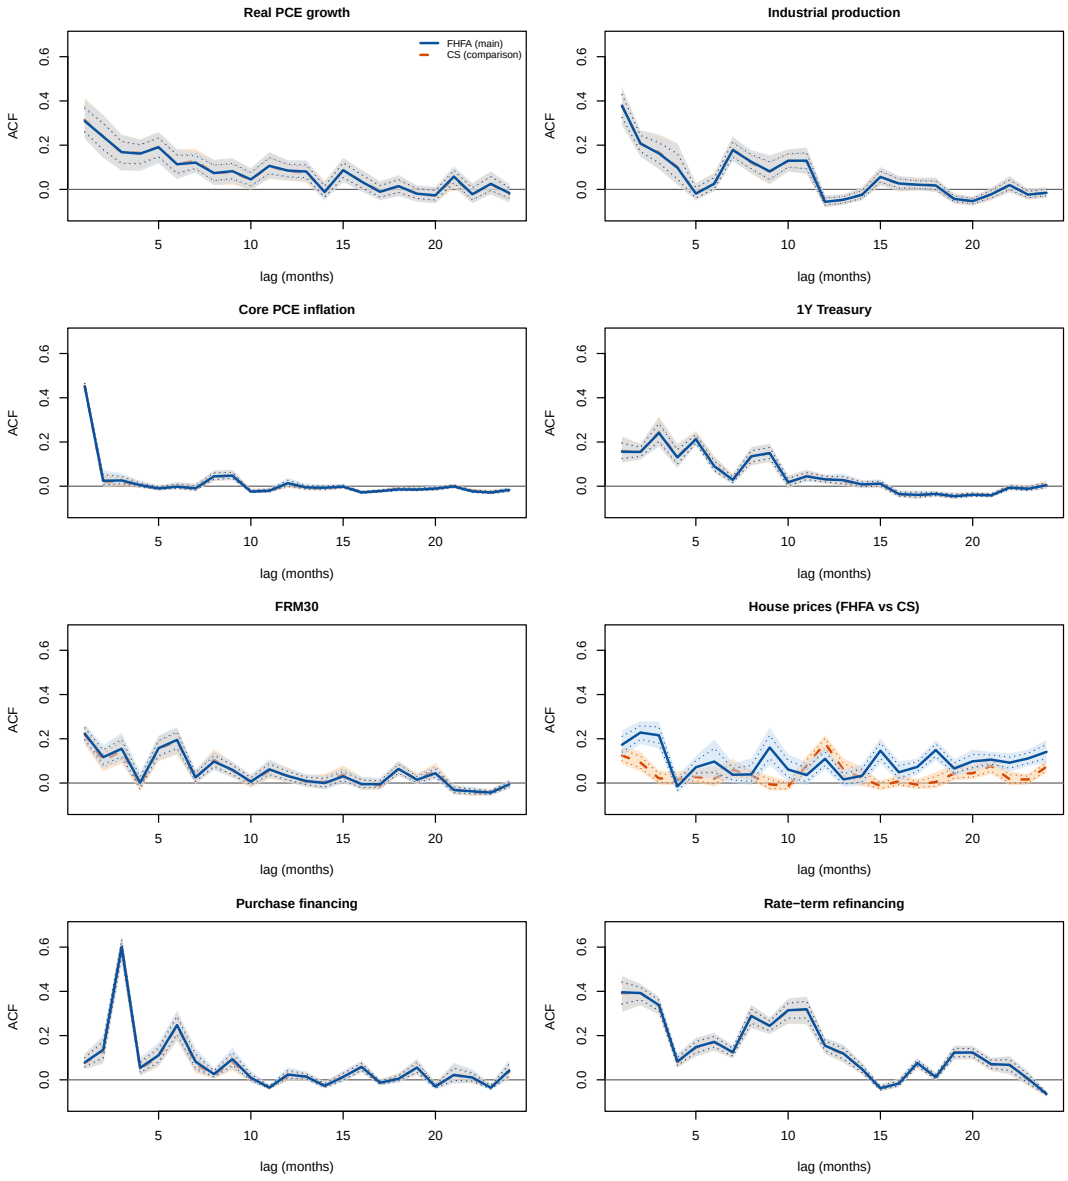


Figure E8. Autocorrelations of native squared residuals. Dimensionless autocorrelations of squared, unstandardized observable innovations at lags 1–24 months, for the FHFA (blue) and Case-Shiller (orange) specifications. Lines show medians, shading shows central 90% pointwise intervals, and dotted limits show central 68% intervals across all 12,000 joint states per specification, measuring persistence in residual magnitudes before volatility standardization.

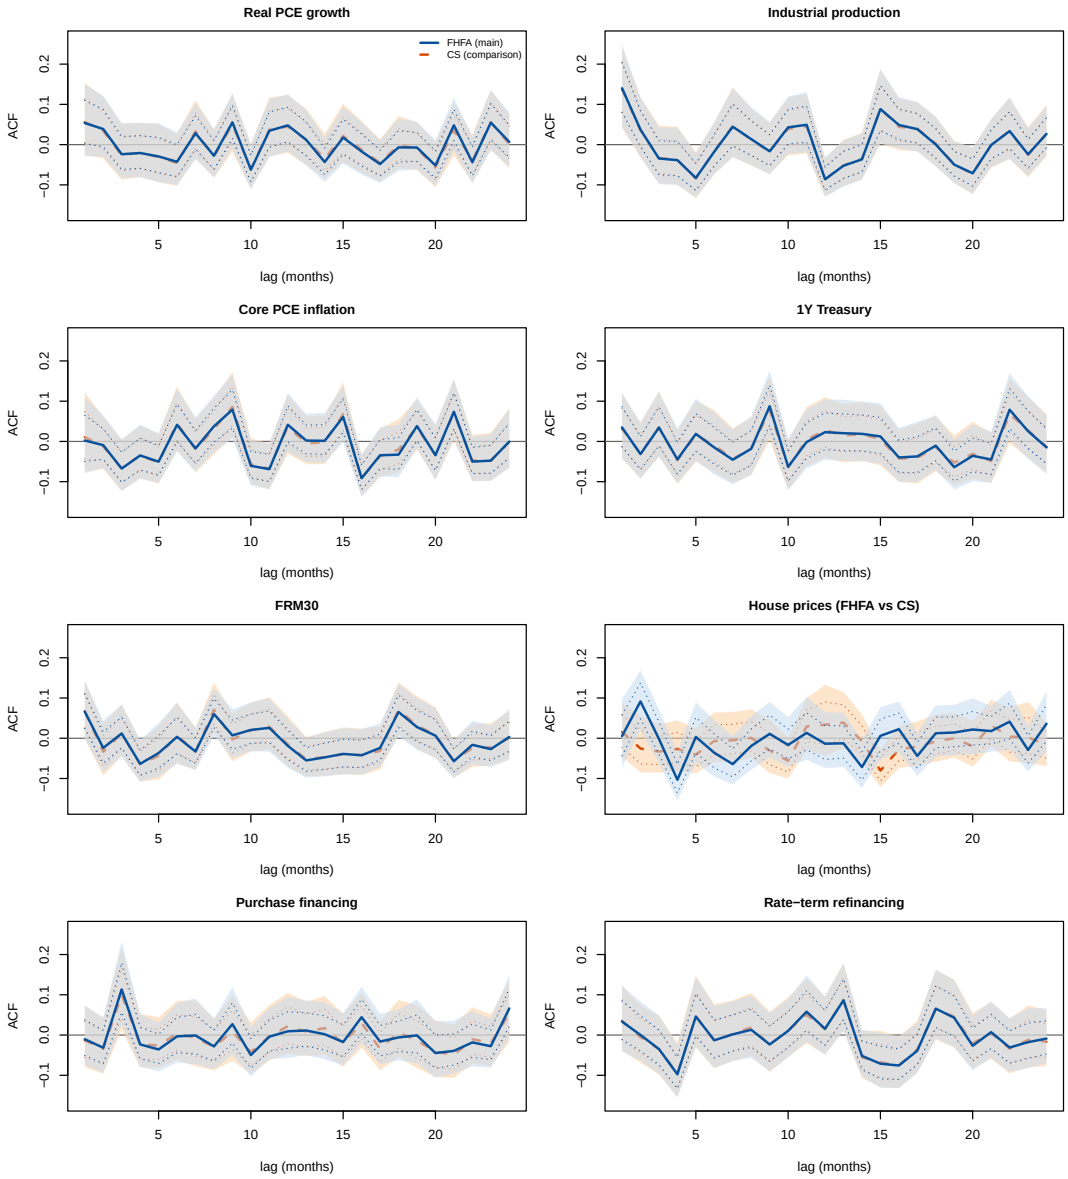


Figure E9. Autocorrelations of marginally standardized squared residuals. Dimensionless autocorrelations at lags 1–24 months after dividing each observable innovation by its time-varying marginal standard deviation and then squaring, for the FHFA (blue) and Case-Shiller (orange) specifications. Lines show medians, shading shows central 90% pointwise intervals, and dotted limits show central 68% intervals across all 12,000 joint states per specification, measuring remaining persistence in squared residuals after marginal volatility standardization.

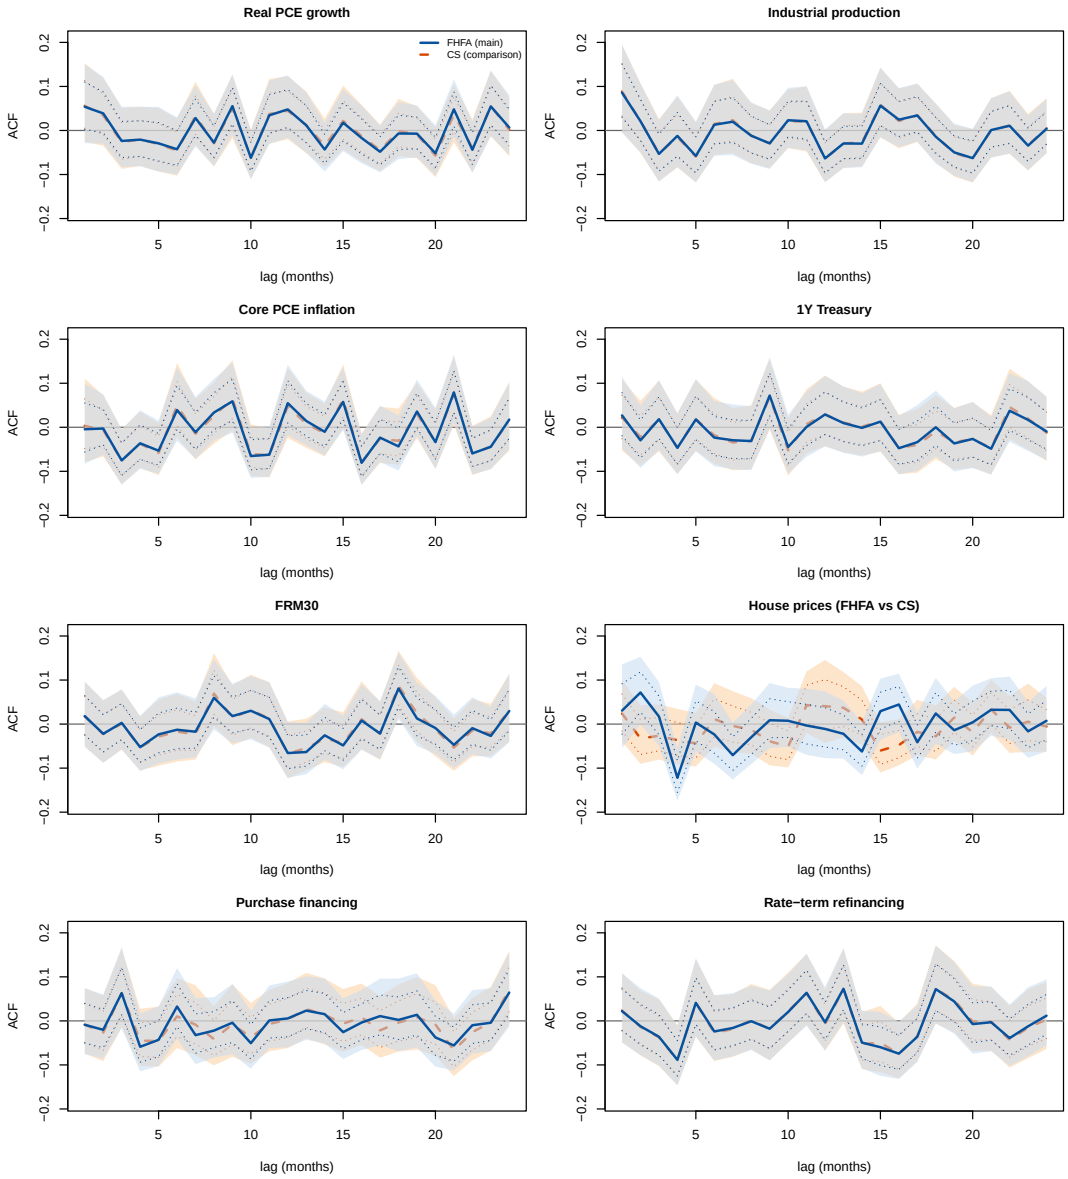


Figure E10. Autocorrelations of structurally standardized squared residuals. Dimensionless autocorrelations at lags 1–24 months for squared structural innovations standardized by their innovation-specific stochastic-volatility paths, for the FHFA (blue) and Case-Shiller (orange) specifications. Lines show medians, shading shows central 90% pointwise intervals, and dotted limits show central 68% intervals across all 12,000 joint states per specification, measuring remaining persistence after structural volatility standardization.

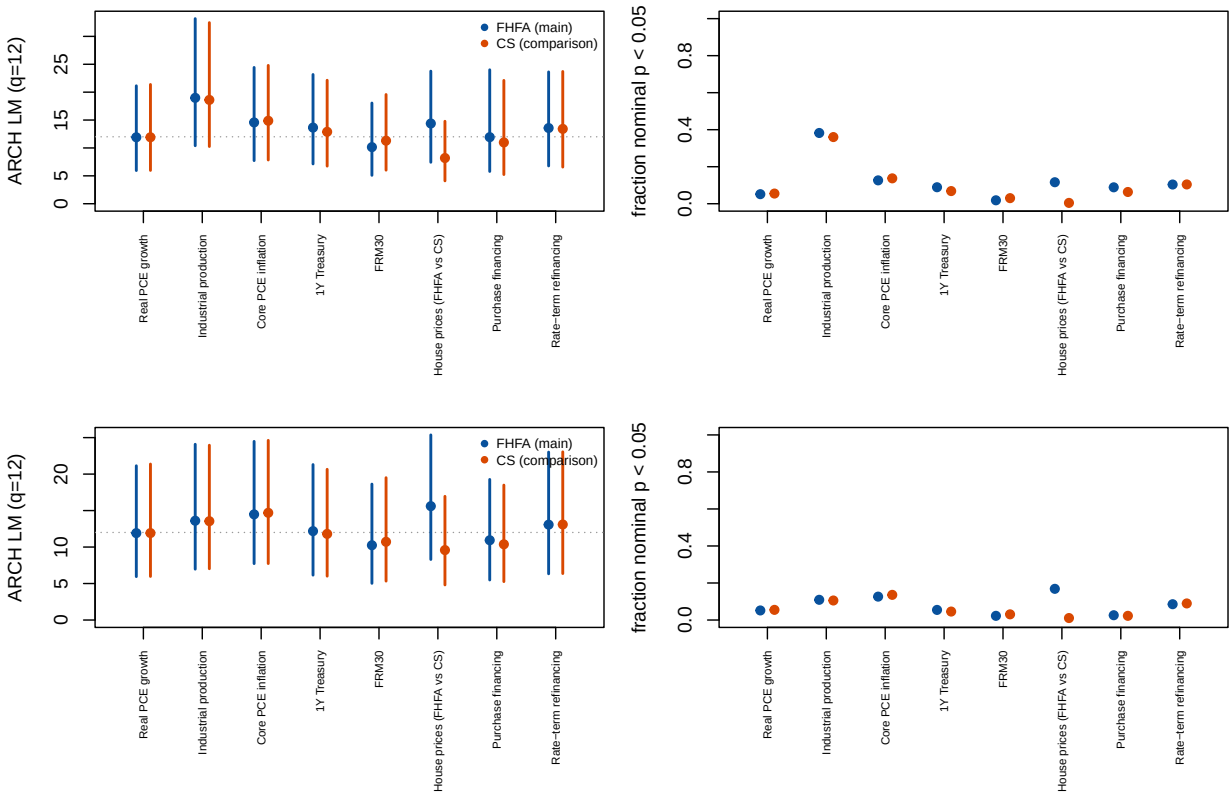


Figure E11. ARCH-type diagnostics at twelve lags. For marginally standardized residuals (top row) and structurally standardized innovations (bottom row), the left panels show dimensionless ARCH LM statistics and the right panels show the fraction of states with nominal $p < 0.05$, using all 12,000 joint states per specification (FHFA in blue; Case-Shiller in orange). Left-panel points are medians and vertical bars are central 90% intervals across states, with the dotted line at 12 marking the mean of the reference $\chi^2(12)$ distribution; the right-panel fractions are descriptive rejection frequencies without uncertainty bars.