

Identifying Policy Causal Effects from Rule Changes[†]

Ed Manuel

LSE

Christian K. Wolf

MIT & NBER

July 7, 2026

Abstract: Recent applied work has used interacted local projections to study how the propagation of macroeconomic shocks changes with the policy regime. We characterize the estimand of such strategies and relate it to the policy shock literature. Our main result is a set of conditions on the regressors and underlying data-generating process under which those two approaches are equivalent in the nature of their estimand, with both identifying slices of the same space of policy dynamic causal effects. Since policy is inherently high-dimensional, however, the two approaches generically recover different slices of that space. For example, for monetary policy, standard shocks tend to deliver the effects of transitory policy rate changes, while looking across policy regimes instead isolates gradual, more forward guidance-like policy treatments.

[†]Email: e.manuel@lse.ac.uk and ckwolf@mit.edu. We received helpful comments from Isaiah Andrews, Gabriel Chodorow-Reich, Marco Del Negro, Sílvia Gonçalves, Lukas Hack, Joe Hazell, Ethan Ilzetzi, Matthias Meier, Geert Mesters, Mikkel Plagborg-Møller, Ricardo Reis, Ludwig Straub, and Jonathan Wright. Wolf acknowledges that this material is based upon work supported by the NSF under Grant #2314736, and thanks the Alfred P. Sloan Foundation for financial support.

1 Introduction

A large literature in applied macroeconomics seeks to learn about the causal effects of changes in macroeconomic policy design. To date, the most common approach to estimating policy causal effects is to regress the policy instrument as well as any macroeconomic outcomes of interest on identified policy shocks (Sims, 1980; Ramey, 2016). In typical structural macroeconomic models, the estimand of such regressions furthermore admits a tight interpretation: it equals the causal effect on the macro-economy of the particular dynamic policy instrument path induced by the shock (see McKay & Wolf, 2023).

Some recent work takes a different approach, leveraging variation, either cross-sectional or across time, in the *systematic* conduct of policy.¹ For example, across time, researchers now often run interacted regressions (“local projections”, Jordà, 2005) of the form

$$y_{t+h} = \text{constant} + \beta_h \varepsilon_t + \gamma_h \zeta_t + \delta_h(\varepsilon_t \times \zeta_t) + \text{controls} + \text{error}, \quad (1)$$

where y_{t+h} is a policy variable or macroeconomic outcome h periods from now, ε_t is a (non-policy) shock, and ζ_t is a measure of the systematic conduct of policy (e.g., of the composition of the monetary policy committee). Intuitively, (1) asks how the propagation of ε_t is shaped by the design of systematic policy. Analogous cross-sectional specifications instead trace out the propagation of ε_t across different regions with different policy regimes.

The principal contribution of this paper is to give conditions on the regressors in (1) and on the data-generating process such that the interaction coefficients δ_h admit a straightforward interpretation analogous to that of regressions on policy shocks. Specifically, running (1) for the policy instrument affected by ζ_t gives the policy “treatment” path, and then doing so for other outcomes gives the associated causal effects. However, since policy is inherently high-dimensional — e.g., policy rates changing today or tomorrow —, any given implementation of (1) generally recovers a different slice of that large space of policy effects than any given policy shock regression. For example, typical monetary shocks identify the effects of transitory rate changes, while looking across policy regimes instead isolates gradual, forward guidance-like treatments, making (1) a useful complement to policy shock regressions.

THE CORE INTUITION. Before discussing our most general identification results, we begin by explaining the core economic intuition through a simple example. Consider two economies,

¹Well-known examples of studies running variants of (1) include Ramey & Zubairy (2018), Cloyne et al. (2023), Miyamoto et al. (2026), Fukui et al. (2025), and Hack et al. (2026). Also see the literature review.

either the same economy at two different points in time, or two different economies at the same time, that are identical in all respects except for the systematic design of a particular policy instrument — e.g., the economies follow different monetary policy rules. Now compare (linearized) impulse responses to the same macroeconomic event, e.g., to the same primitive demand or supply shock, across the two economies. Intuitively, any difference in outcomes is exclusively driven by differences in systematic policy design, and so should reflect causal effects of changes in the policy instrument. This logic is made precise by inverting the policy counterfactual results of McKay & Wolf (2023): differencing impulse responses of the policy instrument across regimes gives the policy “treatment,” and differencing other outcomes yields the associated causal effects — a standard differences-in-differences idea, straightforwardly implementable via local projections like (1). The same logic also delivers an intimate connection between comparisons across policy regimes and regressions on policy shocks. Just like comparisons across regimes, policy shocks induce a revision of the current and expected future policy instrument, and then deliver the causal effects of that “treatment.” Crucially, however, policy treatments, and so also their associated causal effects, are high-dimensional; in linearized economies, the treatment space is every possible *path* of the policy instrument. This yields our second main insight: that any given comparison across regimes and any given policy shock generically identify the effects of *different* policy treatments.

The preceding example was, of course, heavily stylized: policy design changed only once, and permanently, rather than stochastically, and possibly for endogenous reasons; and *only* policy changed, with all the other fundamentals of the macro-economy perfectly stable. Our main identification results concern environments that relax these restrictions.

GENERAL ENVIRONMENT. We address the limitations of the stylized example in two steps, first generalizing assumptions on policy design, and then allowing other, non-policy forms of variation in shock propagation.

The formal setting for our analysis is a general family of linearized macroeconomic models, with the exact same scope and limitations as in McKay & Wolf (2023), but now augmented to allow for essentially unrestricted variation over time in how the policymaker sets her instrument.² Appealingly, this setting is general enough to nest classic structural analyses of stochastically switching policy regimes (e.g., see Davig & Leeper, 2007; Bianchi, 2013),

²The key restriction is that policy shapes private-sector behavior exclusively through current and expected future values of the policy instrument. This is the case in many popular business-cycle models, but typically fails with informational asymmetries between private sector and policymakers.

including endogenous regime switches, thus addressing the first limitation. Key to our identification arguments is the observation that, even in this rich class of environments, equilibrium outcomes admit a simple additive decomposition. At each date t , newly arriving information, in the form of either primitive macroeconomic shocks or policy rule changes, leads to surprise revisions of current as well as expected future macroeconomic outcomes. We prove that these revisions can be written as the sum of two terms: the impulse responses of date- t shocks under some arbitrary “baseline” policy rule, plus the product of policy causal effects times the deviation of actual policy from baseline. This decomposition is exactly analogous to the across-regime intuition described above, just now formalized in a general environment with constantly, and perhaps endogenously, switching policy.

We then generalize the environment further to allow for other sources of variation in the non-policy block of the model, e.g., reflecting permanent differences across countries, or gradual changes in the structure of the economy over time. Returning to our decomposition of equilibrium outcomes, this model extension means that both “baseline” impulse responses as well as policy causal effects are possibly heterogeneous, varying either across countries or across time with a further, possibly endogenous, aggregate state.

RUNNING INTERACTED LPS. Having described the data-generating process, we now turn to conditions on the regressors (ε_t, ζ_t) such that the estimand of δ_h in (1) indeed equals our claimed interpretation. Mechanically, (1) captures how the dynamic propagation of the shock ε_t co-varies with the policy regime variable ζ_t . By our additive decomposition of equilibrium outcomes, this co-movement will capture policy causal effects under two key requirements. First, ε_t needs to be a fixed conditioning event that itself is not shaped by the policy regime; proper structural shocks, or more broadly any macroeconomic news not itself shaped by ζ_t , work, while reduced-form statistical surprises do not.³ Second, the dynamic propagation of ε_t must change with ζ_t *only* because of policy: the interaction $\varepsilon_t \times \zeta_t$ must forecast the policy instrument, and must not co-vary with any other sources of variation, either across time or cross-sectional, in shock propagation. This requirement in principle is consistent with ζ_t being highly endogenous, e.g., systematic policy leaning more aggressively against inflation following past inflationary shocks. What is ruled out, however, is that those past events — to which systematic policy responds — *themselves* shape the propagation of ε_t . Under those

³More generally, just asking how unconditional second moments — even if induced by a stable set of shocks, and given stable non-policy primitives — change with systematic policy is not enough: such exercises yield non-trivial, but typically excessively wide, identified sets for policy causal effects. See Appendix A.2.

two key conditions, the interaction term $\varepsilon_t \times \zeta_t$ co-varies only with the second part of our additive equilibrium decomposition, delivering the desired conclusion: the δ_h 's for the policy instrument are the treatment, and those for endogenous outcomes are the causal effects.

We supplement this headline identification result with several further observations on the estimated policy causal effects. First, as in the earlier simple example, any given interacted local projection (1) and any given policy shock regression will generically identify different policy treatments; indeed, even re-running (1) for different shocks ε_t will in general identify different treatments. Second, if ζ_t changes the design of multiple policy instruments — e.g., interest rate policy and quantitative easing — then (1) will recover the causal effects of a joint policy treatment. Third, if policy effects are heterogeneous over time or in the cross section, then an additional monotonicity condition is required to allow a causal interpretation of the δ_h 's, closely echoing assumptions made in the literature on heterogeneous treatment effects.

ILLUSTRATIVE APPLICATIONS. With our results on required regressor properties and estimand interpretation in hand, we next illustrate our findings by revisiting two sets of applications from prior work, one leveraging policy variation in the time series, one cross-sectional.

Our first and main application is to monetary policy transmission: What kind of policy treatments does recent applied work running variants of the interacted local projection (1) actually identify? And how does that estimand relate to the voluminous literature on identified monetary policy shocks? As is well-known, regressions on textbook monetary policy shocks tend to isolate rather transitory interest rate treatments (e.g., see Caravello et al., 2025). Theory suggests that matters may look quite different for interacted local projections: if the shock ε_t itself has persistent aggregate effects, then it will typically also induce persistent movements in the policy instrument, and so differencing across policy regimes should isolate persistent policy treatments. We confirm this intuition by revisiting and further extending the recent work of Hack et al. (2026). Those authors identify a measure of systematic U.S. monetary policy conduct — the hawk-dove balance of the FOMC, for additional robustness further instrumented based on FOMC voting rights rotations — that plausibly governs shock propagation chiefly through its implications for monetary policy conduct. Interacted with oil as well as fiscal policy shocks, that measure then predicts significant, gradual responses of the policy rate, i.e., forward guidance-like monetary policy treatments. For both of the non-policy shocks ε_t that we look at, the implied delayed interest rate hikes tend to induce similarly delayed contractions in output, with the price responses more muted. Through the lens of our identification results, those findings offer a very useful complement to standard

regressions on monetary policy shocks: while both strategies *in principle* are equivalent in the nature of their estimand, they *in practice*, and for intuitive economic reasons, identify very different slices of that common causal effect space.

For our second, cross-sectional, application we revisit Fukui et al. (2025). Relative to their original analysis, we change the conditioning event ε_t included in the interacted local projection. Doing so illustrates our observation that the implicit policy treatment identified by (1) can change materially as a function of the included shock variable: in some cases we isolate the effects of more transitory exchange rate depreciations, in others we learn about a combination of exchange rate and interest rate treatments.

FURTHER LITERATURE. Our paper first and most obviously adds to a fast-growing literature on interacted local projections of the form (1). A review of that literature, a connection to the Kitagawa-Oaxaca-Blinder decomposition, and an application to fiscal multipliers are all provided in Cloyne et al. (2023). Relative to that paper, and to the broader applied literature, our contribution is to provide a structural interpretation of the econometric estimand, and to prove equivalence of the nature of that estimand with standard policy shock regressions, unifying and connecting two literatures that hitherto have largely been separate. Those same contributions also distinguish our work from Gonçalves et al. (2024), Kolesár & Plagborg-Møller (2025) and Winkler (2026), who all study state-dependent local projections.

Zooming out, our paper is situated in a broader literature that exploits stability restrictions for identification (Lewbel, 2012; Magnusson & Mavroeidis, 2014; Mavroeidis, 2021). We here leverage such ideas in the context of changes in policy design, and provide identification results for a large class of macroeconomic models as data-generating processes. Our arguments build on earlier results in McKay & Wolf (2023), Barnichon & Mesters (2023), and Caravello et al. (2025). In terms of theory, our main incremental contributions are to show when the mappings characterized in those papers are invertible — yes for shocks, no for second moments — and to extend the analysis to models with stochastically time-varying policy design and other sources of state dependence in shock propagation.

Finally, there is a close relationship between (1) and the “shift-share” regression specifications common in the applied microeconomic literature (see Adao et al., 2019; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022). While shift-share regressions compare shock responses across heterogeneously exposed micro units, (1) instead compares shock responses over time or across countries, by policy regime. Our identification results build a bridge from those strategies to the classical time series approach of regressing on policy shocks.

2 A simple example

Echoing the introduction, we begin with a simple two-period example to illustrate the core ideas. Section 2.1 introduces the environment and provides an additive decomposition of equilibrium outcomes. Section 2.2 then previews our main results, contrasting policy shock regressions and interacted local projections as means of recovering policy causal effects.

2.1 Environment

To preview our identification results we use a simple two-period variant of the canonical textbook New Keynesian model (see Galí, 2015; Woodford, 2003).

MODEL. The economy consists of two blocks. The first block collects all private-sector relations, here consisting of the usual IS and NKPC equations:

$$x_t = -\frac{1}{\gamma} (i_t - \mathbb{E}_t [\pi_{t+1}]) + \mathbb{E}_t [x_{t+1}] + (e_{d,t} + \theta_d e_{d,t-1}), \quad (\text{IS})$$

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t [\pi_{t+1}] + (e_{s,t} + \theta_s e_{s,t-1}), \quad (\text{NKPC})$$

where x_t denotes output, i_t is the policy rate, π_t is inflation, and $(e_{d,t}, e_{s,t})$ are demand and supply shocks. The second block is the policy block, here a standard Taylor-type rule:

$$i_t = \phi_\pi \pi_t + \underbrace{v_{0,t}}_{\text{contemp. shock}} + \underbrace{v_{1,t-1}}_{\text{1-period news shock}}, \quad (\text{TR})$$

where $v_{0,t}$ denotes a conventional monetary policy shock that realizes at date t and perturbs policy at date t , while $v_{1,t-1}$ is a one-period forward guidance shock.

We collect endogenous outcomes in the vector $y_t \equiv (x_t, \pi_t, i_t)'$ and shocks in the vector $\varepsilon_t \equiv (e_{d,t}, e_{s,t}, v_{0,t}, v_{1,t})'$, and assume that $\varepsilon_t \stackrel{\text{iid}}{\sim} (0, I_{n_\varepsilon})$. Since the model features no endogenous propagation, and since all shocks at most affect equilibrium dynamics one period in advance, all the shock impulse responses die out after at most one period. We stack impulse responses to the demand shock in the 6×1 vector Θ_d , and do the same for the other shocks as Θ_s , $\Theta_{v,0}$ and $\Theta_{v,1}$, with $\Theta_v \equiv (\Theta_{v,0}, \Theta_{v,1})$ collecting all policy shock impulse responses.

POLICY REGIMES AND EQUILIBRIUM CHARACTERIZATION. We suppose that the economy can be in one of two distinct policy regimes; consistent with the applied literature to which

our analysis relates, the two regimes should be interpreted as two different policy rules over time for the same economy, or for two otherwise identical economies. In the first, “baseline” regime, the rule coefficient ϕ_π is equal to $\bar{\phi}_\pi$, and we indicate all corresponding shock impulse responses with bars. In the hawkish regime, the coefficient is instead $\phi_\pi^H > \bar{\phi}_\pi$, and we denote all impulse responses by H superscripts.

A key building block result — which later will be echoed in our more general economies in Sections 3 to 5 — is that equilibrium outcomes admit a simple additive decomposition into “baseline” impulse responses plus policy (shock) causal effects. Specifically, in the two-regime setting considered here, it follows directly from McKay & Wolf (2023, Proposition 1) that the demand and supply shock impulse responses in the hawkish regime satisfy

$$\Theta_d^H = \bar{\Theta}_d + \bar{\Theta}_{v,0}\omega_{d,0} + \bar{\Theta}_{v,1}\omega_{d,1} \quad (2)$$

as well as

$$\Theta_s^H = \bar{\Theta}_s + \bar{\Theta}_{v,0}\omega_{s,0} + \bar{\Theta}_{v,1}\omega_{s,1} \quad (3)$$

for suitably chosen scalars $\{\omega_{d,0}, \omega_{d,1}\}$ and $\{\omega_{s,0}, \omega_{s,1}\}$. In words, impulse responses to demand and supply shocks in the hawkish regime equal the corresponding impulse responses in the baseline regime plus policy shock impulse responses $\bar{\Theta}_v$ times suitably chosen artificial policy shocks $\omega_\bullet$. The basic idea is simple. In the model considered here, private-sector behavior is shaped by the policymaker’s actions only through current and expected future values of the policy instrument — what McKay & Wolf (2023) have called the “instrument sufficiency” property of a general class of models, including the particular one that we consider here. As a result, it does not matter whether, say, interest rates are high because the rule is hawkish or because a dovish rule was just subject to hawkish shocks; if the resulting level of interest rates is the same, then outcomes are invariably the same. (2) - (3) leverage this logic, with the artificial policy shocks to the baseline policy rule chosen to ensure that, in equilibrium, the policy instrument moves exactly as it would under the hawkish rule.

2.2 Learning about policy causal effects

We now consider an econometrician that lives in this economy and wishes to learn about the effects of monetary policy. We first elaborate on the definition of policy causal effects as the object of interest, before then turning to two empirical strategies — policy shock regressions and interacted policy regime local projections — to learn about them.

THE OBJECT OF INTEREST. Recall that policy shock impulse responses under the baseline rule are denoted by $\bar{\Theta}_v$. For our subsequent analysis, it will prove useful to translate those policy responses from so-called policy shock space to policy instrument space, proceeding as follows. Let

$$\Theta_i \equiv \bar{\Theta}_v \cdot \bar{\Theta}_{i,v}^{-1},$$

where $\bar{\Theta}_{i,v}$ denotes the impulse responses of the policy instrument, i.e., of the short-term rate i_t , to the two monetary policy shocks. While the two columns of $\bar{\Theta}_v$ correspond to the causal effects of monetary policy *shocks* — the standard definition of impulse response functions — Θ_i instead collects causal effects corresponding to *interest rate* paths, with the first column giving the effects of a contemporaneous rate hike without any further rate change tomorrow, and *vice-versa* for the second column. Because of instrument sufficiency, and as we show with the closed-form expressions for Θ_i in Appendix A.1, that causal effect matrix does not actually depend on the monetary policy rule: we can recover the causal effects of an interest rate path by solving (IS) - (NKPC) in isolation, without any reference to the monetary policy rule and thus to ϕ_π .⁴ It follows in particular that Θ_i could have equivalently and without change been defined using the hawkish rule’s Θ_v^H .

For future reference, we also note that policy causal effects in instrument space, and not shock space, are what actually matters for most standard uses of policy causal effects: Θ_i , the effects of interest rate paths, are all that is needed for both model estimation through impulse response matching (as in Christiano et al., 2005) and for policy counterfactual evaluation (as in McKay & Wolf, 2023). Our object of interest throughout this paper will thus be policy causal effects in instrument space, and not shock space.

USING POLICY SHOCKS. The canonical approach, followed in a voluminous literature that goes back to Sims (1980), is to obtain measures of $v_{0,t}$ or $v_{1,t}$, or linear combinations thereof. Specifically, suppose the econometrician observes the policy “instrument”

$$w_t = \omega_{w,0}v_{0,t} + \omega_{w,1}v_{1,t} + \text{noise},$$

and then, for $h = 0, 1$ and generic outcome y_t , runs

$$y_{t+h} = \text{constant} + \beta_{y,h}w_t + \text{error}.$$

⁴More precisely, the policy rule plays two roles: regulating the on-equilibrium path of interest rates, and equilibrium selection. By enforcing *exact* return to steady state after two periods, we remove the selection margin (see Angeletos et al., 2026), leaving only the on-equilibrium role, and so delivering irrelevance of ϕ_π .

If this regression is run in the baseline policy regime (say), then the estimand is equal to a weighted average of columns of $\bar{\Theta}_v$ with weights proportional to $(\omega_{w,0}, \omega_{w,1})$, or equivalently to a weighted average of columns of Θ_i with the rate response $\beta_i \equiv (\beta_{i,0}, \beta_{i,1})'$ as the weights.

We note two important challenges with this canonical “policy shock” approach. First, as is well-known, it is often difficult to come up with credible policy shock measures. Second, and somewhat less appreciated, another non-trivial challenge is that the object of interest, Θ_i , is actually multi-dimensional, corresponding to the dynamic causal effects of changes in policy at all horizons. The available empirical variation, in contrast, is often limited (see the literature review for monetary policy shocks in Caravello et al., 2025).

USING POLICY REGIMES. Suppose now instead that the econometrician observes data from the two policy regimes, and furthermore has access to either or even both of the demand and supply shocks, $e_{d,t}$ and $e_{s,t}$. She then contemplates running the *interacted* local projections

$$y_{t+h} = \text{constant} + \beta_{y,h}\varepsilon_{j,t} + \gamma_{y,h}\zeta_t + \delta_{y,h}(\varepsilon_{j,t} \times \zeta_t) + \text{error}, \quad j = d, s, \quad (4)$$

where again $h = 0, 1$, $\varepsilon_{j,t}$ is either $e_{d,t}$ or $e_{s,t}$, and finally ζ_t is a dummy indicator for the hawkish regime. It is straightforward to see from (2) - (3) that $\delta_y \equiv (\delta_{y,0}, \delta_{y,1})'$ satisfies

$$\delta_y = \Theta_{y,j}^H - \bar{\Theta}_{y,j}$$

where the y subscripts indicate impulse responses of outcome y , and thus

$$\delta_y = \bar{\Theta}_{y,v} \cdot \omega_j,$$

i.e., the interacted local projection recovers policy causal effects, just like regressions on policy shocks. The intuition is straightforward: because of instrument sufficiency, the difference in propagation over time of the same primitive innovation — here either demand or supply shocks — across policy regimes is exclusively a function of policy causal effects. The weights $\omega_j \equiv (\omega_{j,0}, \omega_{j,1})'$ are unknown to the econometrician, but this is no problem since ultimately we anyway are interested in policy causal effects in instrument rather than shock space: by definition of $\Theta_{y,i}$, the local projection estimands satisfy

$$\delta_y = \underbrace{\Theta_{y,i}}_{\text{policy causal effects}} \times \delta_i.$$

In words, first running the interacted local projection (4) for the policy instrument gives the treatment δ_i , and the impulse responses for the other variables then give the associated causal

effects, perfectly analogous to our interpretation above of the usual policy shock regression estimands. We note further that the weights ω_w in any policy shock instrumental variable w_t need not, and in practice generically will not, agree with the weights ω_j isolated by (4). Interacted local projections are thus *both* an alternative *and* a complement to the standard policy shock regressions: an alternative since they achieve identification under very different informational requirements, and a complement since they generally recover the effect of a different sequence of policy shocks, and thus of a different policy instrument path.

Our final observation is that the estimand of (4) not only is generically different from any given policy shock instrumental variable w_t , but also differs by shock $\varepsilon_{j,t}$. For example, if in our current setting the supply shock has an MA(1) component ($\theta_s \neq 0$) but the demand shock does not ($\theta_d = 0$), then the interacted local projection for supply shocks will generally identify a mix of contemporaneous and forward guidance monetary shocks, while the demand shock estimand just equals that of a contemporaneous shock. This again is practically useful: if the econometrician had access to both shocks, then she could run (4) twice, getting a richer picture of how different policy treatments affect the macro-economy.⁵

ASIDE: SECOND MOMENTS. The preceding discussion revealed that across-regime differences in structural shock impulse responses can be used to identify policy causal effects. A natural follow-up question is whether changes in reduced-form second moments — either in impulse responses to Wold innovations, or in autocovariance functions — can do the same. In particular, since policy causal effects are known to be sufficient to map actual into counterfactual reduced-form second moments (see Caravello et al., 2025), the question is whether that mapping is invertible. Perhaps surprisingly, the answer to this turns out to be negative; we postpone a detailed discussion of this point to Section 3.3.

2.3 Outlook

The motivating discussion in this section was subject to three obvious, and material, limitations. First, we considered a highly specific, and two-period, structural model. Second, we simply assumed that the econometrician has access to data that come from two disjoint, permanent policy regimes. In practice, to the extent that the conduct of systematic policy changes over time, it may do so gradually, stochastically, and perhaps endogenously, in ways

⁵In fact, it is immediate that the regression in (4) would also recover valid monetary policy causal effects if any linear combination of the two shocks were included as the “shock” measure $\varepsilon_{j,t}$. As we will discuss at length in the upcoming sections, the key is simply that $\varepsilon_{j,t}$ is an event that is stable across policy regimes.

likely unknown to the econometrician. And third, in this example, changes in policy were the only reason why the propagation of the demand and supply shocks changed across regimes. In practice, the structure of the economy itself may change, also shaping shock propagation.

The next three sections will tackle these challenges step-by-step: first, in Section 3, we consider a much more general class of infinite-horizon structural models, though still with permanent policy regime differences; second, in Section 4, we extend the model to allow for essentially unrestricted variation over time in policy design; and third, in Section 5, we add rich forms of non-policy state dependence in shock propagation. The practical upshot will be conditions on the regressors $\varepsilon_{j,t}$ and ζ_t in local projections like (4) that allow the takeaways of this section — chiefly our proposed interpretation of the δ_y 's — to survive.

3 Permanent policy regimes

We follow the same structure as in Section 2: first the model environment, then identification results for interacted local projections.

3.1 Environment

We will consider a general, though linearized, family of dynamic economies. We first describe the model class and discuss its scope and limitations, before then characterizing equilibrium outcomes across different systematic policy regimes. The model class is identical to that in McKay & Wolf (2023), so our discussion will be brief, though self-contained.

MODEL. The model economy consists of two blocks. The first block is what we refer to as the private-sector block, and it is given as

$$\mathcal{H}_x [\mathbb{E}_t (\mathbf{x}^t) - \mathbb{E}_{t-1} (\mathbf{x}^t)] + \mathcal{H}_z [\mathbb{E}_t (\mathbf{z}^t) - \mathbb{E}_{t-1} (\mathbf{z}^t)] + \mathcal{H}_e e_t = \mathbf{0}. \quad (5)$$

Here x_t is an n_x -dimensional vector of endogenous private-sector outcomes, z_t is a scalar policy instrument, and e_t is an n_e -dimensional vector of non-policy shocks. Boldface denotes time paths, i.e., for any scalar w we have $\mathbf{w}^t = (w_t, w_{t+1}, \dots, w_{t+H-1})'$, and H is the (possibly infinite) maximal horizon. Equation (5) stacks $n_x \times H$ distinct relations, giving current and expected future *revisions* of private-sector outcomes x as a function of date- t shocks e_t as well as the expected time path of the policy instrument z . We assume invertibility of the operator

$\mathcal{H}_x$, delivering a unique mapping from policy and shocks to private-sector outcomes.⁶

The second block is what we refer to as the policy block, and it gives expectation revisions of the policy instrument as a function — either explicit or implicit — of endogenous macroeconomic fluctuations:

$$\mathcal{A}_x [\mathbb{E}_t (\mathbf{x}^t) - \mathbb{E}_{t-1} (\mathbf{x}^t)] + \mathcal{A}_z [\mathbb{E}_t (\mathbf{z}^t) - \mathbb{E}_{t-1} (\mathbf{z}^t)] + \mathcal{A}_v v_t = \mathbf{0}, \quad (6)$$

where v_t is an n_v -dimensional vector of policy shocks, and we assume that the shocks $\varepsilon_t \equiv (e'_t, v'_t)'$ are mutually independent with $\varepsilon_t \stackrel{\text{iid}}{\sim} (0, I_{n_\varepsilon})$. We note that the equilibrium system (5) - (6) is written in sequence-space, expectation revision form. We can translate from impulse response form to actual realized outcomes over time in the stochastic economy by stacking impulse responses as a vector moving average representation, as in Caravello et al. (2025).

SCOPE. For our purposes, the key property of the general environment (5) - (6) is one that we already previewed and leveraged in the simple example in Section 2: independently of how exactly policy is set, the mapping from the policy instrument to private-sector outcomes is always the same — i.e., instrument sufficiency. As discussed in detail in McKay & Wolf (2023), instrument sufficiency is a property consistent with a wide class of structural business-cycle models, including in particular those with many frictions (e.g., Christiano et al., 2005), shocks (e.g., Smets & Wouters, 2007), rich microeconomic heterogeneity (e.g., Kaplan et al., 2018), and also certain kinds of departures from rationality (e.g., as in Gabaix, 2020).⁷ Our quantitative illustrations in Section 3.2 will indeed leverage the model of Smets & Wouters.

POLICY REGIMES AND EQUILIBRIUM CHARACTERIZATION. Exactly as in Section 2, we assume that the economy can be in one of two distinct policy regimes, again interpretable as either variation over time or across (otherwise identical) countries. We begin by characterizing outcomes in one of the policy regimes, the “baseline” rule with coefficients $\{\bar{\mathcal{A}}_x, \bar{\mathcal{A}}_z, \bar{\mathcal{A}}_v\}$, and will write equilibrium outcomes in that particular policy regime as the structural vector

⁶In standard stochastic dynamic economies, it is well-known that policy instrument pegs typically deliver equilibrium indeterminacy, corresponding in the notation of (5) to multiplicity of the mapping from shocks and instruments to outcomes, even when restricting attention to bounded outcome sequences. Our statement about $\mathcal{H}_x$ should thus be understood to also embed the equilibrium selection achieved by the actual in-sample policy (6), exactly as in McKay & Wolf (2023). See also Angeletos et al. (2026) for a discussion of this point.

⁷The perhaps most important class of models not nested are those featuring informational asymmetries between the policymaker and the private sector (as in Lucas, 1972), resulting in the private-sector response coefficients $\mathcal{H}_\bullet$ being shaped by the policy rule itself.

moving average representation

$$\mathbb{E}_t(\bar{\mathbf{y}}^t) - \mathbb{E}_{t-1}(\bar{\mathbf{y}}^t) = \bar{\Theta}\varepsilon_t, \quad (7)$$

where the $(H \cdot n_y) \times n_\varepsilon$ matrix $\bar{\Theta}$ here collects all shock impulse responses under the baseline rule. We also, and again in keeping with our earlier analysis in Section 2, define $\bar{\Theta}_v$ as the impulse responses of shocks to the baseline policy rule *at all horizons*, i.e., $\bar{\Theta}_v$ is $(H \cdot n_y) \times H$ -dimensional, collecting the causal effects over time of all contemporaneous and news policy shocks. In the sequel, y , x , or z subscripts on the $\bar{\Theta}$'s will indicate impulse responses for the corresponding variables. We throughout assume that $\bar{\Theta}_{z,v}$ is invertible, i.e., through policy shocks any possible path of the policy instrument can be implemented.

We now arrive at the following version of our additive equilibrium characterization.

Lemma 1. *For policy regime $\{\mathcal{A}_x, \mathcal{A}_z, \mathcal{A}_v\}$, equilibrium outcomes of the model economy (5) - (6) can be written in expectation revision form as*

$$\mathbb{E}_t[\mathbf{y}^t] - \mathbb{E}_{t-1}[\mathbf{y}^t] = \bar{\Theta}\varepsilon_t + \bar{\Theta}_v\Omega(\varepsilon_t), \quad (8)$$

where the $H \times 1$ vector $\Omega(\varepsilon_t)$ of policy shocks is chosen so that the paths in (8) satisfy the actual policy rule (6).

Lemma 1 generalizes our earlier discussion in Section 2 to the present, richer model setting: impulse responses in any policy regime can be written as the sum of “baseline” impulse responses plus policy (shock) causal effects times suitably chosen artificial policy shocks. We now as before leverage that equilibrium characterization to inform the interpretation of interacted local projections that look across policy regimes.

3.2 Interpreting interacted local projections

As before policy causal effects are the object of interest. Translating from shock space to instrument space just as in the simpler model, we can define those policy causal effects as

$$\Theta_{x,z} \equiv \bar{\Theta}_{x,v} \cdot \bar{\Theta}_{z,v}^{-1},$$

where again, by instrument sufficiency, those causal effects do not depend on the policy rule, and are instead shaped exclusively by the private-sector block (5). Regressions on a policy shock — or on an IV for such policy shocks — identify, yet again just as in the simple model,

a one-dimensional slice of those causal effects. Formally, let

$$w_t = \sum_{\ell=0}^{H-1} \omega_{w,\ell} v_{t+\ell,t} + \text{noise}$$

where $v_{t+\ell,t}$ is a date- t , horizon- ℓ policy news shock (with $\ell = 0$ corresponding to a standard contemporaneous shock), $\omega_{w,\ell}$ is the weight of the policy shock instrumental variable on the horizon- ℓ shock, and the noise term is measurement error orthogonal to current, lagged, and future disturbances. By Plagborg-Møller & Wolf (2021), essentially all time series identification approaches, whether implemented through vector autoregressions or local projections, ultimately amount to regressions on such policy shock instruments. Writing the horizon- h regression for some generic macroeconomic outcome y as

$$y_{t+h} = \text{constant} + \beta_{y,h} w_t + \text{error},$$

it follows that the β_y 's give us the causal effects of the policy shock mixture $\{\omega_{w,\ell}\}$ on the macro-economy, or equivalently, translating to instrument space, that

$$\beta_y = \Theta_{y,z} \times \beta_z.$$

We now contrast this familiar strategy with interacted regime local projections.

INTERACTED LOCAL PROJECTIONS ACROSS REGIMES. We consider an econometrician living in the dynamic economy (5) - (6), observing data from two distinct policy regimes, and running the following regression specification:

$$y_{t+h} = \alpha_{y,h} + \beta_{y,h} \varepsilon_{j,t} + \gamma_{y,h} \zeta_t + \delta_{y,h} (\varepsilon_{j,t} \times \zeta_t) + u_{t,y,h}, \quad (9)$$

where $\varepsilon_{j,t}$ is the j th primitive structural shock, and ζ_t is a dummy indicator for the policy regime, just as before. We again pay particular attention to (9) run for both the policy instrument z_t and macroeconomic outcomes x_t on the left-hand side.

Proposition 1. *Consider the dynamic stochastic economy (5) - (6). The population estimates of (9) satisfy*

$$\delta_x = \underbrace{\Theta_{x,z}}_{\text{policy causal effects}} \times \underbrace{\delta_z}_{H \times 1 \text{ vector of treatments}}.$$




Figure 1: The top left panel shows the identified $H \times 1$ “as-if” policy shock path (defined relative to a standard Taylor rule as the baseline rule), while the other three panels show the interacted local projection coefficients $\{\delta_h\}$ for real interest rates, output, and inflation. All panels show six lines, for each of the primitive non-policy shocks of the model. Responses are scaled so that the peak interest rate response equals one percent. See Appendix B.1 for details on the construction.

With Proposition 1 we have fully dealt with the first important limitation of our simple example in Section 2: the insights of that section extend without any change to a general class of linearized infinite-horizon economies with permanent policy regimes.⁸ We next provide a numerical illustration, showcasing the key applied takeaway that the implicit policy “treatment” isolated by (9) generically differs by shock $\varepsilon_{j,t}$ and is distinct — not in nature, but in precise treatment path — from any given policy shock regression estimand.

ILLUSTRATION. We consider the structural model of Smets & Wouters (2007) as our data-generating process. We assume that the monetary authority initially sets the policy rate according to a first rule, expected by all agents in the economy to be permanent. Then, and

⁸Note that Proposition 1 applies for *any* shock, including in particular shocks to the policy instrument themselves. In our subsequent illustrations and applications, however, we will not further consider such shocks, precisely because we want to consider approaches that complement the usual shock regressions, both in informational requirements and in estimand.

unexpectedly, it changes its conduct to a different rule, yet again expected to be permanent; nothing else changes across the two sub-samples.⁹ We now consider an econometrician that observes long samples pre- and post-policy change, and runs the local projection (9).

Results are reported in Figure 1, which plots the population regression coefficients $\{\delta_h\}$ for three different outcomes of interest, as well as the synthetic “as-if” monetary policy shock paths $\Omega(\varepsilon_{j,t})$ that map the first into the second rule, consistent with Lemma 1. We begin with the top two panels, which show the policy treatment in shock space (left panel) and in instrument space (right panel). There are two main takeaways. First, each non-policy shock $\varepsilon_{j,t}$ identifies a different policy treatment. Intuitively, each shock has different stochastic properties and equilibrium effects, and so looking across the policy regimes will generically identify different changes in policy. Second, the identified treatments are generally persistent, both in shock space (left panel) and in instrument space (right panel), sometimes even changing sign over time. The intuition here is that all of the non-policy shocks have relatively persistent effects, and that policy regimes are expected to be persistent. Putting the two pieces together, it follows that the difference in policy instrument paths, and so the policy treatment, is also persistent — an observation that we will revisit in our empirical applications in Section 6. Finally, the bottom panels reveal that, as expected, the identified contractionary monetary policy treatments lower both output and inflation.

For further practical interpretation, it is useful to note that an econometrician living in this laboratory model economy, but without knowledge of the model’s structural equations, could recover, through her interacted local projections, the impulse responses shown in the top right and bottom panels of Figure 1; she would not, however, know the corresponding “as-if” shock paths in the top left panel. We repeat that this is no problem: her regressions recover policy causal effects in instrument space, and that suffices both for model estimation through impulse response matching as well as policy counterfactual evaluation.

EXTENSION TO MULTIPLE INSTRUMENTS. We conclude our analysis by asking what our interacted local projections identify when the design of *multiple distinct* policy instruments changes over time. Specifically, we consider an extension of the model environment of Section 3.1 in which the policy instrument z_t is now n_z -dimensional; modulo that generalization of the dimensionality, the system equations (5) - (6) remain unchanged. We then obtain the following generalization of Proposition 1.

⁹Equivalently, as discussed before, we can also interpret this experiment cross-sectionally, as data observed from two distinct economies that differ in policy, but are otherwise identical.

Corollary 1. *Consider the dynamic stochastic economy (5) - (6), with the policy instrument z_t of size n_z . The population estimands of (9) satisfy*

$$\boldsymbol{\delta}_x = \underbrace{\Theta_{x,z}}_{\text{policy causal effects}} \times \underbrace{\boldsymbol{\delta}_z}_{(n_z \cdot H) \times 1 \text{ vector of treatments}}.$$

The practical interpretation of this model generalization is straightforward. If the design of multiple policy instruments changes at the same time (e.g., interest rate policy and quantitative easing, or monetary and fiscal policy), then a researcher running (9) will identify a combination of multiple policy treatments at the same time, with those treatments estimable by running (9) with the various policy instruments on the left-hand side. In particular, if the design of an instrument z_k changes over time, but the interaction regressor $\varepsilon_{j,t} \times \zeta_t$ does not revise the setting of that instrument, then $\boldsymbol{\delta}_{z_k} = \mathbf{0}$, and so the estimand does not reflect that particular policy instrument. We will revisit this insight in Section 5, when we allow for other, non-policy-related forms of variation in shock propagation, and in Section 6, where we study applications in which mixtures of policy causal effects are recovered.

3.3 Shocks, macroeconomic news, and second moments

The interacted local projection (9) featured two regressors: a primitive structural shock $\varepsilon_{j,t}$, and the (binary) regime indicator ζ_t . In this section we ask whether the implied formidable informational requirement on $\varepsilon_{j,t}$ — i.e., observing a true structural shock — can be relaxed. Throughout, the core intuition is that, in (9), the fundamental role of $\varepsilon_{j,t}$ is to provide a fixed conditioning event, itself not shaped by policy, but whose propagation varies with ζ_t *only* because of policy. In our current setting the second requirement is trivial, but the first one is not, and the remainder of this section discusses what kind of variables can play this required role. Assumptions on the policy regime variable ζ_t will instead take center stage in the two extensions of the model environment in Sections 4 and 5.

BEYOND PRIMITIVE SHOCKS. It is straightforward to see that running (9) with a *fixed* linear combination of primitive shocks, $b'\varepsilon_t$, in lieu of any single shock $\varepsilon_{j,t}$, will also deliver valid policy causal effects. This is a meaningful qualification of the earlier informational requirements: any date- t macroeconomic news that itself is not shaped by, and so in particular does not change with, the policy regime variable ζ_t suffices for our purposes. We will illustrate the bite of this observation in our later applications in Section 6, where we will

argue that our conclusions apply even if the included regressors are not narrow measures of primitive structural shocks in the textbook sense.

WOLD INNOVATIONS. A natural alternative strategy is to use reduced-form statistical (i.e., “Wold”) innovations in (9) in lieu of the primitive shock $\varepsilon_{j,t}$. This, however, will generally *not* work, simply because those reduced-form innovations themselves tend to be shaped by policy. To make the argument as transparent as possible, assume that the Wold innovations in the policy regimes $k = 1, 2$, denoted $u_{k,t}$, are constructed using an information set large enough to ensure invertibility (see Fernández-Villaverde et al., 2007). The j th Wold innovation in regime k then as usual satisfies

$$u_{j,k,t} = b'_{j,k}\varepsilon_t$$

where $b_{j,k}$ is an n_ε -dimensional vector. The “diff-in-diff” logic underlying the proof of Proposition 1 now breaks down: we are comparing across different policy regimes how *different* mixes of structural shocks, $b'_{j,1}\varepsilon_t$ vs. $b'_{j,2}\varepsilon_t$, propagate, and so the computed difference reflects both policy causal effects as well as the heterogeneous conditioning event.

REDUCED-FORM SECOND MOMENTS. Rather than asking how the change in policy regime shaped the propagation of a given shock or innovation, we may alternatively ask how the regime change shaped unconditional second moments (i.e., the reduced-form autocovariance function), and then whether that change is informative about policy causal effects. An early version of such an analysis is Mussa (1986), who observes that exchange rate regime changes affect the volatility of real outcomes, thus allowing him to reject policy neutrality. Relatedly, the literature on identification through heteroskedasticity (e.g., Rigobon, 2003; Lewis, 2025) has proved that changes in the volatility of policy shocks, together with an otherwise stable economic structure, can identify policy causal effects. And finally, Caravello et al. (2025) establish that policy causal effects together with realized second moments of macroeconomic time series are sufficient to predict counterfactual second moments — thus making it natural to ask whether that mapping is invertible.

Unfortunately, it turns out that policy-induced changes in reduced-form second moments are *not* a usable alternative to the requirement of observing a primitive shock (or a linear combination thereof). We only sketch the high-level intuition here, with a detailed treatment relegated to Appendix A.2.¹⁰ In the literature on identification through heteroskedasticity, a

¹⁰The appendix provides both a formal characterization of the population identified set of policy causal effects, as well as some numerical illustrations.

change in policy shock volatility induces a rank-1 change in the variance-covariance matrix of reduced-form errors, allowing point identification of policy causal effects. Changes in the policy *rule*, in contrast, can be shown to lead to a higher-dimensional change in the reduced-form autocovariance function, leading to *set* identification of policy causal effects. That set may exclude zero, echoing the Mussa (1986) test of policy non-neutrality, but it is generally far too wide to allow informative conclusions. We thus conclude that the headline takeaway of this section — the requirement of a fixed conditioning event $\varepsilon_{j,t}$, to be included in (9), along the lines discussed in the preceding paragraphs — simply cannot be relaxed through informationally less demanding reduced-form arguments.

4 Stochastic policy switches

We now address the second limitation of the simple example in Section 2: we go beyond permanent policy regimes, considering instead stochastically, perhaps endogenously, switching policy conduct. Perhaps surprisingly, we show that this extension does not cause particular difficulties, with our conclusions extending under rather mild conditions on the policy regime variable ζ_t — conditions that in particular allow for endogeneity of that regime measure.

4.1 Extended environment

The non-policy block is exactly as in Section 3.1. The policy block, however, is now different: it gives expectation revisions in the scalar policy instrument z_t as a function of exogenous disturbances to the system:

$$\mathbb{E}_t(z^t) - \mathbb{E}_{t-1}(z^t) = \mathcal{A}(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots), \quad (10)$$

where s_t is a scalar policy regime indicator, whose role will be discussed momentarily. Note that (10) is a solved-out, reduced-form mapping, and not an explicit, primitive policy rule; we will discuss later why we find it convenient to specify policy in this reduced-form way. We throughout assume that the policy regime indicator s_t is a random variable independent of the shocks ε_t , but otherwise unrestricted.

As before, the equilibrium system (5) - (10) is written in sequence-space, expectation revision form. We can then obtain realized equilibrium sequences by recursively summing up solutions to this pair of equations for each date t , exactly as in Wolf (2025, Appendix D). Differently from the analysis in Section 3, this recursive summing does not deliver a simple vector moving average representation like (7).

DISCUSSION OF THE MODEL SPACE. For a simple example of a model environment that can be written in the form (5) - (10), begin with the same textbook New Keynesian private-sector block as in Section 2, but now consider an interest rate rule with exogenous time-variation (following Davig & Leeper, 2007)

$$i_t = \phi_\pi(s_t)\pi_t + v_t, \quad (\text{TR}')$$

where s_t is a regime indicator that follows an exogenous Markov process independent of the other shocks, $\phi_\pi(\bullet)$ gives regime-specific policy rule coefficients, and v_t is the monetary policy shock. Assuming equilibrium existence and uniqueness, we can obtain a mapping from the primitive shocks ε_t and s_t to i_t — an example of the general reduced-form policy mapping (10) in our general model. Alternatively, an example of endogenous policy feedback would be the adjusted Taylor rule (following Davig & Leeper, 2006)

$$i_t = \phi_\pi(\pi_{t-1})\pi_t + v_t, \quad (\text{TR}'')$$

where now the policy rule coefficient depends on past inflation. In equilibrium, policy is again given as a reduced-form mapping from shocks to the policy instrument, and given the policy instrument we can recover all other outcomes from the linear private-sector block.

Going beyond these specific examples, what kind of assumptions on economic primitives are consistent with the assumed structure (5) - (10)? Note that we have materially, in fact largely arbitrarily, generalized the design of policy, but through (5) maintain a simple linear mapping from the policy instrument to private-sector outcomes. Appealingly, such private-sector structures obtain by standard first-order perturbations in aggregate shock volatilities in settings with stochastic regime switches, either exogenous or endogenous, around a fixed deterministic steady state (Farmer et al., 2009; Barthélemy & Marx, 2019). The principal appeal of writing policy not as an explicit rule but simply as a reduced-form equilibrium mapping is that it allows us to seamlessly nest all these distinct cases, plus other possible nonlinearities in policy-setting (e.g., a zero lower bound).¹¹ We have thus extended the model to address the second limitation of our earlier analysis: with a linear time-invariant non-policy block plus policy specified as a general reduced-form mapping from exogenous inputs to policy, our environment nests canonical structural analyses of stochastically switching policy regimes (as in Davig & Leeper, 2007; Bianchi, 2013; Bianchi & Melosi, 2017).

¹¹For example, this extended model setting nests the structural analysis of the Great Recession in Christiano et al. (2015), which features regime switching and a binding lower bound for the policy rate.

EQUILIBRIUM CHARACTERIZATION. We next show that, precisely because the private-sector block (5) is unchanged, equilibrium outcomes continue to admit an additive decomposition into “baseline” impulse responses and policy causal effects. The definitions of both those baseline impulse responses, and of policy causal effects, are entirely unchanged relative to Section 3, as those two objects depend only on the (arbitrary) “baseline” policy rule and on the private-sector block, not on actual policy. The sole purpose of the baseline rule will be to serve as a reference point (i.e., basis) for our decomposition.

Lemma 2. *Equilibrium outcomes of the model economy (5) - (10) can be written in expectation revision form as*

$$\mathbb{E}_t [\mathbf{y}^t] - \mathbb{E}_{t-1} [\mathbf{y}^t] = \bar{\Theta} \varepsilon_t + \bar{\Theta}_v \Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots) \quad (11)$$

where the $H \times 1$ vector $\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots)$ solves

$$\begin{aligned} \mathbb{E}_t [\mathbf{z}^t] - \mathbb{E}_{t-1} [\mathbf{z}^t] &= \bar{\Theta}_z \varepsilon_t + \bar{\Theta}_{z,v} \Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots) \\ &= \mathcal{A}(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots). \end{aligned} \quad (12)$$

We see that, *conditional* on equilibrium policy, moving to a model with rich, stochastic policy regime switches has not materially complicated the characterization of equilibrium outcomes. The intuition is exactly as before: If policy was counterfactually set according to the “baseline” rule, then equilibrium expectation revisions would simply be given as $\bar{\Theta} \varepsilon_t$, by definition of $\bar{\Theta}$. Actual equilibrium outcomes are however generally different, simply because the policy instrument z_t is not actually set in accordance with the baseline rule, but instead satisfies (10). The H -dimensional vector $\Omega(\bullet)$ in (12) is an artificial “shock” to the baseline rule *defined* to ensure that the policy path in (11) satisfies (10), therefore making that part of the lemma tautological. By instrument sufficiency, once z_t is set in that way, the outcomes for x_t stated in (11) must also agree with true equilibrium outcomes.

4.2 Interpreting interacted local projections, again

We wish to learn about policy causal effects, which in instrument space are defined just as before as

$$\Theta_{x,z} \equiv \bar{\Theta}_{x,v} \cdot \bar{\Theta}_{z,v}^{-1}.$$

We also consider the exact same interacted local projection specification as before, equation (9), though now ζ_t is not a binary regime indicator — it is just some date- t random variable,

and we will pay particular attention to its required properties. We then obtain the following generalization of Proposition 1.

Proposition 2. *Consider the dynamic stochastic economy (5) - (10) and suppose that the random variable ζ_t has mean zero and is given as*

$$\zeta_t = \zeta(\varepsilon_{-j,t}, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots). \quad (13)$$

Then the population estimands of (9) satisfy

$$\delta_x = \underbrace{\Theta_{x,z}}_{\text{policy causal effects}} \times \underbrace{\delta_z}_{H \times 1 \text{ vector of treatments}}.$$

Proposition 2 reveals that, under the restriction (13), all of our previous results — the estimand interpretation, the relation to policy shocks, treatment heterogeneity — extend without any change to the present much richer setting, with much more general time-varying policy. In the remainder of this section we first provide intuition for this observation by presenting a brief proof sketch, before then discussing further the economic content and the implications of the assumption (13).

PROOF SKETCH. By Frisch-Waugh-Lovell, and using the assumed independence of ζ_t and $\varepsilon_{j,t}$, the regression estimand for outcome y_{t+h} satisfies

$$\delta_{y,h} \propto \text{Cov}(y_{t+h}, \varepsilon_{j,t} \zeta_t).$$

Since $\varepsilon_{j,t}$ is a date- t shock, and since both $\varepsilon_{j,t}$ and ζ_t are measurable with respect to date- t information, we can re-write that covariance term using date- t expectation revisions, i.e.,

$$\delta_{y,h} \propto \text{Cov}(\mathbb{E}_t[y_{t+h}] - \mathbb{E}_{t-1}[y_{t+h}], \varepsilon_{j,t} \zeta_t).$$

This expectation revision, however, is the object we characterized in Lemma 2, given as

$$\mathbb{E}_t[y_{t+h}] - \mathbb{E}_{t-1}[y_{t+h}] = \bar{\Theta}_{y,h} \varepsilon_t + \bar{\Theta}_{y,v,h} \Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots).$$

The assumed independence of $\varepsilon_{j,t}$ from $(\varepsilon_{-j,t}, \zeta_t)$ then implies that we are left with

$$\delta_{y,h} \propto \bar{\Theta}_{y,v,h} \text{Cov}(\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots), \varepsilon_{j,t} \zeta_t). \quad (14)$$

(14) is the key relation, implying that the interaction coefficients δ_y correspond to the dynamic causal effects of a particular “as-if” $H \times 1$ policy shock vector, given as

$$\text{Cov}(\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots), \varepsilon_{j,t} \zeta_t). \quad (15)$$

The actual statement of Proposition 2 then simply moves from policy shock space to policy instrument space: running (9) for the instrument z_t gives the policy treatment δ_z , and then running it for an outcome x_t gives the associated causal effects. Zooming out, the high-level intuition is that Lemma 2 continues to split equilibrium outcomes into “baseline” expectation revisions and policy causal effects, and that the key regressor $\varepsilon_{j,t} \times \zeta_t$ only co-varies with the policy causal effect term. Our earlier discussion in Section 4.1 reveals that the underlying decomposition of equilibrium outcomes is a quite general feature of structural models, even with rich time variation in policy design; and our assumptions on $\varepsilon_{j,t}$ and ζ_t then do the rest, with the next paragraph elaborating further on the role played by ζ_t .

PROPERTIES OF THE POLICY MEASURE ζ_t . Two properties of ζ_t are crucial in the preceding argument. First, the interaction of $\varepsilon_{j,t}$ with ζ_t needs to *predict* the policy instrument: if it does so, then $\delta_z \neq 0$, and so the regression indeed isolates a non-zero policy treatment. This is simply a standard relevance condition.¹² Second, ζ_t needs to shape the propagation of $\varepsilon_{j,t}$ *only* through policy — the simple “diff-in-diff” intuition that dynamics following the fixed event $\varepsilon_{j,t}$ differ only because of the channel of interest, policy transmission. Condition (13) has to do with this second requirement.

In the data-generating process that we considered in this section, changes over time in policy design are the only source of variation in how a given structural shock $\varepsilon_{j,t}$ propagates. Ruling out any *direct* co-movement between ζ_t and $\varepsilon_{j,t}$, which is what (13) delivers, is thus all that is needed for correct identification of policy causal effects. Importantly, this requirement can be entirely consistent with ζ_t being an *endogenous* correlate of systematic policy. For example, if a past history of inflationary shocks tends to increase the policymaker’s concern for inflation (e.g., along the lines of the simple example (TR”) above), then any correlates of such disturbances will also predict future policy, and yet they will by assumption not further shape the propagation of $\varepsilon_{j,t}$ through any non-policy channels. This discussion however also suggests that, once we allow for other, non-policy forms of time or state dependence in the

¹²Of course the statement of Proposition 2 would strictly speaking still be true even if $\delta_z = 0$, but vacuously so, as then also $\delta_x = \mathbf{0}$.

propagation of $\varepsilon_{j,t}$, the conditions on ζ_t will become more stringent; we will thus next turn to such additional sources of variation in shock propagation, either cross-sectional or over time, in the third, and final, extension of the model environment.

5 Adding other forms of state dependence

As before, we first present the extended environment and equilibrium characterization, before turning to identification results. We close with takeaways for applied work, summarizing the key practical implications of the results of this and the preceding sections.

5.1 Further extended environment

To deal with the third shortcoming of the simple example in Section 2, we replace the general but time-invariant private-sector block (5) with the following state-dependent variant:

$$\mathcal{H}_x(m_t) [\mathbb{E}_t(\mathbf{x}^t) - \mathbb{E}_{t-1}(\mathbf{x}^t)] + \mathcal{H}_z(m_t) [\mathbb{E}_t(\mathbf{z}^t) - \mathbb{E}_{t-1}(\mathbf{z}^t)] + \mathcal{H}_e(m_t)e_t = \mathbf{0}, \quad (16)$$

where m_t is a date- t random variable that is independent of the date- t shocks ε_t but otherwise unrestricted. We will discuss momentarily what assumptions on economic primitives would deliver such a representation of the private-sector block. Next, policy is as before some general reduced-form mapping from histories of shocks into the policy instrument, just now augmented to also allow dependence on m_t :

$$\mathbb{E}_t(\mathbf{z}^t) - \mathbb{E}_{t-1}(\mathbf{z}^t) = \mathcal{A}(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots; m_t). \quad (17)$$

We first discuss the scope of this model class before turning to equilibrium characterization.

DISCUSSION OF THE MODEL SPACE. The extended model (16) - (17) allows for non-policy sources of variation in shock propagation in both cross-sectional as well as time-series designs. For the cross section, it is straightforward to see that the above model structure nests the case of cross-sectional entities with multiple distinct policy regimes (as in Section 3), but now also with the private-sector coefficients $\{\mathcal{H}_x, \mathcal{H}_z, \mathcal{H}_e\}$ changing in arbitrary ways with the policy regime itself. For example, we may now observe data from multiple countries that differ not only in the design of their systematic policy, but also in, say, the severity of nominal rigidities, the average marginal propensity to consume, or the cross-sectional distribution of household

wealth. We provide a detailed discussion of such explicitly cross-sectional environments, and their translation to the system (16) - (17), in Appendix A.3.

For the time series, a model structure like (16) would emerge through the usual linearization in the aggregate shocks ε_t , but now not around a fixed deterministic steady state, but instead around a time-varying non-linear transition path indexed by m_t (as in Auclert et al., 2021).¹³ Just as in the cross-sectional case, this approach allows a time-varying private-sector structure to shape the dynamic propagation of any newly arriving shocks and surprise policy changes, with (16) now giving a time-dependent — but otherwise still linear — mapping of shocks and policy to private-sector outcomes.

EQUILIBRIUM CHARACTERIZATION. The equilibrium in our final model extension yet again delivers an additive decomposition of equilibrium outcomes, just now with shock and policy impulse responses that depend on the private-sector state variable m_t . As before, we assume temporarily that policy counterfactually followed a simple time-invariant “baseline” rule with coefficients $\{\bar{\mathcal{A}}_x, \bar{\mathcal{A}}_z, \bar{\mathcal{A}}_v\}$; combining that rule with the date- t private-sector block (16), we would obtain the state-dependent representation of expectation revisions given as

$$\mathbb{E}_t(\bar{\mathbf{y}}^t) - \mathbb{E}_{t-1}(\bar{\mathbf{y}}^t) = \bar{\Theta}(m_t)\varepsilon_t. \quad (18)$$

Combining this representation with our assumptions on policy in (17) we finally arrive at the following additive decomposition, as expected generalizing our earlier expressions to feature additional state dependence:

Lemma 3. *Equilibrium outcomes of the model economy (16) - (17) can be written as*

$$\mathbb{E}_t[\mathbf{y}^t] - \mathbb{E}_{t-1}[\mathbf{y}^t] = \bar{\Theta}(m_t)\varepsilon_t + \bar{\Theta}_v(m_t)\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots; m_t) \quad (19)$$

where the $H \times 1$ vector $\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots; m_t)$ solves

$$\mathbb{E}_t[\mathbf{z}^t] - \mathbb{E}_{t-1}[\mathbf{z}^t] = \bar{\Theta}_z(m_t)\varepsilon_t + \bar{\Theta}_{z,v}(m_t)\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots; m_t)$$

¹³The private-sector representation (16) omits any changes in the propagation of *past* shocks or policies to private-sector outcomes because of time variation in m_t . If m_t itself changes in response to lagged shocks, such interaction terms would be second-order and thus drop out of the linearization. In addition to linearization around a perfect-foresight transition path, (16) would also obtain from exogenous or endogenous Markov switching in private-sector coefficients, with agents treating drifting parameters as if they would remain constant at the current level going forward in time (as assumed in much of the macro learning literature, see Evans & Honkapohja, 2001, and also Cogley & Sbordone (2008) for an application).

$$= \mathcal{A}(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots; m_t). \quad (20)$$

Just as before, *conditional* on equilibrium policy, we have a simple linear mapping from shocks and policy into endogenous outcomes. Differently from before, however, that mapping is now time-varying, or equivalently, for cross-sectional specifications, differs by unit.

5.2 Interpreting interacted local projections, yet again

Policy causal effects conditional on date- t private-sector state m_t are now defined as

$$\Theta_{x,z}(m_t) \equiv \bar{\Theta}_{x,v}(m_t) \cdot \bar{\Theta}_{z,v}(m_t)^{-1}.$$

Note that these policy causal effects again do not depend at all on policy conduct, though they do depend on m_t and so on the time-varying private-sector structure. We consider the same interacted local projection (9) as before, and now obtain the following result.

Proposition 3. *Consider the dynamic stochastic economy (16) - (17) and suppose that the random variable ζ_t has mean zero and is given as*

$$\zeta_t = \zeta(\varepsilon_{-j,t}, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots), \quad (21)$$

and furthermore independent of the non-policy state m_t . Then the population estimands of (9) satisfy

$$\delta_x = \mathbb{E} \left[\underbrace{\Theta_{x,z}(m_t)}_{\text{policy causal effects}} \times \left(\underbrace{\delta_z}_{\text{av'g treatment}} \odot \underbrace{\mathbf{w}_z(m_t)}_{\text{"weights"}} \right) \right], \quad (22)$$

where the H -dimensional "weights", $\mathbf{w}_z(m_t)$, satisfy $\mathbb{E}(\mathbf{w}_z(m_t)) = \mathbf{1}$, but are not necessarily always positive. Positivity of the weights is ensured if the signs of all entries of the vector

$$\text{Cov}(\mathbf{z}^t, \varepsilon_{j,t} \zeta_t \mid m_t) \quad (23)$$

*are independent of m_t .*¹⁴

Proposition 3 has the same structure as our earlier identification results, but now with two wrinkles: the requirements on the policy measure ζ_t have become more stringent, and we require a further monotonicity condition, (23). We discuss both in turn.

¹⁴Strictly speaking, the representation in (22) furthermore requires that, for all h , $\delta_{z,h} \neq 0$.

TIGHTER RESTRICTIONS ON ζ_t . The interacted local projection (9) asks how the impulse response to $\varepsilon_{j,t}$ covaries with the policy measure ζ_t . Given that now impulse responses are allowed to vary in largely unrestricted fashion with the macroeconomic state m_t , we require ζ_t to not co-vary with m_t . Intuitively, any policy regime changes, and so in particular also their correlate ζ_t , can still be endogenous, but now that endogeneity must not *incrementally* shape the propagation of $\varepsilon_{j,t}$ through any non-policy channels. For example, monetary policy can become more concerned about inflation in response to past inflationary shocks, but those past inflationary shocks must not themselves shape the propagation of future shocks through non-policy channels. If instead such past shocks at the same time steepened the NKPC, say, then the estimated δ 's would also reflect the causal effects of shocks to that model equation, i.e., of cost-push shocks; see Appendix A.4 for further discussion.

To summarize, the required conditions on ζ have become meaningfully more stringent, here in the form of unconditional independence of m_t and ζ_t . As usual, similar conclusions do apply under a weaker condition of mean independence conditional on controls in an extended version of (9). We provide a discussion of that case in Appendix A.5.

MONOTONICITY. The time-varying state m_t adds a further complication, familiar from the microeconometric literature on heterogeneous treatment effects (e.g., see Imbens & Angrist, 1994): that both policy treatment effects, $\Theta_{x,z}(m_t)$, as well as the effective policy treatment itself, can vary with m_t . To see the logic, note that, conditional on a non-policy regime m_t , we could write the regression estimand as

$$\delta_x(m_t) = \sum_{h=0}^{H-1} \underbrace{\Theta_{x,z_h}(m_t)}_{\text{treatment effect}} \times \underbrace{\delta_{z,h}(m_t)}_{\text{treatment}},$$

where $\Theta_{x,z_h}(m_t)$ is column $h+1$ of $\Theta_{x,z}(m_t)$, i.e., the causal effect of z_{t+h} on $\mathbf{x}^t$ conditional on regime m_t . In words, the estimand is just the sum of H treatments of possibly heterogeneous sign and size, for the policy instrument at each horizon. Averaging across regimes to get to the actual estimand, we find

$$\delta_x = \mathbb{E} \left[\sum_{h=0}^{H-1} \Theta_{x,z_h}(m_t) \times \delta_{z,h} \times \underbrace{\frac{\delta_{z,h}(m_t)}{\delta_{z,h}}}_{\equiv w_{z,h}(m_t)} \right].$$

For each policy treatment horizon h , we get a weighted average of treatment effects across regimes m_t , but with weights that are not necessarily positive, because the effective treat-

ments themselves may also be regime-dependent. Assuming that the identified policy treatments $\delta_{z,h}(m_t)$ are always the same for each h across all regimes obviously circumvents that problem; more plausibly, a monotonicity assumption — requiring that, across all regimes, the date- h treatment $\delta_{z,h}(m_t)$ has the same sign — suffices to ensure that the weights are positive, echoing results from microeconomic program evaluation.¹⁵

RELATION TO THE BROADER LITERATURE. Zooming out, our results bear a very close resemblance to those from “shift-share” panel regression specifications. In dynamic, structural general equilibrium settings, the key identifying assumption for panel shift-share regressions is that the “shifter” reflects some (combination of) time- t macroeconomic shocks, and that the “exposure shares” capture the differential effect of the shock across micro units *only* via the specific channel of interest (see Manuel, 2026, for a general treatment). In our case, $\varepsilon_{j,t}$ plays the role of the “shifter” — which, as discussed, can be replaced with any combination of time- t shocks — and the regime indicator ζ_t is then akin to the “exposure shares.” And just like in the panel setting, identification fails when the regime indicator ζ_t is correlated with other sources of state dependence in the dynamic propagation of $\varepsilon_{j,t}$.

5.3 Summary and applied takeaways

The upshot of the preceding sections is a simple recipe for running and interpreting interacted local projections with cross-sectionally or time-varying measures of systematic policy. For the regressors, it is important to ensure that the “shock” variable $\varepsilon_{j,t}$ is not directly shaped by the policy measure ζ_t (e.g., because it is a primitive structural shock), and that ζ_t affects the propagation of $\varepsilon_{j,t}$ exclusively *through* policy. Once that is the case, the estimand of δ_y is, under relatively general assumptions on the data-generating process, a simple weighted average of policy causal effects — i.e., the interacted local projection is “as if” the researcher is regressing on a particular mixture of contemporaneous and news policy shocks. Results can be reported exactly as familiar from the popular shock literature, with the implied policy treatment straightforwardly read off the impulse response of the policy instrument, and the associated causal effects in the other responses.

We will illustrate these conceptual takeaways — and the extent to which such interacted

¹⁵While not explicitly discussed there, we note that variants of the same assumptions are also necessary for the policy shock regression interpretation in Kolesár & Plagborg-Møller (2025): that paper shows that regressions on policy shocks give positive weighted averages of *policy shock* causal effects, but moving to policy instrument space requires additional monotonicity assumptions, exactly as discussed here.

local projections are both an alternative to and complement for standard policy shock regressions — through several applications in Section 6. Before doing so we close here with some brief remarks comparing our identification results with those obtained in more general nonparametric model settings, following recent advances in the literature. Readers focused on the practical applications may skip ahead to Section 6.

5.4 Aside: nonparametric models

For all of our preceding results we imposed non-trivial restrictions on the economic environment; the benefit of doing so was the ability to tightly connect policy shock and interacted local projection estimands, thus setting our analysis apart from other recent contributions on interacted local projections. An alternative path, more general but (as we will see) without these interpretative benefits, is to consider a largely unrestricted, nonparametric, non-linear setting, similar in spirit to the analyses in Rambachan & Shephard (2025) and Kolesár & Plagborg-Møller (2025). The discussion here will be relatively brief, as the below results are minor extensions of that previous work; details are provided in Appendix A.6. Our focus here will instead be on the interpretation of those results, and in particular how they relate to our main identification results.

SETTING. We materially generalize even the extended environment of Section 5.1 and now suppose that we can write any date- $t + h$ macroeconomic outcome x_{t+h} as

$$x_{t+h} = \psi_h(\varepsilon_{j,t}, z(\varepsilon_{j,t}, s_t, u_{t+h}), u_{t+h})$$

where our assumptions on the shock $\varepsilon_{j,t}$ and policy regime s_t are exactly as before, and u_{t+h} is a date- $t + h$ random variable independent of both $\varepsilon_{j,t}$ and s_t , capturing all other sources of stochasticity in the economy. Integrating out u_{t+h} , we define the average structural function for x_{t+h} conditional on $\varepsilon_{j,t}$ and s_t ,

$$\Psi_h(\varepsilon_{j,t}, s_t) = \mathbb{E}_u [\psi_h(\varepsilon_{j,t}, z(\varepsilon_{j,t}, s_t, u_{t+h}), u_{t+h})].$$

In words, this expression gives the average, expected value for x_{t+h} conditional on the shock $\varepsilon_{j,t}$ and policy state s_t , integrating out all other sources of randomness.

INTERACTED LOCAL PROJECTIONS. Now consider running the interacted local projection (9) with $\zeta_t = s_t$, i.e., we assume that the policy measure is orthogonal to both $\varepsilon_{j,t}$ and

to all other sources of stochasticity in the economy. Under regularity conditions stated in Appendix A.6, the interacted local projection estimand δ_x now satisfies

$$\delta_{x,h} = \mathbb{E}_{\varepsilon_j,s} [w(s_t, \varepsilon_{j,t}) \partial_{\varepsilon_j,s} \Psi_h(\varepsilon_{j,t}, s_t)]$$

where $w(s_t, \varepsilon_{j,t}) \geq 0$ and $\mathbb{E}_{\varepsilon_j,s} [w(s_t, \varepsilon_{j,t})] = 1$. In words, the interacted local projection now isolates a weighted average of a *cross partial* of x_{t+h} in the shock $\varepsilon_{j,t}$ and the policy measure $\zeta_t = s_t$. The cross-derivative term furthermore satisfies

$$\partial_{\varepsilon_j,s} \Psi_h(\varepsilon_{j,t}, s_t) = \mathbb{E}_u \left[\psi'_{h,\varepsilon_j z} z_s + z'_{\varepsilon_j} \psi_{h,zz} z_s + \psi'_{h,z} z_{\varepsilon_j s} \right], \quad (24)$$

with the subscripts here denoting partial derivatives. Looking back, the model of Section 5.1 corresponds to the special case of the second derivatives of $\psi_h(\bullet)$ being zero, in which case the first two terms in (24) vanish and so the overall cross partial collapses to a simple derivative with respect to policy, $\psi'_{h,z}$, times the induced movement in the policy instrument, $z_{\varepsilon_j,s}$. If the policy causal effects are furthermore invariant to the other shocks and states u , then we return all the way to our simpler identification result of Section 4.2, with the estimand equal to a weighted average of policy causal effects $\psi_{h,z}$.

To summarize, under assumptions that echo those in our main analysis, interacted local projections of the form (9), perhaps unsurprisingly, identify weighted averages of shock and policy cross-partials (with positive weights), and so admit a valid causal interpretation, very similarly to the results of Kolesár & Plagborg-Møller (2025). The additional, theory-guided structure that we imposed in our main analysis then collapses these cross partials into simple derivatives, and so allows us to connect interacted systematic policy local projections with the voluminous literature on policy shock regressions — our main contribution.

6 Applications

We illustrate our findings by revisiting and extending applications in prior work, both time series and cross-sectional, that have run interacted local projections of the sort we studied here. The applications are selected to showcase our headline practical takeaways on policy treatment heterogeneity: in the first one, that interacted local projections tend to identify very different monetary policy treatments compared to the usual regressions on monetary policy shocks; and in the second, that changing the shock variable in the local projection can materially change the identified policy treatment.

6.1 Variation in FOMC composition

Our first set of applications is based on Hack et al. (2026), who study how the propagation of fiscal policy shocks varies with the composition of the Federal Open Market Committee (FOMC) between hawks and doves. We here extend their analysis to additional shocks and then, importantly, interpret the results through the lens of our novel identification results. The headline takeaway, consistent with our earlier theoretical explorations in Section 3.2, will be that interacted local projections with their policy measure identify a very different, more persistent, slice of policy causal effects than typical regressions on monetary shocks. We at the end briefly show that this takeaway is not tied to the particular Hack et al. measure of systematic monetary policy, but, e.g., also obtains in Miyamoto et al. (2026).

SHOCKS AND POLICY MEASURE. Hack et al. (2026) construct a measure of the hawk-dove balance on the FOMC, and then suggest instrumenting it through variation in that measure induced by the mechanics of the FOMC rotation. We use the resulting series, labeled $Hawk_t^{IV}$ in their paper, as our measure of the monetary stance ζ_t . As it reflects mechanical rotation in the FOMC, this measure is likely to be independent of other (non-policy) forms of state dependence, and thus chiefly affects shock propagation through its effects on monetary policy design, as required.¹⁶ The remaining question is relevance; consistent with Hack et al. (2026), we will indeed find that, when interacted with non-policy disturbances, this policy measure tends to predict future revisions in monetary policy conduct.

For the non-policy shock $\varepsilon_{j,t}$, we consider several options. Consistent with our theoretical discussion, we want shock series available over a relatively long sample, allowing precision when interacted with rule changes. To this end we leverage examples of two kinds of popular non-monetary shocks: oil as well as fiscal policy. For oil shocks, we use the recently updated version of the shock measure of Hamilton (2003); we also experimented with the more recent series of Känzig (2021), but its shorter sample leads to significantly noisier estimates. And for fiscal shocks, as Hack et al. themselves already conduct an analysis based on Ramey (2011), we instead consider the shock series of Ben Zeev & Pappa (2015).

As we discussed in Section 3.3, the key requirement for identification is that the included shocks $\varepsilon_{j,t}$ reflect genuine macroeconomic news that is independent of the FOMC rotation measure ζ_t . Importantly, this is weaker than the traditional requirement of those measures

¹⁶See Hack et al. for further discussion of the construction of this measure, and of its lack of correlation with other plausible determinants of shock propagation (e.g., the state of the economy).

capturing a single, well-identified structural shock. For example, the oil shock measure of Hamilton (2003) captures large upswings in oil prices and thus plausibly reflects shocks to both oil demand and supply. This is however no particular concern for our purposes: either oil demand or supply shocks are valid shock measures in the interacted local projection, and thus so are their combinations. We instead simply require that these moves in oil prices are a surprise to economic agents, and that they do not occur as a result of rotations in FOMC member voting rights. Along the same lines, it is no concern *per se* for our identification strategy if changes in military spending captured by the Ben Zeev & Pappa (2015) series reflect both fiscal and geopolitical news shocks — they just need to be unanticipated changes that do not depend on the FOMC rotation.

REFERENCE POINT: MONETARY SHOCKS. For comparison of our estimated monetary policy treatment effects with the standard literature, we will also run a simple (non-interacted) local projection with the familiar monetary shock series of Romer & Romer (2004).

EMPIRICAL SPECIFICATIONS. Consistent with our theory above and the recommendations of Montiel Olea et al. (2025), we estimate interacted local projections of the form

$$y_{t+h} = \alpha_{y,h} + \beta_{y,h}\varepsilon_{j,t} + \gamma_{y,h}\zeta_t + \delta_{y,h}(\varepsilon_{j,t} \times \zeta_t) + \sum_{\ell=1}^p \xi_{\ell}' q_{t-\ell} + u_{t,y,h}, \quad (25)$$

where the lagged control vector $q_{t-\ell}$ contains both the outcome of interest as well as the three contemporaneous regressors. We estimate the interacted local projections for the federal funds rate as the policy instrument as well as real GDP and inflation as the macroeconomic outcomes of interest. Further details on data construction, sample selection, lag length, and other estimation details are provided in Appendix C.1.

RESULTS. We begin in Figure 2 with the non-interacted local projection on the canonical Romer & Romer (2004) monetary policy shocks. The results are familiar, with the induced interest rate hike lowering output and inflation. For our purposes, the key is the left panel — the identified interest rate “treatment.” We see that the policy rate increase is transitory, lasting for around a year and a half. The more detailed literature review in Caravello et al. (2025) shows that this is in fact quite typical: standard measures of monetary policy shocks identify the causal effects of transitory policy treatments.

Our headline results are displayed in Figure 3 and reveal one of the key applied takeaways





Figure 2: Impulse responses of the policy rate, output, and inflation to a monetary policy shock identified following Romer & Romer (2004). The shaded areas correspond to 68 per cent confidence bands. Quarterly data.

of this paper: that interacted monetary policy rule change regressions tend to identify gradual, more persistent monetary policy treatments, i.e., the effects of forward guidance-type policy. The three panels correspond to interest rate, output, and inflation impulse responses to the identified “as-if” monetary policy shocks — i.e., the estimated δ_h coefficients — for our two non-policy shock measures $\varepsilon_{j,t}$: oil (top panel) and fiscal policy (bottom panel). In each case, the left sub-panel is the policy treatment, while the other two then give the associated dynamic causal effects. We see that in both cases the interest rate signal is relatively precisely estimated, and most importantly much more persistent than the original Romer & Romer policy treatment.¹⁷ In terms of outcomes, the estimated output response is somewhat larger and more gradual than in Figure 2, consistent with the policy treatment here being more delayed and so also overall larger (in cumulative interest rate terms). The picture for inflation is more mixed, with the fiscal shock indicating an initial price puzzle, consistent perhaps with an interest rate cost channel.

The results of this section reveal that, once interpreted through the lens of our identification results, recent work on differential shock propagation by monetary policy regime offers a useful *complement* to the existing monetary policy shock literature: while that literature tends to isolate transitory disturbances, the estimand of interacted regressions is instead a very different slice of the high-dimensional space of possible monetary policy treatments —

¹⁷The intuition is exactly as discussed in Section 3.2. First, the two non-monetary shocks themselves have relatively persistent effects. Second, even though the policy measure $Hawk_t^{IV}$ is only moderately persistent (reflecting annual rotations), the well-documented gradualism in monetary policy conduct means that such changes in systematic policy design will have lasting effects.

OIL SHOCK: HAMILTON (2003)



FISCAL SHOCK: BEN ZEEV & PAPPA (2015)




Figure 3: Impulse responses of the policy rate, output, and inflation to “as-if” monetary policy shocks identified through the hawk-dove measure of Hack et al. (2026) together with an oil shock (top panel) and a fiscal policy shock (bottom panel). The shaded areas correspond to 68 per cent confidence bands. Quarterly data.

a more forward guidance-like treatment effect. This is valuable precisely because the lack of empirical evidence on delayed monetary policy treatments has been a key limitation of the increasingly popular “semi-structural” approach to counterfactual policy evaluation (see McKay & Wolf, 2023; Barnichon & Mesters, 2023).

ALTERNATIVE APPROACH: INTEREST RATE LOWER BOUND. In Appendix C.2, we consider a different measure of time-varying systematic monetary policy: the zero lower bound on the policy rate, e.g., as studied in Miyamoto et al. (2026). Appealingly, those authors provide evidence that, in their setting, shock propagation is not materially different in booms versus

recessions *per se*, suggesting that any differences in shock propagation inside and outside the zero lower bound are chiefly attributable to differential monetary policy conduct.

For our purposes here, it suffices to just state the main takeaway of this second application: interacted monetary policy local projections yet again identify gradual policy treatments, and for the same reason as above — the effects of both shocks and policy regimes are persistent. Interacted local projections with systematic measures of monetary policy conduct thus provide, through the lens of our results, a robust alternative way of recovering an estimand that has so far proved elusive: the causal effects of persistent, forward guidance-like monetary policy treatments.¹⁸

6.2 Different exchange rate treatments

Our second application is based on Fukui et al. (2025), who study the differential responses of dollar peggers and floaters to a depreciation of the U.S. dollar. As shown by those authors, in standard open-economy models, differences in exchange rate regimes shape the propagation of shocks from abroad through two channels: the real rate of interest, and the real effective exchange rate. We first replicate and then extend their applications to illustrate the second key applied takeaway of our analysis: that changing the shock variable can materially change the identified policy treatment.

EMPIRICAL SPECIFICATION. Like Fukui et al. (2025) we estimate the panel regressions

$$\Delta y_{i,t+h} = \alpha_{y,i,h} + \alpha_{y,r(i),t,h} + \gamma_{y,h} \zeta_{i,t} + \delta_{y,h} (\zeta_{i,t} \cdot \varepsilon_{j,t}) + \sum_{\ell=1}^p \xi_{\ell}^{\prime} q_{i,t-\ell} + u_{y,i,t,h}, \quad (26)$$


where i is the country, $\alpha_{y,i,h}$ is a country fixed effect, $\alpha_{y,r(i),t,h}$ is a region-by-time fixed effect, and $q_{i,t-\ell}$ denotes additional control variables. We discuss the choice of controls further in Appendix C.3; our main text discussion will instead focus on the policy and shock variables.¹⁹

SHOCKS AND POLICY MEASURE. In Fukui et al. (2025), the policy measure $\zeta_{i,t}$ is a binary indicator denoting whether country i at time t pegs to the dollar, and the shock $\varepsilon_{j,t}$ is the log

¹⁸Monetary policy conduct however plausibly changed more broadly during the zero lower bound episode, notably including unconventional monetary interventions (like quantitative easing). The estimand is thus a joint rate-plus-balance-sheet policy treatment, illustrating Corollary 1.

¹⁹Note that the shock $\varepsilon_{j,t}$ only varies over time and not across countries, as discussed further below. Thus, because of our included fixed effects, the shock itself (without interactions) does not appear in (26).

U.S. DOLLAR DEPRECIATION SHOCK



GLOBAL FINANCIAL CONDITIONS SHOCK




Figure 4: Impulse responses of the real exchange rate, real interest rate, and output, based on differential responses of peggers and floaters to a U.S. dollar depreciation shock (top panel) and a global financial conditions shock (bottom panel). The shaded areas correspond to 68 per cent confidence bands. Annual data.

change in the dollar nominal effective exchange rate from time $t - 1$ to t . Through the lens of our discussion, the key requirement is that $\varepsilon_{j,t}$ captures macroeconomic news orthogonal to the regime variable $\zeta_{i,t}$, and whose differential propagation in the cross-section of countries is plausibly shaped only by the two policy levers — i.e., the real interest and exchange rates.²⁰ The most obvious threat to identification is that the exchange rate regime is correlated with other sources of country heterogeneity that determine the exposure to shocks from abroad. Fukui et al. provide various pieces of evidence in support of this requirement, in particular

²⁰The set-up thus fits into our extended structure with multiple distinct policy instruments, as discussed earlier in Corollary 1.

noting that, at least conditional on regional fixed effects, there is little difference between pegs and floats across a range of observables.

Our focus in this section will instead be on the shock variable. Nominal exchange rate changes are a very natural candidate for $\varepsilon_{j,t}$ because they are (largely) unpredictable, plausibly unrelated to infrequent changes in exchange rate regimes, and can serve as a common, fixed conditioning event across countries, as is required through the lens of our identification results.²¹ Crucially, however, the exact same logic would also apply to other date- t macroeconomic news — alternative “shock” variables that, when combined with the policy regime, would likely isolate a different mix of policy treatments. To illustrate that point of possible policy treatment heterogeneity, we will consider as a second shock variable (log) changes in the global financial cycle indicator from Miranda-Agrippino & Rey (2020).

RESULTS. Results are displayed in Figure 4. The top panel essentially replicates Fukui et al. (2025): using the dollar depreciation as the shock variable, the policy treatment, displayed in the left and middle panels, is an exchange rate depreciation coupled with (approximately) no change in interest rates; the associated treatment effect, shown in the right panel, is a persistent boom. For our purposes, the key takeaway is the comparison of that top panel with the bottom panel, using instead the change in global financial conditions as the shock variable. We now see a depreciation coupled with a real rate hike, resulting in a much more muted increase in economic activity. Switching the shock variable thus here results in the analyst identifying the effects of a very different policy treatment.

7 Conclusions

Under what assumptions can changes in the systematic design of policy be used to learn about that policy’s causal effects? And how do econometric strategies that leverage such policy design changes relate to the common practice of estimating the dynamic propagation of policy “shocks”? The principal contribution of this paper is a set of identification results that answer these questions. The headline applied takeaway is that local projections on macroeconomic shocks interacted with measures of policy regime changes admit, under suitable assumptions on the policy measure and on the data-generating process, a tight causal interpretation: just

²¹Exchange rates are equilibrium outcomes, and thus correlated with a range of underlying primitive time- t shocks. As stressed earlier, however, this is no problem as long as the weights on the underlying shocks do not change with the policy regime variable (recall Section 3.3) — which here is the case precisely because the conditioning event is *global*, and thus fixed in the cross section of countries.

like regressions on policy shocks, they identify policy causal effects — the impulse response of the policy instrument is the “treatment,” and the other responses then give the associated causal effects. The availability of a primitive, structural non-policy shock to interact with the policy measure, or more generally of some variable that itself is invariant to the change in policy, is furthermore *necessary*: reduced-form Wold innovations typically change with the policy regime, and policy-induced changes in unconditional second moments yield excessively wide identified sets.

More practically, our results suggest that interacted local projections of the sort studied in recent work hold substantial promise as an alternative, or at least as a useful complement, to the hitherto dominant focus on policy shocks. Since policy is inherently high-dimensional, robust model discipline and counterfactual policy evaluation require estimates of a rich space of policy causal effects — richer than is available using the existing policy shocks alone. Appealingly, any given interacted local projection will generically identify a different slice of this high-dimensional space than any given policy shock; for example, in our applications to monetary policy, we find that policy rule changes implicitly isolate gradual interest rate changes, very different from the short-lived interest rate treatments identified by standard monetary policy shocks. We hope that future work will further leverage these insights to provide a more complete picture of policy transmission to the macro-economy.

References

- Adao, R., Kolesár, M. M., & Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4), 1949–2010.
- Almuzara, M. & Sancibrián, V. (2024). Micro responses to macro shocks. Working paper, CEMFI.
- Angeletos, G.-M., Lian, C., Wolf, C. K., & Zhang, D. (2026). Monetary-fiscal interactions: A reappraisal. Working Paper.
- Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5), 2375–2408.
- Barnichon, R. & Mesters, G. (2023). A sufficient statistics approach for macro policy. *American Economic Review*, 113(11), 2809–2845.
- Barthélemy, J. & Marx, M. (2019). Monetary policy switching and indeterminacy. *Quantitative Economics*, 10(1), 353–385.
- Ben Zeev, N. & Pappa, E. (2015). Multipliers of unexpected increases in defense spending: An empirical investigation. *Journal of Economic Dynamics and Control*, 57, 205–226.
- Bianchi, F. (2013). Regime switches, agents’ beliefs, and post-world war ii us macroeconomic dynamics. *Review of Economic Studies*, 80(2), 463–490.
- Bianchi, F. & Melosi, L. (2017). Escaping the great recession. *American Economic Review*, 107(4), 1030–1058.
- Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, 89(1), 181–213.
- Caravello, T. E., McKay, A., & Wolf, C. K. (2025). Evaluating Monetary Policy Counterfactuals: (When) Do We Need Structural Models? Working Paper.
- Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1), 1–45.

- Christiano, L. J., Eichenbaum, M. S., & Trabandt, M. (2015). Understanding the great recession. *American Economic Journal: Macroeconomics*, 7(1), 110–167.
- Cloyne, J., Jordà, Ò., & Taylor, A. M. (2023). State-dependent local projections: Understanding impulse response heterogeneity. Technical report, National Bureau of Economic Research.
- Cogley, T. & Sbordone, A. M. (2008). Trend inflation, indexation, and inflation persistence in the new keynesian phillips curve. *American Economic Review*, 98(5), 2101–2126.
- Darvas, Z. (2021). Timely measurement of real effective exchange rates. Bruegel Working Papers 46497, Bruegel.
- Davig, T. & Leeper, E. M. (2006). Endogenous monetary policy regime change. In *NBER International Seminar on Macroeconomics*, number 1, (pp. 345–391). The University of Chicago Press Chicago, IL.
- Davig, T. & Leeper, E. M. (2007). Generalizing the Taylor Principle. *American Economic Review*, 97(3), 607–635.
- Evans, G. W. & Honkapohja, S. (2001). *Learning and expectations in macroeconomics*. Princeton University Press.
- Farmer, R. E., Waggoner, D. F., & Zha, T. (2009). Understanding markov-switching rational expectations models. *Journal of Economic Theory*, 144(5), 1849–1867.
- Fernández-Villaverde, J., Rubio-Ramírez, J. F., Sargent, T. J., & Watson, M. W. (2007). Abcs (and ds) of understanding vars. *American Economic Review*, 97(3), 1021–1026.
- Fukui, M., Nakamura, E., & Steinsson, J. (2025). The macroeconomic consequences of exchange rate depreciations. *The Quarterly Journal of Economics*, 140(4), 3015–3065.
- Gabaix, X. (2020). A behavioral New Keynesian model. *American Economic Review*, 110(8), 2271–2327.
- Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the New Keynesian framework and its applications*. Princeton University Press.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8), 2586–2624.

- Gonçalves, S., Herrera, A. M., Kilian, L., & Pesavento, E. (2024). State-dependent local projections. *Journal of Econometrics*, 244(2), 105702.
- Hack, L., Istrefi, K., & Meier, M. (2026). Identification of systematic monetary policy. Working Paper.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2), 363–398.
- Herbst, E. P. & Johannsen, B. K. (2024). Bias in local projections. *Journal of Econometrics*, 240(1), 105655.
- Imbens, G. W. & Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2), 467–475.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161–182.
- Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review*, 111(4), 1092–1125.
- Kaplan, G., Moll, B., & Violante, G. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3), 697–743.
- Kolesár, M. & Plagborg-Møller, M. (2025). Dynamic causal effects in a nonlinear world: the good, the bad, and the ugly. *Journal of Business & Economic Statistics*, 43(4), 737–754.
- Lewbel, A. (2012). Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. *Journal of Business & Economic Statistics*, 30(1), 67–80.
- Lewis, D. J. (2025). Identification based on higher moments in macroeconometrics. *Annual Review of Economics*, 17.
- Lucas, R. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2), 103–124.
- Magnusson, L. M. & Mavroeidis, S. (2014). Identification using stability restrictions. *Econometrica*, 82(5), 1799–1851.
- Manuel, E. (2026). Missing intercept, biased slope: Identification in general equilibrium. Unpublished manuscript.

- Mavroeidis, S. (2021). Identification at the zero lower bound. *Econometrica*, 89(6), 2855–2885.
- McKay, A. & Wolf, C. K. (2023). What can time-series regressions tell us about policy counterfactuals? *Econometrica*, 91(5), 1695–1725.
- Miranda-Agrippino, S. & Rey, H. (2020). U.S. monetary policy and the global financial cycle. *The Review of Economic Studies*, 87(6), 2754–2776.
- Miyamoto, W., Nguyen, T. L., & Sergeyev, D. (2026). How oil shocks propagate: Evidence on the monetary policy channel. Working Paper.
- Montiel Olea, J. L. & Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4), 1789–1823.
- Montiel Olea, J. L., Plagborg-Møller, M., Qian, E., & Wolf, C. K. (2025). Local projections or VARs? A primer for macroeconomists. Technical report, National Bureau of Economic Research.
- Mussa, M. (1986). Nominal exchange rate regimes and the behavior of real exchange rates: Evidence and implications. In *Carnegie-Rochester Conference series on public policy*, volume 25, (pp. 117–214). Elsevier.
- Plagborg-Møller, M. & Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2), 955–980.
- Rambachan, A. & Shephard, N. (2025). When do common time series estimands have nonparametric causal meaning? Working Paper.
- Ramey, V. A. (2011). Identifying government spending shocks: It’s all in the timing. *The Quarterly Journal of Economics*, 126(1), 1–50.
- Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. In J. B. Taylor & H. Uhlig (Eds.), *Handbook of Macroeconomics*, volume 2 chapter 2, (pp. 71–162). Elsevier.
- Ramey, V. A. & Zubairy, S. (2018). Government spending multipliers in good times and in bad: evidence from us historical data. *Journal of Political Economy*, 126(2), 850–901.
- Rigobon, R. (2003). Identification through heteroskedasticity. *Review of Economics and Statistics*, 85(4), 777–792.

- Romer, C. D. & Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4), 1055–1084.
- Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica*, 48(1), 1–48.
- Smets, F. & Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3), 586–606.
- Uhlig, H. (2005). What are the effects of monetary policy on output? results from an agnostic identification procedure. *Journal of Monetary Economics*, 52(2), 381–419.
- Wieland, J. F. & Yang, M.-J. (2020). Financial dampening. *Journal of Money, Credit and Banking*, 52(1), 79–113.
- Winkler, V. (2026). When and why state-dependent local projections work. Working Paper.
- Wolf, C. K. (2025). Interest rate cuts versus stimulus payments: An equivalence result. *Journal of Political Economy*, 133(4), 1235–1275.
- Woodford, M. (2003). *Interest and prices: Foundations of a theory of monetary policy*. Princeton University Press.

Appendix for:

Identifying Policy Causal Effects from Rule Changes

This appendix contains supplemental material for the article “Identifying Policy Causal Effects from Rule Changes.” We provide (i) several supplementary theoretical results; (ii) quantitative simulations to illustrate our identification results; (iii) implementation details for the empirical analysis; and (iv) all proofs.

Any references to equations, figures, tables, assumptions, propositions, lemmas, or sections that are not preceded by “A.”—“D.” refer to the main article.

A Supplementary theoretical results

In this appendix we provide further details for several theoretical results referenced or only sketched in our main analysis, in Sections 2 to 5.

A.1 Interest rate causal effects in the two-period model

In Section 2.2 we translated the object of interest, monetary policy causal effects, from shock space to instrument space, i.e., from $\bar{\Theta}_v$ to Θ_i . We here derive closed-form expressions for both representations.

SHOCK SPACE. The entries of the matrix $\bar{\Theta}_v$ can be recovered by solving the triplet of equations (IS) - (TR) as a function of the two monetary policy shocks. Letting $D \equiv \gamma + \kappa\phi_\pi$, this delivers

$$\begin{aligned} x_0 &= -\frac{1}{D}v_0 - \frac{\gamma + \kappa(1 - \beta\phi_\pi)}{D^2}v_1, \\ x_1 &= -\frac{1}{D}v_1, \end{aligned}$$

for output,

$$\begin{aligned} \pi_0 &= -\frac{\kappa}{D}v_0 - \frac{\kappa[\gamma(1 + \beta) + \kappa]}{D^2}v_1, \\ \pi_1 &= -\frac{\kappa}{D}v_1, \end{aligned}$$

for inflation, and

$$\begin{aligned} i_0 &= \frac{\gamma}{D}v_0 - \frac{\kappa\phi_\pi[\gamma(1 + \beta) + \kappa]}{D^2}v_1, \\ i_1 &= \frac{\gamma}{D}v_1, \end{aligned}$$

for nominal interest rates. Stacking, the coefficients on v_0 deliver the first column of $\bar{\Theta}_v$, and those on v_1 deliver the second column.

INSTRUMENT SPACE. The entries of Θ_i can instead be recovered by solving the *pair* of equations (IS) - (NKPC) at dates $t = 0, 1$ as a function of interest rates, and in the absence of any other shocks. Importantly, by enforcing *exact* return of the economy to steady state after

two periods, we implicitly enforce the conventional Taylor principle equilibrium selection (see Angeletos et al., 2026), leaving only the on-equilibrium role of policy — where, by instrument sufficiency, only the implied values of the instrument itself matter. Solving, we recover

$$\begin{aligned}x_0 &= -\frac{1}{\gamma}i_0 - \frac{1}{\gamma}\left(1 + \frac{\kappa}{\gamma}\right)i_1 \\x_1 &= -\frac{1}{\gamma}i_1\end{aligned}$$

for output, and

$$\begin{aligned}\pi_0 &= -\frac{\kappa}{\gamma}\left(i_0 + \left(1 + \frac{\kappa}{\gamma}\right)i_1\right) - \beta\frac{\kappa}{\gamma}i_1 \\ \pi_1 &= -\frac{\kappa}{\gamma}i_1\end{aligned}$$

for inflation. It is immediate that, as claimed, these causal effects do not depend on the policy rule ϕ_π . It is similarly immediate that we can recover the instrument space causal effects from shock space causal effects through the mapping from (v_0, v_1) to (i_0, i_1) , as claimed.

A.2 Set identification through second moments

We provide a detailed analysis of policy causal effect identification through changes in second moments, supplementing the brief discussion in Section 3.3. To build intuition we begin with a simple static analysis, before then turning to dynamics, or equivalently to multiple distinct instruments. Relative to prior work, the key takeaway of the analysis is that the reduced-form policy counterfactual mapping characterized in Caravello et al. (2025) is not invertible, unlike the structural shock mapping in McKay & Wolf (2023).

A.2.1 Static warm-up

We begin with a static analysis based on the general model economy in Section 3.1, setting $H = 1$. The researcher estimates variance-covariance matrices Σ and $\tilde{\Sigma}$ of macroeconomic outcomes y_t across two distinct policy regimes, and then seeks to use the change in those variance-covariance matrices, $\Delta \equiv \tilde{\Sigma} - \Sigma$, to learn as much as possible about policy causal effects. Clearly $\Delta \neq \mathbf{0}$ is informative about policy (non-)neutrality — the main idea of Mussa (1986). We here instead seek to characterize the entire possible set of policy causal effects consistent with the observed second moment change.

THE EFFECT OF POLICY CHANGES ON SECOND MOMENTS. Letting ζ_t as in Section 3.2 denote a regime indicator, the variance-covariance matrix at date t , $\Sigma_t \equiv \text{Var}(y_t)$, is

$$\Sigma_t = \underbrace{\bar{\Theta}\bar{\Theta}'}_{\text{baseline policy}} + \underbrace{U_t\bar{\Theta}'_v + \bar{\Theta}_vU'_t + \bar{\Theta}_vS_t\bar{\Theta}'_v}_{\text{effect of policy change}} \quad (\text{A.1})$$

where $U_t \equiv \bar{\Theta}\mathbb{E}[\varepsilon_t\Omega(\varepsilon_t, \zeta_t) \mid \zeta_t]$ and $S_t \equiv \mathbb{E}[\Omega(\varepsilon_t, \zeta_t)^2 \mid \zeta_t]$. We can see that the variance-covariance matrix consists of two terms: the effect of the shocks ε_t under the baseline policy rule, and how the deviation of the actual policy from the baseline rule alters things. Key to all of the subsequent arguments will be that the second term is generally rank-2, shifting the baseline variance-covariance matrix in two distinct directions: the rank-1 variance-covariance matrix associated with a policy shock $(\bar{\Theta}_vS_t\bar{\Theta}'_v)$, and the rank-1 interaction of policy causal effects with the propagation of other shocks $(U_t\bar{\Theta}'_v)$. This result appears promising: policy rule changes will leave a low-dimensional imprint on reduced-form second moments, and thus changes in those second moments should be informative about policy effects.

Comparing across two policy regimes, with tildes indicating the second regime, we obtain

$$\Delta = (\tilde{U} - U)\bar{\Theta}'_v + \bar{\Theta}_v(\tilde{U} - U)' + \bar{\Theta}_v(\tilde{S} - S)\bar{\Theta}'_v. \quad (\text{A.2})$$

It is immediate that $\text{rank}(\Delta) \leq 2$; for all subsequent arguments we will assume that Δ indeed attains that maximal rank.²²

THE IDENTIFIED SET OF POLICY EFFECTS. The objective now is to characterize the set of possible policy causal effects consistent with the observed variance-covariance matrices.

Definition 1. A vector $\bar{\Theta}_v^\dagger$ lies in the identified set of policy causal effects if it is part of a tuple $\{\bar{\Theta}^\dagger, \bar{\Theta}_v^\dagger, U^\dagger, \tilde{U}^\dagger, S^\dagger, \tilde{S}^\dagger\}$ that rationalizes the observed pair $(\Sigma, \tilde{\Sigma})$ according to (A.1).

In words, a candidate vector $\bar{\Theta}_v^\dagger$ is admissible if and only if it can be used to generate variance-covariance matrices $(\Sigma, \tilde{\Sigma})$ of the form (A.1) and consistent with the observed gap Δ . The following proposition characterizes the set of such vectors.

Proposition A.1. Consider two static economies (5) - (6) (with $H = 1$) that differ only in

²²If the rank instead was 1, then we would recover $\bar{\Theta}_v$ (up to a factor of proportionality) as any basis of the column space of Δ ; and if it was zero, then the policy change had no effect on second-moment properties.




Figure A.1: Identified set for interest rate, output, and inflation impulse responses to a one-off increase in interest rates, in a simple static New Keynesian model. The blue dots indicate members of the identified set, and the orange dots indicate the truth. See Appendix B.2 for model details.

the policy rule. Given Σ and $\tilde{\Sigma}$ with $\text{rank}(\tilde{\Sigma} - \Sigma) = 2$, the identified set for $\bar{\Theta}_v$ is:

$$\left\{ \bar{\Theta}_v^\dagger \mid \bar{\Theta}_v^\dagger = \left(V_+ \Lambda_+^{1/2} \pm V_- \Lambda_-^{1/2} \right) \cdot \alpha, \quad \alpha \neq 0 \right\} \quad (\text{A.3})$$

where $V_\bullet, \Lambda_\bullet$ come from the eigendecomposition $\Delta \equiv \tilde{\Sigma} - \Sigma = \Lambda_+ V_+ V_+' - \Lambda_- V_- V_-'$, with Λ_+ and Λ_- equal to the absolute values of the two eigenvalues (one positive, one negative), and $V_\bullet$ denoting the corresponding unit-length eigenvectors.

The proposition reveals that, up to scale, the identified set for policy causal effects consists of just two points. Intuitively, the fact that Δ itself is rank-2 already reduces candidate policy causal effects to a two-dimensional space; the specific form of the variance-covariance matrix in (A.1) then further reduces that identified set to just two points.

The preceding identification result is intimately connected to existing results on identification through heteroskedasticity, and seems to suggest some promise. In that literature, pure changes in policy shock volatilities induce a rank-1 change of variance-covariance matrices, delivering immediate point identification (up to scale). Here the change is of rank 2, but the further properties of Δ materially shrink the identified set, to the extent that a modicum of additional information (e.g., in the form of sign restrictions, as in Uhlig, 2005) would be sufficient for identification. Figure A.1 visually illustrates this point, showing the

identified set of interest rate causal effects in a static New Keynesian model. Consistent with Proposition A.1, the identified set consists of two points; requiring a monetary contraction to lower inflation would then point-identify the true policy causal effects.

A.2.2 Identified sets with dynamics, or multiple instruments

We now turn to the main result on policy shock causal effects in a general dynamic economy. Unfortunately the takeaway will be that, compared to the simple static case, the identified set widens so much as to make it largely uninformative.

THE EFFECT OF POLICY CHANGES ON SECOND MOMENTS. As before we begin by characterizing unconditional second moments in the economy with stochastic policy rule changes. Almost identically to the static case, we now find the variance-covariance matrix of date- t expectation revisions as

$$\Sigma_t = \underbrace{\bar{\Theta}\bar{\Theta}'}_{\text{baseline policy}} + \underbrace{U_t\bar{\Theta}'_v + \bar{\Theta}_vU'_t + \bar{\Theta}_vS_t\bar{\Theta}'_v}_{\text{effect of policy change}} \quad (\text{A.4})$$

where $U_t \equiv \bar{\Theta}\mathbb{E}_{t-1}[\varepsilon_t\Omega(\bullet)']$ is $(H \cdot n_y) \times H$ and $S_t \equiv \mathbb{E}_{t-1}[\Omega(\bullet)\Omega(\bullet)']$ is $H \times H$. We thus see that second moments have the same shape as in the static case — shock-induced stochasticity under the baseline rule, plus a policy-induced offset — but now the dimensionalities are different, with Σ_t a $(H \cdot n_y) \times (H \cdot n_y)$ matrix, and the offset term generally being of rank- $2H$. Intuitively, in an $(H-1)$ -lag dynamic economy, there are now H distinct policy instruments, reflecting movements of the scalar policy instrument z_t over time. The offset term combines direct stochasticity induced by those policy movements (last term) plus their interaction with baseline shock impulse responses (two middle terms).

We now as in the static case compare across two regimes, giving

$$\Delta = (\tilde{U} - U)\bar{\Theta}'_v + \bar{\Theta}_v(\tilde{U} - U)' + \bar{\Theta}_v(\tilde{S} - S)\bar{\Theta}'_v.$$

It is again immediate that $\text{rank}(\Delta) \leq 2H$; we as above assume that this holds with equality, with detailed arguments provided in the proof of Proposition A.2.

THE IDENTIFIED SET OF POLICY EFFECTS. We define the identified set of possible policy causal effects exactly as in Definition 1. We then arrive at the following characterization of that identified set.

Proposition A.2. *Consider two $(H - 1)$ -lag dynamic economies (5) - (6) that differ only in the policy rule (6). Given Σ and $\tilde{\Sigma}$ with $\text{rank}(\tilde{\Sigma} - \Sigma) = 2H$, the identified set for $\bar{\Theta}_v$ is:*

$$\left\{ \bar{\Theta}_v^\dagger \mid \bar{\Theta}_v^\dagger = \left(V_+ \Lambda_+^{1/2} + V_- \Lambda_-^{1/2} Q \right) \cdot \mathcal{A}, \quad Q \in \mathcal{O}(H) \text{ and } \mathcal{A} \text{ invertible} \right\} \quad (\text{A.5})$$

where $V_\bullet, \Lambda_\bullet$ come from the eigendecomposition $\Delta \equiv \tilde{\Sigma} - \Sigma = V_+ \Lambda_+ V_+' - V_- \Lambda_- V_-'$, and $\mathcal{O}(H)$ denotes the group of H -dimensional orthogonal matrices. The diagonal matrices $\Lambda_\bullet$ collect the (absolute values of the) H positive and H negative eigenvalues of Δ , and the $V_\bullet$'s contain the orthonormal eigenvectors.

The structure of the identified set is similar to the static case, but with one crucial difference. Recall that, in the static case, the identified set only consisted of two distinct points (up to scale), with the negative eigenvalue component added to or subtracted from the positive component in (A.3). One of those two points corresponded to the truth, while the other was an “imposter,” reflecting not true policy causal effects but instead the interaction of those causal effects with non-policy shock impulse responses. In the dynamic case, with H policy instruments, identification is now up to an H -dimensional space — that of H -dimensional rotation matrices Q . Mathematically, there are now more ways to combine true policy causal effects with the “imposter” interaction effects while still maintaining consistency with the observed variance-covariance matrices.²³

How meaningful is this extension of the identified set? There is still useful information in the difference Δ : it directly identifies a candidate $2H$ -dimensional space, and the actual identified set for policy causal effects is then a further restricted subspace. Unfortunately, in practice, the resulting identified set can actually be *very* wide. An illustration is provided in Figure A.2, which shows the identified set corresponding to a one-off interest rate change in a standard two-period New Keynesian model. We see that the resulting identified sets are essentially uninformative, showcasing how the additional freedom of mixing between the two components in (A.5) is so great as to preclude any sharp conclusions.

A.2.3 Discussion

What accounts for the difference between our analysis here and the seemingly similar strategy of identification through (policy) shock heteroskedasticity? As already discussed, policy rule changes induce a higher-dimensional change in reduced-form variance-covariance matrices

²³The invertible matrix $\mathcal{A}$ in (A.5) plays the same role as $\alpha \neq 0$ in the static case — policy *shock* causal effects are only pinned down up to scale and rotations. In instrument space, $\mathcal{A}$ cancels out.




Figure A.2: Identified set for interest rate, output, and inflation impulse responses to a one-off increase in interest rates, in a two-period New Keynesian model. The blue dots indicate members of the identified set, and the orange dots indicate the truth. See Appendix B.2 for model details.

than mere shock volatility changes. In the static case, the associated dimensionality increase of the identified set is manageable, and can be sidestepped through additional, mild identifying information. The problem with the dynamic case is not dynamics *per se*, but rather the fact that the dimensionality of the policy instrument increased. As a result, richer combinations of true policy causal effects and the “imposter” terms become admissible, leading to a substantive widening of the identified set.²⁴

The overall takeaway of this section thus is negative: even though policy regime changes imply meaningful restrictions on how reduced-form second moments change, those restrictions are still insufficient to say (much) more than Mussa’s original conclusions about policy non-neutrality. While policy causal effects and pre-policy change second moments suffice to predict post-change second moments, this mapping unfortunately is not invertible. In other words, the additional restrictions implied by the availability of a fixed conditioning event $\varepsilon_{j,t}$ — leveraged so heavily in our main arguments — are indispensable.

²⁴In principle, the identified set could be tightened further if the researcher were to observe *multiple* regime changes. In practice, however, given practical estimation uncertainty, this is again unlikely to yield tightly identified policy causal effects.

A.3 Environment with cross-sectional units

We here explicitly show how our most general environment in Section 5.1 can be re-cast as a system with heterogeneous cross-sectional units.

ENVIRONMENT. Cross-sectional units are indexed by i ; to connect with the notation of Section 5.1, we assume that each unit i 's private-sector characteristics are summarized by the scalar variable m_i . Outcomes for each i are then described by a system of the same form as in Section 3.1, just now indexed by m_i for the private-sector block:

$$\mathcal{H}_x(m_i) [\mathbb{E}_t(\mathbf{x}_i^t) - \mathbb{E}_{t-1}(\mathbf{x}_i^t)] + \mathcal{H}_z(m_i) [\mathbb{E}_t(\mathbf{z}_i^t) - \mathbb{E}_{t-1}(\mathbf{z}_i^t)] + \mathcal{H}_e(m_i)e_{i,t} = \mathbf{0}, \quad (\text{A.6})$$

as well as

$$\mathcal{A}_{i,x} [\mathbb{E}_t(\mathbf{x}_i^t) - \mathbb{E}_{t-1}(\mathbf{x}_i^t)] + \mathcal{A}_{i,z} [\mathbb{E}_t(\mathbf{z}_i^t) - \mathbb{E}_{t-1}(\mathbf{z}_i^t)] + \mathcal{A}_{i,v}v_{i,t} = \mathbf{0}. \quad (\text{A.7})$$

The system (A.6) - (A.7) repeats the general environment of Section 5.1 for many cross-sectional units, allowing both policy rules as well as private-sector coefficients to be heterogeneous across the units. For all the following arguments we will assume that the random variable m_i is independent of the shocks $\{e_{i,t}, v_{i,t}\}$; we do, however, allow these shocks to be correlated (or even identical) across i 's, reflecting, say, global shocks.

Following the same steps as in our main analysis, it is straightforward to establish that equilibrium outcomes satisfy

$$\mathbb{E}_t(\mathbf{y}_i^t) - \mathbb{E}_{t-1}(\mathbf{y}_i^t) = \bar{\Theta}(m_i)\varepsilon_{i,t} + \bar{\Theta}_v(m_i)\Omega(\varepsilon_{i,t}, \varepsilon_{i,t-1}, \dots; m_i, \mathcal{A}_{i,x}, \mathcal{A}_{i,z}, \mathcal{A}_{i,v}), \quad (\text{A.8})$$

where $\Omega(\varepsilon_{i,t}, \varepsilon_{i,t-1}, \dots; m_i, \mathcal{A}_{i,x}, \mathcal{A}_{i,z}, \mathcal{A}_{i,v})$ is defined implicitly by (A.7). We also note that policy causal effects depend on the cross-sectional variable m_i , and are given as

$$\Theta_{x,z}(m_i) \equiv \bar{\Theta}_{x,v}(m_i) \cdot \bar{\Theta}_{z,v}(m_i)^{-1}.$$

RUNNING INTERACTED LOCAL PROJECTIONS. In keeping with the analysis in Section 5, we suppose that the econometrician has access to the j th primitive shock $\varepsilon_{i,j,t}$ as well as a mean-zero random variable ζ_i that is independent of both $\varepsilon_{i,j,t}$ and m_i . We then consider

interacted local projections of the form

$$y_{i,t+h} = \alpha_{y,h} + \beta_{y,h}\varepsilon_{i,j,t} + \gamma_{y,h}\zeta_i + \delta_{y,h}(\varepsilon_{i,j,t} \times \zeta_i) + u_{i,t,y,h}. \quad (\text{A.9})$$

We now see that the time series analysis of Section 5.2 applies unchanged: under an appropriate monotonicity assumption, δ_y can again be interpreted as capturing a (convex)-weighted average of policy causal effects.

MAPPING TO FUKUI ET AL. (2025). Our second application in Section 6.2 fits into this cross-sectional structure. The “shock” $\varepsilon_{i,j,t}$ is now a common (global) innovation, and ζ_i is an indicator for the exchange rate regime that by assumption does not vary with any country characteristics, at least conditional on controls.²⁵

A.4 Mis-identification with other sources of state dependence

Complementing the brief discussion in Section 5.1, we here elaborate on what happens when the “policy measure” ζ_t co-varies with other sources of state dependence or of time variation in shock propagation. To make the argument as stark as possible, we will assume that there is time variation — e.g., in the form of different long-lasting regimes, or exogenous Markov regime-switching — *only* in one non-policy equation. For example, in the simple model of Section 2, the policy rule would now be time-invariant, but the NKPC may be in one of two regimes, $\kappa \in \{\underline{\kappa}, \bar{\kappa}\}$. The model environment thus fits into the structure of Section 4, but with an important twist in interpretation: the stable “private-sector” block

$$\mathcal{H}_x [\mathbb{E}_t(\mathbf{x}^t) - \mathbb{E}_{t-1}(\mathbf{x}^t)] + \mathcal{H}_z [\mathbb{E}_t(\mathbf{z}^t) - \mathbb{E}_{t-1}(\mathbf{z}^t)] + \mathcal{H}_e e_t = \mathbf{0}$$

now contains the similarly stable policy rule, while the stochastically switching private-sector equation is replaced by a solved-out, reduced-form mapping for one model variable,

$$\mathbb{E}_t(\mathbf{x}_1^t) - \mathbb{E}_{t-1}(\mathbf{x}_1^t) = \mathcal{A}(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots), \quad (\text{A.10})$$

echoing Section 4.1. Proceeding exactly as in our main analysis, we can define $\bar{\Theta}_\xi$ as the dynamic causal effects of shocks to an arbitrary “baseline” variant of the time-varying private-

²⁵In that application, we allow the policy regime variable ζ_i to actually also (slowly) change with time, which does not affect any of the arguments of the present section.

sector equation, and $\bar{\Theta}$ as impulse responses to all shocks ε_t under that baseline equation. We then arrive at an equilibrium representation of the form

$$\mathbb{E}_t [\mathbf{y}^t] - \mathbb{E}_{t-1} [\mathbf{y}^t] = \bar{\Theta} \varepsilon_t + \bar{\Theta}_\xi \Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots),$$

perfectly analogous to Lemma 2, with $\Omega(\bullet)$ exactly as defined there, now to ensure (A.10). It then follows by the same logic as that in the proof of Proposition 2 that interacted local projections of the form (9) identify weighted averages of the dynamic causal effects $\bar{\Theta}_\xi$; for example, with a time-varying NKPC, they would recover the causal effects of a sequence of contemporaneous and news cost-push shocks, as claimed.

A.5 Additional controls

Consider the setting with richer forms of state dependence discussed in Section 5.1. Rather than assuming that the regime-indicator ζ_t is independent of the non-policy state m_t (as well as the shock-measure $\varepsilon_{j,t}$), suppose instead that ζ_t satisfies *conditional* (mean)-independence:

$$\mathbb{E}[\zeta_t \mid \varepsilon_{j,t}, m_t, c_t] = \alpha + \rho' c_t, \quad (\text{A.11})$$

for some n_c -dimensional vector of control variables c_t and $\rho \in \mathbb{R}^{n_c}$. Suppose furthermore, and as before, that the shock measure $\varepsilon_{j,t}$ is mean zero, independent of all other shocks as well as m_t , and additionally independent of c_t , ζ_t :

$$\varepsilon_{j,t} \perp \zeta_t, c_t, m_t, \varepsilon_{-j,t}, \underbrace{\mathbb{E}_{t-1}[y_{t+h}]}_{\text{past shocks}}, \underbrace{y_{t+h} - \mathbb{E}_t[y_{t+h}]}_{\text{future shocks}}. \quad (\text{A.12})$$

We now consider the following extended local projection that also controls for $c_t \cdot \varepsilon_{j,t}$:

$$y_{t+h} = \alpha_{y,h} + \beta_{y,h} \varepsilon_{j,t} + \gamma_{y,h} \zeta_t + \delta_{y,h} (\varepsilon_{j,t} \cdot \zeta_t) + \eta'_{y,h} (\varepsilon_{j,t} \cdot c_t) + u_{t,y,h}. \quad (\text{A.13})$$

Under assumptions (A.11) and (A.12), and by Frisch-Waugh-Lovell, we find that

$$\delta_{y,h} = \frac{\text{Cov}(y_{t+h}, \varepsilon_{j,t} \cdot \tilde{\zeta}_t)}{\text{Var}(\varepsilon_{j,t} \cdot \tilde{\zeta}_t)},$$

where $\tilde{\zeta}_t$ denotes the residual from a regression of ζ_t on a constant and c_t . Next note that, by (A.12), we have

$$\begin{aligned}\mathbb{E}[\mathbb{E}_{t-1}[y_{t+h}](\varepsilon_{j,t} \cdot \tilde{\zeta}_t)] &= 0, \\ \mathbb{E}[(y_{t+h} - \mathbb{E}_t[y_{t+h}])(\varepsilon_{j,t} \cdot \tilde{\zeta}_t)] &= 0, \\ \mathbb{E}[\bar{\Theta}_{-j}(m_t)\varepsilon_{-j,t}(\varepsilon_{j,t} \cdot \tilde{\zeta}_t)] &= 0.\end{aligned}$$

Additionally, by (A.11),

$$\mathbb{E}[\bar{\Theta}_j(m_t)\varepsilon_{j,t}^2\tilde{\zeta}_t] = \mathbb{E}[\bar{\Theta}_j(m_t)\varepsilon_{j,t}^2\mathbb{E}[\tilde{\zeta}_t | \varepsilon_{j,t}, m_t]] = 0.$$

Then, as in the proof of Proposition 3, using (19), we have that

$$\delta_{y,h} \propto \mathbb{E}[\bar{\Theta}_{y,v,h}(m_t)\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots, m_t)\varepsilon_{j,t}\tilde{\zeta}_t].$$

The remaining steps of Proposition 3 then go through replacing ζ_t with $\tilde{\zeta}_t$.

A.6 Extension to a general non-linear environment

Consider the setting introduced in Section 5.4. By Frisch-Waugh-Lovell, we have that

$$\delta_h = \frac{\text{Cov}(x_{t+h}, \varepsilon_{j,t}s_t)}{\text{Var}(\varepsilon_{j,t})\text{Var}(s_t)}. \quad (\text{A.14})$$

The remainder of this section characterizes this estimand further.

We assume that the scalar random variables $\varepsilon_{j,t}$ and s_t have densities denoted by $f_{\varepsilon_j}(\varepsilon_j)$ and $f_s(s)$, respectively, and are mean zero, $\mathbb{E}(\varepsilon_{j,t}) = \mathbb{E}(s_t) = 0$. We furthermore from now on simply impose high-level assumptions on the environment ensuring that all of the below operations are valid and the relevant expectations exist; these assumptions could be weakened along the lines discussed by Kolesár & Plagborg-Møller, but with little return for the purposes of our discussion here. Specifically, we begin with integration by parts to get

$$\text{Cov}(x_{t+h}, \varepsilon_{j,t}s_t) = \mathbb{E}[\Psi_h(\varepsilon_{j,t}, s_t)\varepsilon_{j,t}s_t] = \mathbb{E}[\tau_s(s_t)\tau_{\varepsilon_j}(\varepsilon_{j,t})\partial_{\varepsilon_j s}\Psi_h(\varepsilon_{j,t}, s_t)]$$

where

$$\tau_{\varepsilon_j}(\varepsilon_j) = \frac{1}{f_{\varepsilon_j}(\varepsilon_j)} \int_{-\infty}^{\varepsilon_j} -r f_{\varepsilon_j}(r) dr, \quad \tau_s(s) = \frac{1}{f_s(s)} \int_{-\infty}^s -r f_s(r) dr.$$

We thus have

$$\delta_h = \int \int w(s_t, \varepsilon_{j,t}) \partial_{\varepsilon_j s} \Psi_h(\varepsilon_{j,t}, s_t) ds_t d\varepsilon_{j,t} \quad (\text{A.15})$$

where

$$w(s_t, \varepsilon_{j,t}) = \frac{\tau_s(s_t) f_s(s_t)}{\text{Var}(s_t)} \frac{\tau_{\varepsilon_j}(\varepsilon_{j,t}) f_{\varepsilon_j}(\varepsilon_{j,t})}{\text{Var}(\varepsilon_{j,t})} \quad (\text{A.16})$$

and the cross-derivative term satisfies

$$\partial_{\varepsilon_j s} \Psi_h(\varepsilon_{j,t}, s_t) = \mathbb{E}_u \left[\psi'_{h, \varepsilon_j z} z_s + z'_{\varepsilon_j} \psi_{h, z z} z_s + \psi'_{h, z} z_{\varepsilon_j s} \right], \quad (\text{A.17})$$

which is exactly the expression in Section 5.4. It now remains to note that, by the properties of $\tau_{\varepsilon_j}(\varepsilon_j)$ and $\tau_s(s)$, $w(s_t, \varepsilon_{j,t}) \geq 0$ and $\int \int w(s_t, \varepsilon_{j,t}) ds_t d\varepsilon_{j,t} = 1$.

B Simulations and illustrations

We provide supplementary information for our two model-based illustrations, in Section 3.2 and Appendix A.2.

B.1 Interacted local projections in Smets & Wouters (2007)

Our laboratory data-generating process for the numerical illustrations in Figure 1 is the well-known structural model of Smets & Wouters (2007), but with the monetary authority following rules of the form

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) (\phi_\pi \pi_t + \phi_y y_t + \phi_{dy} (y_t - y_{t-1})), \quad (\text{B.1})$$

which is somewhat simpler than the headline specification considered by Smets & Wouters.

We treat the posterior mode parameterization of that model as the ground truth, and then consider an econometrician that leverages data generated under different assumptions on the monetary rule. Specifically, the econometrician observes data from two long sub-samples: in the first, the monetary policy rule takes the form (B.1) with $\phi_\pi = 2$, $\phi_y = 0.5$, $\phi_{dy} = 0.2$, and $\rho_i = 0.9$; in the second, we have $\phi_\pi = 5$ and $\phi_y = \phi_{dy} = \rho_i = 0$. She furthermore observes each of the model’s six non-monetary policy shocks, and runs (9) with ζ_t equal to a dummy indicator of the policy regime. Figure 1 reports the population estimands of these interacted local projections. For the top left panel, we define the “as-if” monetary shocks relative to a textbook variant of (B.1) with $\phi_\pi = 2$ and $\phi_y = \phi_{dy} = \rho_i = 0$.²⁶

B.2 Second moments in the New Keynesian model

Our two quantitative illustrations in Figures A.1 and A.2 begin with the same textbook New Keynesian model:

$$x_t = \mathbb{E}_t [x_{t+1}] - \frac{1}{\gamma} (i_t - \mathbb{E}_t [\pi_{t+1}]) + u_t^d, \quad (\text{B.2})$$

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t [\pi_{t+1}] + u_t^s, \quad (\text{B.3})$$

$$i_t = \phi_\pi \pi_t + \phi_x x_t + u_t^m. \quad (\text{B.4})$$

²⁶Of course, as we change that reference rule, the associated “as-if” policy shock sequence changes. What does not change, however, is the policy treatment in “instrument space” (top right panel) as well as the associated output and inflation causal effects.

For our static analysis all three shocks are purely transitory; for the two-period variant they follow MA(1) processes. In both cases we consider a standard model calibration: $\gamma = 1$, $\kappa = 0.5$, and $\beta = 0.99$. We then suppose that the econometrician observes large samples generated under two different monetary rules: $\phi_\pi = 1.5$ and $\phi_x = 0.2$, as well as $\phi_\pi = 3$ and $\phi_x = 0.7$. In both cases, all shocks have identical (unit) volatilities; for the dynamic model, the assumed MA(1) coefficients are $\theta_d = 0.7$ for the demand shock, $\theta_s = 0.2$ for the supply shock, and $\theta_m = 0.5$ for the policy shock.

Given the two large samples, the econometrician recovers Σ and $\tilde{\Sigma}$. We then map these objects into the identified sets characterized in Propositions A.1 and A.2. For the dynamic model we simply sample orthogonal rotation matrices Q at random from $\mathcal{O}(H)$, as familiar from the literature on sign-identified vector autoregressions (Uhlig, 2005).

C Empirical applications

We provide supplementary details for the two main applications in Section 6, as well as for the secondary time series application based on Miyamoto et al. (2026).

C.1 Variation in FOMC composition

We begin with the main application in Section 6.1.

DATA. We study three outcome variables.

- *Output.* We take log output per capita from FRED (A939RX0Q048SBEA).
- *Inflation.* We compute the log-differenced GDP deflator (GDPDEF), and then annualize, without further transformations.
- *Federal funds rate.* We obtain the series FEDFUNDS, without further transformations.

For the systematic monetary policy measure ζ_t , we use the $Hawk_t^{IV}$ series kindly shared by Hack et al. (2026). We finally turn to the shock measures $\varepsilon_{j,t}$.

- *Oil.* We use the oil shock series of Hamilton (2003), available as the first shock series in the file `oil_2019.csv` of the recently updated replication files.
- *Fiscal policy.* We take the series `mfev` from the replication files of Ramey (2016).
- *Monetary policy.* We take the Romer & Romer (2004) shock series from the replication and extension of Wieland & Yang (2020) (`rr_2`).

All series are quarterly. For all our applications we consider the exact same sample, from 1969:Q1—2006:Q4, avoiding the zero lower bound on interest rates. We analyze variation induced by the ZLB in Appendix C.2.

ECONOMETRIC IMPLEMENTATION. We estimate the interacted local projections with $p = 4$ lags. We bias-correct the estimates following Herbst & Johannsen (2024), and because of our lag augmentation only use heteroskedasticity-robust standard errors (following Montiel Olea & Plagborg-Møller, 2021).

C.2 Interest rate lower bound

The main takeaway of Section 6.1 — on interacted monetary policy local projections identifying gradual policy treatments — is not limited to the instrument of Hack et al. (2026), but applies more generally. To this end we replicate Miyamoto et al. (2026), just now interpreted through the lens of our identification results. That paper studies shock propagation with and without a binding zero lower bound on the policy rate; the authors furthermore provide evidence that shock propagation is not materially different in booms versus recessions *per se*, suggesting that the estimated differences are chiefly attributable to differential monetary policy conduct. That said, since the required exclusion restriction on ζ_t is less plausibly satisfied than for Hack et al., we only consider this alternative as a brief further check.

DATA. We begin with the three outcome variables. Note that all data is for Japan.

- *Industrial production.* We compute log industrial production using data from the Finaeon Global Financial Database (GFD). The series name is “Japan Industrial Production Index (with GFD Extension),” NDJPNM.
- *Inflation.* We compute 3-month log-differences in the consumer price index, and then annualize, using data from Finaeon GFD. The series name is “Japan Consumer Price Index Inflation Rate (with GFD Extension),” CPJPNM.
- *Interest rate.* Following Miyamoto et al. (2026), the interest rate is the uncollateralized overnight call rate from July 1985, and the collateralized overnight call rate for months prior to that. Both data series are taken from the Bank of Japan website.

For the systematic monetary policy measure ζ_t , we use an indicator variable that takes the value of one if the economy is in the ZLB in period $t - 1$. Following Miyamoto et al. (2026), we define the ZLB periods for Japan as: 1995:M10—2006:M6 and 2009:M1—2019:M12.

As our shock measure $\varepsilon_{j,t}$, we use the oil shock series of Känzig (2021). Specifically, we use the “Oil supply news shock” series from the “Monthly (pre-Covid)” tab available in the 2020M12 update of the series. All series are monthly, and we use a sample from 1975:M1—2019:M12, covering periods with and without a binding interest rate lower bound.

ECONOMETRIC IMPLEMENTATION. We estimate the interacted local projections with $p = 6$ lags. All other settings are as discussed in Appendix C.1.




Figure C.1: Impulse responses of the policy rate, industrial production, and inflation to synthetic “as-if” monetary policy shocks based on comparisons across ZLB and non-ZLB periods. The figure plots differential responses of Japanese variables conditional on an oil shock, following Miyamoto et al. (2026). The shaded areas correspond to 16th and 84th percentile bands. Monthly data.

RESULTS. Figure C.1 re-affirms the main result of Section 6.1: the identified monetary policy treatment is yet again gradual and persistent. As before, the three panels of the figure display impulse responses to the identified “as-if” monetary policy shocks (i.e., the δ_h coefficients), with the left panel showing the policy instrument (“treatment”) path. Similar to the results reported in Figure 3, we see that the interest rate response yet again builds up gradually, here peaking after around one year. In response, industrial production falls significantly in the medium term; we also find evidence of a significant “price puzzle,” with inflation rising sharply in the near term.

C.3 Exchange rate propagation

Our final application closely follows the analysis of Fukui et al. (2025).

DATA. Most of the data we employ comes directly from the replication materials provided by Fukui et al. (2025). We begin with the three outcome variables.

- *Real effective exchange rates.* We use the `lreal` series, defined as the log real effective exchange rate. The data is sourced directly from Darvas (2021).
- *Ex-post real interest rate.* We use the `real_rate_all_bloomberg` series, defined as the nominal interest rate minus realized CPI inflation. The series was constructed by the authors from various sources.

D Proofs

D.1 Proof of Lemma 1 and Proposition 1

Both results are nested by the strictly more general Lemma 2 and Proposition 2, respectively, so we refer to those arguments.

D.2 Proof of Corollary 1

None of the arguments in either the equilibrium characterization of Lemma 1 or the identification result of Proposition 1 actually relied on the policy instrument being a scalar, so we can follow the exact same steps to complete the argument for Corollary 1. $\square$

D.3 Proof of Lemma 2

The existence of a solution to (12) follows from the assumed invertibility of $\bar{\Theta}_{z,v}$. The rest of the argument is then analogous to the proof of Proposition 1 in McKay & Wolf (2023), and leverages the fact that the private-sector block (5) considered here is identical to the one of that paper. $\square$

D.4 Proof of Proposition 2

By Frisch-Waugh-Lovell, and using the assumed independence of ζ_t and $\varepsilon_{j,t}$, the regression estimand for a generic outcome y_t satisfies

$$\delta_{y,h} \propto \text{Cov}(y_{t+h}, \varepsilon_{j,t} \zeta_t).$$

Next re-write y_{t+h} as

$$y_{t+h} = y_{t+h} - \mathbb{E}_t[y_{t+h}] + \mathbb{E}_t[y_{t+h}] - \mathbb{E}_{t-1}[y_{t+h}] + \mathbb{E}_{t-1}[y_{t+h}],$$

and observe that

$$\text{Cov}(y_{t+h} - \mathbb{E}_t[y_{t+h}], \varepsilon_{j,t} \zeta_t) = 0$$

since both $\varepsilon_{j,t}$ and ζ_t are measurable with respect to the history of shocks up to date t . Next, we also have that

$$\text{Cov}(\mathbb{E}_{t-1}[y_{t+h}], \varepsilon_{j,t} \zeta_t) = 0$$

since $\varepsilon_{j,t}$ is jointly independent of both ζ_t and all shocks that realized prior to date t . We are thus left with

$$\delta_{y,h} \propto \text{Cov}(\mathbb{E}_t[y_{t+h}] - \mathbb{E}_{t-1}[y_{t+h}], \varepsilon_{j,t}\zeta_t).$$

Now recall from (11) that

$$\mathbb{E}_t[y_{t+h}] - \mathbb{E}_{t-1}[y_{t+h}] = \bar{\Theta}_{y,h}\varepsilon_t + \bar{\Theta}_{y,v,h}\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots).$$

By the mutual independence of $\varepsilon_{j,t}$ from $(\varepsilon_{-j,t}, \zeta_t)$ it follows that

$$\delta_{y,h} \propto \bar{\Theta}_{y,v,h} \text{Cov}(\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots), \varepsilon_{j,t}\zeta_t).$$

With $\bar{\Theta}_{z,v}$ invertible, this finally delivers

$$\delta_{x,h} = \bar{\Theta}_{x,v,h}\bar{\Theta}_{z,v}^{-1}\delta_z.$$

Stacking across horizons h , the conclusion follows. $\square$

D.5 Proof of Lemma 3

The existence of a solution to (20) follows from the assumed invertibility of $\bar{\Theta}_{z,v}(m_t)$ for each possible value of m_t . The rest of the argument is again analogous to the proof of Proposition 1 in McKay & Wolf (2023), and leverages the fact that, conditional on the fixed m_t , the private-sector block (16) considered here is identical to the one of that paper. $\square$

D.6 Proof of Proposition 3

Using (19), we can proceed exactly as in the proof of Proposition 2 to get

$$\delta_{y,h} \propto \mathbb{E}[\bar{\Theta}_{y,v,h}(m_t)\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots, m_t)\varepsilon_{j,t}\zeta_t].$$

Re-write this as

$$\delta_{y,h} \propto \mathbb{E}[\bar{\Theta}_{y,v,h}(m_t)\mathbb{E}[\Omega(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots, m_t)\varepsilon_{j,t}\zeta_t \mid m_t]].$$

Denoting the interior expectation by $\omega_{\varepsilon_j,\zeta}(m_t)$, we can write this as

$$\delta_{y,h} \propto \mathbb{E}[\bar{\Theta}_{y,v,h}(m_t)\omega_{\varepsilon_j,\zeta}(m_t)].$$

Now define the weights $\mathbf{w}_z(m_t)$ via

$$\boldsymbol{\delta}_z \odot \mathbf{w}_z(m_t) = \frac{\bar{\Theta}_{z,v}(m_t) \boldsymbol{\omega}_{\varepsilon_j, \zeta}(m_t)}{\text{Var}(\varepsilon_{j,t} \zeta_t)}$$

It then follows that

$$\delta_{y,h} = \mathbb{E} [\bar{\Theta}_{y,v,h}(m_t) \bar{\Theta}_{z,v}(m_t)^{-1} \times (\boldsymbol{\delta}_z \odot \mathbf{w}_z(m_t))],$$

as claimed. Next, since

$$\boldsymbol{\delta}_z = \frac{\mathbb{E} [\bar{\Theta}_{z,v}(m_t) \boldsymbol{\omega}_{\varepsilon_j, \zeta}(m_t)]}{\text{Var}(\varepsilon_{j,t} \zeta_t)},$$

it also follows that $\mathbb{E} [\mathbf{w}_z(m_t)] = \mathbf{1}$. Finally, note that, by construction, the sign of $\mathbf{w}_z(m_t)$ equals the sign of

$$\frac{\text{Cov}(\mathbf{z}^t, \varepsilon_{j,t} \zeta_t \mid m_t)}{\text{Cov}(\mathbf{z}^t, \varepsilon_{j,t} \zeta_t)},$$

where the fraction here denotes an elementwise ratio. The sign of all entries of (23) being independent of m_t is sufficient to ensure positivity of this weight vector for all m_t . $\square$

D.7 Proof of Proposition A.1

We begin with some preliminary definitions. The variance-covariance matrices in the two policy regimes are given as

$$\Sigma = \bar{\Theta} \bar{\Theta}' + U \bar{\Theta}'_v + \bar{\Theta}_v U' + \bar{\Theta}_v S \bar{\Theta}'_v, \quad (\text{D.1})$$

$$\tilde{\Sigma} = \bar{\Theta} \bar{\Theta}' + \tilde{U} \bar{\Theta}'_v + \bar{\Theta}_v \tilde{U}' + \bar{\Theta}_v \tilde{S} \bar{\Theta}'_v, \quad (\text{D.2})$$

where tildes indicate the second regime, and U , S , $\tilde{U}$ and $\tilde{S}$ are defined as stated in Appendix A.2. Next, letting $B \equiv (\tilde{U} - U) + \frac{1}{2} \bar{\Theta}_v (\tilde{S} - S)$, we can write the variance-covariance gap Δ as

$$\Delta = \bar{\Theta}_v B' + B \bar{\Theta}'_v. \quad (\text{D.3})$$

(D.3) implies that $\text{col}(\Delta) \subseteq \text{span} \{ \bar{\Theta}_v, B \}$, and thus $\text{rank}(\Delta) \leq 2$, with equality being the generic case. It follows that we can write Δ in the following eigendecomposition form,

$$\Delta = V_+ \Lambda_+ V_+' - V_- \Lambda_- V_-' \quad (\text{D.4})$$

where $\{V_+, V_-\}$ are orthonormal, $\Lambda_+ > 0$ is the positive eigenvalue, and $\Lambda_- > 0$ is the absolute value of the negative eigenvalue.²⁷

Next note that, since $\bar{\Theta}_v$ lies in the column space of Δ , we can write it as

$$\bar{\Theta}_v = \alpha V_+ + \beta V_-$$

for some scalars α and β . Now let $z \in \mathbb{R}^{n_y}$ be any vector orthogonal to $\bar{\Theta}_v$, i.e., $z' \bar{\Theta}_v = 0$. It then follows from (A.2) that

$$z' \Delta z = 0$$

We now look for such a vector inside the column space of Δ . Consider the candidate

$$z = \beta V_+ - \alpha V_-.$$

By construction $z \perp \bar{\Theta}_v$, and so it follows that

$$z' \Delta z = \beta^2 \Lambda_+ - \alpha^2 \Lambda_- = 0.$$

This equation provides restrictions on $\{\alpha, \beta\}$. Specifically, the true $\bar{\Theta}_v$ will be proportional to one of the following two candidates:

$$\begin{aligned} \bar{\Theta}_{v,1} &= V_+ \Lambda_+^{1/2} + V_- \Lambda_-^{1/2}, \\ \bar{\Theta}_{v,2} &= V_+ \Lambda_+^{1/2} - V_- \Lambda_-^{1/2}. \end{aligned}$$

It remains to show that, for any $\bar{\Theta}_v^\dagger \propto \bar{\Theta}_{v,i}$ for $i = 1, 2$, we can find a tuple $\{\bar{\Theta}^\dagger, \bar{\Theta}_v^\dagger, \Omega^\dagger, \tilde{\Omega}^\dagger\}$ so that (D.1) - (D.2) hold. To this end define the projector on the column space of $\bar{\Theta}_v^\dagger$ as

$$\Pi_{\bar{\Theta}_v^\dagger} = \bar{\Theta}_v^\dagger [(\bar{\Theta}_v^\dagger)' \bar{\Theta}_v^\dagger]^{-1} (\bar{\Theta}_v^\dagger)'$$

and let

$$P_{\bar{\Theta}_v^\dagger} = I - \Pi_{\bar{\Theta}_v^\dagger}.$$

²⁷Since Δ is rank-2 by assumption and since the quadratic form defined by Δ vanishes on an $(n_y - 1)$ -dimensional subspace it follows that the quadratic form is indefinite, so one eigenvalue is positive, and the other negative.

Now note that, by definition of $\bar{\Theta}_v^\dagger$, we have that

$$P_{\bar{\Theta}_v^\dagger} \Delta P'_{\bar{\Theta}_v^\dagger} = 0$$

and thus

$$P_{\bar{\Theta}_v^\dagger} \tilde{\Sigma} P'_{\bar{\Theta}_v^\dagger} = P_{\bar{\Theta}_v^\dagger} \Sigma P'_{\bar{\Theta}_v^\dagger}.$$

It thus follows that, for some orthogonal matrix $R^* \in O(n_y)$, we have that

$$P_{\bar{\Theta}_v^\dagger} \text{chol}(\tilde{\Sigma}) R^* = P_{\bar{\Theta}_v^\dagger} \text{chol}(\Sigma).$$

Now consider setting

$$\bar{\Theta}^\dagger = \text{chol}(\Sigma), \quad \Omega^\dagger = \mathbf{0}', \quad U^\dagger = \bar{\Theta}^\dagger \Omega^{\dagger'}, \quad S^\dagger = \Omega^\dagger \Omega^{\dagger'}$$

as well as

$$\tilde{\Theta}^\dagger = \text{chol}(\tilde{\Sigma}) R^*.$$

Since by construction $\tilde{\Theta}^\dagger - \bar{\Theta}^\dagger$ lies in the column space of $\bar{\Theta}_v^\dagger$ we can find $\tilde{\Omega}^\dagger$, given as

$$\tilde{\Omega}^\dagger = [(\bar{\Theta}_v^\dagger)' \bar{\Theta}_v^\dagger]^+ (\bar{\Theta}_v^\dagger)' [\tilde{\Theta}^\dagger - \bar{\Theta}^\dagger]$$

such that

$$\tilde{\Theta}^\dagger = \bar{\Theta}^\dagger + \bar{\Theta}_v^\dagger \tilde{\Omega}^\dagger.$$

But this means that, with

$$\tilde{U}^\dagger = \bar{\Theta}^\dagger \tilde{\Omega}^{\dagger'}, \quad \tilde{S}^\dagger = \tilde{\Omega}^\dagger \tilde{\Omega}^{\dagger'}$$

we have found the desired tuple $\{\bar{\Theta}^\dagger, \bar{\Theta}_v^\dagger, U^\dagger, \tilde{U}^\dagger, S^\dagger, \tilde{S}^\dagger\}$, completing the argument. $\square$

D.8 Proof of Proposition A.2

We again begin with preliminary definitions. Forecasting variance-covariance matrices in the two regimes are now

$$\Sigma_t = \bar{\Theta} \bar{\Theta}' + U_t \bar{\Theta}_v' + \bar{\Theta}_v U_t' + \bar{\Theta}_v S_t \bar{\Theta}_v', \quad (\text{D.5})$$

$$\tilde{\Sigma}_t = \bar{\Theta} \bar{\Theta}' + \tilde{U}_t \bar{\Theta}_v' + \bar{\Theta}_v \tilde{U}_t' + \bar{\Theta}_v \tilde{S}_t \bar{\Theta}_v'. \quad (\text{D.6})$$

For all subsequent arguments we for simplicity drop the t subscripts. Again letting $B \equiv (\tilde{U} - U) + \frac{1}{2}\bar{\Theta}_v(\tilde{S} - S)$, we can also re-write this as

$$\Delta = \bar{\Theta}_v B' + B \bar{\Theta}_v'. \quad (\text{D.7})$$

Note that generically the rank of Δ here is even (equal to some $2m$, where $m \leq \min(n_\varepsilon, H)$), and half ($= m$) of the non-zero eigenvalues are positive, while the remaining ones ($= m$) are negative.²⁸ We will from now on assume that $m = H$, which in actual primitive structural models will generally be the case if n_ε is larger than H , i.e., economies that are subject to a large number of distinct shocks. Now as before we write the spectral decomposition of Δ as

$$\Delta = V_+ \Lambda_+ V_+' - V_- \Lambda_- V_-', \quad (\text{D.8})$$

where Λ_+, Λ_- are positive diagonal matrices, stacking the positive and negative (in absolute value) eigenvalues of Δ .

Since $\bar{\Theta}_v$ lies in the column space of Δ , we can write it as

$$\bar{\Theta}_v = V_+ A + V_- C$$

for $H \times H$ matrices A and C . Now let $Z \in \mathbb{R}^{(H \cdot n_y) \times H}$ be a rank- H matrix such that $\bar{\Theta}_v' Z = 0$, and thus

$$Z' \Delta Z = 0.$$

This is again the restriction that we will leverage for further tightening of the identified set of policy causal effects.²⁹ We now as in the static case look for candidate Z 's inside $\text{col}(\Delta)$, which we can write as

$$Z = V_+ X + V_- Y.$$

Since Z is rank- H we can furthermore without loss of generality restrict $X = I_H$. The

²⁸To see this, write

$$\Delta = C J C'$$

where $C \equiv (\bar{\Theta}_v \ B)$ and $J \equiv \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix}$, so J has H positive and H negative eigenvalues. It follows from Sylvester's law of inertia that the inertia of Δ equals the inertia of J restricted to $\text{col}(C)$, so Δ has an even number of non-zero eigenvalues ($2m$), and generically m are positive, and m are negative.

²⁹The rank- H assumption is without loss of generality since the below arguments are supposed to apply to all $z \in \mathbb{R}^{H \cdot n_y}$ in $\text{col}(\Delta)$ and such that $\bar{\Theta}_v' z = 0$; a rank- H Z allows us to span all such z 's.

requirement that $\bar{\Theta}'_v Z = 0$ is now equivalent to the matrix equation

$$A' + C'Y = 0.$$

We will from now on use the fact that A and C must be invertible.³⁰ It follows that the above matrix equation has the unique solution

$$Y = -(C')^{-1}A'.$$

Plugging this into $Z'\Delta Z = 0$ we find that

$$Z'\Delta Z = \Lambda_+ - Y'\Lambda_-Y = \Lambda_+ - (-(C')^{-1}A')'\Lambda_-(-(C')^{-1}A') = 0$$

and so

$$\Lambda_+ = AC^{-1}\Lambda_-(C')^{-1}A'$$

or

$$A'\Lambda_+^{-1}A = C'\Lambda_-^{-1}C.$$

Let $\mathcal{A} = \Lambda_+^{-1/2}A$ and $\tilde{C} = \Lambda_-^{-1/2}C$. Defining $Q = \tilde{C}\mathcal{A}^{-1}$, it is immediate that $Q \in O(H)$. Since $\tilde{C} = Q\mathcal{A}$ we have

$$C = \Lambda_-^{1/2}Q\mathcal{A}$$

and so

$$\bar{\Theta}_v = V_+A + V_-C = \left(V_+\Lambda_+^{1/2} + V_-\Lambda_-^{1/2}Q\right)\mathcal{A}$$

It now again remains to show that for any $\bar{\Theta}_v^\dagger$ that can be written in this way for some invertible $\mathcal{A}$ we can find a tuple $\{\bar{\Theta}^\dagger, \bar{\Theta}_v^\dagger, U^\dagger, \tilde{U}^\dagger, S^\dagger, \tilde{S}^\dagger\}$ so that (D.5) - (D.6) hold. To this end proceed as in the static case to note that

$$P_{\bar{\Theta}_v^\dagger}\Delta P'_{\bar{\Theta}_v^\dagger} = 0$$

and thus

$$P_{\bar{\Theta}_v^\dagger}\tilde{\Sigma}P'_{\bar{\Theta}_v^\dagger} = P_{\bar{\Theta}_v^\dagger}\Sigma P'_{\bar{\Theta}_v^\dagger}.$$

³⁰Suppose A is not invertible. Then there exists $u \neq 0$ such that $A'u = 0$. Now take $z = V_+u$, and note that $\bar{\Theta}'_v z = (V_+A + V_-C)'V_+u = A'u = 0$, so $z \perp \bar{\Theta}_v$ and $z'\Delta z = 0$. But $z'\Delta z = u'\Lambda_+u > 0$, so we have a contradiction. An analogous argument works for C .

It follows that, for some orthogonal matrix $R^* \in O(H \cdot n_y)$, we have that

$$P_{\bar{\Theta}_v^\dagger} \text{psd}(\tilde{\Sigma}) R^* = P_{\bar{\Theta}_v^\dagger} \text{psd}(\Sigma)$$

where $\text{psd}(\bullet)$ denotes any positive semi-definite factorization. Now consider setting

$$\bar{\Theta}^\dagger = \text{psd}(\Sigma), \quad \Omega^\dagger = \mathbf{0}', \quad U^\dagger = \bar{\Theta}^\dagger \Omega^{\dagger'}, \quad S^\dagger = \Omega^\dagger \Omega^{\dagger'}$$

as well as

$$\tilde{\Theta}^\dagger = \text{psd}(\tilde{\Sigma}) R^*.$$

Since by construction $\tilde{\Theta}^\dagger - \bar{\Theta}^\dagger$ lies in the column space of $\bar{\Theta}_v^\dagger$ we can find $\tilde{\Omega}^\dagger$, given as

$$\tilde{\Omega}^\dagger = [(\bar{\Theta}_v^\dagger)' \bar{\Theta}_v^\dagger]^+ (\bar{\Theta}_v^\dagger)' [\tilde{\Theta}^\dagger - \bar{\Theta}^\dagger]$$

such that

$$\tilde{\Theta}^\dagger = \bar{\Theta}^\dagger + \bar{\Theta}_v^\dagger \tilde{\Omega}^\dagger$$

But this means that, with

$$\tilde{U}^\dagger = \bar{\Theta}^\dagger \tilde{\Omega}^{\dagger'}, \quad \tilde{S}^\dagger = \tilde{\Omega}^\dagger \tilde{\Omega}^{\dagger'}$$

we have found the desired tuple $\{\bar{\Theta}^\dagger, \bar{\Theta}_v^\dagger, U^\dagger, \tilde{U}^\dagger, S^\dagger, \tilde{S}^\dagger\}$, completing the argument. $\square$