Identifying Policy Causal Effects from Rule Changes[†]

Ed Manuel — LSE

Christian K. Wolf — MIT & NBER

Preliminary and incomplete – comments welcome!

March 26, 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 this strategy and relate it to the classical policy shock literature. Our main result is a set of conditions on regressors and underlying data-generating process under which the two approaches are equivalent, i.e., the interacted local projection identifies the causal effects of a weighted average of policy shocks. In applications to monetary policy, this identified "as-if" policy shock does not, however, look like the literature's usual monetary shocks: while the standard shocks tend to yield transitory policy rate responses, the interacted local projection instead differences across more persistent rate paths, identifying the causal effects of gradual, forward guidance-like policy.

[†]Email: e.manuel@lse.ac.uk and ckwolf@mit.edu. We received helpful comments from Isaiah Andrews, Gabriel Chodorow-Reich, Marco Del Negro, Joe Hazell, Geert Mesters, 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 dynamic causal effects of changes in macroeconomic policy design. The most common approach to estimating policy causal effects is to regress the policy instrument as well as any macroeconomic outcomes of interest on so-called "policy shocks" (see Sims, 1980; Ramey, 2016). In typical structural macroeconomic models, the econometric estimand of such regressions furthermore admits a tight interpretation (by McKay & Wolf, 2023): it equals the causal effect on the macroeconomy of the particular dynamic policy instrument path induced by the shock.

Some recent work takes a different approach, leveraging time variation in the *systematic* conduct of policy. Here researchers run interacted regressions ("local projections", following Jordà, 2005) of the form

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

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., a zero lower bound on policy rates, or the composition of the monetary policy committee).[1] Intuitively, (1) asks how the dynamic propagation of the non-policy shock ε_t is shaped by systematic policy.

The principal contribution of this paper is to connect the two strategies, giving conditions under which (1) can be interpreted as a regression on a synthetic, "as-if" policy shock. A key assumption is that the propagation of ε_t changes with the regime measure ζ_t solely because of the conduct of policy, and not for other, extraneous reasons. If that is the case, then the coefficients δ_h in (1) identify policy causal effects: running the regression for the instrument identifies the policy "treatment," and doing so for other outcomes gives the associated causal effects—exactly like in policy shock regressions. Our second, more applied, contribution then concerns the nature of this synthetic shock in applications to monetary policy. Since typical non-policy shocks have gradual effects, and since measures of monetary regimes tend to be persistent, the identified "as-if" monetary shock moves policy rates gradually—a forward guidance "treatment" very different from the literature's typical shocks.

Environment. Our "as-if" identification result is derived under specific assumptions on the properties of the two regressors ε_t and ζ_t in (1), as well as on the underlying economic

[1] Recent well-known examples of studies running variants of (1) include Ramey & Zubairy (2018), Cloyne et al. (2023), Miyamoto et al. (2024), and Hack et al. (2024). See also the literature review.

environment. The baseline setting for our analysis is a general family of linearized macroeconomic models, with the 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), making it a natural starting point for our analysis. That said, a key restriction of this setting is that time variation in the propagation of non-policy shocks occurs *only* because of the changing design of policy. We later relax this assumption, discussing when and how our conclusions generalize beyond it.

We begin by showing that equilibrium outcomes in this model family admit a very simple characterization. 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 an arbitrary but time-invariant "baseline" policy rule, plus the product of policy causal effects times the deviation of actual policy from baseline. Importantly, these policy causal effects are invariant to any assumptions about policy: they are simply the uniquely defined effects of any current and expected future path of the policy instrument on the macro-economy. This characterization reflects two key properties of our environment: first, that policy shapes the economy only through current and expected future paths of the policy instrument (i.e., "instrument sufficiency", following McKay & Wolf, 2023); and second, that policy is the only source of time variation in the propagation of any given shock ε_t. We now consider an econometrician that lives in this economy and runs variants of the interacted local projections (1).

Interpreting the interacted LP (1). Our first main identification result gives conditions on the regressors (ε_t, ζ_t) such that (1) can be interpreted as a regression on a synthetic, "as-if" policy shock. The formal result supposes that ε_t is a primitive shock, that the policy measure ζ_t is independent of ε_t (but otherwise possibly endogenous), and finally that their interaction $\varepsilon_t \cdot \zeta_t$ forecasts future values of the policy instrument, i.e., that the interaction coefficients δ_h for the policy instrument are non-zero. Under those assumptions, the δ_h's for the instrument are the (non-zero) treatment, and the interaction coefficients for all other

[2]The 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 (see McKay & Wolf, 2023).

macroeconomic outcomes are the associated causal effects. Importantly, the synthetic "as-if" policy shock identified by the interacted local projection will generically change if (1) is run for two distinct non-policy shocks $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$: intuitively, ζ_t is likely to differentially revise the policy path following $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$, thus isolating two distinct policy treatments.

The proof of the identification result highlights the role played by the shock ε_t: it serves as a conditioning event that itself is not shaped by the policy regime, and yet whose propagation changes with the regime. In our environment, the first of those two requirements can be violated; for example, using a reduced-form Wold innovation in lieu of a true structural shock in (1) will not work, as that innovation itself may be shaped directly by policy. Insisting that ε_t is a shock—or a combination of shocks, or more generally any macroeconomic news not shaped by the regime measure—ensures the first requirement. The second requirement, in contrast, is automatic, as policy by assumption is the *only* channel generating time or state dependence in shock propagation. To speak to this margin we generalize the environment.

Richer forms of state dependence. In practice the propagation of the shock ε_t may vary with the policy measure ζ_t not only because of changes in policy, but also because of possible other channels that are correlated with ζ_t. To explore this threat, we generalize the model environment to allow for the two components in the date-t expectation revisions of macroeconomic outcomes to depend on a further measure of an aggregate state.

Our second identification result states that the synthetic shock logic extends to this richer environment, but only under more stringent assumptions on the policy measure ζ_t: beyond the conditions stated above, we now also require that ζ_t is orthogonal to the additional time-varying aggregate state. This means that ζ_t can still be endogenous, e.g., with systematic policy design changing in response to past economic conditions; importantly, however, those other factors must not further shape the propagation of ε_t through non-policy channels. If so, and under a further monotonicity assumption that closely echoes those made in the literature on heterogeneous treatment effects (e.g., see Imbens & Angrist, 1994), the interacted local projection (1) is again equivalent with an "as-if" policy shock regression.[3]

Synthetic monetary shocks in practice. With our results on required regressor properties and estimand interpretation in hand, we next turn to a more practical question:

[3]Finally, we also provide identification results in an even more general, essentially unrestricted non-linear data-generating process (following Kolesár & Plagborg-Møller, 2025). Under assumptions on the regressors (ε_t, ζ_t) close to those for our main results, the estimand of (1) is a positive weighted average of cross partials of the potential outcomes function. Our main results are an informative special case of this characterization.

In the context of monetary policy, what “as-if” policy shock 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?

We begin with theoretical arguments for why the identified synthetic monetary policy shock will in practice likely look quite different from the textbook notion of a one-off monetary shock. First, typical non-policy shocks ε_t tend to have persistent aggregate effects and thus, for standard monetary policy rules, induce persistent movements in the policy rates. Second, monetary regimes ζ_t themselves are persistent and often feature gradualism in interest rate setting. Putting the two together, the *difference* in rate movements in response to ε_t across regimes is also persistent—i.e., a gradual, forward guidance-like policy treatment. We provide an illustration through simulations in models with policy regime changes, based on Smets & Wouters (2007) as well as Davig & Leeper (2007): for each non-policy shock ε_t, running (1) isolates a different policy “treatment,” but in all cases that treatment is persistent.

We next confirm these predictions in empirical applications revisiting Hack et al. (2024) and Miyamoto et al. (2024). Hack et al. identify a measure of systematic U.S. monetary policy conduct—the Hawk-Dove balance of the FOMC, for additional robustness instrumented based on FOMC voting rights rotations—that plausibly governs shock propagation chiefly through its implications for policy conduct. Interacted with oil, fiscal, and technology shocks, that measure predicts significant, gradual responses of the policy rate, i.e., forward guidance-like policy treatments. Across all the non-policy shocks ε_t that we look at, delayed interest rate hikes tend to induce similarly delayed contractions in output, while price responses are more muted. Through the lens of our identification results, those findings provide a useful addition to our understanding of monetary policy causal effects: while existing shocks are mostly informative about the causal effects of transitory policy rate revisions (see Ramey, 2016), the estimands here instead are the causal effects of interest rate forward guidance. Finally, we report similar conclusions replicating Miyamoto et al., who compare oil shock propagation with and without a binding lower bound on policy rates.

Do we actually require ε_t? Our results so far leveraged two distinct inputs: a conditioning event ε_t that is invariant to policy, plus a policy regime measure ζ_t that shapes the propagation of ε_t only through policy. A natural question is whether the first of those requirements can be relaxed, asking instead how ζ_t shapes the unconditional second moments, i.e., autocovariances, of macroeconomic aggregates. An early antecedent of such a strategy is Mussa (1986), who uses the effect of changes in the exchange rate regime on macroeconomic volatilities to reject policy neutrality. More recently, the literature on shock identification through heteroskedasticity (Rigobon, 2003; Lewis, 2025) has demonstrated that changes in the volatility of policy shocks alone are enough to point-identify those shocks' causal effects through the induced time variation in unconditional second moments.

Unfortunately, our final identification result is negative, establishing that the effects of policy rule changes on second moments yield non-trivial, but often excessively wide, identified sets for policy causal effects. Intuitively, while identification through heteroskedasticity works because shock causal effects are just scaled up or down, rule changes affect second moments by altering the dynamic propagation of every single non-policy shock. This increases the complexity of the change in implied second moments, precluding (point) identification.

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 well-known Kitagawa-Blinder-Oaxaca 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 prove equivalence with conventional policy shock regressions. Those same contributions also distinguish our work from Gonçalves et al. (2024) and Kolesár & Plagborg-Møller (2025), who both study state-dependent local projections.

Zooming out, our paper is situated in a broader literature that exploits stability restrictions for identification (e.g., see Lewbel, 2012; Magnusson & Mavroeidis, 2014; Mavroeidis, 2021, in addition to the heteroskedasticity references above). We leverage such ideas in the specific context of changes in policy design, and characterize identified sets for a large class of structural macroeconomic models as candidate data-generating processes. Our arguments use, and in fact can be seen as an inversion of, the policy counterfactual identification results of McKay & Wolf (2023), Barnichon & Mesters (2023), and Caravello et al. (2025).

Finally, there is a close relationship between (1) and the "shift-share" regression specifications common in the applied microeconometric 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, by policy regime. Studying the effects of exchange rate depreciations and terms of trade changes, Broda (2004) and Fukui et al. (2025) similarly leverage variation in policy regimes, but they look across countries, not time. Our identification results build a bridge from those strategies to the classical time series approach of regressing on policy shocks.

2 A simple example

We begin with a simple two-period example that illustrates the core economic logic of our identification results. Section 2.1 describes the economy and provides a constructive equilibrium characterization, previewing the logic of our general infinite-horizon equilibrium characterization in Section 3. Then, in Section 2.2, we show how knowledge of non-policy shock propagation across policy regimes identifies policy causal effects. This is the core insight at the heart of our later, most general identification results in Section 4.

2.1 Model and equilibrium characterization

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

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

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

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

where x_t denotes output, i_t is the policy rate, π_t is inflation, and (e_t^d, e_t^s) denote demand and supply shocks, respectively. The second block is the policy block, which here is simply a standard Taylor-type rule:

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

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 $\varepsilon_t \overset{\text{iid}}{\sim} (0, I_{n_\varepsilon})$. 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. In the baseline policy 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 $\phi_\pi^H > \phi_\pi$, and we denote all impulse responses by H superscripts. It now follows directly from McKay & Wolf (2023, Proposition 1) that those impulse responses satisfy

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

as well as

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

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.

We emphasize that the decompositions (2) - (3) are constructed by comparing across two policy regimes that are each expected to be in place forever. Our more general equilibrium characterization results in Section 3 will arrive at similar additive decompositions, but there in a richer environment with stochastic, potentially endogenous policy regime switches.

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.

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 (see Caravello et al., 2025), proceeding as follows. Let

$$\Theta_i \;\equiv\; \bar{\Theta}_v \cdot \bar{\Theta}_{i,v}^{-1}, \tag{4}$$

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. Importantly, because of instrument sufficiency, the policy causal effect matrix Θ_i in (4) 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 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.

We also note that, even if the econometrician is unable to recover the underlying impulse response functions in shock space—i.e. $\bar{\Theta}_v$—knowledge of policy causal effects in instrument space—i.e. Θ_i—is all she needs for both model estimation through impulse response matching (as in Christiano et al., 2005) and for policy counterfactual evaluation (McKay & Wolf, 2023). The object of interest is thus from now on policy effects in instrument 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"

$$w_t = \omega_{w,0} v_{0,t} + \omega_{w,1} v_{1,t} + \text{noise}, \tag{5}$$

and then, for $h = 0, 1$ and (with slight abuse of notation) generic outcome y_t, runs

$$y_{t+h} = \text{constant} + \beta_{y,h} w_t + \text{error}. \tag{6}$$

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

[4] To be precise, the policy rule does play one additional role: equilibrium selection. Just like in McKay & Wolf (2023), our arguments here apply for switches across policy rules with identical equilibrium selection.

to a weighted average of columns of Θ_i with the policy rate impulse response as the weights.

We note two important challenges with this familiar "policy shock" approach. First, as is well-known, it is often difficult to come up with credible policy shock measures. Second, and somewhat less widely appreciated, another important problem 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 at best limited (see the literature review for monetary policy shocks in Caravello et al., 2025). We thus now turn to an alternative strategy for learning about Θ_i.

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

$$y_{t+h} = \text{constant} + \beta^j_{y,h}\varepsilon_{j,t} + \gamma^j_{y,h}\zeta_t + \delta^j_{y,h}\left(\varepsilon_{j,t}\cdot\zeta_t\right) + \text{error}, \quad j = d, s, \tag{7}$$

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), and also follows from the much more general identification result in Proposition 2 below, that $\boldsymbol{\delta}^j_y \equiv (\delta^j_{y,0}, \delta^j_{y,1})'$ satisfies

$$\boldsymbol{\delta}^j_y = \Theta^H_{y,j} - \bar{\Theta}_{y,j} \tag{8}$$

and thus

$$\boldsymbol{\delta}^j_y = \bar{\Theta}_{y,v}\cdot\boldsymbol{\omega}_j. \tag{9}$$

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 $\boldsymbol{\omega}_j \equiv (\omega_{j,0}, \omega_{j,1})'$ 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

$$\boldsymbol{\delta}^j_y \quad = \quad \underbrace{\Theta_{y,i}}_{\text{policy causal effects}} \quad \times \quad \boldsymbol{\delta}^j_i.$$

In words, first running the interacted local projection (7) for the policy instrument gives the "treatment" $\boldsymbol{\delta}^j_i$, and the impulse responses for the other variables then give the associated causal effects. This is the core idea of our "as-if" identification results: just like regressions on policy shocks, interacted local projections deliver policy causal effects. We note further that the weights $\boldsymbol{\omega}_w$ in a policy instrumental variable w_t need not, and in practice generally will not, agree with the "as-if" weights $\boldsymbol{\omega}_j$ isolated by (7). 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 (7) 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 (7) twice, getting a richer picture of how different policy treatments affect the macro-economy.[5]

Discussion and outlook. The preceding analysis has previewed two of this paper's key messages: that interacted local projections like (7) can, in a particular sense, be informative about policy causal effects in exactly the same way as is the case for standard policy shocks; and that such regressions may offer useful complementary information, as they could feasibly isolate different policy "treatments."

That said, our discussion was rather restrictive, in three key respects. 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 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 two sections will tackle all of these challenges by generalizing our analysis to a general class of

[5]In fact, it is immediate that (7) would also recover 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 Section 4.1, the key is simply that $\varepsilon_{j,t}$ is an event that is stable across policy regimes.

infinite-horizon structural models with essentially unrestricted variation over time in policy design, together with rich forms of non-policy state dependence in shock propagation. The practical upshot will be conditions on the regressors ε_t and ζ_t in local projections like (7) that allow the takeaways of this section to survive.

3 Main environment

We begin in Section 3.1 with an outline of the general model environment, including a discussion of its scope and limitations. We then in Section 3.2 show that equilibrium outcomes can be characterized using the exact same two-part additive decomposition as in our simple toy example in Section 2.

3.1 Model outline

Our general environment features a linear but otherwise very general non-policy block, and essentially unrestricted variation over time in how the policymaker sets her instrument. Relative to the simple example in Section 2, we thus extend the analysis in two key respects: to a richer *family* of non-policy blocks, and to stochastic, perhaps endogenous, time variation in policy, rather than two permanent regimes. The final extension to other, non-policy forms of state and time dependence in shock propagation will follow later, in Section 4.3.

We proceed in three steps: we first describe the environment, then for illustrative purposes give an explicit example of a particular nested model economy, and finally provide a more general discussion of the kind of primitive models our setting nests (and does not nest). Even though we will throughout consider stochastic economies subject to recurring shocks, our formulation of the equilibrium conditions and dynamics will heavily leverage insights from recent work on so-called "sequence-space" methods (e.g., see Boppart et al., 2018; Auclert et al., 2021). Throughout this section, date-t expectation operators are defined with respect to the filtration induced by the history of shocks and policy regimes up to date t.

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

$$\mathcal{H}_x \left[\mathbb{E}_t \left(\boldsymbol{x}^t\right) - \mathbb{E}_{t-1} \left(\boldsymbol{x}^t\right)\right] + \mathcal{H}_z \left[\mathbb{E}_t \left(\boldsymbol{z}^t\right) - \mathbb{E}_{t-1} \left(\boldsymbol{z}^t\right)\right] + \mathcal{H}_e e_t = \mathbf{0}. \tag{10}$$

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 a n_e-dimensional vector of non-policy shocks. Boldface denotes time paths, i.e., for any scalar w we have $\boldsymbol{w}^t = (w_t, w_{t+1}, \dots, w_{t+H-1})'$, and H is the (possibly infinite) maximal horizon. Equation (10) 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 in the scalar policy instrument z_t as a function of exogenous disturbances to the system:

$$\mathbb{E}_t\left(\boldsymbol{z}^t\right) - \mathbb{E}_{t-1}\left(\boldsymbol{z}^t\right) = \mathcal{A}\left(e_t, e_{t-1}, \dots, v_t, v_{t-1}, \dots, s_t, s_{t-1}, \dots\right), \tag{11}$$

where v_t is a n_v-dimensional vector of policy shocks, and s_t is a scalar policy regime variable. Note that (11) is a solved-out, reduced-form mapping, and not an explicit, primitive policy rule; we will discuss later why we specify policy in this way. We assume that the shocks $\varepsilon_t \equiv (e_t', v_t')'$ are distributed as $\varepsilon_t \overset{\text{iid}}{\sim} (0, I_{n_\varepsilon})$, and the policy regime indicator s_t is a random variable independent of those shocks, but otherwise unrestricted.

We are now in a position to define an equilibrium in this economy.

Definition 1. *Given initial conditions $\mathbb{E}_{-1}\left(\boldsymbol{z}^0\right) = \mathbf{0}$ and $\mathbb{E}_{-1}\left(\boldsymbol{x}^0\right) = \mathbf{0}$, and given any realization of the exogenous disturbances $\{\varepsilon_t\}_{t=0}^T$ and the policy regime $\{s_t\}_{t=0}^T$, an equilibrium of the model is given by sequences of expectation revisions $\mathbb{E}_t\left(\boldsymbol{z}^t\right) - \mathbb{E}_{t-1}\left(\boldsymbol{z}^t\right)$ and $\mathbb{E}_t\left(\boldsymbol{x}^t\right) - \mathbb{E}_{t-1}\left(\boldsymbol{x}^t\right)$ such that (10) - (11) hold for all dates $t = 0, \dots, T$. The realized sequences $\{z_t, x_t\}_{t=0}^T$ follow, for each date t, from summing expectation revisions up to t.*

Note that the equilibrium system (10) - (11) is written in sequence-space, expectation revision form. Definition 1 then uses the implied equilibrium expectation revisions to construct realized outcomes, translating from impulse response form to those actual outcomes by proceeding exactly as in Wolf (2025, Appendix D).

An example. For a simple example of a model environment that can be written in the form (10) - (11), consider as in Section 2 a variant of the textbook New Keynesian model, but

[6] In standard stochastic dynamic economies, it is well-known that policy instrument pegs typically deliver equilibrium indeterminacy, corresponding in the notation of (10) 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 (11), exactly as in McKay & Wolf (2023).

now augmented to feature stochastic policy rule switches. Beginning with the private-sector block, the IS and NKPC relations are

$$y_t = -\frac{1}{\gamma}\left(i_t - \mathbb{E}_t\left[\pi_{t+1}\right]\right) + \mathbb{E}_t\left[y_{t+1}\right] + e_t^d, \tag{IS}$$

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t\left[\pi_{t+1}\right] + e_t^s. \tag{NKPC}$$

At any given date t, stacking those relations from dates t to $t + H - 1$, we obtain the private-sector block (10). Turning to policy, an example of an explicit interest rate rule with exogenous time-variation is (following Davig & Leeper, 2007)

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

where s_t as in our general environment is a regime indicator that is independent of other shocks, $\phi_\pi(\bullet)$ gives regime-specific policy rule coefficients, and v_t is the monetary policy shock. Assuming equilibrium existence and uniqueness of the system (IS) - (TR'), we obtain a mapping from $\varepsilon_t \equiv (e_t^d, e_t^s, v_t)$ and s_t to i_t—an example of the general reduced-form policy mapping (11) in our general model. Alternatively, an example of endogenous policy feedback would be the adjusted Taylor rule (following Davig & Leeper, 2008)

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

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 private-sector outcomes from the private-sector block.

Scope and limitations. For our purposes, the key property of the general environment (10) - (11) 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.

What kind of assumptions on economic primitives are consistent with this assumed model property? McKay & Wolf, in the context of time-invariant policy rules, establish that 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 certain kinds of departures from rationality (e.g., as in Gabaix, 2020).[7] In richer settings with stochastic regime switches, standard first-order perturbations in aggregate shock volatilities yield linearized private-sector blocks that do not depend on the policy rule as long as all possible policy rules are consistent with the same deterministic steady state (e.g., Farmer et al., 2009; Barthélemy & Marx, 2019). As in the two simple examples above, our setting is general enough to nest models of both exogenous and endogenous rule switches. 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 non-linearities in policy-setting (e.g., a zero lower bound).

To summarize, our environment (10) - (11) is a natural generalization of the rule-invariant system studied in McKay & Wolf (2023): we make the same instrument sufficiency assumption on the private sector, but now allow for largely arbitrary stochasticity in policy design. To make such stochasticity consistent with invariance of the private-sector block, the key and meaningful additional assumption we require is that the policy rule changes leave the system's deterministic steady state unchanged. This makes good on our claims in the introduction: 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 (Davig & Leeper, 2007; Bianchi, 2013), including also regime changes induced by non-linearities such as a binding lower bound on interest rates (Guerrieri & Iacoviello, 2015).

3.2 Equilibrium characterization

The objective of this section is to arrive at an additive equilibrium decomposition analogous to that of (2) - (3) in the simple example of Section 2. To this end we begin with a definition of what we will call a "baseline" policy rule and of policy causal effects.

Baseline policy and policy causal effects. Consider the following time-invariant, linear, baseline policy rule:

$$\bar{\mathcal{A}}_x \left[\mathbb{E}_t \left(\boldsymbol{x}^t\right) - \mathbb{E}_{t-1} \left(\boldsymbol{x}^t\right)\right] + \bar{\mathcal{A}}_z \left[\mathbb{E}_t \left(\boldsymbol{z}^t\right) - \mathbb{E}_{t-1} \left(\boldsymbol{z}^t\right)\right] + \bar{\mathcal{A}}_\nu \nu_t = \mathbf{0}. \tag{11'}$$

[7] The 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 $\mathcal{H}_\bullet$ to be shaped by the policy rule itself.

The sole purpose of this baseline rule will be to serve as a reference point (i.e., basis) for our equilibrium decomposition. If policy in our economy were (counterfactually) set following (11') rather than (11), and further assuming that the rule (11') implies existence of a unique equilibrium, equilibrium outcomes could be written as

$$\mathbb{E}_t\left(\bar{\boldsymbol{y}}^t\right) - \mathbb{E}_{t-1}\left(\bar{\boldsymbol{y}}^t\right) = \bar{\Theta}\varepsilon_t. \tag{12}$$

(12) is a vector moving average representation, with the $(H \cdot n_y) \times n_\varepsilon$ matrix $\bar{\Theta}$ collecting the impulse responses under the baseline rule. Finally we follow McKay & Wolf and define $\bar{\Theta}_v$ as the impulse responses of shocks to the baseline policy rule (11') *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 indicate impulse responses for the corresponding variables. We will throughout assume that $\bar{\Theta}_{z,v}$ is invertible, i.e., that, through policy shocks, any possible path of the policy instrument can be implemented.

The equilibrium decomposition. We now arrive at the following equilibrium characterization, generalizing our earlier additive decomposition in (2) - (3).

Proposition 1. *Equilibrium outcomes of the model economy* (10) - (11) *can be written as*

$$\mathbb{E}_t\left[\boldsymbol{y}^t\right] - \mathbb{E}_{t-1}\left[\boldsymbol{y}^t\right] = \bar{\Theta}\varepsilon_t + \bar{\Theta}_v\Omega(\varepsilon_t, \varepsilon_{t-1}, \ldots, s_t, s_{t-1}, \ldots) \tag{13}$$

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

$$\begin{aligned} \mathbb{E}_t\left[\boldsymbol{z}^t\right] - \mathbb{E}_{t-1}\left[\boldsymbol{z}^t\right] &= \bar{\Theta}_z\varepsilon_t + \bar{\Theta}_{z,v}\Omega(\varepsilon_t, \varepsilon_{t-1}, \ldots, s_t, s_{t-1}, \ldots) \\ &= \mathcal{A}\left(e_t, e_{t-1}, \ldots, v_t, v_{t-1}, \ldots, s_t, s_{t-1}, \ldots\right). \end{aligned} \tag{14}$$

Proposition 1 decomposes equilibrium outcomes into date-t revisions under the baseline policy plus the product of policy causal effects $\bar{\Theta}_v$ times a policy gap. The intuition is simple but instructive. If policy was counterfactually set according to the baseline rule (11'), 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 (11'), but instead satisfies (11). The H-dimensional vector $\Omega(\bullet)$ in (14) is an artificial "shock" to the baseline rule (11') *defined* to ensure that the policy path in (13) satisfies (11), thus making that part of the proposition tautological. By instrument sufficiency, once z_t is set in that way, the outcomes for x_t stated in (13) must also agree with true equilibrium outcomes. This is the same logic as in Section 2, now extended to a richer economy with stochastically, and perhaps endogenously, changing policy rules.

A key and obvious limitation of (13) is that, just as in our simple example, variation over time in the design of policy is the *only* source of time or state dependence in the propagation of the macroeconomic shocks ε_t. We will for now continue working under this assumption, but later, in Section 4.3 further generalize the environment to allow for a rich set of further, non-policy reasons for such time or state dependence.

Outlook. The upshot of the analysis in this section is that a general family of linearized macroeconomic models with stochastic policy rule switches admits a very simple equilibrium characterization—a characterization in which policy causal effects $\bar{\Theta}_v$ feature prominently. The remainder of this paper will use this characterization, as well as its extensions in further generalized model variants, to provide a structural interpretation of the estimand of interacted local projections like (1), tying that estimand to $\bar{\Theta}_v$ and thus also to textbook policy shock regressions, exactly as in the earlier simple example.

4 Identification with interacted local projections

We now turn to our main results on interpretation of the estimand of interacted local projections of the form (1). Section 4.1 gives the conditions on the included regressors so that the estimand admits a tight "as-if" policy shock interpretation. In Section 4.2 we then argue that, in practice, this synthetic policy shock is likely to correspond to a gradual, persistent policy treatment. Finally, extensions of our results to multiple policy treatments as well as to more general data-generating processes follow in Section 4.3.

4.1 Identification results

To set the stage, just as in Section 2, we begin by reviewing the estimand of policy shock regressions. We then turn to our main identification results on interacted local projections.

The estimand of policy shock regressions. Suppose a researcher observes, or through her identifying assumptions is able to isolate, an instrumental variable w_t that satisfies

$$w_t = \sum_{\ell=0}^{H-1} \omega_{w,\ell} v_{t+\ell,t} + \text{noise} \tag{15}$$

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

$$y_{t+h} = \text{constant} + \beta_{y,h} w_t + \text{error}, \tag{16}$$

it follows that the β_y's give us the causal effects of the policy shock mixture $\{\omega_{w,\ell}\}$ on the macro-economy.[8] In particular, if w_t is a contemporaneous policy shock (i.e., if $\omega_{w,0} = 1$, $\omega_{w,\ell} = 0$ for $\ell > 0$), the econometric estimand is simply the first column of $\bar{\Theta}_v$. Equivalently, expressing things not in terms of policy shocks but policy instruments, running (16) for the policy instrument z_t gives the "treatment" $\boldsymbol{\beta}_z \equiv (\beta_{z,0}, \beta_{z,1}, \dots)'$, and doing the same for macroeconomic outcomes x_t gives the associated causal effects. All of this discussion is exactly analogous to Section 2, just here presented for a more general model class.

Interacted local projections. We now turn to our actual object of interest: interacted local projections. Specifically, we consider an econometrician living in the dynamic economy (10) - (11), and running the following regression specification:

$$y_{t+h} = \alpha_{y,h} + \beta_{y,h}\varepsilon_{i,t} + \gamma_{y,h}\zeta_t + \delta_{y,h}\left(\varepsilon_{i,t} \cdot \zeta_t\right) + u_{t,y,h}, \tag{17}$$

where $\varepsilon_{i,t}$ is the ith primitive structural shock, and ζ_t is a (currently otherwise unrestricted) date-t random variable. As before we will pay particular attention to (17) run for both the policy instrument z_t and macroeconomic outcomes x_t on the left-hand side.

Proposition 2. *Consider the dynamic stochastic economy* (10) - (11) *and suppose that the random variable* ζ_t *has mean zero and is given as*

$$\zeta_t = \zeta\left(\varepsilon_{-i,t}, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots\right) \tag{18}$$

[8] To be precise, relative to the model setting in Section 3, the informal discussion here implicitly supposes that policy is set according to the time-invariant rule (11'). Nothing of essence would change, however, with stochastic rule switches: the regression estimand would now reflect policy causal effects because of both the first and second parts in the additive decomposition in Proposition 1.

Then the population estimands of (17) satisfy

$$\boldsymbol{\delta}_x = \underbrace{\bar{\Theta}_{x,v} \cdot \bar{\Theta}_{z,v}^{-1}}_{\text{policy causal effects}} \times \underbrace{\boldsymbol{\delta}_z}_{H \times 1 \text{ vector of treatments}} \tag{19}$$

The remainder of this section is dedicated to interpretation of this proposition. We first discuss the estimand and its connection to policy shock regressions, before then digging deeper into the required assumptions on the regressors $\varepsilon_{i,t}$ and ζ_t.

Estimand interpretation. To understand the estimand of (17) and relate it to policy shock regressions, it will prove useful to follow some of the steps of the proof of Proposition 2. By Frisch-Waugh-Lovell, and using the assumed independence of ζ_t and $\varepsilon_{i,t}$, the regression estimand for outcome y_{t+h} satisfies

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

Since $\varepsilon_{i,t}$ is a date-t shock, and since both $\varepsilon_{i,t}$ 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.,

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

This expectation revision, however, is exactly the object we characterized in Section 3, given as

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

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

$$\delta_{y,h} \propto \bar{\Theta}_{y,v,h}\,\text{Cov}\left(\Omega\left(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots\right), \varepsilon_{i,t}\zeta_t\right) \tag{20}$$

(20) is the key relation, implying that the interaction coefficients $\boldsymbol{\delta}_y$ correspond to the dynamic causal effects of a particular synthetic, "as-if" $H \times 1$ policy shock vector, given as

$$\text{Cov}\left(\Omega\left(\varepsilon_t, \varepsilon_{t-1}, \dots, s_t, s_{t-1}, \dots\right), \varepsilon_{i,t}\zeta_t\right) \tag{21}$$

The actual statement of Proposition 2 then simply moves from policy shock space to policy instrument space: running (17) for the instrument z_t gives the policy "treatment" $\boldsymbol{\delta}_z$, and then running it for an outcome x_t gives the associated causal effects. Interpretation of that interacted local projection is therefore, under our stated assumptions, isomorphic to how we interpret regressions on policy shocks, as claimed.[9] This thus generalizes all insights from Section 2: interacted local projections are an *alternative* to policy shock regressions because they share the same kind of estimand but rely on different informational requirements; and a *complement* because the policy treatment they isolate may differ from typical policy shocks. We will provide theoretical guidance on this question of possible differences in the identified policy treatment in Section 4.2, before turning to actual applications in Section 5.

Properties of the shock measure $\varepsilon_{i,t}$. We next turn to regressor properties, beginning with the "shock" variable $\varepsilon_{i,t}$. Under the assumptions stated so far, this regressor indeed is one of the various primitive structural shocks hitting the economy (10) - (11). Intuitively, the role of $\varepsilon_{i,t}$ in Proposition 2 is as a "conditioning event" that itself is invariant to the other regressor ζ_t, but whose dynamic propagation is shaped by ζ_t. This logic suggests—and it is straightforward to prove—that linear combinations of date-t shocks ε_t *with fixed weights*, or more generally any date-t news independent of ζ_t would also yield interacted local projection estimands that equal "as-if" policy shocks. We will further leverage these observations in Section 5, where we argue that our conclusions apply even if our included regressors are not narrow measures of true macroeconomic shocks.

By the same token, however, date-t innovations that themselves are shaped by the other regressor ζ_t would not work, as here both dynamic propagation and the conditioning event itself vary with ζ_t. The easiest way to see the problem is to consider a simpler model variant in which there are two very long-lasting policy regimes (akin to Section 2), and the researcher runs (17) with ζ_t indicating the policy regime, and the other regressor equal to the ith Wold innovation $u_{i,t}$ in each regime. Assuming for simplicity invertibility (see Fernández-Villaverde et al., 2007) in each of the two regimes, the regressor $u_{i,t}$ in regime j satisfies

$$u_{i,t} = p'_{i,j}\varepsilon_t$$

where $p_{i,j}$ is an n_ε-dimensional vector. The "diff-in-diff" logic underlying the proof of Proposition 2 now breaks down: we are comparing across different policy regimes how *different* mixes of structural shocks, $p'_{i,1}\varepsilon_t$ vs. $p'_{i,2}\varepsilon_t$, propagate, and so the computed difference reflects

[9] Just like for policy shocks, interpreting regression estimands in policy shock space is necessarily dependent on the assumed (baseline) policy rule, here reflected in the presence of $\bar{\Theta}_v$ in (20). Moving to instrument space, because of instrument sufficiency, drops that dependence of the baseline rule: the product $\bar{\Theta}_{x,v} \cdot \bar{\Theta}_{z,v}^{-1}$ is independent of the rule, simply giving us the rule-invariant causal effects of arbitrary policy paths.

both policy causal effects as well as the heterogeneous conditioning event.

Properties of the policy measure ζ_t. We next consider the policy measure ζ_t. Two properties of that random variable are crucial to the interacted local projection indeed teaching us about policy causal effects. First, the interaction of $\varepsilon_{i,t}$ with ζ_t needs to *predict* the policy instrument: if it does so, then $\boldsymbol{\delta}_z \neq 0$, and so the regression indeed isolates a non-zero policy treatment. This is simply a standard relevance condition. Second, ζ_t needs to shape the propagation of $\varepsilon_{i,t}$ *only* through policy. This second condition is automatically satisfied in our family of data-generating processes, simply because policy by assumption is the only source of time variation in the propagation of $\varepsilon_{i,t}$.

An important takeaway of our analysis is that, in the environment considered so far, the two stated requirements on ζ_t can be entirely consistent with ζ_t being an *endogenous* correlate of systematic policy. For example, if some past history of inflationary shocks increases the policymaker's concern for inflation, then correlates of such shocks will predict future policy, and yet they will by construction not shape the propagation of $\varepsilon_{i,t}$ 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 propagation of $\varepsilon_{i,t}$, the conditions on ζ_t will become more stringent; we will provide that analysis in Section 4.3. Before doing so, however, we will first provide a visual illustration of our "as-if" policy shock result, and of how the identified synthetic shock is likely to relate to more standard measures of policy shocks.

4.2 What does the synthetic shock look like?

This section plots the estimand of the interacted local projection (17) in a simple environment based on the structural model of Smets & Wouters (2007). We make two main points: first, as we change the non-policy shock $\varepsilon_{i,t}$, the identified policy "treatment" changes; and second, across all the $\varepsilon_{i,t}$'s, this identified treatment is gradual and persistent, and thus meaningfully different from textbook transitory policy shocks. Analogous results for the structural model of Davig & Leeper (2007), which we also use as the laboratory for a finite-sample simulation study, are reported in Appendix B.2.

Data-generating process. We consider the structural model of Smets & Wouters, but change assumptions on policy. The monetary authority initially sets the policy rate according to a first rule, expected by all agents in the economy to be permanent. Then, unexpectedly, it changes its conduct to a different rule, again expected to be permanent; nothing else changes across the two sub-samples. We consider an econometrician that observes long samples pre- and post-policy change, and runs the interacted local projection (17) with $\varepsilon_{i,t}$ equal to each of the model's six primitive non-monetary shocks, and ζ_t equal to a dummy indicator for the policy regime. This setting is thus similar to our simple example in Section 2, just now obviously in a much richer infinite-horizon economy.

We note that, compared to our general model in Section 3, the experiment here is simpler: we just consider a one-off permanent rule change rather than recurring stochastic switches. We do so because such a simpler environment is actually sufficient for our purposes here, to illustrate in population the likely nature of identified "as-if" policy shocks. Our finite-sample simulation study in Appendix B.2 considers a more challenging, empirically relevant setting with recurring rule switches.

Plotting the "as-if" shocks. 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" policy shock paths, for each of the model's six shocks. 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 $\varepsilon_{i,t}$ identifies a different policy treatment. Intuitively, each shock has different stochastic properties and equilibrium effects, and so looking across policy regimes generically identifies 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 (here by assumption) persistent. Putting the two together, it follows that the difference in policy instrument paths is also persistent. Finally, the bottom two panels show 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 stress 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 in Christiano et al., 2005) as well as policy counterfactual evaluation (McKay & Wolf, 2023).

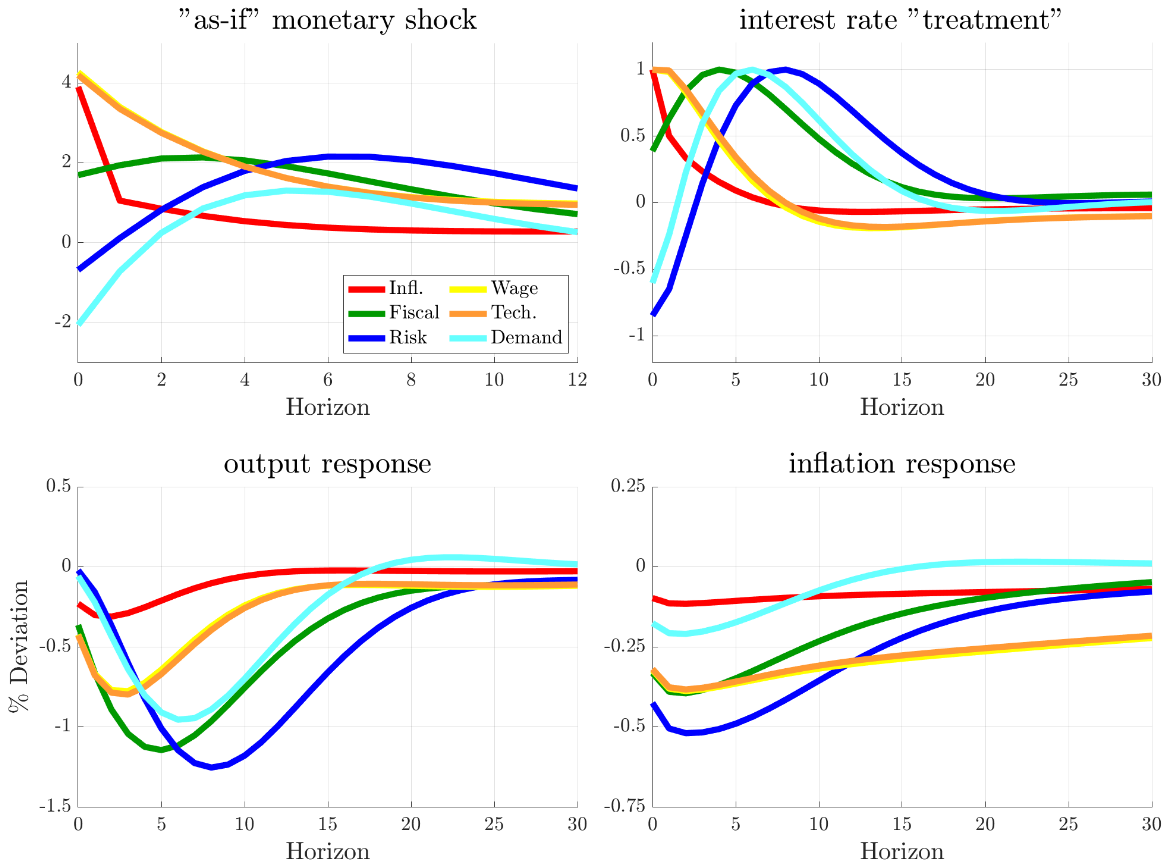

Figure 1: The top left panel shows the identified $H \times 1$ "as-if" policy shock path (defined relative to a standard Taylor rule as the baseline rule), while the other three panels show the interacted local projection coefficients $\{\delta_h\}$ for real interest rates, output, and inflation. All 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.

Empirical preview. The results in Figure 1 plus the accompanying discussion suggest that the estimand of interacted local projections like (17) is likely to correspond to persistent policy treatments; in the specific context of monetary policy, this would translate to forward guidance-like policy perturbations. Our empirical applications in Section 5 will confirm this.

4.3 Extensions

We now extend the headline identification result in Proposition 2 in three directions. First, we ask what our interacted local projections identify when the design of *multiple distinct* policy instruments changes over time. This analysis provides the segue to our second and key extension, allowing shock propagation to vary over time for non-policy reasons. Finally, we extend the model environment even further, to a largely unrestricted non-linear system akin to those studied in Rambachan & Shephard (2025) and Kolesár & Plagborg-Møller (2025).

Multiple policy instruments. We consider an extension of the model environment of Section 3 in which the policy instrument z_t is now n_z-dimensional; modulo that generalization of the dimensionality, the system equations (10) - (11) are unchanged. We then obtain the following straightforward generalization of Proposition 2.

Proposition 3. *Consider the dynamic stochastic economy* (10) - (11)*, with the policy instrument z_t of size n_z. Then, under the same assumptions as for Proposition 2, the population estimands of* (17) *satisfy*

$$\boldsymbol{\delta}_x = \underbrace{\bar{\Theta}_{x,v} \cdot \bar{\Theta}_{z,v}^{-1}}_{\text{policy causal effects}} \times \underbrace{\boldsymbol{\delta}_z}_{(n_z \cdot H) \times 1 \text{ vector of weights}} \tag{22}$$

The interacted local projection now generically recovers simultaneous treatments to all n_z policy instruments. In particular, following the same steps as in the sketch argument for Proposition 2, we now have that

$$\delta_{y,h} \propto \sum_{j=1}^{n_z} \bar{\Theta}_{y,v_j,h} \operatorname{Cov}\left(\Omega_j\left(\varepsilon_t, \varepsilon_{t-1}, \ldots, s_t, s_{t-1}, \ldots\right), \varepsilon_{i,t}\zeta_t\right) \tag{23}$$

where v_j is the shock to instrument j and $\Omega_j(\bullet)$ indicates the arguments of $\Omega(\bullet)$ corresponding to policy instrument j, defined exactly as in Proposition 1. As before, (23) is in policy shock space, while (22) translates to instrument space.

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 (17) will identify a combination of multiple policy treatments at the same time, with those treatments estimable by running (17) with the various policy instruments on the left-hand side. In particular, if the design of an instrument z_j changes over time, but the interaction regressor $\varepsilon_{i,t} \cdot \zeta_t$ does not revise the setting of that instrument, then $\boldsymbol{\delta}_{z_j} = \mathbf{0}$, and so the estimand does not reflect that particular policy instrument. This insight is intimately connected to our next extension, generalizing the model environment to allow for other forms of state dependence.

Richer forms of state dependence. Our second model extension allows for relatively general forms of non-policy state dependence in the propagation of macroeconomic shocks. Specifically, we replace the simple time-invariant private-sector block (10) by the following state-dependent generalization,

$$\mathcal{H}_x(m_t)\left[\mathbb{E}_t\left(\boldsymbol{x}^t\right)-\mathbb{E}_{t-1}\left(\boldsymbol{x}^t\right)\right]+\mathcal{H}_z(m_t)\left[\mathbb{E}_t\left(\boldsymbol{z}^t\right)-\mathbb{E}_{t-1}\left(\boldsymbol{z}^t\right)\right]+\mathcal{H}_e(m_t)e_t=\mathbf{0}, \tag{24}$$

where m_t is a date-t random variable that is independent of the date-t shocks ε_t but otherwise unrestricted. For example, the dependence of the private-sector block on m_t could capture time variation in the severity of nominal frictions (i.e., the slope of the Phillips curve), or in the fraction of households that are at or close to a binding borrowing constraint.

Following the same steps as in our earlier equilibrium characterization in Proposition 1, we now have

$$\mathbb{E}_t\left[\bar{\boldsymbol{y}}^t\right]-\mathbb{E}_{t-1}\left[\bar{\boldsymbol{y}}^t\right]=\bar{\Theta}(m_t)\varepsilon_t+\bar{\Theta}_v(m_t)\Omega(\varepsilon_t,\varepsilon_{t-1},\ldots,s_t,s_{t-1},\ldots;m_t) \tag{25}$$

where the $H\times 1$ vector $\Omega(\bullet)$ is defined exactly as before, just now conditional on the non-policy regime m_t. Proposition 2 then generalizes as follows.

Proposition 4. *Consider the dynamic stochastic economy with private-sector block* (24) *and policy* (11). *Suppose that the mean-zero random variable* ζ_t *is given as*

$$\zeta_t=\zeta\left(\varepsilon_{-i,t},\varepsilon_{t-1},\ldots,s_t,s_{t-1},\ldots\right) \tag{26}$$

and furthermore independent of the non-policy state m_t. *Then the population estimands of* (17) *satisfy*

$$\boldsymbol{\delta}_x=\mathbb{E}\Big[\underbrace{\bar{\Theta}_{x,v}(m_t)\cdot\bar{\Theta}_{z,v}(m_t)^{-1}}_{\text{policy causal effects}}\times\Big(\underbrace{\boldsymbol{\delta}_z}_{\text{av'g treatment}}\odot\underbrace{\boldsymbol{w}_z(m_t)}_{\text{“weights”}}\Big)\Big], \tag{27}$$

where the H*-dimensional “weights”,* $\boldsymbol{w}_z(m_t)$*, satisfy* $\mathbb{E}\left(\boldsymbol{w}_z(m_t)\right)=\mathbf{1}$*, but are not necessarily always positive. Positivity additionally requires that all entries of*

$$\operatorname{Cov}(\boldsymbol{z}^t,\varepsilon_{i,t}\zeta_t\mid m_t)$$

have the same sign independently of m_t.[10]

[10] Strictly speaking, the representation in (27) furthermore requires that, for all h, $\delta_{z,h}\neq 0$.

Intuitively, the interacted local projection (17) asks how the impulse response to the shock $\varepsilon_{i,t}$ covaries with ζ_t. Given that now impulse responses are allowed to vary in unrestricted fashion with the macroeconomic state m_t, we require ζ_t to not co-vary with m_t; once that is the case, $\boldsymbol{\delta}_x$ yet again admits a tight structural interpretation, as revealed by (27). There is, however, one additional wrinkle: in principle both treatment as well as treatment effect can change with the macroeconomic state m_t, which can cause weighting issues familiar from the microeconometric literature on heterogeneous treatment effects (e.g., see Imbens & Angrist, 1994). To see the logic, note that, conditional on a regime m_t, we could write the regression estimand as

$$\boldsymbol{\delta}_x(m_t) = \sum_{h=0}^{H-1} \underbrace{\bar{\Theta}_{x,z_h}(m_t)}_{\text{treatment effect}} \times \underbrace{\delta_{z,h}(m_t)}_{\text{treatment}},$$

where $\bar{\Theta}_{x,z_h}(m_t)$ is column $h+1$ of $\bar{\Theta}_{x,v}(m_t) \cdot \bar{\Theta}_{z,v}(m_t)^{-1}$, i.e., the causal effect of z_{t+h} on $\boldsymbol{x}^t$ conditional on regime m_t. Averaging across regimes to get to the actual estimand, we find

$$\boldsymbol{\delta}_x = \mathbb{E}\Big[\sum_{h=0}^{H-1} \bar{\Theta}_{x,z_h}(m_t) \ \times \ \delta_{z,h} \ \times \ \underbrace{\frac{\delta_{z,h}(m_t)}{\delta_{z,h}}}_{\equiv w_{z,h}(m_t)}\Big].$$

We thus get a weighted average of treatment effects across regimes m_t, but with weights that are not necessarily positive, because the effective treatments may also be regime-dependent. Assuming that the identified policy treatments are the same across regime 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 microeconometric program evaluation.[11]

We finally note that, while the above discussion assumed unconditional independence of m_t and ζ_t, similar conclusions apply under mean independence conditional on controls in an extended version of (17). We provide a discussion of that case in Appendix A.1.

Broader relation to the literature. Our identification results on causal effect identification through variation over time in policy design bear a close resemblance to those from "shift-share" panel regression specifications. In dynamic, general equilibrium settings, the

[11]While 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.

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 units *only* via the specific channel of interest (see Manuel, 2026, for a general theoretical treatment). In our case, $\varepsilon_{i,t}$ 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 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 propagation of $\varepsilon_{i,t}$.

A general non-linear setting. As our last extension we discuss how the interpretation of (17) and thus our identification results change in a largely unrestricted 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.2. Our focus here will instead be the interpretation of those results, in particular of how they relate to the foregoing conclusions about "as-if" shock equivalence.

We materially generalize the environment of Section 3 and now suppose that we can write any date-$t+h$ macroeconomic outcome x_{t+h} as

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

where our assumptions on the shock $\varepsilon_{i,t}$ and policy regime s_t are exactly as before, and u_{t+h} is a date-$t+h$ random variable independent of both $\varepsilon_{i,t}$ 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 $\varepsilon_{i,t}$ and s_t,

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

Now consider running the interacted local projection (17) with $\zeta_t = s_t$, i.e., we assume that the policy measure is orthogonal to both $\varepsilon_{i,t}$ and to all other sources of stochasticity in the economy. Under regularity conditions stated in Appendix A.2, the interacted local projection estimand $\boldsymbol{\delta}_x$ now satisfies

$$\delta_{x,h} = \mathbb{E}_{\varepsilon_i, s} \left[w(s_t, \varepsilon_{i,t}) \partial_{\varepsilon_i, s} \Psi_h(\varepsilon_{i,t}, s_t)\right] \tag{28}$$

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

$$\partial_{\varepsilon_i,s}\Psi_h(\varepsilon_{i,t}, s_t) = \mathbb{E}_u\left[\psi'_{h,\varepsilon_i z} z_s + z'_{\varepsilon_i}\psi_{h,zz} z_s + \psi'_{h,z} z_{\varepsilon_i s}\right], \tag{29}$$

with the subscripts here denoting partial derivatives. The baseline model of Section 3 and its extension discussed earlier in this section then correspond to the special case of the second derivatives of $\psi_h(\bullet)$ being zero, in which case the first two terms in (29) 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, $z_{\varepsilon_i,s}$. If the policy causal effects are furthermore invariant to other shocks and states u, then we return all the way to our headline identification result, with the estimand equal to a weighted average of policy causal effects $\psi_{h,z}$.

4.4 Summary and takeaways

The analysis of this section suggests a simple recipe for running and interpreting interacted local projections of the popular form

$$y_{t+h} = \alpha_{y,h} + \beta_{y,h}\varepsilon_{i,t} + \gamma_{y,h}\zeta_t + \delta_{y,h}(\varepsilon_{i,t} \cdot \zeta_t) + u_{t,y,h}.$$

In terms of the included regressors, it is important to ensure that $\varepsilon_{i,t}$ 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 $\varepsilon_{i,t}$ chiefly through policy. Once that is the case, the estimand of $\boldsymbol{\delta}_y$ is, under relatively weak assumptions on the underlying data-generating process, a simple weighted average of policy causal effects—i.e., the interacted local projection is an "as-if" policy shock regression. In the next section we leverage these insights in practical applications to monetary policy and ask how the identified synthetic shock compares to the standard monetary policy shocks familiar from the literature.

5 Applications to monetary policy causal effects

This section uses empirical applications to monetary policy to illustrate our identification results, building on the contributions of Hack et al. (2024) and Miyamoto et al. (2024). The headline takeaway will be that, consistent with the discussion in Section 4.2, interacted local projections with time-varying monetary policy tend to identify forward guidance-like "as-if" monetary shocks, and thus recover an estimand meaningfully different from that of standard monetary shock regressions.

We supplement the applications in this section with a simulation study in Appendix B.2. The empirical applications here suggest that, even in only moderately long samples, interacting shocks with rule changes can isolate a meaningful policy treatment signal, and thus allow estimation of policy causal effects. The simulations, based on the model of Davig & Leeper (2007), confirm this finding in a setting with recurring, stochastic policy regime switches.

5.1 Variation in FOMC composition

Our first set of applications is based on Hack et al. (2024), 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 interpret the results through the lens of our novel identification results. We first briefly recap their estimation strategy, before then presenting the headline results on the interacted local projection coefficients.

Shocks and policy measure. Hack et al. (2024) 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.[12] The remaining question then of course is relevance; consistent with Hack et al. (2024), we will indeed find that, when interacted with non-policy disturbances, this policy measure will tend to predict future revisions in monetary policy conduct.

For the non-policy shock $\varepsilon_{i,t}$, 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—and with relatively persistent macroeconomic effects. To this end we will leverage examples of three kinds of popular non-monetary shocks: oil, fiscal, and technology. For oil shocks, we use the recently updated version of the shock measure

[12] See Hack et al. (2024) 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).

of Hamilton (2003); we also experimented with the more recent series of Känzig (2021), but its relatively short sample leads to significantly noisier estimates. For fiscal shocks, as Hack et al. (2024) themselves already conduct an analysis based on Ramey (2011), we here instead consider the shock series of Ben Zeev & Pappa (2015). Finally, for technology shocks, we use the Beaudry & Portier (2006) news shock.

As we discussed in Section 4.1, the key requirement for identification is that the included shocks $\varepsilon_{i,t}$ 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 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 resulting "as-if" monetary shock estimands 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

$$y_{t+h} = \alpha_{y,h} + \beta_{y,h}\varepsilon_{i,t} + \gamma_{y,h}\zeta_t + \delta_{y,h}\left(\varepsilon_{i,t} \cdot \zeta_t\right) + \sum_{\ell=1}^{p} \xi_\ell q_{t-\ell} + u_{t,y,h}, \tag{30}$$

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. Finally, for the monetary policy shock reference point based on Romer & Romer (2004), we run a non-interacted local projection.

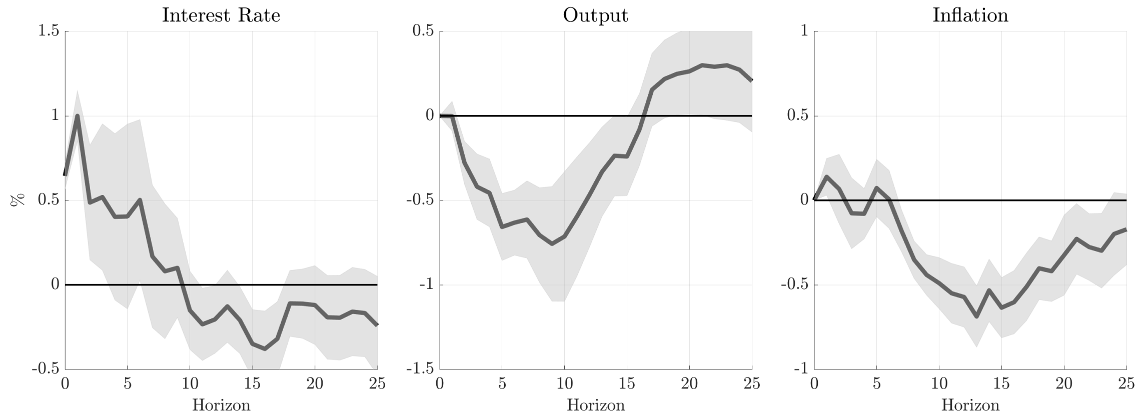

Figure 2: Impulse responses of the policy rate, output, and inflation to a monetary shock identified as in Romer & Romer (2004). The shaded areas correspond to 16th and 84th percentile bands. Quarterly data.

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 the key applied takeaway of this paper: that interacted monetary policy rule change regressions identify gradual, persistent "as-if" monetary policy shocks, i.e., forward guidance-type disturbances. The three panels correspond to interest rate, output, and inflation impulse responses to the identified synthetic monetary shocks—i.e., the δ_h coefficients—for our three non-policy shock measures $\varepsilon_{i,t}$: oil (top panel), fiscal (middle panel), and technology (bottom panel). Results are starkest for the top two panels, where the interest rate signal is relatively precisely estimated, and in both cases much more persistent than the Romer & Romer treatment.[13] 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 overall larger (in cumulative interest rate terms). The picture for inflation is more mixed, with the fiscal shock indicating an initial

[13]The intuition is exactly as discussed in Section 4.2. First, the three 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.

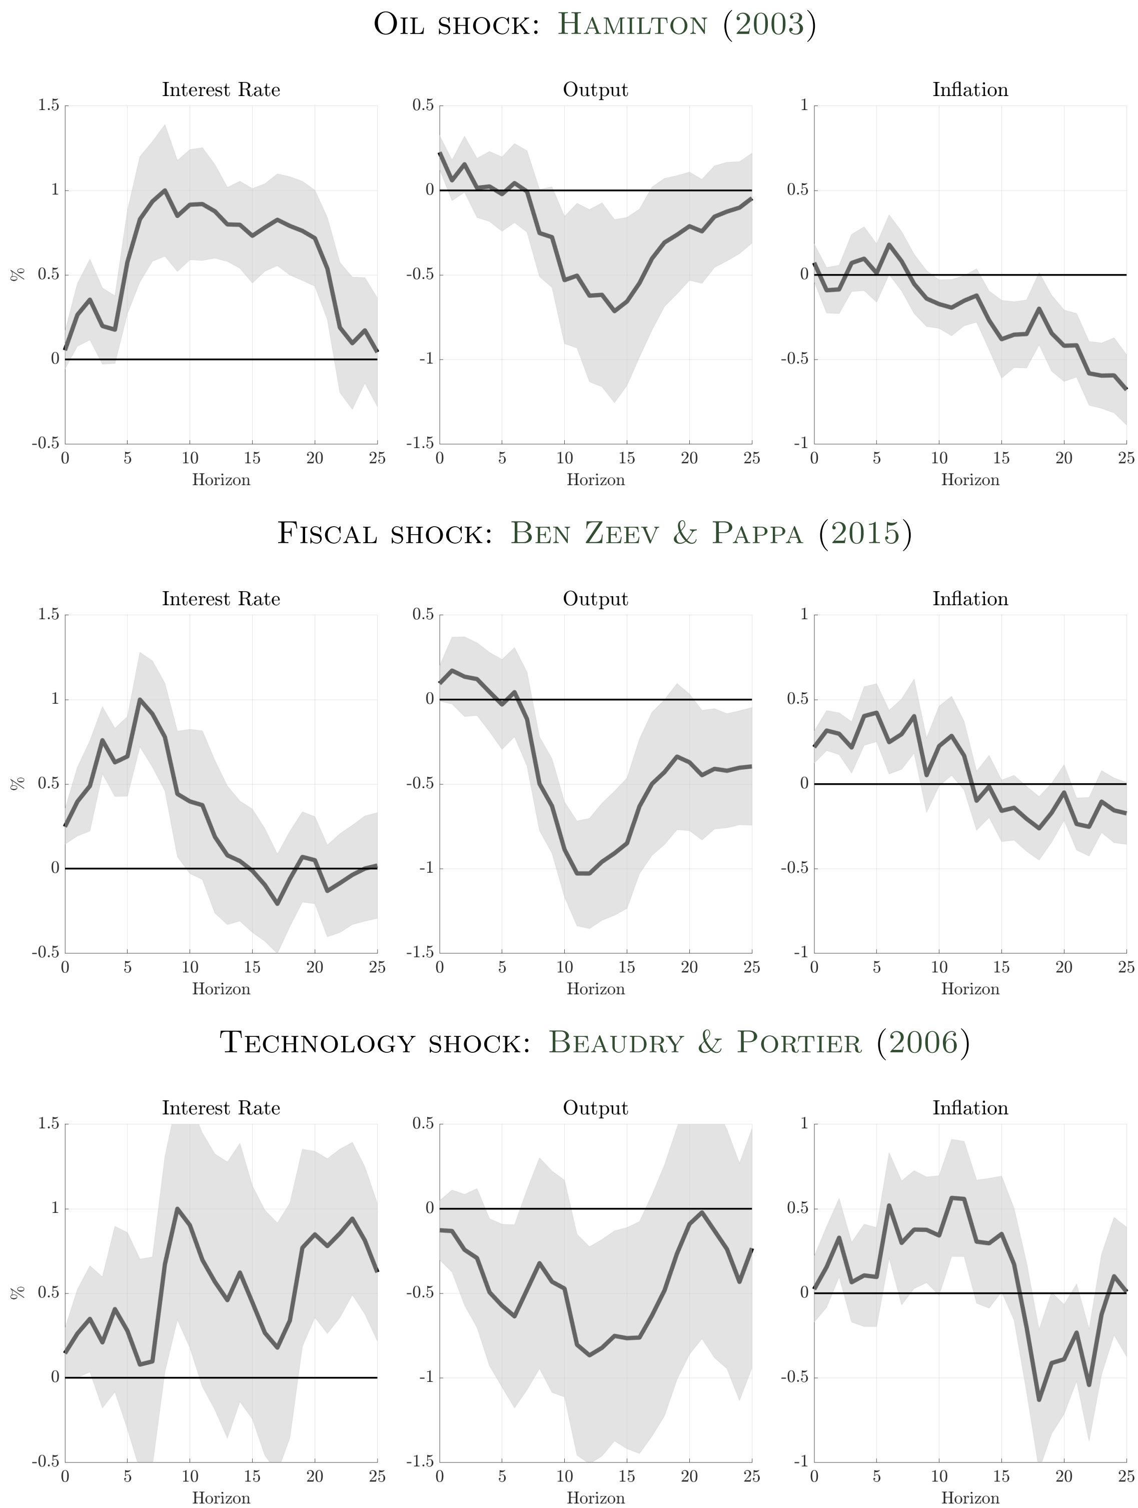

Figure 3: Impulse responses of the policy rate, output, and inflation to synthetic "as-if" monetary policy shocks based on the hawk-dove measure of Hack et al. (2024) together with an oil shock (top panel), fiscal shock (middle panel), and technology shock (bottom panel). The shaded areas correspond to 16th and 84th percentile bands. Quarterly data.

price puzzle, consistent perhaps with an interest rate cost channel. Finally, the results for the technology shock are broadly similar, but a bit noisier: the interest rate treatment is also gradual but much less precisely estimated, the output decline is of comparable magnitude, and inflation again shows an initial price puzzle.

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 an "as-if" more forward guidance-like disturbance. 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).

5.2 Interest rate lower bound

We conclude the empirical analysis by arguing that our headline applied takeaway—on interacted monetary policy local projections identifying gradual policy treatments—is not limited to the instrument of Hack et al. (2024), but applies more generally. To this end we replicate Miyamoto et al. (2024), 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 the policy measure ζ_t is less plausibly satisfied than for Hack et al., we only consider this alternative as a brief further robustness check.

Empirical specification. We yet again consider interacted local projections of the form (30). Our outcomes of interest are now the policy rate as the policy instrument, industrial production as our measure of economic activity, and as before inflation. The policy measure ζ_t is an indicator for whether the zero lower bound on the policy rate is binding in period $t-1$, and the shock $\varepsilon_{i,t}$ is equal to the oil shock series from Känzig (2021); all of these choices mirror Miyamoto et al. (2024). For further estimation details see Appendix C.2.

Results. Figure 4 re-affirms the main result of this section: the identified monetary policy treatment is yet again gradual and persistent. As before, the three panels of the figure display

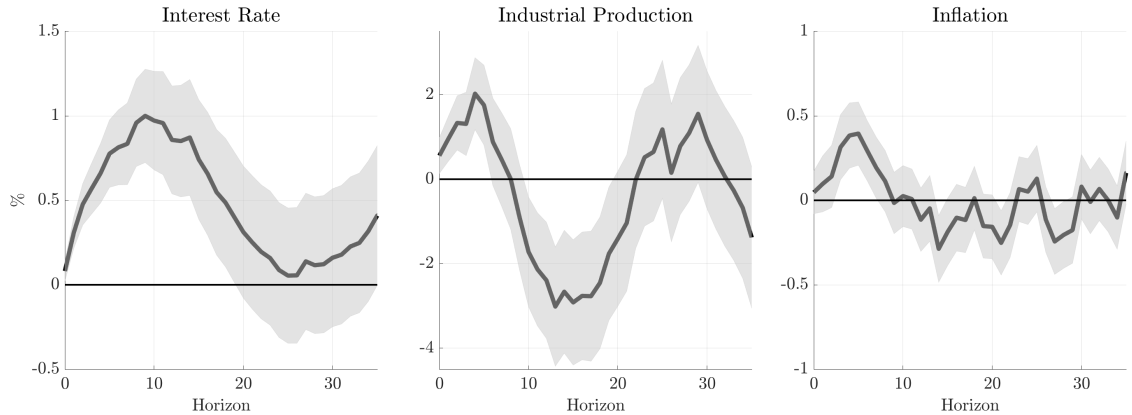

Figure 4: Impulse responses of the policy rate, industrial production, and inflation to synthetic "as-if" monetary policy shocks based on comparisons across ZLB and non-ZLB periods. The figure plots differential responses of Japanese variables conditional on an oil shock, following Miyamoto et al. (2024). The shaded areas correspond to 16th and 84th percentile bands. Monthly data.

impulse responses to the identified synthetic monetary shocks (i.e., the δ_h coefficients), with the left panel showing the policy instrument path. Similar to the results reported in Figure 3, we see that the interest rate response 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.

6 Identification through second moments

We conclude our analysis by asking whether the informational requirements on the "conditioning event" $\varepsilon_{i,t}$ can be meaningfully weakened, asking what we can learn from the effects of policy regime changes on unconditional second moments. There are at least three reasons to expect such an exercise to be informative. First, an early version of such an analysis is Mussa (1986), who observes that exchange rate regime changes affect the volatility of real outcomes, thus rejecting policy neutrality. We simply ask the stronger question of whether we can learn more than a mere test of neutrality. Second, the literature on identification through heteroskedasticity (e.g., Rigobon, 2003; Lewis, 2025) shows that the effects of changes in policy shock volatility on variance-covariance matrices are sufficient to allow identification of policy causal effects. We here check whether that same logic extends to policy rule changes. And third, 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. An open question is whether that mapping is invertible.

The headline takeaway of this section is that, unfortunately, the answer to all of these questions is "no." We show that, in the setting of Section 3, the effect of policy rule changes on unconditional second moments implies non-trivial, but typically excessively wide, identified sets for policy causal effects—policy neutrality can be tested, but not much more. Since the takeaways of this section are negative, it may be skipped by readers focused on practical applications, but may be relevant to those with interest in policy rule change identification results *per se*, or coming from the tradition of identification through heteroskedasticity.

6.1 Static warm-up

To build intuition we begin with a simple static analysis, based on the general model economy in Section 3 but setting $H = 1$.[14] 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 neutrality—the basic 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. Conditional on a policy regime s_t today, the variance-covariance matrix at date t, $\Sigma_t \equiv \text{Var}(y_t)$, is given as

$$\Sigma_t = \underbrace{\bar{\Theta}\bar{\Theta}'}_{\text{baseline policy}} + \underbrace{U_t\bar{\Theta}_v' + \bar{\Theta}_v U_t' + \bar{\Theta}_v S_t \bar{\Theta}_v'}_{\text{effect of policy change}} \tag{31}$$

where $U_t \equiv \bar{\Theta}\mathbb{E}\left[\varepsilon_t \Omega(\varepsilon_t, s_t) \mid s_t\right]$ and $S_t \equiv \mathbb{E}\left[\Omega(\varepsilon_t, s_t)^2 \mid s_t\right]$. 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 ($\bar{\Theta}_v S_t \bar{\Theta}_v'$), 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

[14] Specifically, the model environment here specializes to $\mathcal{H}_x x_t + \mathcal{H}_z z_t + \mathcal{H}_e e_t = \mathbf{0}$ for the private-sector block and $z_t = \mathcal{A}(e_t, v_t, s_t)$ for policy.

changes in those second moments could be informative about policy effects.

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

$$\Delta = (\tilde{U} - U)\bar{\Theta}_v' + \bar{\Theta}_v(\tilde{U} - U)' + \bar{\Theta}_v(\tilde{S} - S)\bar{\Theta}_v'. \tag{32}$$

It is immediate that $\text{rank}(\Delta) \leq 2$; for all subsequent arguments we will assume that Δ indeed assumes that maximal rank.[15]

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 2. *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 (31).*

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 (31) and consistent with the observed gap Δ. The following proposition characterizes the set of such vectors.

Proposition 5. *Consider two static economies (10) - (11) (with $H = 1$) that differ only in the policy rule. Given Σ and $\tilde{\Sigma}$ with $\text{rank}(\tilde{\Sigma} - \Sigma) = 2$, the identified set for $\bar{\Theta}_v$ is:*

$$\left\{ \bar{\Theta}_v^\dagger \ \mid \ \bar{\Theta}_v^\dagger = \left(V_+ \Lambda_+^{1/2} \ \pm \ V_- \Lambda_-^{1/2} \right) \cdot \ \alpha, \ \ \alpha \neq 0 \right\} \tag{33}$$

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 (31) 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

[15] If 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.

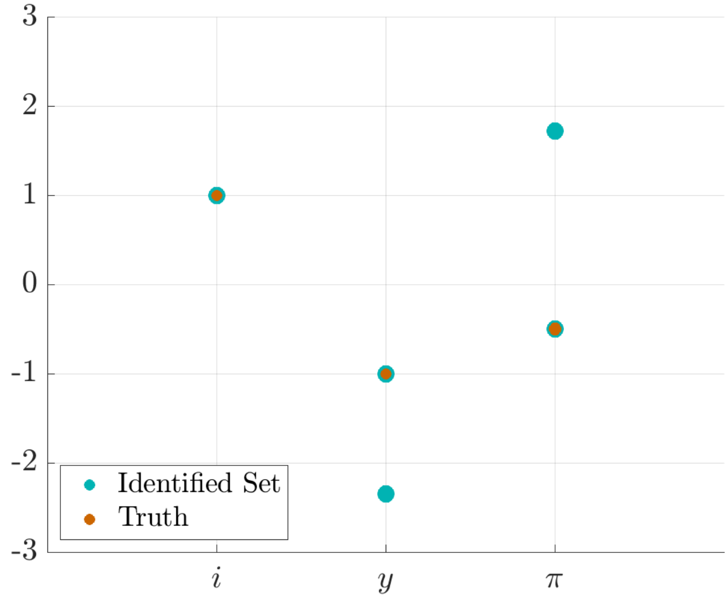

Figure 5: Identified set for interest rate, output, and inflation impulse responses to a one-off increase in interest rates, in a simple static New Keynesian model. The blue dots indicate members of the identified set, and the orange dots indicate the truth. See Appendix B.3 for model details.

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 5 visually illustrates this point, showing the identified set of interest rate causal effects in a static New Keynesian model. Consistent with Proposition 5, the identified set consists of two points; requiring a monetary contraction to lower inflation would then point-identify the true policy causal effects.

6.2 Identified sets with 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

$$\Sigma_t = \underbrace{\bar{\Theta}\bar{\Theta}'}_{\text{baseline policy}} + \underbrace{U_t\bar{\Theta}_v' + \bar{\Theta}_v U_t' + \bar{\Theta}_v S_t \bar{\Theta}_v'}_{\text{effect of policy change}} \tag{34}$$

where $U_t \equiv \bar{\Theta}\mathbb{E}_{t-1}\left[\varepsilon_t \Omega(\bullet)'\right]$ is $(H \cdot n_y) \times H$ and $S_t \equiv \mathbb{E}_{t-1}\left[\Omega(\bullet)\Omega(\bullet)'\right]$ 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. The supplementary discussion in Appendix A.3 argues that this is the typical case, and that furthermore H of Δ's eigenvalues are positive, and H negative.

The identified set of policy effects. We define the identified set of possible policy causal effects exactly as in Definition 2. We then arrive at the following characterization of that identified set.

Proposition 6. *Consider two $H-1$-lag dynamic economies (10) - (11) that differ only in the policy rule (11). Given Σ and $\tilde{\Sigma}$ with $\text{rank}(\tilde{\Sigma} - \Sigma) = 2H$, the identified set for $\bar{\Theta}_v$ is:*

$$\left\{\bar{\Theta}_v^\dagger \;\middle|\; \bar{\Theta}_v^\dagger = \left(V_+\Lambda_+^{1/2} + V_-\Lambda_-^{1/2}Q\right) \cdot \mathcal{A}, \;\; Q \in \mathcal{O}(H) \;\; \textit{and } \mathcal{A} \textit{ invertible}\right\} \tag{35}$$

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 (33). 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.[16]

How meaningful is this extension of the identified set? Clearly there is still much 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 of that. Unfortunately, in practice, the resulting identified set can be *very* wide. A simple illustration is provided in Figure 6, 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 (35) is too great to preclude any sharp conclusions.

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 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.[17]

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

[16] The invertible matrix $\mathcal{A}$ is in (35) 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.

[17] In 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.

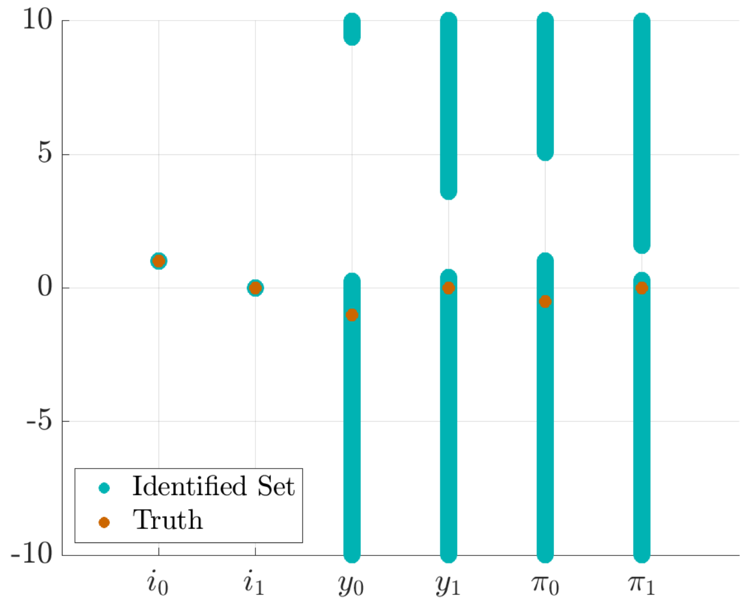

Figure 6: Identified set for interest rate, output, and inflation impulse responses to a one-off increase in interest rates, in a two-period New Keynesian model. The blue dots indicate members of the identified set, and the orange dots indicate the truth. See Appendix B.3 for model details.

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 $\varepsilon_{i,t}$—leveraged so heavily in prior sections—are indispensable for our purposes here.

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: they are equivalent to a regression on a weighted average of synthetic, "as-if" policy shocks. The availability of a primitive non-policy shock to interact with the policy measure—or more generally of a variable itself invariant to the policy change—is furthermore *necessary*:

policy-induced changes in unconditional second moments are not enough.

More practically, our results suggest that interacted local projections of the sort studied in recent work hold substantial promise as an alternative to the hitherto dominant focus on policy shocks. Since policy is high-dimensional, counterfactual policy evaluation requires estimates of a rich set of policy causal effects—richer than is available using the available policy shocks alone. Appealingly, interacted local projections promise to isolate a different policy treatment for each included non-policy shock; in practical applications to monetary policy, we furthermore find such treatments to be relatively gradual and thus quite materially different from standard monetary 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.

Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. Working Paper.

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.

Beaudry, P. & Portier, F. (2006). Stock prices, news, and economic fluctuations. *American economic review*, *96*(4), 1293–1307.

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.

Boppart, T., Krusell, P., & Mitman, K. (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, *89*, 68–92.

Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, *89*(1), 181–213.

Broda, C. (2004). Terms of trade and exchange rate regimes in developing countries. *Journal of International economics*, *63*(1), 31–58.

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.

Cloyne, J., Jordà, Ò., & Taylor, A. M. (2023). State-dependent local projections: Understanding impulse response heterogeneity. Technical report, National Bureau of Economic Research.

Davig, T. & Leeper, E. M. (2007). Generalizing the taylor principle. *American Economic Review*, *97*(3), 607–635.

Davig, T. & Leeper, E. M. (2008). *Endogenous Monetary Policy Regime Change*. National Bureau of Economic Research.

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.

Guerrieri, L. & Iacoviello, M. (2015). Occbin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics*, *70*, 22–38.

Hack, L., Istrefi, K., & Meier, M. (2024). 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.

Miyamoto, W., Nguyen, T. L., & Sergeyev, D. (2024). 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.

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 "A."—"D." refer to the main article.

A Supplementary theoretical results

We begin in Appendix A.1 by extending our identification results to policy measure orthogonality conditional on controls. Appendix A.2 provides the detailed discussion for identification results in a general non-linear setting, building heavily on Kolesár & Plagborg-Møller (2025). Finally, in Appendix A.3, we give missing details for the second-moment identification discussion of Section 6.

A.1 Additional controls

Consider the setting with richer forms of state dependence discussed in Section 4.3. Rather than assuming that the regime-indicator ζ_t is independent of the non-policy state m_t (as well as the shock-measure $\varepsilon_{i,t}$), suppose instead that ζ_t satisfies *conditional* (mean)-independence:

$$\mathbb{E}[\zeta_t \mid \varepsilon_{i,t}, m_t, c_t] = \alpha + \rho' c_t, \tag{A.1}$$

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 $\varepsilon_{i,t}$ is mean zero, independent of all other shocks as well as m_t, and additionally independent of c_t, ζ_t:

$$\varepsilon_{i,t} \perp \zeta_t, c_t, m_t, \varepsilon_{-i,t}, \underbrace{\mathbb{E}_{t-1}[y_{t+h}]}_{\text{past shocks}}, \underbrace{y_{t+h} - \mathbb{E}_t[y_{t+h}]}_{\text{future shocks}}. \tag{A.2}$$

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

$$y_{t+h} = \alpha_{y,h} + \beta_{y,h}\varepsilon_{i,t} + \gamma_{y,h}\zeta_t + \delta_{y,h}\left(\varepsilon_{i,t} \cdot \zeta_t\right) + \eta'_{y,h}\left(\varepsilon_{i,t} \cdot c_t\right) + u_{t,y,h}. \tag{A.3}$$

Under assumptions (A.2) and (A.1), and by Frisch-Waugh-Lovell, we find that

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

where $\tilde{\zeta}_t$ denotes the residual from a regression of ζ_t on a constant and c_t. Next note that, by (A.2), we have

$$\mathbb{E}\left[\mathbb{E}_{t-1}\left[y_{t+h}\right]\left(\varepsilon_{i,t} \cdot \tilde{\zeta}_t\right)\right] \;\; = \;\; 0,$$

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

Additionally, by (A.1),

$$\mathbb{E}\left[\bar{\Theta}_i(m_t)\varepsilon_{i,t}^2\tilde{\zeta}_t\right] = \mathbb{E}\left[\bar{\Theta}_i(m_t)\varepsilon_{i,t}^2\mathbb{E}[\tilde{\zeta}_t \mid \varepsilon_{i,t}, m_t]\right] = 0.$$

Then, as in the proof of Proposition 4, using (25), we have that

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

The remaining steps of Proposition 4 then go through replacing ζ_t with $\tilde{\zeta}_t$.

A.2 Extension to a general non-linear environment

Consider the setting introduced in the last part of Section 4.3. By Frisch-Waugh-Lovell, we have that

$$\delta_h = \frac{\operatorname{Cov}(y_{t+h}, \varepsilon_{i,t}s_t)}{\operatorname{Var}(\varepsilon_{i,t})\operatorname{Var}(s_t)}. \tag{A.4}$$

The remainder of this section characterizes this estimand further.

We assume that the scalar random variables $\varepsilon_{i,t}$ and s_t have densities denoted by $f_{\varepsilon_i}(\varepsilon_i)$ and $f_s(s)$, respectively, and are mean zero, $\mathbb{E}(\varepsilon_{i,t}) = \mathbb{E}(s_t) = 0$. 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

$$\operatorname{Cov}(y_{t+h}, \varepsilon_{i,t}s_t) = \mathbb{E}\left[\Psi_h(\varepsilon_{i,t}, s_t)\varepsilon_{i,t}s_t\right] = \mathbb{E}\left[\tau_s(s_t)\tau_{\varepsilon_i}(\varepsilon_{i,t})\partial_{\varepsilon_i s}\Psi_h(\varepsilon_{i,t}, s_t)\right]$$

where

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

We thus have

$$\delta_h = \int\int w(s_t, \varepsilon_{i,t})\partial_{\varepsilon_i s}\Psi_h(\varepsilon_{i,t}, s_t)ds_t d\varepsilon_{i,t} \tag{A.5}$$

where

$$w(s_t, \varepsilon_{i,t}) = \frac{\tau_s(s_t) f_s(s_t)}{\text{Var}(s_t)} \frac{\tau_{\varepsilon_i}(\varepsilon_{i,t}) f_{\varepsilon_i}(\varepsilon_{i,t})}{\text{Var}(\varepsilon_{i,t})} \tag{A.6}$$

and the cross-derivative term satisfies

$$\partial_{\varepsilon_i s} \Psi_h(\varepsilon_{i,t}, s_t) = \mathbb{E}_u \left[\psi'_{h,\varepsilon_i z} z_s + z'_{\varepsilon_i} \psi_{h,zz} z_s + \psi'_{h,z} z_{\varepsilon_i s} \right], \tag{A.7}$$

which is exactly the expression in Section 4.3. It now remains to note that, by the properties of $\tau_{\varepsilon_i}(\varepsilon_i)$ and $\tau_s(s)$, $w(s_t, \varepsilon_{i,t}) \geq 0$ and $\int\int w(s_t, \varepsilon_{i,t}) ds_t d\varepsilon_{i,t} = 1$.

A.3 Further discussion of second-moment identification

This appendix provides some supplementary information for the set-up of the identification problems in Section 6. We here follow the same structure as in that section, beginning with the static analysis and then turning to dynamics.

Static model. We begin with variance-covariance matrices in the two policy regimes, with the second regime indicated using tildes. We have

$$\Sigma = \bar{\Theta}\bar{\Theta}' + U\bar{\Theta}_v' + \bar{\Theta}_v U' + \bar{\Theta}_v S \bar{\Theta}_v', \tag{A.8}$$

$$\tilde{\Sigma} = \bar{\Theta}\bar{\Theta}' + \tilde{U}\bar{\Theta}_v' + \bar{\Theta}_v \tilde{U}' + \bar{\Theta}_v \tilde{S} \bar{\Theta}_v', \tag{A.9}$$

with U, S, $\tilde{U}$ and $\tilde{S}$ defined as stated in Section 6.1. 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

$$\Delta = \bar{\Theta}_v B' + B\bar{\Theta}_v'. \tag{A.10}$$

(A.10) implies that $\text{col}(\Delta) \subseteq \text{span}\left\{\bar{\Theta}_v, B\right\}$, 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,

$$\Delta = V_+ \Lambda_+ V_+' - V_- \Lambda_- V_-' \tag{A.11}$$

where $\{V_+, V_-\}$ are orthonormal, $\Lambda_+ > 0$ is the positive eigenvalue, and $\Lambda_- > 0$ is the absolute value of the negative eigenvalue.[18] This completes the set-up and initial discussion

[18]Since 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

of environment properties; the formal proof of Proposition 5 is provided in Appendix D.

Dynamic model/multiple instruments. Forecasting variance-covariance matrices in the two regimes are now

$$\Sigma_t = \bar{\Theta}\bar{\Theta}' + U_t\bar{\Theta}_v' + \bar{\Theta}_v U_t' + \bar{\Theta}_v S_t \bar{\Theta}_v', \tag{A.12}$$

$$\tilde{\Sigma}_t = \bar{\Theta}\bar{\Theta}' + \tilde{U}_t\bar{\Theta}_v' + \bar{\Theta}_v \tilde{U}_t' + \bar{\Theta}_v \tilde{S}_t \bar{\Theta}_v'. \tag{A.13}$$

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

$$\Delta = \bar{\Theta}_v B' + B\bar{\Theta}_v'. \tag{A.14}$$

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.[19] 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

$$\Delta = V_+\Lambda_+V_+' - V_-\Lambda_-V_-', \tag{A.15}$$

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

other negative.

[19]To see this, write

$$\Delta = CJC'$$

where $C \equiv \begin{pmatrix} \bar{\Theta}_v & B \end{pmatrix}$ 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 $\mathrm{col}(C)$, so Δ has an even number of eigenvalues ($2m$), and generically m are positive, and m are negative.

B Simulations and illustrations

We illustrate our identification results with several model-based simulations. In Appendix B.1 we give details on our computation of the "as-if" monetary policy shocks in the Smets & Wouters (2007) model. Appendix B.2 conducts a simple simulation study based on the model of Davig & Leeper (2007). Finally, in Appendix B.3, we explain how we construct identified sets based on second moments in simple one- and two-period New Keynesian models.

B.1 "As-if" shocks 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

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

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 (17) 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 $\phi_y = \phi_{dy} = \rho_i = 0$.[20]

B.2 Simulation study

We use simulations from a simple variant of the model of Davig & Leeper (2007) for two purposes: first, to illustrate our identification results in a setting with stochastic rule switches; and second, to showcase the performance of the proposed empirical strategy in finite samples

[20] Of 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.

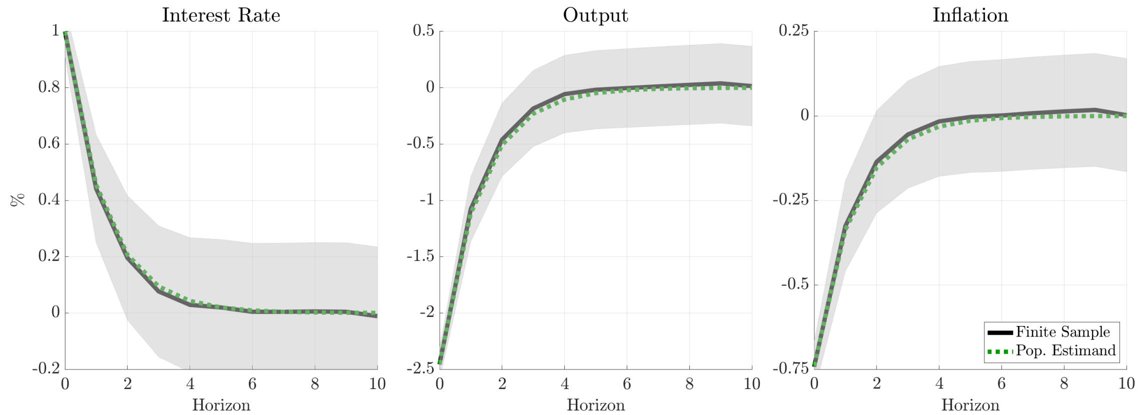

Figure B.1: Impulse responses of the policy rate, output, and inflation to synthetic "as-if" monetary policy shocks in population (green) and finite-sample simulations (black and gray) in the structural model of Davig & Leeper (2007), for demand shock interacted local projections. The shaded areas correspond to 16th and 84th percentile bands.

in a controlled environment. We first introduce the model and then present our results.

Model. We parameterize the model of Davig & Leeper (2007) as follows. Beginning with the private sector, we set $\sigma = 1$ (EIS), $\beta = 0.99$ (discount factor), as well as $\kappa = 0.166$—all relatively standard values. Next, for the two disturbances, we consider AR(1) processes, with persistence $\rho_d = 0.7$ for demand and $\rho_s = 0.2$ for supply and volatilities $\sigma_d = 1$ and $\sigma_s = 0.4$, respectively. Finally, for monetary policy, we allow for switches between two Taylor-type rules, both of the form

$$i_t = \phi_\pi \pi_t + \phi_y y_t,$$

with i_t denoting the policy rate, π_t inflation, and y_t output. For the first rule we assume $\phi_\pi = 3$ and $\phi_y = 1$ (i.e., an active rule), while for the second rule $\phi_\pi = 0.5$ and $\phi_y = 0$ (a passive rule). The probability of staying in the first regime is 0.85, while that of staying in the second regime is 0.8. We note that all of these parameter choices are purely illustrative and that the conclusions reported below are robust to a wide range of changes in parameterization.

Results. We simulate $1,000$ samples of size $T = 240$ from this economy, and consider an econometrician that runs the interacted local projection (30) with the demand shock as $\varepsilon_{i,t}$, ζ_t equal to a regime indicator, and $p = 1$ lags. For each sample we store point estimates as well as 16th and 84th percentile bands, and then average across all $1,000$ samples.

The main results of the simulation exercise are summarized in Figure B.1. The green line first of all shows the population estimand, which by Proposition 2 equals the causal effects of a particular mix of monetary policy shocks. The black lines then reveal that, even for $T = 240$, the interacted local projection is essentially unbiased, recovering correctly (on average) the target causal effects. Finally, the shaded areas indicate that those effects are reasonably precisely estimated.

B.3 Second moments in the New Keynesian model

Our two quantitative illustrations in Figures 5 and 6 begin with the same textbook New Keynesian model:

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

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t\left[\pi_{t+1}\right] + u_t^s, \tag{B.3}$$

$$i_t = \phi_\pi \pi_t + \phi_x x_t + u_t^m. \tag{B.4}$$

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 5 and 6. 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 present some supplementary information for the two empirical applications in Section 5: first the application based on Hack et al. (2024) in Appendix C.1, and second that based on Miyamoto et al. (2024) in Appendix C.2.

C.1 Variation in FOMC composition

We provide further details on our empirical applications in Section 5.1.

DATA. We begin with the 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. (2024). We finally turn to the shock measures $\varepsilon_{i,t}$.

- *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).
- *Technology.* We use Valerie Ramey's update of the Beaudry & Portier (2006) TFP news shock based on long-run restrictions, available as the series `bp_tfp_news_lr` in 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 our second application.

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

We provide further details on the second application in Section 5.2.

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. (2024), 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. (2024), we define the ZLB periods for Japan as: 1995:M10—2006:M6 and 2009:M1—2019:M12.

As our shock measure $\varepsilon_{i,t}$, 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 1960: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.

D Proofs

D.1 Proof of Proposition 1

The existence of a solution to (14) follows from the assumed invertibility of $\bar{\Theta}_{z,v} \neq 0$. 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 (10) considered here is identical to the one of that paper. □

D.2 Proof of Proposition 2

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

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

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

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

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

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

since $\varepsilon_{i,t}$ is jointly independent of both ζ_t and all shocks that realized prior to date $t-1$. We are thus left with

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

Now recall from (13) 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 $\varepsilon_{i,t}$ of $(\varepsilon_{-i,t}, \zeta_t)$ it follows that

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

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

$$\delta_{x,h} = \bar{\Theta}_{x,v,h}\bar{\Theta}_{z,v}^{-1}\boldsymbol{\delta}_z$$

Stacking across horizons h, the conclusion follows. Finally, it follows that $\boldsymbol{\delta}_z \neq \mathbf{0}$ as long as $\text{Cov}\left(\Omega\left(\varepsilon_t, \varepsilon_{t-1}, \ldots, s_t, s_{t-1}, \ldots\right), \varepsilon_{i,t}\zeta_t\right) \neq \mathbf{0}$. □

D.3 Proof of Proposition 3

None of the arguments in either the equilibrium characterization of Proposition 1 or the identification result of Proposition 2 actually relied on the policy instrument being a scalar, so we can follow the exact same steps to complete the argument for Proposition 3. □

D.4 Proof of Proposition 4

Using (25), we can proceed exactly as in the proof of Proposition 2 to get

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

Re-write this as

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

Denoting the interior expectation by $\boldsymbol{\omega}_{\varepsilon_i,\zeta}(m_t)$, we can write this as

$$\delta_{y,h} \propto \mathbb{E}\left[\bar{\Theta}_{y,v,h}(m_t)\boldsymbol{\omega}_{\varepsilon_i,\zeta}(m_t)\right].$$

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

$$\boldsymbol{\delta}_z \cdot \boldsymbol{w}_z(m_t) = \frac{\bar{\Theta}_{z,v}(m_t)\boldsymbol{\omega}_{\varepsilon_i,\zeta}(m_t)}{\text{Var}(\varepsilon_{i,t}\zeta_t)}$$

It then follows that

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

as claimed. Finally noting that

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

it also follows that $\mathbb{E}\left[\boldsymbol{w}_z(m_t)\right] = \mathbf{1}$, completing the argument. $\square$

D.5 Proof of Proposition 5

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 (32) 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 (A.8) - (A.9) hold. To this end define the projector on the column space of $\bar{\Theta}_v^\dagger$ as

$$\Pi_{\bar{\Theta}_v^\dagger} = \bar{\Theta}_v^\dagger\left[(\bar{\Theta}_v^\dagger)'\bar{\Theta}_v^\dagger\right]^+(\bar{\Theta}_v^\dagger)'$$

and let

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

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

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

Now consider setting

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

as well as

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

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

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

such that

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

But this means that, with

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

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. □

D.6 Proof of Proposition 6

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.[21] 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 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.[22] 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

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

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 (A.12) - (A.13) hold. To this

[21] The 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.

[22] Suppose 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.

end proceed as in the static case to note that

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

and thus

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

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

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

where we use the fact that, under our assumptions on the rank of Δ, we will generally also have that $\operatorname{rank}(\Sigma) = \operatorname{rank}(\tilde{\Sigma}) = H \cdot n_y$. Now consider setting

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

as well as

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

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

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

such that

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

But this means that, with

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

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. □