Econometrica Supplementary Material

SUPPLEMENT TO “WHAT CAN TIME-SERIES REGRESSIONS TELL US ABOUT POLICY COUNTERFACTUALS?”

(*Econometrica*, Vol. 91, No. 5, September 2023, 1695–1725)

ALISDAIR MCKAY
Federal Reserve Bank of Minneapolis

CHRISTIAN K. WOLF
MIT and NBER

THIS SUPPLEMENT CONTAINS ADDITIONAL MATERIAL FOR THE ARTICLE “What Can Time-Series Regressions Tell Us About Policy Counterfactuals?”. We provide (i) supplementary results complementing our theoretical identification analysis in Section 2, (ii) implementation details for our empirical methodology in Section 3, and (iii) several supplementary findings and alternative experiments complementing our empirical applications in Section 4.

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

APPENDIX A: SUPPLEMENTARY THEORETICAL RESULTS

This appendix provides several results complementing our theoretical identification analysis of Section 2. Appendix A.1 discusses examples of macro models that are nested by our results, Appendix A.2 gives an example of a model that is not, Appendix A.3 extends our optimal policy arguments to more general loss functions, Appendix A.4 works out what happens when the policy rule itself is changing, Appendix A.5 provides the details for unconditional second-moment counterfactuals, Appendix A.6 studies optimal monetary policy in our illustrative HANK model, Appendix A.7 shows how we construct counterfactuals with a limited number of policy shocks (as displayed in Figure 1), Appendix A.8 provides a global identification analysis with even higher informational requirements, and finally, Appendix A.9 extends our results to the case where only the policy rule is nonlinear.

A.1. *Examples of Nested Models*

We provide further details on three sets of models: the three-equation New Keynesian model of Section 2.1, a general class of behavioral models, and the HANK model of Section 2.4.

Alisdair McKay: alisdair.mckay@mpls.frb.org
Christian K. Wolf: ckwolf@mit.edu

Three-Equation NK Model

We here state the three-equation model of Section 2.1 in the form of our general matrix system (6)–(7). We begin with the non-policy block. The Phillips curve can be written as

$$\begin{pmatrix} 1 & -\beta & 0 & \dots \\ 0 & 1 & -\beta & \dots \\ 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \boldsymbol{\pi} - \kappa \boldsymbol{y} - \boldsymbol{\varepsilon}^s = 0,$$

while the Euler equation can be written as

$$-\sigma \begin{pmatrix} 0 & 1 & 0 & \dots \\ 0 & 0 & 1 & \dots \\ 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \boldsymbol{\pi} + \begin{pmatrix} 1 & -1 & 0 & \dots \\ 0 & 1 & -1 & \dots \\ 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \boldsymbol{y} + \sigma \boldsymbol{i} = 0.$$

Letting $\boldsymbol{x} \equiv (\boldsymbol{\pi}', \boldsymbol{y}')'$, we can stack these linear maps into the form (6). Finally, the policy rule can be written as

$$\phi_\pi \boldsymbol{\pi} - \boldsymbol{i} + \boldsymbol{\nu} = 0,$$

which directly fits into the form of (7) with $\boldsymbol{z} = \boldsymbol{i}$.

Behavioral Model

Our general framework (6)–(7) is rich enough to nest popular behavioral models such as the cognitive discounting framework of Gabaix (2020) or the sticky information setup of Mankiw and Reis (2002). We here provide a sketch of the argument for a particular example—the consumption-savings decision of behavioral consumers. Our discussion leverages sequence-space arguments as in Auclert, Bardóczy, Rognlie, and Straub (2021).

Let the linear map $\mathcal{E}$ summarize the informational structure of the consumption-savings problem, with entry (t, s) giving the expectations of consumers at time t about shocks at time s. Here an entry of 1 corresponds to full information and rational expectations, while entries between 0 and 1 can capture behavioral discounting or incomplete information. For example, cognitive discounting at rate θ would correspond to

$$\mathcal{E} = \begin{pmatrix} 1 & \theta & \theta^2 & \dots \\ 1 & 1 & \theta & \dots \\ 1 & 1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix},$$

while sticky information with a fraction $1 - \theta$ receiving the latest information could be summarized as

$$\mathcal{E} = \begin{pmatrix} 1 & 1-\theta & 1-\theta & \dots \\ 1 & 1 & 1-\theta^2 & \dots \\ 1 & 1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Let p denote an input to the household consumption-savings problem (e.g., income or interest rates). In sequence space, we can use the matrix $\mathcal{E}$ to map derivatives of the aggregate consumption function with respect to p, denoted $\mathcal{C}_p$, into their behavioral analogues $\tilde{\mathcal{C}}_p$ via

$$\tilde{\mathcal{C}}_p(t,s) = \sum_{q=1}^{\min(t,s)} \big[\mathcal{E}(q,s) - \mathcal{E}(q-1,s)\big]\mathcal{C}_p(t-q+1, s-q+1).$$

Typical behavioral frictions thus merely affect the matrices that enter our general non-policy block (6), but do not affect the separation of policy and non-policy blocks at the heart of our identification result.

Quantitative HANK Model

The HANK model used for our quantitative illustration in Section 2.4 is a simplified version of that in Wolf (2021). Since the model is standard, our discussion here will be relatively brief.

- **Demand block.** The economy is populated by a unit continuum of households that can save in a nominally risk-free, liquid asset. We set the steady-state quarterly real return on the asset to $\bar{r} = 0.01$, and its supply as a share of quarterly output to 1.04 (Kaplan, Moll, and Violante (2018)). Households have time-separable log preferences over consumption, with discount factor β. Their total non-asset income is $(1-\tau_y)e_{it}y_t + \tau_t$. Here y_t is aggregate income and e_{it} is idiosyncratic household productivity, with $\int_0^1 e_{it}\,di = 1$ at all t and where e_{it} follows the income process of Kaplan, Moll, and Violante (2018), ported to discrete time. Households also receive lump-sum transfers τ_t from the government; we set the steady-state level of transfers as a share of quarterly output to 0.05. We recover the time-invariant income tax rate τ_y to balance the government budget.

  The consumption block of the model is solved in two steps. First, we iterate over the household discount factor β to clear the liquid asset market in steady state. Second, we differentiate the aggregate consumption function around the deterministic steady state, giving the linearized relation

  $$\boldsymbol{c} = \mathcal{C}_y\boldsymbol{y} + \mathcal{C}_i\boldsymbol{i} + \mathcal{C}_\pi\boldsymbol{\pi} + \mathcal{C}_\tau\boldsymbol{\tau},$$

  where c is consumption, y is aggregate income, i is the nominal rate on the liquid asset, π is inflation, and τ denotes the uniform lump-sum transfer. Without uninsurable income risk, this demand block would collapse to the familiar Euler equation.
- **Supply block.** The supply relation of our economy is a standard NKPC:[28]

  $$\pi_t = \kappa y_t + \frac{1}{1+\bar{r}}\pi_{t+1} + \varepsilon_t,$$

  where ε_t is a contractionary cost-push shock and we set $\kappa = 0.1$. We assume that ε_t follows a standard AR(1) process with persistence 0.8.
- **Policy.** The fiscal authority fixes the total amount of outstanding government debt at its steady-state level and adjusts lump-sum transfers to balance the budget. The monetary authority follows the policy rules described in Section 2.4.

[28]See McKay and Wolf (2022) for a discussion of the assumptions on primitives necessary to derive such a supply relation from a combination of nominal rigidities and consumer labor supply in a HANK model.

A.2. *Filtering Problems*

To illustrate how an asymmetry in information between the private sector and the policy authority can break our separation of the policy and non-policy blocks in (6)–(7) even for a linear model, we consider a standard Lucas (1972) island model with a slightly generalized policy rule. The policy authority sets nominal demand x_t according to the rule

$$x_t = \phi_y y_t + x_{t-1} + \varepsilon_t^m,$$

where y_t denotes real aggregate output and ε_t^m is a policy shock with volatility σ_m. The private sector of the economy as usual yields an aggregate supply curve of the form

$$y_t = \theta(p_t - \mathbb{E}_{t-1} p_t),$$

where the response coefficient θ follows from a filtering problem and is given as

$$\theta = \frac{\sigma_z^2}{\sigma_z^2 + \sigma_p^2},$$

with σ_z denoting the (exogenous) volatility of idiosyncratic demand shocks and σ_p denoting the (endogenous) volatility of the surprise component of prices, $p_t - \mathbb{E}_{t-1} p_t$. A straightforward guess-and-verify solution of the model gives

$$p_t = \frac{1}{1+\theta} x_t + \frac{\theta}{1+\theta} x_{t-1}$$

and so

$$\sigma_p^2 = \left(\frac{1}{1+\theta}\right)^2 \mathrm{Var}\big(\phi_y y_t + \varepsilon_t^m\big).$$

But since

$$y_t = \frac{1}{1 - \dfrac{\theta}{1+\theta}\phi_y} \frac{\theta}{1+\theta} \varepsilon_t^m,$$

it follows that θ depends on the policy rule coefficient ϕ_y, breaking our separation assumption.

A.3. *More General Loss Functions*

Proposition 2 can be generalized to allow for a non-separable quadratic loss function. Suppose the policymaker's loss function takes the form

$$\mathcal{L} = \frac{1}{2} \boldsymbol{x}' Q \boldsymbol{x}, \tag{A.1}$$

where Q is a weighting matrix. Following the same steps as the proof of Proposition 2, we can formulate the policy problem as minimizing the loss function (A.1) subject to (25). The first-order conditions of this problem are

$$\Theta'_{\nu,x,A}\big(Q + Q'\big)\boldsymbol{x} = 0,$$

so we can recover the optimal policy rule as

$$\mathcal{A}_x^* = \Theta'_{\nu,x,A}(Q+Q'),$$
$$\mathcal{A}_z^* = \mathbf{0}.$$

Even outside of the quadratic case, the causal effects of policy shocks on $\boldsymbol{x}$ are still enough to formulate a set of necessary conditions for optimal policy, but in this general case, the resulting optimal policy rule will not fit into the linear form (7).

A.4. *Changing Policy Rules*

Our results apply without change to economies in which the policy rule is changing over time; instead, for our purposes, the key requirement is that the non-policy block (6) remains stable. To see why, consider an econometrician who observes data generated from an economy described first by the pair of equations

$$\mathcal{H}_w \boldsymbol{w} + \mathcal{H}_x \boldsymbol{x} + \mathcal{H}_z \boldsymbol{z} + \mathcal{H}_\varepsilon \boldsymbol{\varepsilon} = \mathbf{0}, \tag{A.2}$$

$$\mathcal{A}_{x,1} \boldsymbol{x} + \mathcal{A}_{z,1} \boldsymbol{z} + \boldsymbol{\nu}_1 = \mathbf{0}, \tag{A.3}$$

before then changing to

$$\mathcal{H}_w \boldsymbol{w} + \mathcal{H}_x \boldsymbol{x} + \mathcal{H}_z \boldsymbol{z} + \mathcal{H}_\varepsilon \boldsymbol{\varepsilon} = \mathbf{0}, \tag{A.4}$$

$$\mathcal{A}_{x,2} \boldsymbol{x} + \mathcal{A}_{z,2} \boldsymbol{z} + \boldsymbol{\nu}_2 = \mathbf{0}. \tag{A.5}$$

That is, while the private-sector block remains stable throughout, the policy rule has changed over time. In keeping with Assumption 1, we assume that both policy rules induce a unique bounded equilibrium for any bounded sequence of $\{\boldsymbol{\varepsilon}, \boldsymbol{\nu}_1\}$ or $\{\boldsymbol{\varepsilon}, \boldsymbol{\nu}_2\}$.

Now consider an econometrician who separately studies the propagation of policy shocks on the two subsamples. In particular, suppose she has successfully estimated the causal effects of some bounded policy shock vector $\boldsymbol{\nu}_1$ under the first rule. Then, by the exact same logic as in the proof of Proposition 1, we know that the similarly bounded shock sequence

$$\boldsymbol{\nu}_2 = (\mathcal{A}_{x,1} - \mathcal{A}_{x,2})\boldsymbol{x}_{\mathcal{A}_1}(\boldsymbol{\nu}_1) + (\mathcal{A}_{z,1} - \mathcal{A}_{z,2})\boldsymbol{z}_{\mathcal{A}_1}(\boldsymbol{\nu}_1) + \boldsymbol{\nu}_1$$

under the new policy rule will induce the exact same impulse responses—that is, we have

$$\boldsymbol{x}_{\mathcal{A}_1}(\boldsymbol{\nu}_1) = \boldsymbol{x}_{\mathcal{A}_2}(\boldsymbol{\nu}_2), \qquad \boldsymbol{z}_{\mathcal{A}_1}(\boldsymbol{\nu}_1) = \boldsymbol{z}_{\mathcal{A}_2}(\boldsymbol{\nu}_2).$$

The same argument also works in reverse, mapping any bounded $\boldsymbol{\nu}_2$ into a similarly bounded $\boldsymbol{\nu}_1$ with the exact same causal effects. This in particular implies that our econometrician could use her evidence on policy shock propagation from either subsample to implement our identification results in Section 2.3. Intuitively, since the private-sector block is unchanged, policy shocks in both cases identify the exact same space of dynamic causal effects, just with respect to two different reference points—the rule $\{\mathcal{A}_{x,1}, \mathcal{A}_{z,1}\}$ in the first subsample, and the rule $\{\mathcal{A}_{x,2}, \mathcal{A}_{z,2}\}$ in the second subsample.

As the final step in the argument, suppose now that the econometrician would estimate policy shock causal effects across the two (large) subsamples; in particular, suppose that the relative sizes of her two subsamples are $(\omega, 1-\omega)$. In that case, direct projection on an instrumental variable that correlates with a given shock sequence $\boldsymbol{\nu}$ would asymptotically recover a weighted average of subsample causal effects, for example, $\omega\Theta_{x,\nu,\mathcal{A}_1} \times \boldsymbol{\nu} + (1-\omega)\Theta_{x,\nu,\mathcal{A}_2} \times \boldsymbol{\nu}$. The overall estimated dynamic causal effect matrices would thus just be

$$\Theta_{x,\nu,\mathcal{A}} \equiv \omega\Theta_{x,\nu,\mathcal{A}_1} + (1-\omega)\Theta_{x,\nu,\mathcal{A}_2},$$
$$\Theta_{z,\nu,\mathcal{A}} \equiv \omega\Theta_{z,\nu,\mathcal{A}_1} + (1-\omega)\Theta_{z,\nu,\mathcal{A}_2}.$$

Importantly, post-multiplying those dynamic causal effects matrices by *any* policy shock vector $\boldsymbol{\nu}$ yields sequences of private-sector outcomes x and policy instruments z that are consistent with the common (sample-invariant) private-sector block (A.2) = (A.4).

Now consider using those causal effects matrices to implement Proposition 1; that is, given baseline non-policy shock impulse responses $\{\boldsymbol{x}(\boldsymbol{\varepsilon}), \boldsymbol{z}(\boldsymbol{\varepsilon})\}$ (which could come from either subsample, or also from a weighted average), find $\tilde{\boldsymbol{\nu}}$ such that

$$\tilde{\mathcal{A}}_x\big[\boldsymbol{x}(\boldsymbol{\varepsilon}) + \Theta_{x,\nu,\mathcal{A}} \times \tilde{\boldsymbol{\nu}}\big] + \tilde{\mathcal{A}}_z\big[\boldsymbol{z}(\boldsymbol{\varepsilon}) + \Theta_{z,\nu,\mathcal{A}} \times \tilde{\boldsymbol{\nu}}\big] = \boldsymbol{0},$$

and then compute the impulse responses

$$\boldsymbol{x}(\boldsymbol{\varepsilon}) + \Theta_{x,\nu,\mathcal{A}} \times \tilde{\boldsymbol{\nu}},$$
$$\boldsymbol{z}(\boldsymbol{\varepsilon}) + \Theta_{z,\nu,\mathcal{A}} \times \tilde{\boldsymbol{\nu}}.$$

Since the counterfactual policy rule holds by construction of $\tilde{\boldsymbol{\nu}}$, and since the averaged policy shock causal effects $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ embed the common (sample-invariant) private-sector block, it follows by the same arguments as in the proof of Proposition 1 that those impulse responses equal $\boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})$ and $\boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})$, respectively.

A.5. *Counterfactual Second-Moment Properties*

Our analysis is largely focused on constructing counterfactuals conditional on particular non-policy shock paths $\boldsymbol{\varepsilon}$. This is in keeping with much of the empirical policy counterfactual literature that followed the lead of Sims and Zha (1995) (e.g., Bernanke, Gertler, Watson, Sims, and Friedman (1997), Eberly, Stock, and Wright (2020), Antolin-Diaz, Petrella, and Rubio-Ramírez (2021)). However, under some additional assumptions, our results can also be used to construct *unconditional* counterfactual second-moment properties—that is, predict how variances and covariances of macroeconomic aggregates would change under a counterfactual rule. This section provides the detailed argument.

Setting

We consider a researcher that observes and is interested in the counterfactual properties of some vector of aggregates $y = (x, z)$—the endogenous outcomes and policy instruments of our main analysis. We assume that, under the prevailing policy rule, this vector of macro aggregates follows a standard structural vector moving average representation:[29]

$$y_t = \sum_{\ell=0}^{\infty} \Theta_\ell \varepsilon_{t-\ell} = \Theta(L)\varepsilon_t, \quad \varepsilon_t \sim N(0, I). \tag{A.6}$$

[29]Given our focus on second moments, the normality restriction is purely for notational convenience (see, e.g., Plagborg-Møller and Wolf (2021)).

We would like to predict the second-moment properties of the macroeconomic aggregates y_t under some counterfactual policy rule (8).

If the researcher can estimate the causal effects of all shocks ε_t on the outcomes y_t, then the identification argument is trivial: she simply applies Proposition 1 for each individual shock, stacks the resulting impulse responses into a new vector moving average representation $\tilde{\Theta}(L)$, and from here computes the counterfactual second-moment properties. This approach may, however, not be feasible, as it requires the researcher to be able to correctly disentangle all of the structural shocks driving the macro-economy.

Procedure

Our proposed procedure has three steps. First, the researcher estimates the Wold representation of the observables y_t. Second, using Proposition 1, she maps the impulse responses to the Wold errors into new impulse responses corresponding to the counterfactual policy rule. Third, she stacks those new impulse responses to arrive at a new vector moving average representation, and from this representation constructs a new set of second-moment properties. Our identification result states that, if the vector moving representation (A.6) under the baseline rule is invertible, then this procedure correctly recovers the desired counterfactual second moments.

Identification Result

Let $\tilde{\Theta}_\ell$ denote the lag-ℓ impulse responses of the observables y_t to the shocks ε_t under the counterfactual policy rule. The process for y_t under the counterfactual policy rule thus becomes

$$y_t = \sum_{\ell=0}^{\infty} \tilde{\Theta}_\ell \varepsilon_{t-\ell} = \tilde{\Theta}(L)\varepsilon_t,$$

and so the second moments of the true counterfactual process are given by

$$\Gamma_y(\ell) = \sum_{m=0}^{\infty} \tilde{\Theta}_m \tilde{\Theta}'_{m+\ell}. \tag{A.7}$$

Now consider instead the output of our proposed procedure. Let u_t denote the Wold errors under the observed policy rule, and let ε_t^* denote any unit-variance orthogonalization of these Wold errors (e.g., $\varepsilon_t^* = \text{chol}(\text{Var}(u_t))^{-1} \times u_t$). Then y_t under the observed policy rule satisfies

$$y_t = \Psi(L)\varepsilon_t^* = \sum_{\ell=0}^{\infty} \Psi_\ell \varepsilon_{t-\ell}^*,$$

where $\varepsilon_t^* \sim N(0, I)$. Under invertibility—that is, $\Theta(L)$ has a one-sided inverse—we in fact know that $\varepsilon_t^* = P\varepsilon_t$, $\Psi(L) = \Theta(L)P'$ for some orthogonal matrix P. The second step of our procedure gives the counterfactual vector moving average representation

$$y_t = \tilde{\Psi}(L)\varepsilon_t^*.$$

Now consider the two lag polynomials $\tilde{\Theta}(L)$ and $\tilde{\Psi}(L)$. Since $\varepsilon_t^* = P\varepsilon_t$, applying our counterfactual mapping to the jth Wold innovation $\varepsilon_{j,t}^*$ gives causal effects $\tilde{\Psi}_j(L) = \tilde{\Theta}(L) \times p_j'$, where p_j is the jth row of P.[30] Thus, overall, we have

$$\tilde{\Psi}(L) = \tilde{\Theta}(L)P'.$$

But then the second-moment properties of y_t implied by our proposed procedure are given as

$$\sum_{m=0}^{\infty} \tilde{\Psi}_m \tilde{\Psi}'_{m+\ell} = \sum_{m=0}^{\infty} \tilde{\Theta}_m P' P \tilde{\Theta}'_{m+\ell} = \sum_{m=0}^{\infty} \tilde{\Theta}_m \tilde{\Theta}'_{m+\ell} = \Gamma_y(\ell), \tag{A.8}$$

which is exactly equal to (A.7), completing the argument.

Finally, we emphasize that this identification result inherently relies on invertibility. Under invertibility, there is a static one-to-one mapping between true shocks ε_t and Wold errors ε_t^*; thus, if we can predict the propagation of the Wold errors under the counterfactual rule, then we also match the propagation of the true shocks, and so we correctly recover second-moment properties. Under non-invertibility, however, there is no analogous one-to-one mapping, and so it is not guaranteed that second moments will be matched.

A.6. *Optimal Policy Counterfactual in HANK*

Section 2.4 used a HANK model to illustrate the logic of Proposition 1—the general counterfactual rule identification result. We here do the same for the analogous optimal policy identification result in Proposition 2.

We consider a policymaker with a standard dual mandate loss function

$$\mathcal{L} = \lambda_\pi \boldsymbol{\pi}' W \boldsymbol{\pi} + \lambda_y \boldsymbol{y}' W \boldsymbol{y} \tag{A.9}$$

with $\lambda_\pi = \lambda_y = 1$. As in Section 2.4, we start by solving for the optimal policy using conventional methods: that is, we first derive the policy rule corresponding to the first-order conditions (18)–(20), then solve the model given that policy rule, and finally report the result as the solid lines in the left and middle panels of Figure A.1. We see that, at the optimum, the cost-push shock moves inflation by much more than output, consistent with the assumed policy weights and the relatively flat Phillips curve. Compared to this optimal policy, the simple baseline rule of the form (28) tightens too much (dotted).

We then instead use Proposition 2 to equivalently recover the optimal policy rule and the corresponding impulse responses. We begin with the optimal rule itself. By (26), the optimal rule is given as

$$\lambda_\pi \Theta'_{\pi,\nu,\mathcal{A}} W \boldsymbol{\pi} + \lambda_y \Theta'_{y,\nu,\mathcal{A}} W \boldsymbol{y} = \mathbf{0}.$$

A researcher with knowledge of the effects of monetary policy shocks on inflation and output, $\{\Theta_{\pi,\nu,\mathcal{A}}, \Theta_{y,\nu,\mathcal{A}}\}$, is able to construct this optimal policy rule. We can then create a counterfactual response to the cost-push shock using (11)–(13), again requiring only knowledge of the causal effects of policy shocks as well as the impulse responses to the cost-push shock under the baseline rule. As expected, the resulting impulse responses—the dashed lines—are identical to those obtained by explicitly solving the optimal policy problem. Finally, the right panel of Figure A.1 shows the optimal policy as a deviation $\tilde{\boldsymbol{\nu}}$

[30] To see this, let $\tilde{\boldsymbol{\nu}}_i$ denote the sequence of policy shocks that maps the true shock ε_i into its counterfactual causal effects. Then the sequence of policy shocks $\tilde{\boldsymbol{\nu}}_j^*$ that implements our counterfactual mapping for the reduced-form shock ε_j^* is given as $\tilde{\boldsymbol{\nu}}_j^* = \sum_i \tilde{\boldsymbol{\nu}}_i p_{ij}$, and thus we have $\tilde{\Psi}_j(L) = \sum_i \tilde{\Theta}_i(L) p_{ij}$, as claimed.

Optimal Policy, HANK Model

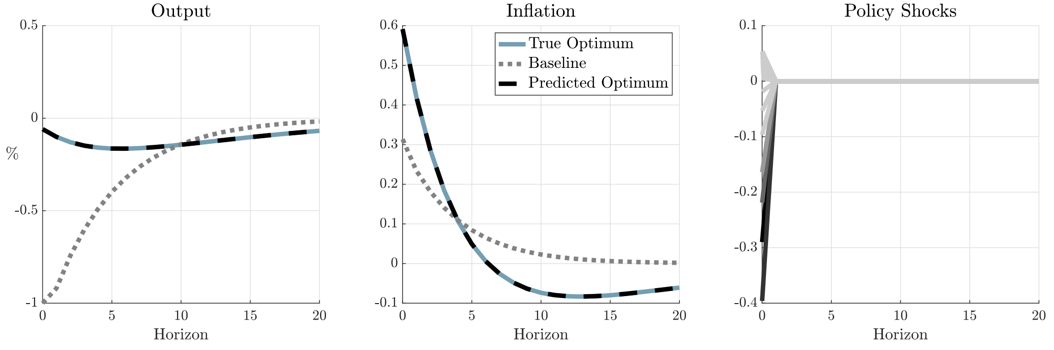

FIGURE A.1.—The dotted and solid lines in the left and middle panels show output and inflation responses to the cost-push shock ε_t for the HANK model with policy rule (28) and the optimal rule for the loss function (A.9). The dashed lines give output and inflation counterfactuals constructed through the policy shocks on the right, set in line with Proposition 2. Lighter shades correspond to news about policy at longer horizons.

from the prevailing rule. The optimal rule accommodates the inflationary cost-push shock more than the baseline rule (28), so the required policy "shock" is persistently negative (i.e., expansionary). Consistent with our discussion in Figure 1, we choose to display those shocks $\tilde{\boldsymbol{\nu}}$ in a way that emphasizes that the optimum is achieved through a sequence of *date-0* policy shocks.

A.7. *Counterfactuals With a Limited Number of Shocks*

In Figure 1, we constructed counterfactuals using a limited number n_s of policy shocks. We here provide the computational details for this construction. We discuss the general case of a researcher with access to n_s shocks (which converges to our identification result for $n_s \to \infty$), with the original proposal of Sims and Zha (1995) nested as the $n_s = 1$ special case.

The approach of Sims and Zha leverages the idea that evidence on one policy shock—that is, any single fixed path $\boldsymbol{\nu}$—is sufficient to enforce any given counterfactual ex post. With n_s distinct shocks, the counterfactual rule can be implemented ex post as well as in ex ante expectation for the next $n_s - 1$ time periods. To compute the counterfactuals corresponding to this multi-shock case, we proceed as follows. First, at $t = 0$, we solve for the n_s-dimensional vector of policy shocks $\boldsymbol{\nu}^0_{1:n_s} \equiv (\nu^0_0, \dots, \nu^0_{n_s-1})'$ such that, in response to $\boldsymbol{\varepsilon}$ and $\boldsymbol{\nu}^0_{1:n_s}$, the counterfactual rule holds at $t = 0$ and is expected to hold for $t = 1, \dots, n_s - 1$. Output and inflation at $t = 0$ are simply given as the thus-derived impulse responses to $\boldsymbol{\varepsilon}$ and $\boldsymbol{\nu}^0_{1:n_s}$. Second, at $t = 1$, we solve for the n_s-dimensional vector of shocks $\boldsymbol{\nu}^1_{1:n_s} \equiv (\nu^1_0, \dots, \nu^1_{n_s-1})'$ such that, in response to the time-0 shocks $\{\boldsymbol{\varepsilon}, \boldsymbol{\nu}^0_{1:n_s}\}$ and the time-1 shocks $\boldsymbol{\nu}^1_{1:n_s}$, the counterfactual policy rule holds at $t = 1$ and in expectation for $t = 2, \dots, n_s$. These impulse responses then give us output and inflation at $t = 1$. Continuing iteratively, we then obtain the entire output and inflation impulse responses, as plotted in the left and middle panels of Figure 1. The corresponding policy shock paths are shown in the right panel.

A.8. *Global Identification Argument*

We here extend our identification results to a general non-linear model with aggregate risk.

Setting

We consider an economy that runs for T periods overall. As in our main analysis, the economy consists of a private block and a policy block. Differently from our main analysis, there is no exogenous non-policy shock sequence $\boldsymbol{\varepsilon}$; rather, there is a stochastic event ω_t each period, with stochastic events drawn from a finite (n_ω-dimensional) set. Let $x_t(\omega^t)$ be the value of the endogenous variables after history $\omega^t \equiv \{\omega_0, \omega_1, \dots, \omega_t\}$ and let $z_t(\omega^t)$ be the realization of the policy instruments after history ω^t. Let $\boldsymbol{x}$ and $\boldsymbol{z}$ be the full contingent plans for all $t \in \{0, 1, \dots, T\}$ and all histories. $\boldsymbol{x}$ and $\boldsymbol{z}$ are vectors in $\mathbb{R}^{n_x \times N}$ and $\mathbb{R}^{n_z \times N}$, respectively, where $N = n_\omega + n_\omega^2 + \cdots + n_\omega^{T+1}$.

We can write the private-sector block of the model as the non-linear equation

$$\mathcal{H}(\boldsymbol{x}, \boldsymbol{z}) = \mathbf{0}. \tag{A.10}$$

Similarly, we can write the policy block corresponding to a baseline policy rule as

$$\mathcal{A}(\boldsymbol{x}, \boldsymbol{z}) + \boldsymbol{\nu} = \mathbf{0}, \tag{A.11}$$

where the vector of policy shocks $\boldsymbol{\nu}$ is now $n_z \times N$-dimensional. We assume that, for any $\boldsymbol{\nu} \in \mathbb{R}^{n_z \times N}$, the system (A.10)–(A.11) has a unique solution. We write this solution as

$$\boldsymbol{x} = x(\boldsymbol{\nu}), \qquad \boldsymbol{z} = z(\boldsymbol{\nu}).$$

We want to construct counterfactuals under the alternative policy rule

$$\tilde{\mathcal{A}}(\boldsymbol{x}, \boldsymbol{z}) = \mathbf{0} \tag{A.12}$$

replacing (A.11). We again assume that the system (A.10) and (A.12) has a unique solution, now written as $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}})$. If we are interested in the counterfactual following a particular path of exogenous events, then we are interested in selections from these vectors.

PROPOSITION A.1: *For any alternative policy rule $\tilde{\mathcal{A}}$, we can construct the desired counterfactuals as*

$$x(\tilde{\boldsymbol{\nu}}) = \tilde{\boldsymbol{x}}, \qquad z(\tilde{\boldsymbol{\nu}}) = \tilde{\boldsymbol{z}}, \tag{A.13}$$

where $\tilde{\boldsymbol{\nu}}$ solves

$$\tilde{\mathcal{A}}\big(x(\tilde{\boldsymbol{\nu}}), \qquad z(\tilde{\boldsymbol{\nu}})\big) = \mathbf{0}. \tag{A.14}$$

The solution $\tilde{\boldsymbol{\nu}}$ to this system exists and any such solution generates the counterfactual $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}})$.

PROOF: We construct the solution $\tilde{\boldsymbol{\nu}}$ as

$$\tilde{\boldsymbol{\nu}} \equiv \tilde{\mathcal{A}}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}}) - \mathcal{A}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}}).$$

By the definition of the functions of $x(\bullet)$ and $z(\bullet)$, we know that

$$\mathcal{H}\big(x(\tilde{\boldsymbol{\nu}}), z(\tilde{\boldsymbol{\nu}})\big) = \mathbf{0}, \tag{A.15}$$

$$\mathcal{A}\big(x(\tilde{\boldsymbol{\nu}}), z(\tilde{\boldsymbol{\nu}})\big) + \tilde{\mathcal{A}}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}}) - \mathcal{A}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}}) = \mathbf{0}. \tag{A.16}$$

Similarly, by the definition of the functions $\tilde{x}(\bullet)$ and $\tilde{z}(\bullet)$, we also know that

$$\mathcal{H}\big(\tilde{x}(\mathbf{0}), \tilde{z}(\mathbf{0})\big) = \mathcal{H}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}}) = \mathbf{0}, \tag{A.17}$$

$$\tilde{\mathcal{A}}\big(\tilde{x}(\mathbf{0}), \tilde{z}(\mathbf{0})\big) = \tilde{\mathcal{A}}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}}) = \mathbf{0}. \tag{A.18}$$

Since the system (A.15)–(A.16) by assumption has a unique solution for any $\tilde{\boldsymbol{\nu}}$, it thus follows that we must have $\{x(\tilde{\boldsymbol{\nu}}) = \tilde{\boldsymbol{x}}, z(\tilde{\boldsymbol{\nu}}) = \tilde{\boldsymbol{z}}\}$.

We now show that any solution to (A.14) must generate $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}})$. Proceeding by contradiction, consider any other $\tilde{\boldsymbol{\nu}}$ that solves (A.14) and suppose that either $x(\tilde{\boldsymbol{\nu}}) \neq \tilde{\boldsymbol{x}}$ and/or $z(\tilde{\boldsymbol{\nu}}) \neq \tilde{\boldsymbol{z}}$. By definition of the functions $x(\bullet)$ and $z(\bullet)$ together with the property (A.14), we know that

$$\mathcal{H}\big(x(\tilde{\boldsymbol{\nu}}), z(\tilde{\boldsymbol{\nu}})\big) = \mathbf{0},$$

$$\tilde{\mathcal{A}}\big(x(\tilde{\boldsymbol{\nu}}), z(\tilde{\boldsymbol{\nu}})\big) = \mathbf{0},$$

and so $(x(\tilde{\boldsymbol{\nu}}), z(\tilde{\boldsymbol{\nu}}))$ is a solution of (A.10) and (A.12) that is distinct from $(\tilde{x}, \tilde{z})$. But by assumption, only one such solution exists, so we have a contradiction. *Q.E.D.*

Informational Requirements

To construct the desired policy counterfactual for all possible alternative policy rules, we in general need to be able to evaluate the functions $x(\bullet)$ and $z(\bullet)$ for every possible $\boldsymbol{\nu} \in \mathbb{R}^{n_z \times N}$. That is, we need to know the effects of policy shocks of all possible sizes at all possible dates and all possible histories.

To understand how our baseline analysis relaxes these informational requirements, it is useful to proceed in two steps: first removing uncertainty (but keeping non-linearity), and then moving to a linear system.

1. **Non-linear perfect foresight.** For a non-linear perfect foresight economy, we replace our general $(n_x + n_z) \times N$-dimensional system with an $(n_x + n_z) \times T$-dimensional one:

$$\mathcal{H}(\boldsymbol{x}, \boldsymbol{z}, \boldsymbol{\varepsilon}) = \mathbf{0},$$

$$\mathcal{A}(\boldsymbol{x}, \boldsymbol{z}) + \boldsymbol{\nu} = \mathbf{0}.$$

   Because of the lack of uncertainty, other possible realizations of the exogenous events do not matter—only the particular time path, now denoted $\boldsymbol{\varepsilon}$, is relevant. Proceeding exactly in line with the analysis above, we can conclude that now we need the causal effects of all possible policy shocks $\boldsymbol{\nu} \in \mathbb{R}^{n_z \times T}$ *at the equilibrium path induced by* $\boldsymbol{\varepsilon}$. Thus, since we only care about the actual realized history of the exogenous inputs, the dimensionality of the informational requirements has been reduced substantially.
2. **Linear perfect foresight/first-order perturbation.** Linearity further reduces our informational requirements in two respects. First, because of linearity, to know the effects of every possible $\boldsymbol{\nu} \in \mathbb{R}^{n_z \times T}$, it suffices to know the effects of $n_z \times T$ distinct paths $\boldsymbol{\nu}$ that together span $\mathbb{R}^{n_z \times T}$. Second, estimates given any possible exogenous state path of the economy suffice, simply because the effects of policy and non-policy shocks are additively separable. We have thus reduced the problem to the (still formidable) one of finding the effects of $n_z \times T$ distinct policy shock paths.

A.9. *Non-Linear Policy Rules*

We emphasize that the simplicity of our baseline identification results in Section 2.3 relative to the much more involved discussion in Appendix A.8 hinges on linearity of the non-policy block (6). Non-linear policy blocks, on the other hand, are straightforward to handle with the same informational requirements as in our baseline analysis. Specifically, suppose the object of interest is the counterfactual propagation of a shock sequence $\boldsymbol{\varepsilon}$ under the following example of a non-linear policy rule:

$$\boldsymbol{z} = \max\{\tilde{\mathcal{A}}_x \boldsymbol{x}, \underline{z}\}. \tag{A.19}$$

Compared to (8), this policy rule is simplified to set $\tilde{\mathcal{A}}_z = -I$, but then enriched to allow for a canonical form of non-linearity—a kink in the policy rule, for example, a zero lower bound constraint on nominal interest rates. Proceeding perfectly analogously to the proof of Proposition 1, we see that we can construct counterfactuals corresponding to (A.19) by solving for a sequence of policy shocks $\tilde{\boldsymbol{\nu}}$ such that (A.19) holds for all $t = 0, 1, \ldots$. Intuitively, by our unchanged assumptions on the non-policy block (6), it is still just the time path of the policy instrument that matters for private-sector behavior, irrespective of whether this time path is generated by a linear rule like (7) or a non-linear rule like (A.19).

APPENDIX B: DETAILS FOR EMPIRICAL METHOD

This appendix provides econometric implementation details for our empirical method. Appendix B.1 begins with counterfactuals for a given alternative policy rule, while Appendix B.2 discusses optimal policy counterfactuals.

B.1. *Policy Rule Counterfactuals*

The solution to the problem (31) is given as

$$\begin{aligned} \boldsymbol{s} = &-\big[(\tilde{\mathcal{A}}_x \Omega_{x,\mathcal{A}} + \tilde{\mathcal{A}}_z \Omega_{z,\mathcal{A}})'(\tilde{\mathcal{A}}_x \Omega_{x,\mathcal{A}} + \tilde{\mathcal{A}}_z \Omega_{z,\mathcal{A}})\big]^{-1} \\ &\times \big[(\tilde{\mathcal{A}}_x \Omega_{x,\mathcal{A}} + \tilde{\mathcal{A}}_z \Omega_{z,\mathcal{A}})'\big(\tilde{\mathcal{A}}_x \boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}) + \tilde{\mathcal{A}}_z \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon})\big)\big]. \end{aligned}$$

The final step is simply to compute impulse responses to the combination of (i) the original non-policy shock $\boldsymbol{\varepsilon}$ and (ii) the derived policy shocks $\boldsymbol{s}$. For counterfactual second-moment properties (as discussed in Appendix A.5), the only change is that these steps are applied separately for each innovation in the Wold representation of observed macro aggregates.

B.2. *Optimal Policy Counterfactuals*

For our optimal policy counterfactual, we analogously consider the following constrained optimal policy problem:

$$\min_{\boldsymbol{s}} \frac{1}{2} \sum_{i=1}^{n_x} \lambda_i \boldsymbol{x}_i' W \boldsymbol{x}_i \tag{B.1}$$

such that

$$\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{\varepsilon}) + \Omega_{x,\mathcal{A}} \boldsymbol{s}.$$

This gives the optimality conditions:

$$(W \otimes \Lambda)\boldsymbol{x} + \boldsymbol{\varphi}_x = \mathbf{0},$$
$$\Omega'_{x,\mathcal{A}}\boldsymbol{\varphi}_x = \mathbf{0},$$

where $\Lambda = \mathrm{diag}(\lambda_1, \lambda_2, \dots)$.

Solving this system gives our optimal policy counterfactual. The solution is given as

$$\boldsymbol{s} = -\left[\Omega'_{x,\mathcal{A}}(\Lambda \otimes W)\Omega_{x,\mathcal{A}}\right]^{-1} \times \left[\Omega'_{x,\mathcal{A}}(\Lambda \otimes W)\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon})\right].$$

As before, for counterfactual second-moment properties, the analysis is repeated for impulse responses to all Wold innovations separately.

APPENDIX C: SUPPLEMENTARY DETAILS FOR MONETARY APPLICATIONS

This appendix provides further results supplementing the discussion in Section 4 on our systematic monetary policy rule counterfactual applications. Appendices C.1 and C.2 begin by describing the data and our baseline monetary policy shock causal effect estimates. Results for the policy counterfactuals omitted in the main text are presented in Appendix C.3, and we investigate the robustness of our results to the use of other monetary policy shock measures in Appendix C.4. Finally, we in Appendix C.5 illustrate how to use our monetary shock estimates to construct counterfactual second-moment properties.

C.1. *Data*

Our analysis of investment-specific technology shocks follows Ben Zeev and Khan (2015), while our monetary policy shock identification closely mimics that of (i) Romer and Romer (2004) and (ii) Gertler and Karadi (2015).

Outcomes

We are interested in impulse responses of three outcome variables: the output gap, inflation, and the policy rate. For the output gap, we use the series `ygap_hp` of Barnichon and Mesters (2020).[31] For inflation, we compute annual changes in the GDP deflator (using the series `pgdp` from the replication files of Ramey (2016)). Finally, we consider the federal funds rate as our measure of the policy rate, obtained from the St. Louis Federal Reserve FRED database. In keeping with much prior work, we also additionally control for commodity prices, with our measure obtained from the replication files of Ramey (2016) (`lpcom`). All series are quarterly.

Shocks & Identification

We take the investment-specific technology shock series from Ben Zeev and Khan (2015) (`bzk_ist_news` in the replication files of Ramey (2016)), the Romer and Romer (2004) shock series from the replication and extension of Wieland and Yang (2020) (`rr_3`), and the high-frequency monetary policy surprise series from Gertler and Karadi

[31]All results are essentially unchanged if we use a measure of log real GDP instead (`rgdp` scaled by `pop`, taken from the replication files of Ramey (2016)).

(2015) (mp1_tc in the replication files of Ramey (2016)).[32] When applicable, the shock series are aggregated to quarterly frequency through simple averaging.

In Appendix C.4, we examine the robustness of our conclusions to other policy shock series—those of Aruoba and Drechsel (2022) and Miranda-Agrippino and Ricco (2021). For the former, we obtain the shock series directly from their replication files (shock). For the latter, we use the publicly available replication files to construct the SVAR-IV shock series for the full sample (from 1979:M1 onwards), with the shocks constructed at the posterior mode of the estimated reduced-form VAR (the specification for their Figure 3).

C.2. *Shock and Policy Dynamic Causal Effects*

For maximal consistency, we try to estimate all impulse responses within a common empirical specification. For the investment-specific technology shocks, we order the shock measure first in a recursive VAR containing our outcomes of interest (following Plagborg-Møller and Wolf (2021)), estimated on a sample from 1969:Q1 to 2007:Q4. For our two monetary policy shocks, we estimate a single VAR in the two shock series, our three outcomes of interest, as well as commodity prices, also estimated from 1969:Q1 to 2007:Q4.[33] For identification, we order the Gertler and Karadi shock first (again consistent with the results in Plagborg-Møller and Wolf (2021)) and the Romer and Romer shock second-to-last, before the federal funds rate (the additional "exogeneity insurance" as in Romer and Romer (2004)).

We use three lags in the technology shock specification, and four lags in the joint monetary policy VAR. We furthermore estimate all VARs with a constant as well as a deterministic linear trend. For the baseline investment-specific technology shock, we fix the OLS point estimates. We construct policy counterfactuals using our identified monetary policy shocks, taking into account their estimation uncertainty. Since the transmission of both shocks is estimated within a single VAR, we can draw from the posterior and compute the counterfactuals for each draw, thus taking into account joint estimation uncertainty.

Results

The OLS point estimates for the technology shocks of Ben Zeev and Khan (2015) are reported as the dotted lines in Figure 3. For monetary policy, the estimated causal effects for our two outcomes of interest as well as the policy instrument are displayed in Figure C.1. The results are in line with prior work: both policy shocks induce the expected signs of the output gap and inflation responses, though the response shapes are quite distinct, consistent with the differences in the induced interest rate paths. We also note that the magnitudes of the estimated responses are at the lower end of empirical estimates (cf. Table 2 and Figures 1 and 2 in Ramey (2016)).

C.3. *Results for Omitted Monetary Policy Counterfactuals*

In Section 4, we presented detailed results for only three of our policy rule counterfactuals—strict output gap targeting, the Taylor rule, and optimal average inflation targeting policy. We here provide the remaining results.

[32]Results are very similar if we use the alternative surprise series ff4_tc instead.

[33]The Gertler and Karadi shock series is only available from 1988 onwards. We thus follow prior work in the macro IV literature (e.g., Känzig (2021)) and set the missing values to zero.

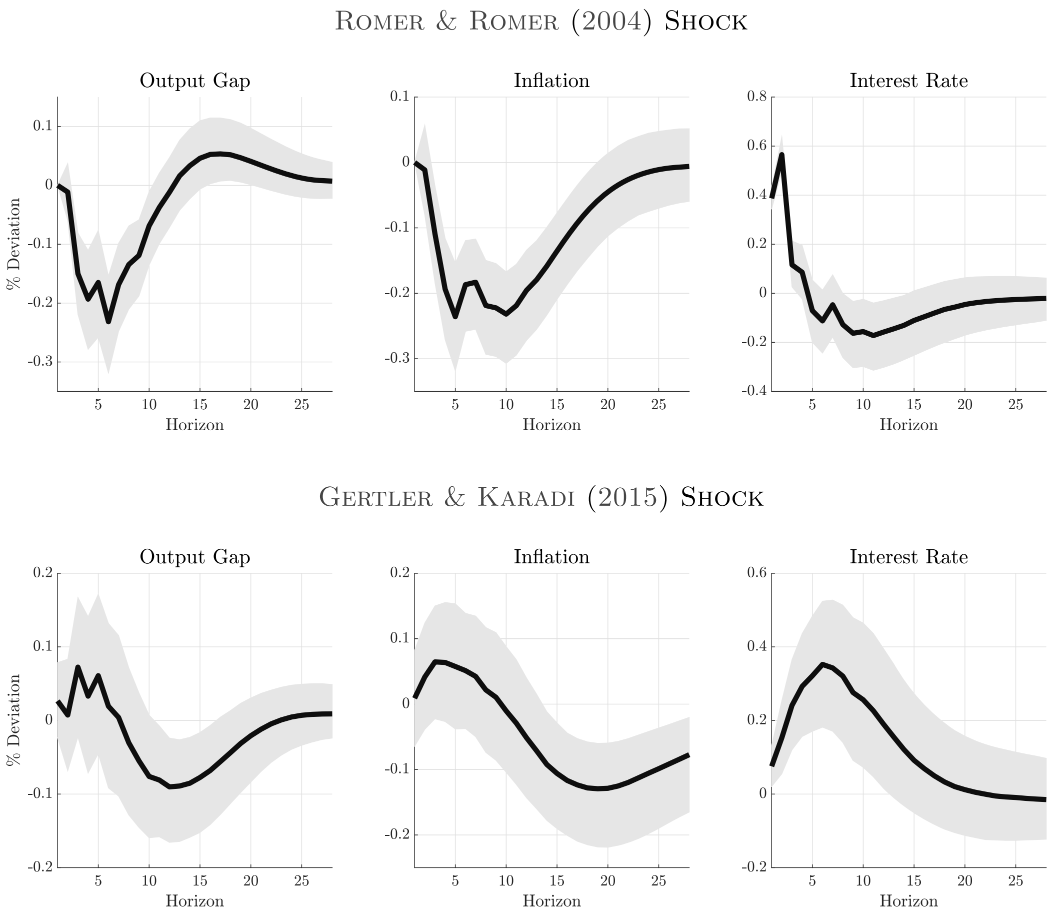

FIGURE C.1.—Impulse responses after the Romer and Romer monetary policy shock (top panel) and the Gertler and Karadi monetary policy shock (bottom panel). The gray areas correspond to 16th and 84th percentile confidence bands, constructed using 10,000 draws from the posterior distribution of the reduced-form VAR parameters.

Nominal Interest Rate Peg

Results for the nominal interest rate peg are presented in Figure C.2. We see that the desired counterfactual policy is implemented well from a couple of quarters out onwards, but that nominal rates are still cut by quite a bit too much immediately after the shock. Since rates are cut by less than in the baseline, the output gap and inflation remain marginally lower for a longer period of time. Compared with the policy counterfactuals discussed in Section 4, we see that a nominal rate peg is a counterfactual policy that is not spanned particularly well by our available monetary shock evidence.

Nominal GDP Targeting

Results for nominal GDP targeting are presented in Figure C.3. The counterfactual policy is implicitly defined by the targeting rule

$$\widehat{\pi}_t + (\widehat{y}_t - \widehat{y}_{t-1}) = 0, \quad \forall t = 0, 1, \ldots$$

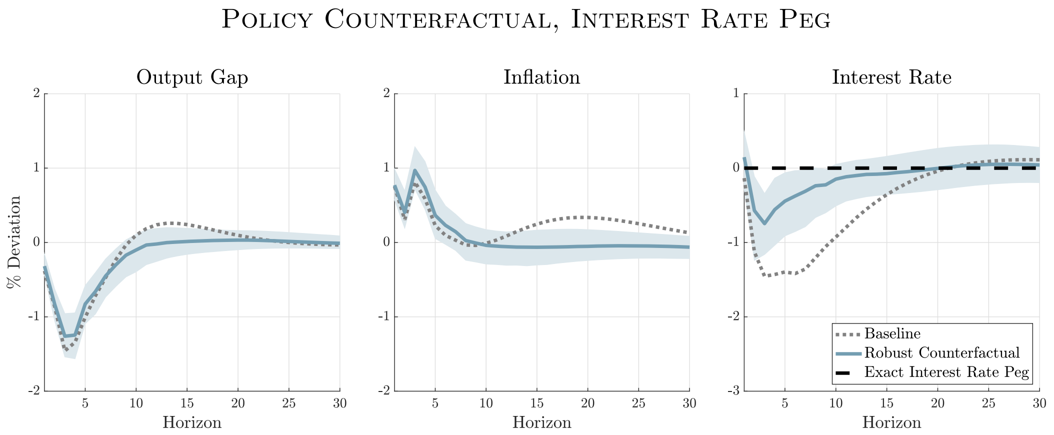

FIGURE C.2.—Output gap, inflation, and interest rate impulse responses to a contractionary investment-specific technology shock under the prevailing baseline rule (dotted) and the best feasible approximation to a nominal rate peg (solid), computed following (31). The shaded areas correspond to 16th and 84th percentile confidence bands. Perfect nominal rate peg (i.e., $\widehat{i}_t = 0$ for all t) is displayed as the black dashed line.

We find that implementation errors are quite small throughout (black dashed). Interestingly, the policy instrument path is quite close to the estimated baseline (dotted gray), indicating that nominal GDP is already stabilized quite well under the prevailing rule.

C.4. *Counterfactuals With Alternative Shock Measures*

Some recent work has questioned the validity of the canonical monetary policy shocks of Romer and Romer and Gertler and Karadi (see Ramey (2016), Nakamura and Steinsson (2018), and the references therein). To examine the robustness of our conclusions to

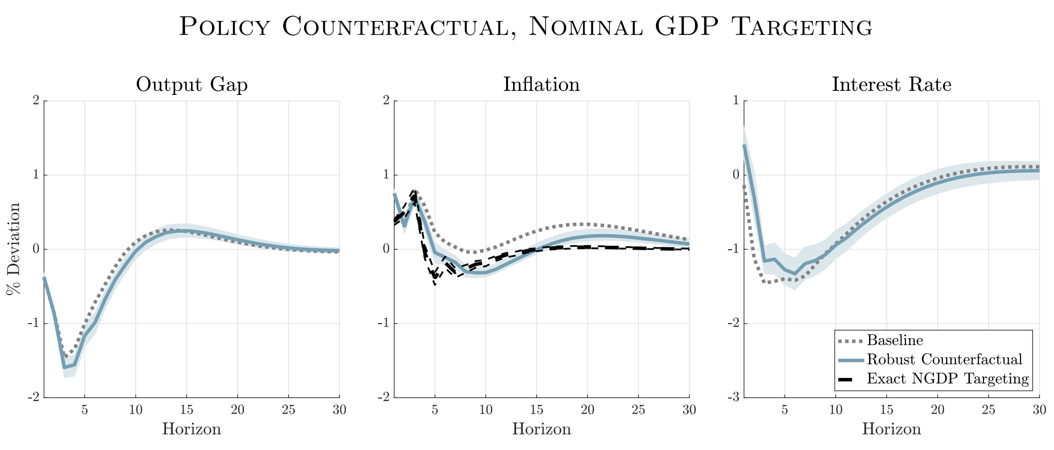

FIGURE C.3.—Output gap, inflation, and interest rate impulse responses to a contractionary investment-specific technology shock under the prevailing baseline rule (dotted) and the best feasible approximation to nominal GDP targeting (solid), computed following (31). The shaded areas correspond to 16th and 84th percentile confidence bands. Given each counterfactual draw for the output gap, the sequence of inflation corresponding to perfect nominal GDP targeting (i.e., so that $\widehat{\pi}_t + (\widehat{y}_t - \widehat{y}_{t-1}) = 0$ for all t) is displayed as the black dashed line.

the use of alternative measures of monetary policy shocks, we now use the policy shock series of Miranda-Agrippino and Ricco (2021) and Aruoba and Drechsel (2022). These shock series are constructed using methods similar to those of Gertler and Karadi and Romer and Romer, but use a richer set of controls for the state of the economy as perceived by the Federal Reserve.

We study the propagation of these shocks in a single integrated VAR, exactly as in our baseline analysis. We find that the two shocks differ in the implied interest rate movements, with the shock of Miranda-Agrippino and Ricco (2021) mirroring the transitory rate movement of Romer and Romer (2004), and the shock of Aruoba and Drechsel (2022) similar to the gradual interest rate movement of Gertler and Karadi (2015). We then leverage these shock estimates to construct monetary policy rule counterfactuals, proceeding exactly as in Section 4. Results for our two main systematic policy rule counterfactuals—output gap targeting and the Taylor rule—are displayed in Figure C.4.

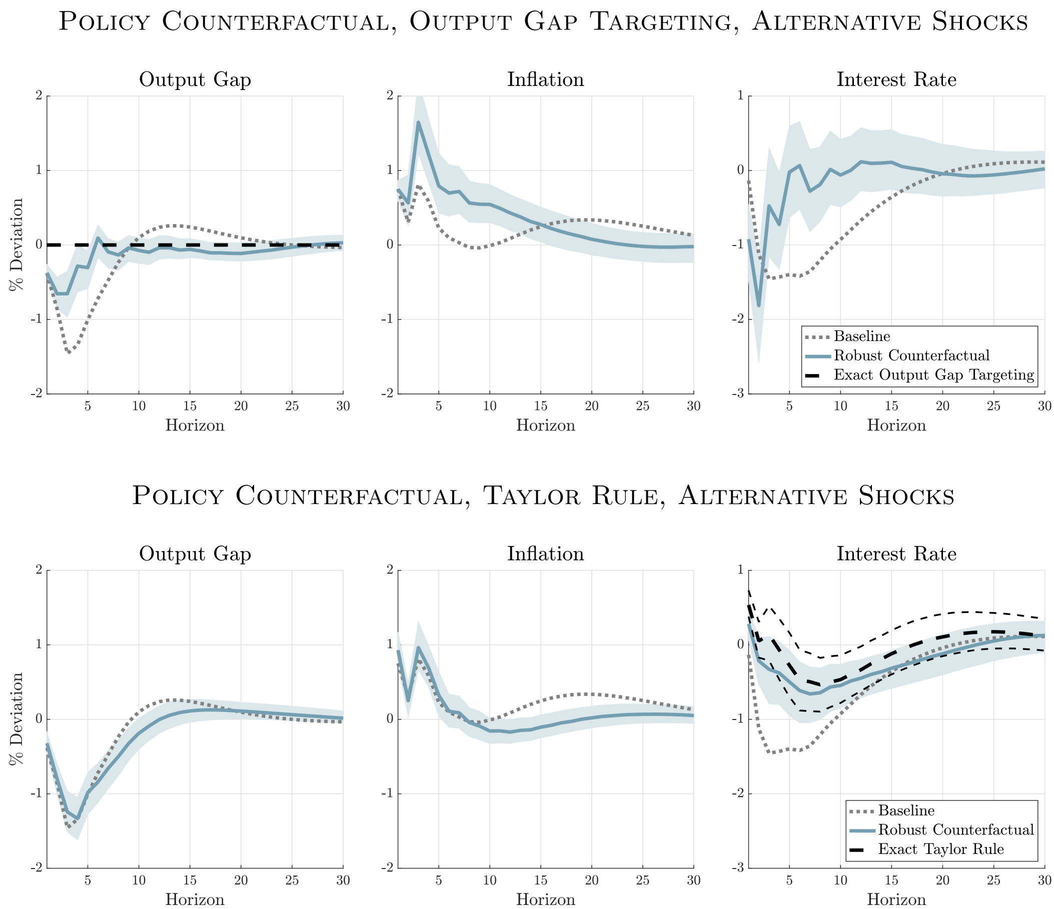

FIGURE C.4.—Output gap, inflation, and interest rate impulse responses to a contractionary investment-specific technology shock under the prevailing baseline rule (dotted gray) and the best feasible approximation to output gap targeting (solid, top panel) and a simple Taylor-type rule $\widehat{i}_t = 0.5\widehat{i}_{t-1} + 0.5 \times (1.5\widehat{\pi}_t + \widehat{y}_t)$ (solid, bottom panel) computed following (31) and using the monetary shocks of Miranda-Agrippino and Ricco (2021) and Aruoba and Drechsel (2022). The shaded areas correspond to 16th and 84th percentile confidence bands.

The main takeaway is that the systematic monetary policy rule counterfactuals are very similar to our headline results. The underlying reason is simply that the impulse responses to the Miranda-Agrippino and Ricco and Aruoba and Drechsel shocks are quite similar to those displayed in Figure C.1 for Romer and Romer and Gertler and Karadi. The perhaps most notable difference is that the shocks of Miranda-Agrippino and Ricco and Aruoba and Drechsel have somewhat *larger* effects on output and inflation (for a given peak interest rate response), so the interest rate cut for the output gap targeting counterfactual is somewhat less steep, and the inflation spike is somewhat more pronounced.

C.5. *Counterfactual Second-Moment Properties*

In this section, we illustrate how estimates of monetary policy shock causal effects can also be used to construct counterfactual *average* business-cycle statistics. Specifically, we construct optimal policy counterfactuals for the average inflation targeting loss function (32).

Our procedure follows the steps outlined in Appendices A.5 and B.2. First, we estimate the Wold representation for our three macroeconomic observables (output gap, inflation, policy rate), giving us impulse responses to the three reduced-form Wold innovations. Then, for each of these three reduced-form shocks, we find the linear combination of date-0 monetary policy shocks that minimizes the policymaker loss function. We then stack these three sets of impulse responses in a new, counterfactual Wold representation, and finally use it to construct counterfactual second-moment properties. We do so for 10,000 draws of monetary policy shock causal effects from our reduced-form VAR.

Results are reported in Table C.I, where the top panel shows business-cycle statistics under observed policy conduct while the bottom panel presents our optimal policy counterfactual. We see that the empirically available subspace of two identified monetary shock paths suffices to somewhat lower the standard deviation of the output gap and inflation. However, the reported gains in aggregate volatility are not substantial, suggesting that rather little policy improvement was feasible within our identified space of policy shock causal effects. These conclusions echo our *conditional* shock conclusions in Section 4.2.

TABLE C.I

COUNTERFACTUAL BUSINESS-CYCLE STATISTICS, OPTIMAL AIT POLICY RULE

	Standard Deviation (per cent)	**Correlation With Output Gap**
Baseline		
Output Gap	1.52	–
Inflation	1.69	0.14
Nominal Rate	2.68	0.27
Counterfactual		
Output Gap	1.27	–
	(1.06, 1.44)	
Inflation	1.38	−0.02
	(1.27, 1.53)	(−0.18, 0.11)
Nominal Rate	2.12	0.39
	(1.73, 2.62)	(0.24, 0.53)

Note: Baseline and counterfactual business-cycle statistics for the best Lucas critique-robust approximation to an optimal average inflation targeting monetary policy rule, computed as discussed in Appendix B.2 applied to each of the three reduced-form innovations in the Wold representation of output gap, inflation, and policy rate. The values in brackets correspond to 16th and 84th percentile confidence bands.

REFERENCES

ANTOLIN-DIAZ, JUAN, IVAN PETRELLA, AND JUAN F. RUBIO-RAMÍREZ (2021): "Structural Scenario Analysis With SVARs," *Journal of Monetary Economics*, 117, 798–815. [6]

ARUOBA, S. BORAGAN, AND THOMAS DRECHSEL (2022): "Identifying Monetary Policy Shocks: A Natural Language Approach," CEPR Discussion Paper No. DP17133. [14,17,18]

AUCLERT, ADRIEN, BENCE BARDÓCZY, MATTHEW ROGNLIE, AND LUDWIG STRAUB (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89 (5), 2375–2408. [2]

BARNICHON, RÉGIS, AND GEERT MESTERS (2020): "Identifying Modern Macro Equations With Old Shocks," *The Quarterly Journal of Economics*, 135 (4), 2255–2298. [13]

BEN ZEEV, NADAV, AND HASHMAT KHAN (2015): "Investment-Specific News Shocks and US Business Cycles," *Journal of Money, Credit and Banking*, 47 (7), 1443–1464. [13,14]

BERNANKE, BEN S., MARK GERTLER, MARK WATSON, CHRISTOPHER A. SIMS, AND BENJAMIN M. FRIEDMAN (1997): "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings papers on economic activity*, 1997 (1), 91–157. [6]

EBERLY, JANICE C., JAMES H. STOCK, AND JONATHAN H. WRIGHT (2020): "The Federal Reserve's Current Framework for Monetary Policy: A Review and Assessment," *International Journal of Central Banking*, 16 (1), 5–71. [6]

GABAIX, XAVIER (2020): "A Behavioral New Keynesian Model," *American Economic Review*, 110 (8), 2271–2327. [2]

GERTLER, MARK, AND PETER KARADI (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7 (1), 44–76. [13-18]

KÄNZIG, DIEGO R. (2021): "The Macroeconomic Effects of Oil Supply News: Evidence From OPEC Announcements," *American Economic Review*, 111 (4), 1092–1125. [14]

KAPLAN, GREG, BENJAMIN MOLL, AND GIANLUCA VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108 (3), 697–743. [3]

LUCAS, ROBERT (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4 (2), 103–124. [4]

MANKIW, N. GREGORY, AND RICARDO REIS (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117 (4), 1295–1328. [2]

MCKAY, ALISDAIR, AND CHRISTIAN K. WOLF (2022): "Optimal Policy Rules in HANK," Technical report, Working Paper, FRB Minneapolis. [3]

MIRANDA-AGRIPPINO, SILVIA, AND GIOVANNI RICCO (2021): "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 13 (3), 74–107. [14,17,18]

NAKAMURA, EMI, AND JÓN STEINSSON (2018): "Identification in Macroeconomics," *Journal of Economic Perspectives*, 32 (3), 59–86. [16]

PLAGBORG-MØLLER, MIKKEL, AND CHRISTIAN K. WOLF (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89 (2), 955–980. [6,14]

RAMEY, VALERIE A. (2016): "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2, 71–162. [13,14,16]

ROMER, CHRISTINA D., AND DAVID H. ROMER (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94 (4), 1055–1084. [13-18]

SIMS, CHRISTOPHER A., AND TAO ZHA (1995): "Does Monetary Policy Generate Recessions?" Working Paper. [6,9]

TAYLOR, JOHN B. (1999): "A Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules*. University of Chicago Press, 319–348. [17]

WIELAND, JOHANNES F., AND MU-JEUNG YANG (2020): "Financial Dampening," *Journal of Money, Credit and Banking*, 52 (1), 79–113. [13]

WOLF, CHRISTIAN K. (2021): "Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result," Working Paper. [3]

Co-editor Charles I. Jones handled this manuscript.

Manuscript received 8 August, 2022; final version accepted 14 April, 2023; available online 18 April, 2023.