What Can Time-Series Regressions Tell Us About Policy Counterfactuals?[†]

Alisdair McKay
FRB Minneapolis

Christian K. Wolf
MIT & NBER

Preliminary and incomplete – comments welcome!

July 2021

Abstract: We show that, in a general family of linearized structural macroeconomic models, knowledge of the dynamic causal effects of contemporaneous and news shocks to the prevailing policy rule can be used to: (a) construct counterfactuals under alternative policy rules; and (b) recover the optimal policy rule corresponding to a given loss function. Under our assumptions, the derived counterfactuals and optimal policies are robust to the Lucas critique. We then discuss strategies for implementing these methods in the empirically relevant case of a limited amount of evidence on policy shock transmission.

[†]Email: alisdair.mckay@mpls.frb.org and ckwolf@mit.edu.

We received helpful comments from Martin Beraja. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

A core principle of modern macroeconomics is that valid policy counterfactuals require a fully-specified structural model. This methodological rule dates back to Lucas (1976), who argued that reduced-form estimates are unlikely to be invariant to changes in policy and so will invariably lead to invalid counterfactuals.

In this paper we revisit the Lucas critique in the context of the linearized structural models that are typically used for macroeconomic policy analysis. We argue that, through the lens of such models, the information contained in empirical time-series regressions – most commonly in the form of structural vector autoregressions or local projections – is in principle sufficient to sidestep the Lucas critique, in the sense that it: (a) allows researchers to construct counterfactuals for arbitrary alternative policy rules; and (b) fully characterizes the optimal policy rule corresponding to a policymaker's loss function. The first part of the paper shows precisely what information is needed for (a) and (b). A key challenge for our approach is that existing empirical evidence falls somewhat short of those high informational requirements. In the second part of the paper we thus present several strategies that allow researchers to leverage the actually available, limited empirical evidence. We argue that, at least for monetary policy counterfactuals, the existing evidence is already sufficient to yield sharp conclusions about several policy counterfactuals and optimal policy rules.

Our results apply to structural models in which private sector behavior depends on policy *only* through the expected path of the policy instrument, and not the coefficients of the policy rule.[1] This property is a feature shared by essentially all linearized business-cycle models, from the analysis of fiscal policy in the real business-cycle model (e.g., Baxter & King, 1993), to monetary policy in the New Keynesian model (e.g., Woodford, 2011; Galí, 2015), to richer models with many frictions and shocks (e.g., Christiano et al., 2005; Smets & Wouters, 2007), to the more recent models with rich consumer and firm heterogeneity (e.g., Kaplan et al., 2018; Ottonello & Winberry, 2020). The key implication of this model property is that we can equivalently re-interpret any arbitrary policy rule as the prevailing baseline rule together with a suitable set of shocks to current policy and expectations of future policy. Intuitively, there always exists *some* sequence of shocks to the baseline rule that changes expectations of the policy instrument just like a given change in the policy rule

[1] More precisely, the policy rule matters through the path of the instrument and through equilibrium selection. Our assumptions on equilibrium existence and uniqueness for the different rules that we consider address equilibrium selection.

itself. But then it follows that, if an econometrician can estimate the dynamic causal effects of a sufficiently rich menu of shocks to the baseline rule – something that is in principle feasible using semi-structural time-series methods (see Ramey, 2016) – then she in fact has enough information to construct counterfactuals for any alternative policy rule. For example, given impulse responses for some shock of interest (e.g., to TFP) under the prevailing rule, she can use her evidence on policy shocks relative to that prevailing rule to construct counterfactual impulse responses that would correspond to a different policy rule. By the same logic, given a policymaker objective function, she can derive a forecast target criterion. Such a forecast target criterion then fully characterizes optimal policy (Svensson, 1997; Woodford, 2011).[2]

The main challenge to implementing these insights is that existing empirical evidence on policy shocks is limited. Formally, the identification result requires that the econometrician can estimate the causal effects of the full menu of possible contemporaneous and news shocks to the prevailing policy rule. For example, in the context of monetary policy, she would need to know the effects of shocks to interest rates at every single point along the yield curve. Such fine-grained, maturity-by-maturity evidence is not available. We present two complementary approaches to dealing with this lack of data.

The first approach realizes that, in the face of incomplete empirical evidence, our results apply without any change to the subspace of dynamic causal effects spanned by that evidence. Intuitively, we can still construct counterfactuals for alternative policy rules that deviate from the prevailing baseline in a way consistent with the empirically identified shock paths. By the same token, we show that any single policy shock estimate provides one additional restriction on the optimal policy rule. Infinitely many restrictions are in principle needed to fully pin down the optimal rule, but a few observed shock paths may already materially restrict it. We illustrate these insights in the context of monetary policy transmission. For this, we use three different identification schemes to identify distinct monetary shocks with information about three different mixtures of current and future changes in interest rates: the shocks identified by Gertler & Karadi (2015) are primarily informative about the effects of changes in contemporaneous rates, while those of Romer & Romer (2004) and Christiano et al. (1999) contain information about joint changes of contemporaneous and future rates. We then apply these estimates to construct output and inflation counterfactuals following a contractionary investment-specific technology shock under a counterfactual rule that aggressively stabilizes

[2] To be clear, our identification results are silent on the shape of the objective function. Explicit, fully specified structural models are one way to arrive at such objective functions. However, given that objective functions in practice are often derived from a legislated mandate rather than economic theory, we believe it is useful to have a method of calculating optimal policy for an objective function that is taken as given.

aggregate output.

Our second approach uses additional information to extrapolate evidence on the policy shocks that we did observe into impulse responses for the ones that we did not. Our starting point is that economic theory imposes restrictions on how impulse responses across variables and across shocks at different horizons are tied together. For example, a dynamic Phillips curve relationship – one block of many modern macro models – implies strong restrictions on the joint behavior of output and inflation. We propose to use those restrictions in conjunction with evidence on monetary policy shock transmission to characterize *optimal* monetary policy, proceeding in two steps. First, we leverage the monetary policy shock evidence as a way to pin down the dynamic Phillips curve (as in Barnichon & Mesters, 2020a). Then, we use our identification results together with the restrictions implied by this Phillips curve to (i) fully characterize the optimal policy rule and (ii) recover the optimal policy response to investment-specific technology shocks. We emphasize that this derived optimal policy response is semi-structural in the sense that any model consistent with our estimated Phillips curve relationship will agree on this policy response – independent of any details of consumer behavior, firm financial frictions, expectation formation, and so on.

Literature. Our analysis connects with two main strands of literature.

First, we study the relationship between fully structural policy counterfactuals and semi-structural estimates of policy shock transmission (Lucas, 1976; Sims, 1980). Previous work has used the dynamic causal effects corresponding to a single policy shock to construct policy counterfactuals (Bernanke et al., 1997; Sims & Zha, 2006). Since this approach relies on a sequence of *unanticipated* shocks, it is well-known to be subject to the Lucas critique. Our approach instead relies on the full menu of contemporaneous *and news* shocks, thus sidestepping those concerns.[3] Our work is also related to Barnichon & Mesters (2020b), who use the dynamic causal effects of policy shocks to evaluate the change in a policymaker's objective function following a small change in policy. They argue that such a perturbation is not subject to the Lucas critique because it is small and therefore would not be considered a change in policy regime by the private sector. The logic underlying our results is quite different: we argue that the Lucas critique can be sidestepped not because expectations do not change, but because suitably chosen (news) shocks change expectations just like a counterfactual change in the policy rule. This argument – which is made possible by the restrictions

[3] These concerns are similarly sidestepped in Beraja (2020). His analysis has weaker informational requirements than ours, but relies on stronger theoretical restrictions. The two approaches are thus complementary.

implied by our model environment – is what allows us to form policy counterfactuals and compute optimal policy rules.

Second, our identification results build on recent advances in solution methods for dynamic general equilibrium models. As in Auclert et al. (2021), we characterize model equilibria using sequence-space methods; and as in Guren et al. (2021) and Wolf (2020), we connect this sequence-space representation to empirically estimable objects. In contemporaneous and independent work, Hebden & Winker (2021) show how to use the same arguments as in our identification results to efficiently *compute* policy counterfactuals in structural models with occasionally binding constraints. Our focus is not computational – we aim to calculate policy counterfactuals directly from empirical evidence, forcing us to confront the fact that such evidence is limited.

Outline. The remainder of the paper proceeds as follows. Section 2 presents our main identification results. Section 3 then discusses our two approaches to dealing with realistic data limitations, and applies our results to construct monetary policy counterfactuals and an optimal dual mandate rule. Section 4 concludes.

2 Dynamic causal effects & policy counterfactuals

This section contains our core identification results. We begin in Section 2.1 by first describing the class of models to which our arguments apply and then defining our two objects of interest: (a) outcomes under alternative policy rules and (b) optimal policy for a given loss function. Section 2.2 shows that (rich) evidence on the dynamic causal effects of shocks to the prevailing rule is sufficient to recover both (a) and (b). Section 2.3 then concludes with a brief numerical illustration.

2.1 Model & objects of interest

We consider a linearized perfect foresight model economy. By certainty equivalence, all results below extend without change to models with aggregate risk linearized around their deterministic steady state.[4] Throughout, boldface denotes time paths for $t = 0, 1, 2, \ldots$, and all variables are expressed in deviations from the model's deterministic steady state.

[4] See Auclert et al. (2021) for a detailed discussion of this point.

The model economy is summarized by the equilibrium system

$$\mathcal{H}_w \boldsymbol{w} + \mathcal{H}_x \boldsymbol{x} + \mathcal{H}_z \boldsymbol{z} + \mathcal{H}_\varepsilon \boldsymbol{\varepsilon} = \mathbf{0} \tag{1}$$

$$\mathcal{A}_x \boldsymbol{x} + \mathcal{A}_z \boldsymbol{z} + \boldsymbol{\nu} = \mathbf{0} \tag{2}$$

w and x are n_w- and n_x-dimensional vectors of endogenous variables, z is a n_z-dimensional vector of policy variables, ε is a n_ε-dimensional vector of exogenous shock sequences, and ν is an n_z-dimensional vector of policy shocks. The distinction between w and x is that all variables in x are observable to the policymaker and econometrician alike, while the variables in w are not. The infinite-dimensional linear maps $\{\mathcal{H}_w, \mathcal{H}_x, \mathcal{H}_z, \mathcal{H}_\varepsilon\}$ summarize the non-policy block of the economy, yielding $n_w + n_x$ restrictions for each t. Our key assumption is that these linear maps do not depend in any way on the coefficients of the policy rule $\{\mathcal{A}_x, \mathcal{A}_z\}$; instead, policy only matters through the path of the instrument z, with the rule (2) giving n_z restrictions for each t.

Given $\{\boldsymbol{\varepsilon}, \boldsymbol{\nu}\}$, an equilibrium is a set $\{\boldsymbol{w}, \boldsymbol{x}, \boldsymbol{z}\}$ that solves (1) - (2). We assume that the baseline policy rule $\{\mathcal{A}_x, \mathcal{A}_z\}$ is such that an equilibrium exists and is unique for any $\{\boldsymbol{\varepsilon}, \boldsymbol{\nu}\}$.

Assumption 1. *The policy rule in (2) induces a unique and determinate equilibrium. That is, the infinite-dimensional linear map*

$$\mathcal{B} \equiv \begin{pmatrix} \mathcal{H}_w & \mathcal{H}_x & \mathcal{H}_\varepsilon \\ \mathbf{0} & \mathcal{A}_x & \mathcal{A}_z \end{pmatrix}$$

is invertible.

Given $\{\boldsymbol{\varepsilon}, \boldsymbol{\nu}\}$, we write that unique solution as $\{\boldsymbol{w}_\mathcal{A}(\boldsymbol{\varepsilon}, \boldsymbol{\nu}), \boldsymbol{x}_\mathcal{A}(\boldsymbol{\varepsilon}, \boldsymbol{\nu}), \boldsymbol{z}_\mathcal{A}(\boldsymbol{\varepsilon}, \boldsymbol{\nu})\}$. Most interest will center on impulse responses to exogenous shocks $\boldsymbol{\varepsilon}$ when the policy rule is followed perfectly ($\boldsymbol{\nu} = \mathbf{0}$); with some slight abuse of notation we will simply write those impulse responses as $\{\boldsymbol{w}_\mathcal{A}(\boldsymbol{\varepsilon}), \boldsymbol{x}_\mathcal{A}(\boldsymbol{\varepsilon}), \boldsymbol{z}_\mathcal{A}(\boldsymbol{\varepsilon})\}$

SCOPE. Our results in the remainder of this paper will apply to *any* structural model that can be written in the general form (1) - (2). As emphasized, the key property of the model is that policy matters for the non-policy block *only* through the realized path of the policy variables $\boldsymbol{z}$; equivalently, in the linearized economy with aggregate risk, policy matters only through its effects on the expected future path of the instrument z. While restrictive, the separation between policy rule and non-policy model blocks that our results require is a common feature of the types of structural models that are typically used for (quantitative) counterfactual policy analysis. To illustrate this point, we show here how to map both a simple three-equation New Keynesian model as well as a much richer Heterogeneous Agent New Keynesian (HANK) model into the general structure (1) - (2).

The canonical small-scale New Keynesian model consists of a New Keynesian Phillips Curve (NKPC),

$$\boldsymbol{\pi} = \kappa \boldsymbol{y} + \beta \boldsymbol{\pi}_{+1} + \boldsymbol{\varepsilon}^s, \tag{3}$$

a consumer demand block (or IS relation),

$$-\boldsymbol{y} = -\boldsymbol{y}_{+1} + \sigma(\boldsymbol{i}_b - \boldsymbol{\pi}_{+1}) + \boldsymbol{\varepsilon}^d, \tag{4}$$

and a monetary policy rule,

$$\boldsymbol{i}_b = \phi_{i_b} \boldsymbol{i}_{b,-1} + (1 - \phi_{i_b})(\phi_\pi \boldsymbol{\pi} + \phi_y \boldsymbol{y}) + \boldsymbol{\nu}, \tag{5}$$

where π is inflation, y is output, i_b is the nominal rate of interest, $(\varepsilon^s, \varepsilon^d)$ are supply and demand shocks, respectively, and ν is the policy shock. The parameters β and σ are deep parameters that come from household preferences, while the parameter κ is a composite parameter that reflects household preferences and nominal rigidities. Crucially, none of them depend on the policy rule. This simple model fits into the structure (1) - (2) for $w = \{\}$, $x = (\pi, y)$, $z = i_b$ and $\varepsilon = (\varepsilon^s, \varepsilon^d)$, with (3) - (4) constituting the policy-invariant non-policy block (1), and (5) as the policy rule (2).

For a second example consider the HANK model of Wolf (2021). This model also features an NKPC of the form (3), and the monetary policy rule is unchanged and given as (5). The fiscal authority passively adjusts lump-sum taxes and transfers to balance the budget, so we have that

$$\mathbf{0} = \bar{\tau} \boldsymbol{\tau} + \bar{b}(1 + \bar{i}_b)(\boldsymbol{i}_{b,-1} - \boldsymbol{\pi}) \tag{6}$$

where τ denotes lump-sum taxes/transfers, bars indicate steady state, and $\bar{b}$ is the steady-state level of government debt. Finally, the consumer block (or IS relation) now becomes

$$\boldsymbol{y} = \mathcal{C}_y \boldsymbol{y} + \mathcal{C}_\pi \boldsymbol{\pi} + \mathcal{C}_{i_b} \boldsymbol{i_b} + \mathcal{C}_\tau \boldsymbol{\tau} + \boldsymbol{\varepsilon}^d. \tag{7}$$

The coefficient matrices in (7) are derived from aggregating the partial equilibrium behavior of household consumption decisions and thus again do not depend on policy rules.[5] This

[5]Decisions do not depend on dynamic policy rules of the sort (3), applied around a given deterministic

model fits into the structure (1) - (2) with $w = \tau$, $x = (\pi, y)$, $z = i_b$ and $\varepsilon = (\varepsilon^s, \varepsilon^d)$, and (6) - (7) added to the block (1).[6]

Policy counterfactuals. In the context of the general structural model represented by (1) - (2) we now want to learn about the following two sets of policy counterfactuals.

1. Arbitrary alternative rules. Consider an alternative policy rule

$$\tilde{\mathcal{A}}_x \boldsymbol{x} + \tilde{\mathcal{A}}_z \boldsymbol{z} = \mathbf{0} \tag{8}$$

   Just like the baseline rule, this alternative policy rule is also assumed to induce a unique, determinate equilibrium.

   Assumption 2. *The policy rule in (8) induces a unique and determinate equilibrium. That is, the infinite-dimensional linear map*

$$\tilde{\mathcal{B}} \equiv \begin{pmatrix} \mathcal{H}_w & \mathcal{H}_x & \mathcal{H}_\varepsilon \\ \mathbf{0} & \tilde{\mathcal{A}}_x & \tilde{\mathcal{A}}_z \end{pmatrix}$$

   is invertible.

   Given this alternative rule, we now ask: what are the dynamic response paths $\boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})$ and $\boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})$ to the exogenous shock path $\boldsymbol{\varepsilon}$?

2. Optimal policy. Consider a policymaker with a quadratic loss function of the form

$$\mathcal{L} = \sum_{i=1}^{n_x} \lambda_i \boldsymbol{x}_i' W \boldsymbol{x}_i \tag{9}$$

   where i indexes the n_x distinct macroeconomic aggregates collected in x, λ_i denotes policy weights, and W is a symmetric positive-definite discounting matrix.[7] We assume that the optimal policy problem has a unique solution.

---

steady state. Permanent changes in fiscal policy (e.g., the tax-and-transfer system) of course do affect (7).

[6] The only actual policy choice here is the nominal rate i_b. Lump-sum taxes – which passively adjust to balance the budget – are thus part of the policy-invariant block (1).

[7] Our focus on a separable quadratic loss functions is in line with standard optimal policy analysis, but not essential. As shown in Appendix A.1, our results extend to the non-separable quadratic case, where the loss is now given by $\boldsymbol{x}' Q \boldsymbol{x}$ for a weighting matrix Q. While our approach in principle also applies to even richer loss functions, the resulting policy rule will generally not fit into the form (2).

Assumption 3. *Given any vector of exogenous shocks* $\boldsymbol{\varepsilon}$*, the problem of choosing the policy variable* $\boldsymbol{z}$ *to minimize the loss function* (9) *subject to the non-policy constraint* (1) *has a unique solution.*

We are interested in two questions. First, what policy rule is optimal for such a policy-maker? Second, given that optimal rule $(\mathcal{A}_x^*, \mathcal{A}_z^*)$, what are the corresponding dynamic response paths $\boldsymbol{x}_{\mathcal{A}^*}(\boldsymbol{\varepsilon})$ and $\boldsymbol{z}_{\mathcal{A}^*}(\boldsymbol{\varepsilon})$?

The objective of the remainder of this section is to characterize the information required to recover these desired policy counterfactuals. Crucially, all of the information we require will be available from data generated under the baseline policy rule.

2.2 Identification

We begin by defining the dynamic causal effects that lie at the heart of our identification results. By Assumption 1, we can write the solution to the system (1) - (2) as

$$\begin{pmatrix} \boldsymbol{w} \\ \boldsymbol{x} \\ \boldsymbol{z} \end{pmatrix} = \underbrace{-\mathcal{B}^{-1} \times \begin{pmatrix} \mathcal{H}_\varepsilon & \mathbf{0} \\ \mathbf{0} & I \end{pmatrix}}_{\equiv \Theta_\mathcal{A}} \times \begin{pmatrix} \boldsymbol{\varepsilon} \\ \boldsymbol{\nu} \end{pmatrix}$$

The linear map $\Theta_\mathcal{A}$ collects the impulse responses of $\boldsymbol{w}$, $\boldsymbol{x}$ and $\boldsymbol{z}$ to the non-policy and policy shocks $(\boldsymbol{\varepsilon}, \boldsymbol{\nu})$ under the prevailing, baseline policy rule (2).[8] We will partition it as

$$\Theta_\mathcal{A} \equiv \begin{pmatrix} \Theta_{w,\varepsilon,\mathcal{A}} & \Theta_{w,\nu,\mathcal{A}} \\ \Theta_{x,\varepsilon,\mathcal{A}} & \Theta_{x,\nu,\mathcal{A}} \\ \Theta_{z,\varepsilon,\mathcal{A}} & \Theta_{z,\nu,\mathcal{A}} \end{pmatrix}. \tag{10}$$

All of our identification results will require knowledge of $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ – the *full* sets of dynamic causal effects for policy shocks. In words, the researcher needs to know the dynamic causal effects of every possible current and future (announced) deviation from the prevailing policy rule on the policy instruments z as well as the (observable) endogenous variables x. Furthermore, to construct counterfactual paths that correspond to a given shock sequence

[8] Note that, in the analogous linearized economy with aggregate risk, $\Theta_\mathcal{A}$ collects the coefficients of the model's vector moving average representation for the full menu of contemporaneous and news policy and non-policy shocks.

$\boldsymbol{\varepsilon}$, the researcher also needs to know the dynamic causal effects of that shock sequence under the baseline policy, $\{\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon})\}$.

Alternative policy rules. We begin with the first object of interest – policy counterfactuals after a shock sequence $\boldsymbol{\varepsilon}$ under an arbitrary alternative policy rule.

Proposition 1. *Suppose that $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ and $\{\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon})\}$ are known. Then, for any alternative policy rule $\{\tilde{\mathcal{A}}_x, \tilde{\mathcal{A}}_z\}$ that induces a unique, determinate equilibrium, we can recover the policy counterfactuals $\boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})$ and $\boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})$ as the unique solution to the system*

$$\begin{pmatrix} I & \mathbf{0} & -\Theta_{x,\nu,\mathcal{A}} \\ \mathbf{0} & I & -\Theta_{z,\nu,\mathcal{A}} \\ \tilde{\mathcal{A}}_x & \tilde{\mathcal{A}}_z & \mathbf{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{z} \\ \boldsymbol{\nu} \end{pmatrix} = \begin{pmatrix} \boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}) \\ \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon}) \\ \mathbf{0} \end{pmatrix}. \tag{11}$$

Proof. The equilibrium system under the new policy rule can be written as

$$\begin{pmatrix} \mathcal{H}_w & \mathcal{H}_x & \mathcal{H}_z \\ \mathbf{0} & \tilde{\mathcal{A}}_x & \tilde{\mathcal{A}}_z \end{pmatrix} \begin{pmatrix} \boldsymbol{w} \\ \boldsymbol{x} \\ \boldsymbol{z} \end{pmatrix} = \begin{pmatrix} -\mathcal{H}_\varepsilon \\ \mathbf{0} \end{pmatrix} \boldsymbol{\varepsilon} \tag{12}$$

By Assumption 2 we know that (12) has a unique solution $\{\boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})\}$. To characterize this solution as a function of observables, consider instead the alternative system (11). Since (1) also holds under the initial policy rule, and since the last line of (11) imposes the new policy rule, it follows that any $(\boldsymbol{x}, \boldsymbol{z})$ that are part of a solution of (11) are also part of a solution of (12). Since by assumption (12) has a unique solution, it follows that the system (11) is solved by at most one set of paths $(\boldsymbol{x}, \boldsymbol{z})$.

It remains to establish that the system (11) has a solution. For this consider the candidate tuple $\{\boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}), \boldsymbol{\nu} = (\tilde{\mathcal{A}}_x - \mathcal{A}_x)\boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}) + (\tilde{\mathcal{A}}_z - \mathcal{A}_z)\boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})\}$. Since the paths $\{\boldsymbol{w}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}), \boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})\}$ solve (12), it follows that they are also a solution to the system

$$\begin{pmatrix} \mathcal{H}_w & \mathcal{H}_x & \mathcal{H}_z \\ \mathbf{0} & \mathcal{A}_x & \mathcal{A}_z \end{pmatrix} \begin{pmatrix} \boldsymbol{w} \\ \boldsymbol{x} \\ \boldsymbol{z} \end{pmatrix} = - \begin{pmatrix} \mathcal{H}_\varepsilon \boldsymbol{\varepsilon} \\ (\tilde{\mathcal{A}}_x - \mathcal{A}_x)\boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}) + (\tilde{\mathcal{A}}_z - \mathcal{A}_z)\boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}) \end{pmatrix} \tag{13}$$

But by Assumption 1 we know that the system (13) has a unique solution, so indeed the paths $\{\boldsymbol{w}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}), \boldsymbol{x}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\tilde{\mathcal{A}}}(\boldsymbol{\varepsilon})\}$ are that solution. It then follows from the definition of $\Theta_{\mathcal{A}}$ in (10) that the candidate tuple also solves (11), completing the argument. □

The logic of the argument is simple: since we know the effects of all possible perturbations $\boldsymbol{\nu}$ of the baseline rule, we can always construct a perturbation that mimics the equilibrium instrument path under the new instrument rule. But since the first model block (1) depends on the policy rule *only* via the expected instrument path, the equilibrium allocations under the new counterfactual rule and the perturbed prevailing rule are the same.

It is instructive to compare this argument to that of Bernanke et al. (1997) and Sims & Zha (2006). In those analyses, the counterfactual policy rule is enforced along the equilibrium impulse response path by subjecting the economy to a new, suitably scaled *contemporaneous* policy shock ν_0 at each time $t = 0, 1, 2, \ldots$. These surprises were not expected at $t = 0$ (unlike the actual counterfactual policy rule), so the approach is vulnerable to the Lucas critique. Our argument circumvents this problem by using a sufficiently rich menu of contemporaneous *and news* shocks, all realized at $t = 0$.

Optimal policy. The second identification result concerns optimal policy. The logic of the argument is identical to that for the case of arbitrary alternative policy rules.

Proposition 2. *Consider a policymaker loss function (9), and suppose that $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ are known. Then we can recover the optimal policy rule $\{\mathcal{A}_x^*, \mathcal{A}_z^*\}$ as*

$$\mathcal{A}_x^* = \left(\lambda_1 \Theta'_{x_1,\nu,\mathcal{A}} W, \lambda_2 \Theta'_{x_2,\nu,\mathcal{A}} W, \ldots, \lambda_{n_x} \Theta'_{x_{n_x},\nu,\mathcal{A}} W\right) \tag{14}$$

$$\mathcal{A}_z^* = \mathbf{0}. \tag{15}$$

If $\{\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon})\}$ are also known, then we can furthermore recover counterfactuals for the shock path $\boldsymbol{\varepsilon}$ under the optimal policy rule, $\boldsymbol{x}_{\mathcal{A}^*}(\boldsymbol{\varepsilon})$ and $\boldsymbol{z}_{\mathcal{A}^*}(\boldsymbol{\varepsilon})$, by using (11).

Proof. The solution to the optimal policy problem is characterized by the following first-order conditions:

$$\mathcal{H}_w'(I \otimes W)\boldsymbol{\varphi} = \mathbf{0} \tag{16}$$

$$(\Lambda \otimes W)\boldsymbol{x} + \mathcal{H}_x'(I \otimes W)\boldsymbol{\varphi} = \mathbf{0} \tag{17}$$

$$\mathcal{H}_z' W \boldsymbol{\varphi} = \mathbf{0} \tag{18}$$

where $\Lambda = \text{diag}(\lambda_1, \lambda_2, \ldots)$ and φ is the multiplier on (1). By Assumption 3 we know that the system (16) - (18) together with (1) has a unique solution $\{\boldsymbol{x}^*(\boldsymbol{\varepsilon}), \boldsymbol{z}^*(\boldsymbol{\varepsilon}), \boldsymbol{\varphi}^*(\boldsymbol{\varepsilon})\}$.

Now consider the alternative problem of choosing deviations $\boldsymbol{\nu}$ from the prevailing rule to minimize (9) subject to the system (1) - (2). This second problem gives the FOCs

$$\mathcal{H}_w'(I \otimes W)\boldsymbol{\varphi} = \mathbf{0} \tag{19}$$

$$(\Lambda \otimes W)\boldsymbol{x} + \mathcal{H}_x'(I \otimes W)\boldsymbol{\varphi} + \mathcal{A}_x' W \boldsymbol{\varphi}_z = \mathbf{0} \tag{20}$$

$$\mathcal{H}_z'(I \otimes W)\boldsymbol{\varphi} + \mathcal{A}_z' W \boldsymbol{\varphi}_z = \mathbf{0} \tag{21}$$

$$W\boldsymbol{\varphi}_z = \mathbf{0} \tag{22}$$

where φ_z is the multiplier on (2). It follows from (22) that $\boldsymbol{\varphi}_z = \mathbf{0}$. But then (19) - (21) together with (1) determine the same unique solution as before, and $\boldsymbol{\nu}$ adjusts residually to satisfy (2). The original problem and the alternative problem are thus equivalent.

Next note that, by Assumption 1, we can re-write the alternative problem's constraint set as

$$\begin{pmatrix} \boldsymbol{w} \\ \boldsymbol{x} \\ \boldsymbol{z} \end{pmatrix} = \Theta_{\mathcal{A}} \times \begin{pmatrix} \boldsymbol{\varepsilon} \\ \boldsymbol{\nu} \end{pmatrix} \tag{23}$$

The problem of minimizing (9) subject to (23) gives the optimality condition

$$\sum_{i=1}^{n_x} \lambda_i \Theta_{x_i,\nu,\mathcal{A}}' W \boldsymbol{x}_i = 0 \tag{24}$$

By the equivalence of the policy problems, it follows that (24) is an optimal policy *rule*, taking the form (14) - (15). Finally, the second part of the result follows from Proposition 1 since (24) is just a special example of a policy rule.

□

The characterization of the optimal policy rule in (14)-(15) uses the fact that, because we know the dynamic causal effects of every possible policy perturbation $\boldsymbol{\nu}$ on the targets $\boldsymbol{x}$, we know the space of those targets that is implementable through policy actions. At an optimum, we must be at the point of this space that minimizes the policymaker loss (9).

It is also useful to note that, by mapping our perfect foresight economy to a linearized economy with aggregate risk, we can re-write that optimal policy rule as a forecasting targeting rule (Svensson, 1997):

$$\sum_{i=1}^{n_x} \lambda_i \Theta_{x_i,\nu,\mathcal{A}}' W \mathbb{E}_t \left[\boldsymbol{x}_i\right] = 0 \tag{25}$$

where now $\boldsymbol{x}_i = (x_{it}, x_{it+1}, \ldots)'$. In words, expectations of future targets must always minimize the policymaker loss within the space of (expected) allocations that are implementable via changes in the policy stance.[9]

2.3 Quantitative illustration

We use the simple HANK model of Wolf (2021), sketched in Section 2.1, to provide a numerical illustration of Propositions 1 and 2. Details of the model parameterization are relegated to Appendix A.2.

Throughout, our analysis relies on policy causal effects $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ and shock impulse responses $\{\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon})\}$ obtained under a simple baseline rule of the form

$$i_{b,t} = \phi_\pi \pi_t \tag{26}$$

for $\phi_\pi = 1.5$. Our object of interest are counterfactuals under alternative policy rules following a contractionary cost-push shock ε^s to the model's Phillips curve (3).

Alternative policy rules. For our first experiment, we would like to learn about the behavior of output and inflation under an alternative policy rule

$$i_{b,t} = \phi_i i_{b,t-1} + (1-\phi_i)(\phi_\pi \pi_t + \phi_y y_t) \tag{27}$$

for $\phi_i = 0.9$, $\phi_\pi = 2$ and $\phi_y = 0.5$.

To begin, we construct the impulse responses of output and inflation under the two policy rules (26) and (27) – the orange and grey lines in the left and middle panels of Figure 1 – by simply solving the full model twice, once for each rule.

We now use Proposition 1 to equivalently construct the counterfactual under the new rule (27) without actually re-solving the model. We do so using $\{\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon})\}$ and $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ – the dynamic causal effects of the fundamental shock and of policy shocks generated under the prevailing baseline rule (26). We feed these inputs into (11) to solve for $\boldsymbol{x}$, $\boldsymbol{z}$ and $\boldsymbol{\nu}$. The dark blue lines in the left and middle panels of Figure 1 show that, as expected, the solution is identical to that implied by re-solving the model. The right panel then shows the sequence of shocks $\boldsymbol{\nu}$ that maps the baseline prevailing rule into the new rule. Since the new rule is

[9] (25) is the optimal policy rule for a time-0 policy problem. For a timeless perspective, (25) must apply to *revisions* of policymaker expectations at each t.

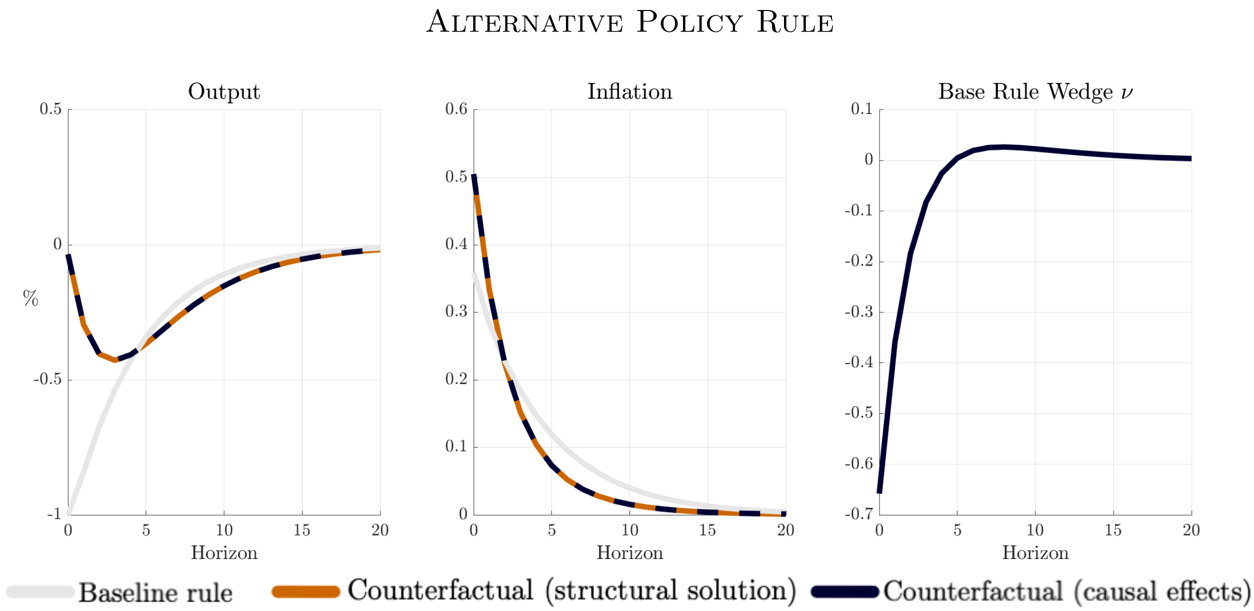

Figure 1: Output and inflation impulse responses together with the equivalence shock wedge $\boldsymbol{\nu}$ (see (11)) for the HANK model with policy rules (26) and (27). The impact output contraction under the prevailing baseline rule is normalized to -1%.

more accommodating, the sequence of shocks is persistently negative.

OPTIMAL POLICY. Our second experiment studies optimal policy under a dual mandate loss function

$$\mathcal{L} = \lambda_\pi \boldsymbol{\pi}' W \boldsymbol{\pi} + \lambda_y \boldsymbol{y}' W \boldsymbol{y} \tag{28}$$

with $W = \text{diag}(1, \beta, \beta^2, \dots)$ and $\lambda_\pi = \lambda_y = 1$. We again start by solving for the optimal policy using conventional methods: we derive the policy rule corresponding to the first-order conditions (16) - (18), solve the model given that policy rule, and report the result as the orange lines in the left and middle panels of Figure 2. 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 (26) tightens too much.

We now instead use Proposition 2 to equivalently recover the optimal policy rule and the corresponding supply shock impulse responses. We begin with the optimal rule itself. By (24), the optimal rule is given as

$$\lambda_\pi \Theta'_{\pi,\nu,\mathcal{A}} W \boldsymbol{\pi} + \lambda_y \Theta'_{y,\nu,\mathcal{A}} W \boldsymbol{y} = 0 \tag{29}$$

Optimal Policy

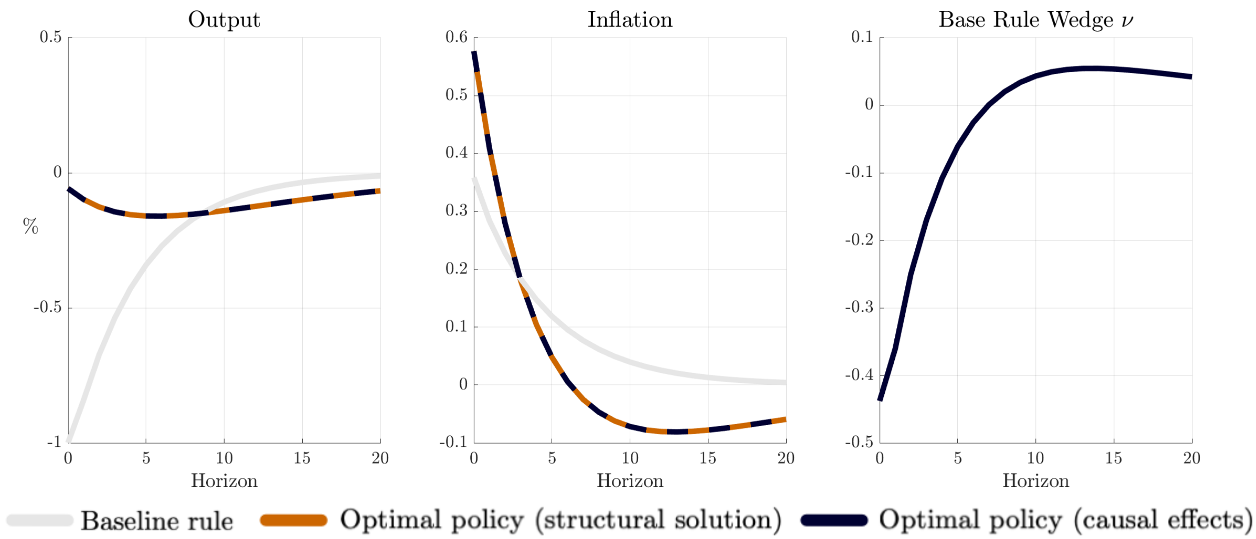

Figure 2: Output and inflation impulse responses together with the equivalence shock wedge $\boldsymbol{\nu}$ (see (11)) for the HANK model with policy rules (26) and the optimal policy given by (28). The impact output contraction under the prevailing baseline rule is normalized to -1%.

A researcher with knowledge of the dynamic causal effects of monetary policy shocks on output and inflation, $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$, is able to construct this optimal policy rule.

Next, we use those same policy causal effects together with the known shock impulse responses $\{\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}), \boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon})\}$ under the baseline rule to recover output and inflation paths under the optimal rule, using (11). Again as expected, the resulting impulse responses – the dark blue lines – are identical to those obtained by explicitly solving the optimal policy problem. Finally, the right panel of Figure 2 again represents this optimal policy as a deviation $\boldsymbol{\nu}$ from the prevailing baseline rule.

3 Estimating policy counterfactuals

The identification results of Section 2 show that, for structural models that fit within the general class (1) - (2), policy counterfactuals can *in principle* be obtained from a large enough number of time-series regressions – the VARs and local projections that have long been used to estimate the dynamic causal effects of policy shocks (Sims, 1980; Ramey, 2016). Unfortunately, while the policy shock literature has undoubtedly made a lot of progress in studying the propagation of policy interventions with different time profiles (e.g., contemporaneous vs. forward guidance monetary shocks), we are far short of knowing the required full menu of policy causal effects $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$.

In light of this challenge, how can we make the identification results in Propositions 1 and 2 useful in practice? In this section we present two ways forward. First, in Section 3.1, we show that the limited information that we do have: (a) can be used to construct counterfactuals for a restricted set of alternative rules that lies in the space spanned by the available evidence on policy shocks; and (b) provides material restrictions on – while falling short of fully identifying – the optimal policy rule. Second, in Section 3.2, we discuss conditions under which the restrictions across impulse responses at different horizons implied by popular *families* of structural models can be used to translate the available partial empirical evidence into an estimate of the required full set of policy causal effects.

3.1 Partial measurement

We first consider the case of a researcher that was only able to estimate the dynamic causal effects of a finite (small) number of particular shock vectors.

Let $\mathcal{V}$ be a linear map whose columns collect the responses of the policy shock $\boldsymbol{\nu}$ to $n_{\tilde{\nu}}$ identified sources of variation in policy. That is, we assume that the researcher knows the dynamic causal effects of shocks to the baseline rule $\boldsymbol{\nu}$ that take the form $\boldsymbol{\nu} = \mathcal{V} \times \tilde{\boldsymbol{\nu}}$, where $\tilde{\boldsymbol{\nu}}$ is $n_{\tilde{\nu}}$-dimensional, and write those causal effects as $\{\Theta_{x,\tilde{\nu},\mathcal{A}}, \Theta_{z,\tilde{\nu},\mathcal{A}}\}$, where now both maps have a finite number of columns $n_{\tilde{\nu}}$. For example, the researcher may know impulse responses for contemporaneous and several forward guidance monetary shocks (e.g., as in the functional shock approach of Inoue et al. (2021)), but not for all possible forward guidance shocks. We first present extensions of Propositions 1 and 2 to this partial measurement case, and then discuss an application to policy counterfactuals after technology news.[10]

Policy counterfactuals in identified subspaces. For both general alternative policy rules as well as the characterization of optimal policy, our results from Section 2 extend naturally to the smaller subspace spanned by the finite set of observed shocks.

Consider first the case of alternative policy rules. With the researcher now only observing $\{\Theta_{x,\tilde{\nu},\mathcal{A}}, \Theta_{z,\tilde{\nu},\mathcal{A}}\}$, the proof of Proposition 1 in general fails because it now may not be possible

[10]Our discussion in this section focusses on the finite-shock case, so $\{\Theta_{x,\tilde{\nu},\mathcal{A}}, \Theta_{z,\tilde{\nu},\mathcal{A}}\}$ have a small number of columns. In any empirical application, those linear maps of course also have a finite number T of rows. We do not pay much attention to this limitation since we consider shocks and counterfactual policies with sufficiently short-lived dynamics, making the maximal truncation horizon immaterial.

to set the shock vector $\boldsymbol{\nu}$ in a way that perfectly mimics any desired alternative policy rule $\{\tilde{\mathcal{A}}_x, \tilde{\mathcal{A}}_z\}$. However, the proof continues to apply without change for *particular* alternative policy rules that satisfy the restriction

$$\tilde{\mathcal{A}}_x(\boldsymbol{x}_{\mathcal{A}}(\boldsymbol{\varepsilon}) + \Theta_{x,\tilde{\nu},\mathcal{A}} \times \tilde{\boldsymbol{\nu}}) + \tilde{\mathcal{A}}_z(\boldsymbol{z}_{\mathcal{A}}(\boldsymbol{\varepsilon}) + \Theta_{z,\tilde{\nu},\mathcal{A}} \times \tilde{\boldsymbol{\nu}}) = 0 \tag{30}$$

for some $\tilde{\boldsymbol{\nu}} \in \mathbb{R}^{n_{\tilde{\nu}}}$. Thus, while the researcher now cannot characterize counterfactuals for all possible policy rules, she can still explore those that move away from the baseline prevailing rule $\{\mathcal{A}_x, \mathcal{A}_z\}$ in an $n_{\tilde{\nu}}$-dimensional subspace spanned by the observed shocks $\boldsymbol{\nu} = \mathcal{V} \times \tilde{\boldsymbol{\nu}}$. Our application to investment technology shocks provides an illustration of this insight.

For optimal policy, we follow the same steps as the proof of Proposition 2 to now consider the problem of minimizing the policymaker loss function (9) within the identified subspace of policy changes. This problem gives the optimality condition

$$\sum_{i=1}^{n_x} \lambda_i \Theta'_{x_i,\tilde{\nu},\mathcal{A}} W \boldsymbol{x}_i = 0. \tag{31}$$

Equation (31) fully characterizes the optimal policy rule within the identified space of policy rules – i.e., the prevailing baseline rule plus deviations that are spanned by the identified policy shocks. Furthermore, it also provides $n_{\tilde{\nu}}$ restrictions that *any* solution to the full optimal policy problem must satisfy.[11] While characterization of the full rule of course requires infinitely many restrictions, the restrictions embedded in (31) may already be meaningful.

Application: counterfactual rules after technology shocks. We use our results to estimate the effects of contractionary investment-specific technology shocks under a counterfactual policy rule that aggressively stabilizes fluctuations in aggregate output. We present the results here, and relegate details to Appendices A.3 and B.1.

Our approach requires two inputs. First, we need the dynamic causal effects of the shock of interest under the baseline policy rule. For this application we consider the investment technology shock identified in Ben Zeev & Khan (2015). Second, we need the dynamic causal effects of monetary policy shocks at several different horizons. To this end we estimate impulse responses using the popular identification approaches of Christiano et al. (1999), Romer

[11]Equation (31) is related to Barnichon & Mesters (2020b), who propose to use a condition of this form to test the optimality of a given policy. Since their analysis relies on fixed private sector expectations, they do not draw any implications for optimal policy *rules*, unlike our approach.

Policy Counterfactual for Investment Shocks

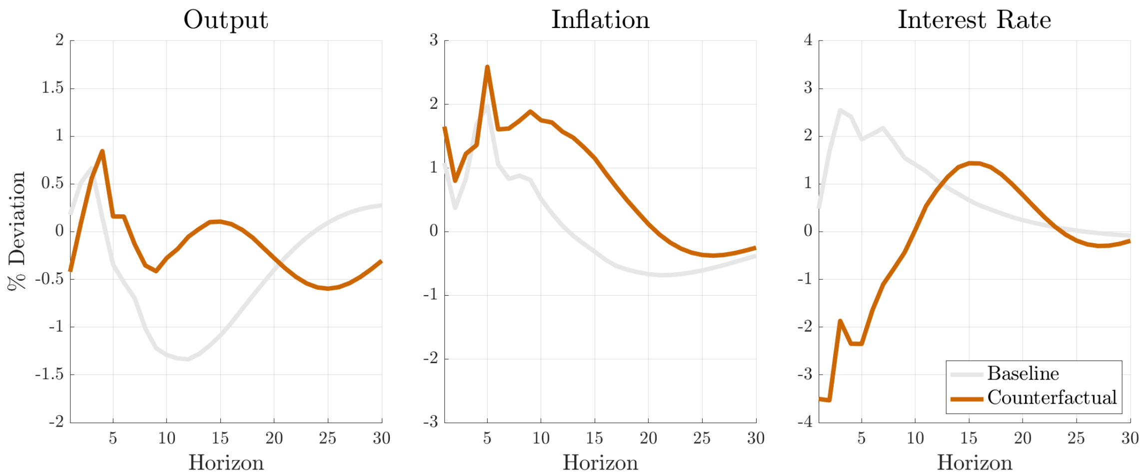

Figure 3: Output, inflation and interest rate impulse responses to a contractionary investment technology shock under the prevailing baseline rule (grey) and the best feasible approximation to a rule that stabilizes output (orange).

& Romer (2004) and Gertler & Karadi (2015). Importantly, since the dynamic response of interest rates differs across these three identifications schemes, they reflect information about the effects of different kinds of current and future interest rate changes. We focus on OLS point estimates for all three, and then construct the best possible approximation to a rule that aims to perfectly stabilize aggregate output.[12]

Figure 3 presents our results. Under the prevailing baseline rule (grey), the policymaker leans against the inflationary effects of the negative technology shock, further pushing down aggregate real activity. Under our counterfactual (orange), on the other hand, interest rates are cut aggressively, keeping output relatively close to trend throughout, but of course at the cost of persistently elevated inflation. By our identification results, any structural model of the general form (1) - (2) and consistent with our empirical estimates of monetary transmission will invariably agree with those alternative rule counterfactuals.

3.2 Imposing cross-column restrictions

Instead of restricting our analysis within the subspace of dynamic causal effects identified in previous work, we may alternatively leverage additional outside information to translate the

[12]That is, we find the best possible fit to a policy rule with $\mathcal{A}_y = I$ and $\mathcal{A}_q = \mathbf{0}$ for all $q \neq y$. For details on the construction see Appendix A.3.

limited empirical evidence into the required full sets of policy causal effects $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$. We first discuss the general promise of this approach and then present an application to optimal dual mandate policy.

The key building block underlying our approach is the idea that the different columns of the causal effect maps $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ are not completely independent objects – for example, impulse responses to a forward guidance shock eight quarters out should be very similar to forward guidance shocks nine quarters out, just shifted by one period. A first, very weak way in which to formalize this intuition is to require that the dynamic causal effect matrices $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ are asymptotically time-invariant, in the sense that, for all $s \in \mathbb{N}$,

$$\lim_{t\to\infty} \Theta_{x,\nu,\mathcal{A}}(t+s,t) = \bar{\Theta}_{x,\nu,\mathcal{A}}(s), \quad \lim_{t\to\infty} \Theta_{z,\nu,\mathcal{A}}(t+s,t) = \bar{\Theta}_{z,\nu,\mathcal{A}}(s) \tag{32}$$

where $\bar{\Theta}_{x,\nu,\mathcal{A}}$ and $\bar{\Theta}_{z,\nu,\mathcal{A}}$ are two sequences.[13] Imposing (32) after some finite horizon H reduces the problem of dynamic causal effect identification from an infinite-dimensional one to an H-dimensional one.

A second, more restrictive approach is to use cross-column restrictions on the linear maps $\{\Theta_{x,\nu,\mathcal{A}}, \Theta_{z,\nu,\mathcal{A}}\}$ implied by families of explicit structural models. In the remainder of this section we follow this second approach, and show that it can be a powerful way to construct counterfactuals and characterize optimal policy for dual mandate policymakers.[14]

Cross-column restrictions for the dual mandate. For any model in the general class (1) - (2), we can write the optimal policy rule of a dual mandate policymaker with preference weights $\{\lambda_\pi, \lambda_y\}$ as

$$\lambda_\pi \Theta'_{\pi,\nu,\mathcal{A}} W \boldsymbol{\pi} + \lambda_y \Theta'_{y,\nu,\mathcal{A}} W \boldsymbol{y} = 0 \tag{33}$$

The key property of optimal policy implicit in (33) is that the dual mandate policymaker cares only about the joint space of output and inflation that she can implement by varying her policy instrument. We summarize the relationship between the two sets of impulse

[13] Auclert et al. (2021) show that such asymptotic time-invariance is a relatively general property of linearized perfect foresight macroeconomic models; to illustrate, we provide a numerical example in the context of the HANK model of Section 2.3 in Appendix A.4.

[14] An even more restrictive version of this second approach is model estimation through impulse response matching (e.g., as in Christiano et al., 2005). As it turns out, for the purposes of optimal policy under a dual mandate, such full-blown model estimation is not necessary.

responses in the linear map Π_y, defined implicitly via

$$\Theta_{\pi,\nu,\mathcal{A}} = \Pi_y \times \Theta_{y,\nu,\mathcal{A}} \tag{34}$$

The map Π_y gives inflation impulse responses *relative* to output impulse responses; for example, the nth column of Π_y gives the impulse response of inflation to a policy shock that changes output by one unit at horizon $n-1$. At this point, of course, Π_y could be arbitrarily complicated, so we have not yet imposed any cross-column restrictions.

Our key insight for dual mandate policy analysis is that, in *any* structural model that features a Phillips curve relationship, that curve fully pins down the linear map Π_y. For example, in the simple New Keynesian Phillips Curve in (3), we would find

$$\Pi_y = \begin{pmatrix} \kappa & \kappa\beta & \kappa\beta^2 & \dots \\ 0 & \kappa & \kappa\beta & \dots \\ 0 & 0 & \kappa & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \tag{35}$$

More generally, any researcher willing to commit to a particular functional form for Π_y can estimate that map using empirical evidence on the two original impulse response maps $\Theta_{\pi,\nu,\mathcal{A}}$ and $\Theta_{y,\nu,\mathcal{A}}$, and then use the result to recover the optimal policy rule as[15]

$$\lambda_\pi W \boldsymbol{\pi} + \lambda_y (\Pi_y')^{-1} W \boldsymbol{y} = 0 \tag{36}$$

This rule would then be optimal for any model that does indeed feature a Phillips curve relationship of the assumed form, independent of any further details of preferences, technology, expectation formation, and so on.

Application: optimal dual mandate & technology shock. We apply this insight to study the *optimal* dual mandate policy response to the investment technology shock of Ben Zeev & Khan (2015) studied in Section 3.1. Throughout, we consider a dynamic Phillips

[15] Strictly speaking, going from (33) to (36) requires that $\Theta_{\pi,\nu,\mathcal{A}}$ and $\Theta_{y,\nu,\mathcal{A}}$ are invertible. This model property is satisfied in the standard business-cycle models routinely used for policy analysis. For example, invertibility is easy to verify in the baseline New Keynesian model, as the two maps are upper triangular. In standard HANK models like that in Section 2.1, it is ensured by invertibility of the linear maps $\mathcal{C}_{i_b}$ and $\mathcal{C}_\tau$ (see Wolf (2021)), as long as prices are not perfectly sticky.

Optimal Policy for Investment Shocks

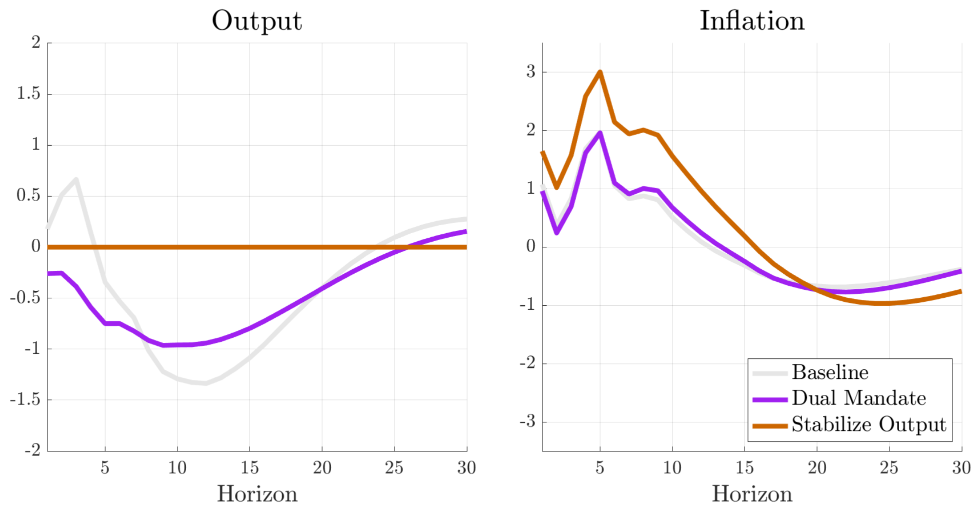

Figure 4: Output and inflation impulse responses to a contractionary investment technology shock under the optimal policy rule (38) for a standard dual mandate ($\lambda_\pi = \lambda_y = 1$, purple) and a policymaker that only cares about output ($\lambda_\pi = 0$, $\lambda_y > 0$, orange).

curve relationship of the form

$$\pi_t = \gamma_b \pi^4_{t-1} + \gamma_f \mathbb{E}_t \left[\pi^4_{t+4}\right] + \kappa y_t \tag{37}$$

where $\pi^4_{t-1} = \frac{1}{4} \times (\pi_{t-1} + \pi_{t-2} + \pi_{t-3} + \pi_{t-4})$. Appendix A.5 shows the linear map Π_y corresponding to this Phillips curve specification. Going from this linear map to the optimal policy rule (36), and setting the weight matrix as $W = I$, straightforward algebra gives a very simple (forecast) targeting rule

$$\lambda_\pi \pi_t + \lambda_y \frac{1}{\kappa} \left[y_t - \frac{1}{4}\gamma_b \sum_{h=1}^{4} y_{t+h} - \frac{1}{4}\gamma_f \sum_{h=1}^{4} y_{t-h} \right] = 0, \quad \forall t = 0, 1, 2, \ldots \tag{38}$$

Barnichon & Mesters use the dynamic causal monetary policy effects of Gertler & Karadi (2015) to estimate a Phillips curve of the form (37), and report the parameter estimates $\{\gamma_b = 0.71, \gamma_f = 0.29, \kappa = 0.12\}$. Using those estimates in (38), we get a fully-specified, implicit targeting rule – a policy rule that we immediately know to be optimal for any dual mandate policymaker facing an economy described by a structural model of the form (1) - (2) that (i) is consistent with the empirical monetary policy shock estimates of Gertler & Karadi (2015) and (ii) features a Phillips curve relationship of the form (37).

We now use the estimated optimal policy rule (38) to solve for the optimal policy response to the investment technology shock identified by Ben Zeev & Khan (2015). Results for two possible loss functions – a conventional equal-weight dual mandate (purple), and a policymaker that only cares about output (orange) – are reported in Figure 4.[16] We emphasize two main takeaways. First, for a policymaker that only cares about output, the result is the expected smoothed-out version of the policy rule approximation in Figure 3. Second, for a dual mandate policymaker, the optimal inflation and output paths are quite close to those under the baseline rule. Thus, given the available estimate of the dynamic Phillips curve, the prevailing baseline policy rule actually appears to be close to optimal.

4 Conclusions & next steps

The standard approach to counterfactual policy analysis relies on fully specified, dynamic general equilibrium models. Our identification results instead point in a different direction: for valid policy counterfactuals, researchers can estimate dynamic causal effects of policy shocks and combine them to form policy counterfactuals and optimal policy responses. These counterfactuals are valid in a large class of models that encompasses the majority of structural business cycle models that are currently used for policy analysis.

The main challenge in implementing this strategy is that evidence on an infinitely large set of policy perturbations is unattainable. One natural reaction is simply to try get as close as possible; viewed in this light, further empirical evidence on policy news shocks would be extremely valuable. Every additional piece of empirical evidence will allow researchers to (a) expand the space of alternative, counterfactual policy rules that we can analyze and (b) find further restrictions on optimal policy rules. The obvious alternative path forward is to use outside prior information to map the available empirical evidence into estimates of the entire space of policy-relevant causal effects. In future versions of this paper, we plan to go further in that direction, with a focus on the asymptotic time invariance of impulse responses as a natural starting point. Clearly, refining and adding to those two approaches remains an important avenue for further research.

[16] Following the steps in the proof of Proposition 2, we construct those optimal inflation and output paths as the solution to the problem of minimizing the policymaker loss function subject to the constraint that $\boldsymbol{\pi} = \Pi_y \times \boldsymbol{y}$. See Appendix A.5 for further details.

References

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. (2020a). Identifying modern macro equations with old shocks. *The Quarterly Journal of Economics*, *135*(4), 2255–2298.

Barnichon, R. & Mesters, G. (2020b). Testing Macroeconomic Policies with Sufficient Statistics. Working Paper.

Baxter, M. & King, R. G. (1993). Fiscal policy in general equilibrium. *The American Economic Review*, 315–334.

Ben Zeev, N. & Khan, H. (2015). Investment-specific news shocks and us business cycles. *Journal of Money, Credit and Banking*, *47*(7), 1443–1464.

Beraja, M. (2020). Counterfactual equivalence in macroeconomics. Working Paper.

Bernanke, B. S., Gertler, M., Watson, M., Sims, C. A., & Friedman, B. M. (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings papers on economic activity*, *1997*(1), 91–157.

Christiano, L. J., Eichenbaum, M., & Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of macroeconomics*, *1*, 65–148.

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.

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gertler, M. & Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, *7*(1), 44–76.

Guren, A., McKay, A., Nakamura, E., & Steinsson, J. (2021). What do we learn from cross-regional empirical estimates in macroeconomics? *NBER Macroeconomics Annual*, *35*(1), 175–223.

Hebden, J. & Winker, F. (2021). Impulse-Based Computation of Policy Counterfactuals. Working Paper.

Inoue, A., Rossi, B., et al. (2021). A new approach to measuring economic policy shocks, with an application to conventional and unconventional monetary policy. Working Paper.

Kaplan, G., Moll, B., & Violante, G. (2018). Monetary Policy according to HANK. *American Economic Review*, *108*(3), 697–743.

Lucas, R. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester Conference Series on Public Policy*, volume 1, (pp. 19–46). Elsevier.

Ottonello, P. & Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, *88*(6), 2473–2502.

Plagborg-Møller, M. & Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, *89*(2), 955–980.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. *Handbook of macroeconomics*, *2*, 71–162.

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: journal of the Econometric Society*, *48*(1), 1–48.

Sims, C. A. & Zha, T. (2006). Were there regime switches in us monetary policy? *American Economic Review*, *96*(1), 54–81.

Smets, F. & Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, *97*(3), 586–606.

Svensson, L. E. (1997). Inflation forecast targeting: Implementing and monitoring inflation targets. *European economic review*, *41*(6), 1111–1146.

Wieland, J. F. & Yang, M.-J. (2020). Financial dampening. *Journal of Money, Credit and Banking*, *52*(1), 79–113.

Wolf, C. K. (2020). The missing intercept: A demand equivalence approach. Working Paper.

Wolf, C. K. (2021). Interest rate cuts vs. stimulus payments: An equivalence result. Working Paper.

Woodford, M. (2011). *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.

A Further results

This appendix presents several supplementary theoretical results. Appendix A.1 extends Proposition 2 to a more general loss function. Appendix A.2 gives further details for the HANK model of Sections 2.1 and 2.3, Appendix A.3 shows how to construct the best fit to any target policy rule on the estimated policy shock subspace, Appendix A.4 discusses the asymptotic time invariance of impulse response maps, and Appendix A.5 provides additional information for our Phillips curve analysis.

A.1 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

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

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 (23). The first-order conditions of this problem are

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

so we can recover the optimal policy rule as

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

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 (2), so we do not consider this case here.

A.2 Model details & parameterization

The HANK model sketched in Section 2.1 and used for our quantitative illustration in Section 2.3 is exactly the same as in Wolf (2021) (including the parameterization), with only one change: rather than imposing uniform hours worked $\ell_{it} = \ell_t$ for all households i, we consider a labor rationing rule that ensures that

$$w_t \ell_{it} e_{it} + d_{it} = e_{it} y_t$$

for all households i. This rationing rule is feasible since $\int_0^1 w_t \ell_{it} e_{it} + \int_0^1 d_{it} = w_t \ell_t + d_t = y_t$, and it allows us to write the perfect foresight consumer demand block as

$$\boldsymbol{c} = \mathcal{C}\left(\boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{i}_b, \boldsymbol{\tau}, \boldsymbol{\varepsilon}^d\right)$$

Linearizing, we recover (7).

A.3 Best-fitting policy counterfactual

For identified policy shocks corresponding to a subspace $\mathcal{V}$ with $\Theta_{x,\tilde{\nu},\mathcal{A}} \equiv \Theta_{x,\nu,\mathcal{A}} \times \mathcal{V}$ and $\Theta_{z,\tilde{\nu},\mathcal{A}} \equiv \Theta_{z,\nu,\mathcal{A}} \times \mathcal{V}$, any policy rule consistent with (30) can be replicated perfectly. For all alternative policy rules $\{\tilde{\mathcal{A}}_x, \tilde{\mathcal{A}}_z\}$ outside of this class, we can construct a best-fitting policy counterfactual within the spanned subspace. Using a simple quadratic loss function for deviations from the target rule, this gives

$$\tilde{\boldsymbol{\nu}} = -\left[\left(\tilde{\mathcal{A}}_w \Theta_{w,\nu,\mathcal{A}} \mathcal{V}\right)' \times \left(\tilde{\mathcal{A}}_w \Theta_{w,\nu,\mathcal{A}} \mathcal{V}\right)\right]^{-1} \times \left[\left(\tilde{\mathcal{A}}_w \Theta_{w,\nu,\mathcal{A}} \mathcal{V}\right)' \times \mathcal{A}_w \times \boldsymbol{w}_{\mathcal{A}}(\boldsymbol{\varepsilon})\right] \tag{A.2}$$

where $\boldsymbol{w} = (\boldsymbol{x}', \boldsymbol{z}')'$. In (3) we use (A.2) to construct the best approximation to a policy rule that perfectly stabilizes aggregate output.

A.4 Asymptotic time invariance

Auclert et al. (2021) show that, under relatively general conditions on economic primitives, the private-sector linear maps $\{\mathcal{H}_w, \mathcal{H}_x, \mathcal{H}_z, \mathcal{H}_\varepsilon\}$ in typical business-cycle models are asymptotically time-invariant, in the sense that

$$\lim_{t \to \infty} \mathcal{H}_q(t+s, t) = \bar{\mathcal{H}}_q(s), \quad q \in \{w, x, z, \varepsilon\} \tag{A.3}$$

where $\bar{\mathcal{H}}_q$ is a sequence. It follows that, for suitable policy rules $\{\mathcal{A}_x, \mathcal{A}_z\}$, impulse response functions will also be asymptotically time-invariant. Figure A.1 provides an illustration in the quantitative HANK model of Section 2.3, showing output impulse responses to various different contemporaneous and forward guidance monetary shocks. We see that, for forward guidance shocks far into the future (large shock horizon h), the output impulse responses are left- and right-translations of each other, exactly as expected.

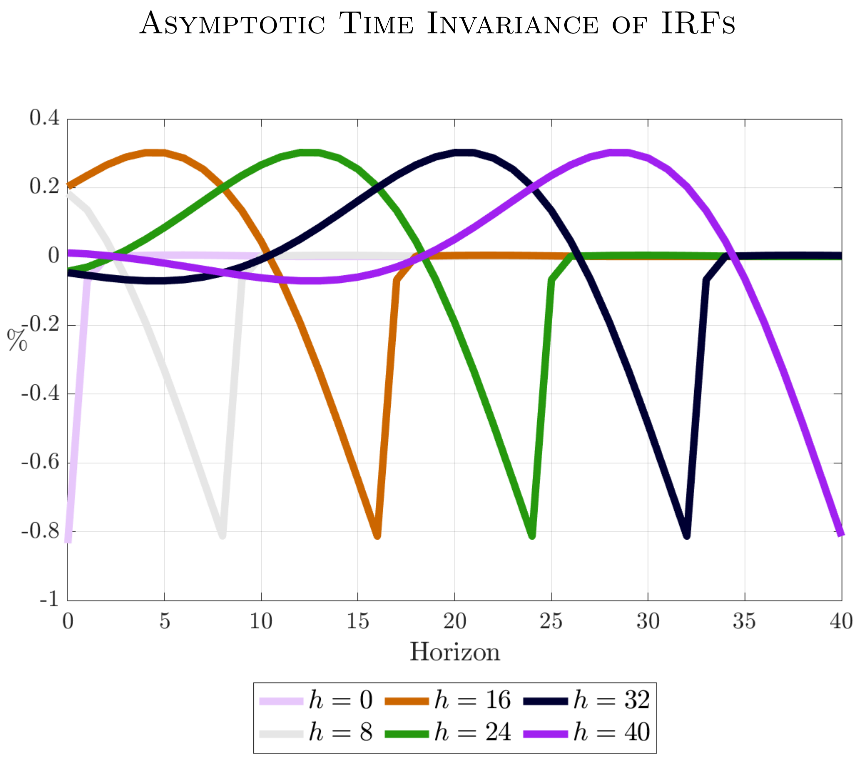

Figure A.1: Output impulse responses to contemporaneous and forward guidance monetary policy shocks in the HANK model of Appendix A.2.

A.5 Phillips curve & optimal policy

Consider the augmented Phillips curve (37). Along a perfect foresight transition path, we can write this relationship as

$$\underbrace{\begin{pmatrix} 1 & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & 0 & \dots \\ -\frac{1}{4}\gamma_b & 1 & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & \dots \\ -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & 1 & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & \dots \\ -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & 1 & -\frac{1}{4}\gamma_f & -\frac{1}{4}\gamma_f & \dots \\ -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & 1 & -\frac{1}{4}\gamma_f & \dots \\ 0 & -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & -\frac{1}{4}\gamma_b & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}}_{\equiv \Pi_\pi} \times \boldsymbol{\pi} \quad = \quad \kappa \times \boldsymbol{y}$$

It thus follows that

$$\Pi_y = \Pi_\pi^{-1} \quad \times \quad \kappa$$

For $W = I$, the optimal policy rule is then given as

$$\lambda_\pi \boldsymbol{\pi} + \lambda_y (\Pi_y')^{-1} \boldsymbol{y} = 0$$

Solving this we find that, for all $t \geq 0$, (38) must hold.

Finally, we show how to use empirical estimates of Π_y to recover optimal output and inflation paths for a given aggregate shock. Following the steps in the proof of Proposition 2, we can write the optimal policy problem as

$$\min_{\boldsymbol{\pi},\boldsymbol{y},\boldsymbol{\nu}} \lambda_\pi \boldsymbol{\pi}' W \boldsymbol{\pi} + \lambda_y \boldsymbol{y}' W \boldsymbol{y} \tag{A.4}$$

subject to

$$\boldsymbol{\pi} \quad = \quad \boldsymbol{\pi}_{\mathcal{A}}(\boldsymbol{\varepsilon}) + \Pi_y \times \boldsymbol{\nu} \tag{A.5}$$

$$\boldsymbol{y} \quad = \quad \boldsymbol{y}_{\mathcal{A}}(\boldsymbol{\varepsilon}) + \boldsymbol{\nu} \tag{A.6}$$

This problem gives the optimality condition

$$\boldsymbol{\nu} = -\left[\lambda_\pi \Pi_y' W \Pi_y + \lambda_y W\right]^{-1} \times \left[\lambda_\pi \Pi_y' W \boldsymbol{\pi}_{\mathcal{A}}(\boldsymbol{\varepsilon}) + \lambda_y W \boldsymbol{y}_{\mathcal{A}}(\boldsymbol{\varepsilon})\right] \tag{A.7}$$

Plugging (A.7) into (A.5) - (A.6), we recover the optimal output and inflation time paths. Figure 4 plots those time paths for the investment news shock of Ben Zeev & Khan (2015).

B Empirical appendix

This appendix elaborates on our empirical estimation. First, in Appendix B.1, we discuss the dynamic causal effects estimated for the counterfactuals in Section 3.1. Second, in Appendix B.2, we briefly review the approach of Barnichon & Mesters (2020a) to Phillips curve estimation.

B.1 Shock & policy dynamic causal effects

Our analysis of investment-specific technology shocks follows Ben Zeev & Khan (2015), while our monetary policy shock identification mimics that of (i) Christiano et al. (1999), (ii) Romer & Romer (2004) and (iii) Gertler & Karadi (2015).

Outcomes. We are interested in impulse responses of three outcome variables: output, inflation, and the policy rate. For output, we follow Ramey (2016) and deflate per-capita nominal GDP by the GDP deflator. For inflation, we consider the annualized growth rate of the GDP deflator. All series are obtained from the replication files for 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. All series are quarterly.

Shocks & identification. We take the investment-specific technology shock series from Ben Zeev & Khan (2015), the updated series of the Romer-Romer shocks from Wieland & Yang (2020), and the high-frequency monetary policy surprise series from Gertler & Karadi (2015), aggregated to quarterly frequency through simple averaging. Recursive shocks are identified through the estimated VAR itself.

Estimation details. For maximal consistency, we try to estimate all impulse responses within a common empirical specification. For the investment-specific technology shocks and the Romer-Romer monetary policy shock, we order the two shock measures first in a recursive VAR containing our outcomes of interest (following Plagborg-Møller & Wolf, 2021), estimated on a sample from 1969:Q1 – 2007:Q4. Our estimation of the Gertler-Karadi shock is identical, except for the fact that – because of constraints on when the shock is actually available – the sample runs from 1990:Q1 – 2012:Q4. Finally, for a recursive shock, we return to our baseline long sample period, and now identify a monetary shock as the *last* recursively ordered shock in a system containing output, inflation, and the nominal rate.

We estimate all VARs using four lags, a constant, and deterministic linear and quadratic trends, and for Figure 3 rely on OLS point estimates.

B.2 NKPC estimation

Barnichon & Mesters (2020a) show how to use estimates of monetary policy impulse responses to identify a Phillips curve relationship of the form (37). We rely on the point estimates corresponding to their Gertler & Karadi analysis and with the additional constraint that $\gamma_f + \gamma_b = 1$, reported in their Table IV.