Evaluating Policy Institutions*

−150 Years of US Monetary Policy−

Régis Barnichon[(a)] *and Geert Mesters*[(b)]

[(a)] Federal Reserve Bank of San Francisco and CEPR

[(b)] Universitat Pompeu Fabra, Barcelona School of Economics and CREI

June 30, 2023

Abstract

How should we evaluate and compare the performances of policy institutions? In this work, we show that two statistics are sufficient to evaluate policy institutions based on their reaction function: (i) the impulse responses of the policy objectives to non-policy shocks (e.g., energy price shocks or financial shocks), which capture what the institution did on average to counteract these shocks, and (ii) the impulse responses of the same policy objectives to policy shocks, which capture what the institution could have done to counteract the shocks. A regression on (i) on (ii) —a regression in impulse response space— allows to compute the distance to the optimal reaction function, and thereby evaluate or rank policy makers. We use our methodology to evaluate US monetary policy over the past 150 years; from the Gold standard period, the early Fed years and the Great Depression, to the post World War II period, and the post-Volcker regime.

JEL classification: C14, C32, E32, E52.

Keywords: optimal policy, policy evaluation, reaction function, monetary history.

*Preliminary. We thank our discussant Frank Smets as well as Christian Brownlees, Ryan Chahrour, Kirill Evdokimov, Sylvain Leduc, Adam Shapiro and seminar participants at Paris School of Economics, Mannheim, UC Berkeley, INSPER, PUC Rio, FGV, Bank of Italy, ECB and Riksbank for helpful comments. Mesters acknowledge support from the Spanish Ministry of Economy and Competitiveness, and the 2021 ERC Starting Grant POLICYMETRICS. The views expressed in this paper are the sole responsibility of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

1 Introduction

How should we evaluate policy institutions? How should we evaluate policy makers after their term in office? These questions are of central importance to the good functioning of democratic and accountable institutions, but there is little consensus on a method for evaluating and comparing performance.[1]

A naive approach could consist in measuring performance based on realized macroeconomic outcomes. For instance, we could assess a central banker based on average inflation and unemployment outcomes over her term. Unfortunately, that approach suffers from three problems: (i) different policy makers may face different initial conditions, e.g. a central banker can inherit a strong or weak economy from her predecessor, (ii) different policy makers may face different economic disturbances, e.g., a central banker may experience a financial crisis or an energy price shock that will affect her ability to stabilize inflation and unemployment, and (iii) different policy makers may live in different economic environments, e.g., a steeper or flatter Phillips curve will affect a central banker's ability to control inflation.

This triplet of confounding factors coming from different initial conditions, different disturbances and different economic environments severely limits our ability to evaluate policy makers based on realized outcomes.[2]

To make progress it is instructive to consider an ideal, yet infeasible, approach for comparing policy makers: an experimental setting. Consider setting up a laboratory, in which different policy makers are given the same mandate —minimizing a loss function involving some policy objectives— and are subjected to the same initial conditions and the same economic environment. The different policy makers are then exposed to the same sequence of shocks, and they each make decisions that aim to achieve their mandate. Afterward, we can compare performance from the realized losses and conclude which policy maker performed better. Such comparison would be on equal grounds as the only source of variation would come from the different ways each policy maker *reacted* to the same of sequence shocks, i.e., from the different reaction functions.

In this paper, we propose an empirical method that aims to mimic this ideal "reaction function comparison" experiment while making minimal structural assumptions on the underlying economic model and underlying policy rule. Our approach exploits a simple idea: while different policy makers are never exposed to the same *sequences* of non-policy shocks, they are often exposed to the same *types* of shocks; for instance energy price shocks, financial shocks, or even pandemic shocks. By comparing how well different policy makers performed

[1]Take the example of modern central banks, which are granted a substantial degree of independence to achieve their mandates. That independence comes with a need for accountability, and central bank governors are regularly called to testify in front of elected officials to justify their actions. Yet, there is no systematic quantitative evaluation of these actions.

[2]See Fair (1978) for an early discussion of these points.

in response to such common shocks, we can approach the ideal empirical setting sketched above: assessing performance from the different ways each policy makers reacted to the same types of shocks.

Geometrically speaking, our strategy amounts to projecting realized macroeconomic outcomes on a space spanned by well chosen non-policy shocks and to study policy performance in that space. In fact, in that subspace policy evaluation reduces to a simple optimization problem that only involves two well-known (and estimable) sufficient statistics: (i) the impulse responses of the policy objectives to a (well-chosen) non-policy shock, and (ii) the same impulse responses to a policy shock.

The first set of impulse responses —the impulse responses to a specific non-policy shock— capture the average effects of that non-policy shock under the policy maker's reaction function and allow to compute a *conditional* loss; a loss conditional on that non-policy shock. For instance, with a quadratic loss function the conditional loss is simply the sum-of-squares of that impulse response. While it is tempting to assess and compare performance based on that impulse response alone, this is not enough since other factors beyond a policy maker's reaction function could generate a lower conditional loss, i.e., a more stable impulse response. For instance, a different economic environment could make the economy more stable independently of the policy makers' reaction function. To assess how well a policy maker reacted to that specific non-policy shock, we need to know the outcome of a policy rule counterfactual: how a different reaction would have affected the economy. That counter-factual is precisely given by the second set of impulse responses —the impulse responses to policy shocks—, which allow to compute how a different reaction function would have affected the conditional loss —what the policy maker could have done to counteract the non-policy shock—.

We show that for a large class of models and loss functions the *distance to the optimal reaction*, or Optimal Reaction Adjustment (ORA), can be computed from a simple regression in "impulse response space": a regression of the impulse responses to the non-policy shock on the impulse responses to policy shocks.[3]

The ORA measures by how much more or less a policy maker should have reacted in response to a given non-policy shock, and it provides a direct measure of policy performance *conditional* on a specific type of non-policy shock. Overall policy performance can then be assessed by measuring the ORAs for different types of non-policy shocks. Further, for non-policy shocks that are common across different policy makers, the ORAs are portable

[3]The correlation between these two sets of impulse responses directly captures the performance of the policy institution: A correlation of zero indicates best performance —the policy institution could not have reacted any better to the macro shocks that affected the economy—, while a correlation of one (in absolute value) indicates worst performance —the institution could have (but did not) perfectly met its objectives by undoing the effects of these macro shocks—.

moments in the sense of Nakamura and Steinsson (2018): we can use the ORAs to compare policy makers or policy institutions across time (say the Fed in 1930s vs the Fed in the 2000s) or across space (say the Fed vs the ECB).

It may seem surprising to be able to assess a reaction function without specifying or estimating any policy rule. This reason this is possible, and the key insight underlying our approach, is that the effects of any reaction function are *encoded* in the impulse responses to policy and non-policy shocks. Even if we do not know the specific form of some past policy rule, that reaction function left a footprint on the effects of policy and non-policy shocks, and that footprint is sufficient to evaluate the reaction function. This is the essence of our sufficient statistics approach.

The evaluation and comparison of policy makers thus reduces to estimating structural impulse responses, and this realization opens a number of important avenues for policy evaluation, as one can draw on a large macro-econometric literature to evaluate policy institutions. A recent discussion of structural shock identification can be found in Ramey (2016) and modern impulse response estimation methods are discussed and compared in e.g., Stock and Watson (2016), Kilian and Lütkepohl (2017) and Li, Plagborg-Mller and Wolf (2022).

We then apply our methodology to study the performance of US monetary policy over the past 150 years. Our method allows us to address and revisit many interesting questions regarding the conduct of monetary policy. To name a few, (i) did the founding of the Federal Reserve in 1913 led to superior macro outcomes than during the Gold standard period (e.g., Bordo and Kydland, 1995)? And if so, by how much? (ii) While many people would agree that monetary policy was superior during the 2007-2009 financial crisis than during the 1929-1933 financial crisis (e.g., Wheelock et al., 2010), can we confirm and quantify this improvement? (iii) To what extent is the stable and low inflation environment of 2000s versus the 1970s the outcome of good policy or simply good luck (e.g., Clarida, Galí and Gertler, 2000)?

To assess and compare monetary policy performance across historical periods, we evaluate how monetary policy responded to five types of non-policy shocks that are common across periods: (i) financial shocks, (ii) government spending shocks, (iii) energy price shocks, (iv) inflation expectation shocks and (v) productivity shocks, and we evaluate US monetary policy over four distinct periods: (a) 1879-1912 covering the Gold standard period until the founding of the Federal Reserve, (b) 1913-1941 covering the early Fed years to the US entering World War II, (c) 1954-1984 covering the post World War II period until the beginning of the Great Moderation, and (d) 1990-2019 covering the Great Moderation period, the financial crisis and up to the COVID crisis.

Over these historical periods US monetary policy was confronted with different sequences of shocks and possibly very different economic environments. With our ORA-based policy

evaluation, we can compare policy performance over these four periods while by-passing the many confounding factors that have plagued previous comparisons. The (still substantial) empirical challenge is then to consistently estimate the impulse responses to monetary shocks and the five non-policy shocks over each sub-period. Fortunately, we can leverage on a large empirical literature on structural shocks identification, and we will use as much as possible the state of the art in each setting: Hamilton (2003) for energy price shocks, Ramey and Zubairy (2018) for government spending shocks, Leduc, Sill and Stark (2007) for inflation expectation shocks, Gali (1999) for productivity shocks and Reinhart and Rogoff (2009) for banking panics. To identify monetary shocks, we rely on Romer and Romer (2004*b*) and Gürkaynak, Sack and Swanson (2005) for the post WWII periods. Identification is more challenging (and less developed) for the pre WWII periods, and we propose a new identification of monetary shocks for the Gold Standard period. Specifically, we exploit the specificity of the Gold Standard, in that the monetary base depends on the amount of gold in circulation, and we use unanticipated large Gold mine discoveries (discoveries that led to Gold rush) as an instrument for movements in the monetary base. For monetary shocks during the early Fed years, we use the Friedman and Schwartz (1963) narrative dates extended by Romer and Romer (1989).

Evaluating and comparing policy makers require to take a stand on a set of objectives, i.e., on a loss function. In our empirical application, we consider a quadratic loss function with equal weights on inflation and unemployment.[4] Given that loss function, our results point to overall improvements in the conduct of monetary policy, though improvements have not been monotonic. In response to bank runs, we do find that the policy response is substantially better after 1951 than during the Gold Standard or the early Fed period. During the Early Fed period the monetary response was much too passive in the face of financial disturbances; not lowering the discount rate enough, in fact running a contractionary monetary policy. In contrast, the monetary response was much closer to optimal in the post Volcker period, though the zero-lower bound did constrain partially the Fed's response.

That said, improvements in the conduct of monetary policy have not been monotonic, and the post WWII period saw the worst performance in response to inflationary pressures shocks: energy price shocks, inflation expectation shocks and TFP shocks. For all these shocks, the Fed reaction was much too weak over 1951-1984, particularly in response to inflation expectation shocks. This result echoes a large literature on the Fed's failure to satisfy the Taylor principle in the 1970s, but also goes further by allowing to compute the distance to the optimal reaction —by how much more the Fed should have responded to supply shocks—. In addition, we find that the Fed's excess passivity carries to aggregate

[4]Importantly, our approach could accommodate other loss functions, for instance different loss functions across time periods, or even micro-founded welfare-based loss functions.

demand shocks and financial shocks as well. For instance, the Fed's interest rate reaction was too weak in the face of the military buildup shocks of the Vietnam war.

Last, we find that performance is universally superior during the post Volcker period: the distances to the optimal reaction coefficients are smallest (and non-significantly different from zero) for all types of shocks that we considered. The only (mild) exception is the reaction to the financial shock for which the zero-lower bound did constrain partially the Fed's response, though the distance to an optimal reaction is still substantially smaller than during the early Fed period.

Related literature

Perhaps surprisingly, the literature has produced few methods for evaluating and comparing policy makers over time or over space.

An early contribution is Fair (1978) who highlights the distortions stemming from different initial conditions and economic environments. His solution was to adopt optimal control methods to compare policy makers. This approach amounts to specifying a structural model, calculating what would have been the optimal policy based on the model and comparing the loss under such optimal policy to the loss under the implemented policy. This general approach has been used within the context of other structural models, notably New Keynesian models (e.g. Gali and Gertler, 2007; Galı, López-Salido and Vallés, 2003; Blanchard and Galí, 2007). Unfortunately, specifying the correct model for (i) the policy rule and (ii) the macroeconomic non-policy block is a very difficult task (e.g., Svensson, 2003; Mishkin, 2010). Using our approach we can greatly reduce the risk of model mis-specification as impulse responses can be estimated with *reduced form* econometric methods that are more robust to model mis-specification (e.g. Montiel Olea and Plagborg-Møller, 2021). Loosely speaking, by assessing policy performance in response to *specific* non-policy shocks, we can break the optimal policy assessment problem into smaller parts, which do not require a fully-fledged structural model. That said, compared to a model-based assessment, our approach will only evaluate performance from the reaction to a subset of all possible non-policy shocks.

In the context of monetary policy, the literature has studied the performance of the Fed pre- and post-Volcker;[5] specifically by assessing whether the Taylor principle —a central bank should react more than one-to-one to inflation movements— was satisfied.[6] However, beyond that Taylor principle, that literature can say little about the optimality of the reaction function, whether the Fed was reacting too much or too little after 1984. Overall, that

[5]See Judd and Rudebusch (1998); Taylor (1999); Clarida, Galí and Gertler (2000); Boivin (2005); Coibion and Gorodnichenko (2011) for policy rules estimates.

[6]Based on Taylor rule estimates, these papers found that the parameters of the Taylor rule shifted around 1984 and that the Fed responded more vigorously to inflation variations after 1984, though this conclusion has not gone unchallenged (e.g., Orphanides, 2003).

approach can only provide a coarse evaluation of reaction functions.

Closer to our approach, Galı, López-Salido and Vallés (2003) and Blanchard and Galí (2007) study the effect of technology shocks or oil shocks to assess the performance of the Fed pre- and post-Volcker. Different from our approach however, their assessment of good monetary policy focuses on the response of the real interest rate to a technology shock and on its distance to a specific New-Keynesian model. In contrast, our approach allows to evaluate the distance to optimality for the Fed's reaction to technology or oil price shocks without relying on any specific structural model. Instead, impulse responses are sufficient statistics.

In the context of fiscal policy Blinder and Watson (2016) improve on the naive approach of policy evaluation —measuring performance based on unconditional realized outcomes— by *projecting out* specific macro shocks, i.e., by trying to control for good luck or bad luck. In contrast, our approach *projects on* the space spanned by specific non-policy shocks and study performance in that space: comparing policy makers by studying how well they reacted to the same shock. In the context of fiscal policy Blinder and Watson (2016) improve on the naive approach of policy evaluation —measuring performance based on unconditional realized outcomes— by *projecting out* specific macro shocks and try to explain the unconditional difference between policy makers, i.e., to explain the difference from "good/bad" luck. In contrast, our approach *projects on* the space spanned by specific non-policy shocks and study performance in that space: comparing policy makers by studying how well they reacted to the same shock.

A less structural literature has proposed ways to study policy rule counter-factuals (e.g., Sims and Zha, 2006*a*; Bernanke et al., 1997; Leeper and Zha, 2003), though those approaches are not fully robust to the Lucas critique. Instead, our approach builds on recent work that shows how robustness to the Lucas critique is possible in a large class of macroeconomic models, see McKay and Wolf (2022). The key underlying assumption is that the underlying (unknown) model has a linear structure, and more specifically that the coefficients of the non-policy block are independent of the coefficients of the policy block. The present paper builds on these insights to evaluate policy institutions with minimal assumptions on the underlying (and possibly time-varying) economic structure.

Last, our paper relates to Barnichon and Mesters (2022) sufficient statistics approach to policy evaluation, though they focus on a different policy problem: the time t optimal policy problem —how to set the policy path today given the state of the economy—, instead of the unconditional policy problem that we consider here —how to set up the policy rule to minimize the unconditional loss—. Barnichon and Mesters (2022) show that the characterization of the time t optimal policy path can be reduced to the estimation of two sufficient statistics (i) forecasts for the policy objectives conditional on some baseline policy choice, (ii) the impulse responses of the policy objectives to policy shocks. However, these two statistics

are not sufficient to evaluate the optimality of the underlying policy rule. The present paper shows that a sufficient statistics approach to rule evaluation is possible, but it requires a different set of statistics, and notably additional identifying restrictions: the identification of (at least some) non-policy shocks.

The remainder of this paper is organized as follows. The next section illustrates our method for a simple New Keynesian model. Section 3 presents the general environment. Section 4 provide the key population results for evaluating and ranking policy makers. The results from the empirical study for monetary policy are discussed in Section 5. Section 6 concludes.

2 Illustrative example

To provide the intuition for our approach we informally present an example to illustrate how we can evaluate and compare policy makers based on their reaction function without having access to the underlying economic model nor the policy rule. We take the baseline New Keynesian (NK) model as the underlying economy and postulate that the researcher is interested in evaluating a central bank based on its ability to control inflation and the output gap. In this setting the optimal policy is well understood and analytically tractable, allowing us to highlight the main mechanisms of our approach and to link back to the broad NK literature (e.g. Galí, 2015).

Specifically, we consider evaluating a central bank based on the loss function

$$\mathcal{L}_t = \frac{1}{2}(\pi_t^2 + x_t^2) \ , \tag{1}$$

with π_t the inflation gap and x_t the output gap.

The log-linearized Phillips curve and intertemporal (IS) curve of the baseline New-Keynesian model (Galí, 2015) are given by

$$\pi_t = \mathbb{E}_t\pi_{t+1} + \kappa x_t + \xi_t \ , \tag{2}$$

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}(i_t - \mathrm{E}_t\pi_{t+1}) \ , \tag{3}$$

with i_t the nominal interest rate set by the central bank and ξ_t an iid cost-push shock.

To illustrate our approach suppose that the policy maker decides on the interest rate by responding to the economy according to

$$i_t = \phi_\pi \pi_t + \phi_\xi \xi_t + \epsilon_t \ , \tag{4}$$

where $\phi = (\phi_\pi, \phi_\xi)$ is the reaction function which captures the systematic response of the

central bank and ϵ_t is a policy shock. We assume that the structural shocks ξ_t and ϵ_t have mean zero and are mutually uncorrelated.

For $\phi_\pi > 1$ we can solve the model and express the endogenous variables $Y_t = (\pi_t, x_t)'$ as functions of the exogenous shocks:

$$Y_t = \Gamma(\phi)\xi_t + \mathcal{R}(\phi)\epsilon_t \ , \quad \text{with} \quad \Gamma(\phi) = \begin{bmatrix} \frac{1-\kappa\phi_\xi/\sigma}{1+\kappa\phi_\pi/\sigma} \\ \frac{-\phi_\pi/\sigma-\phi_\xi/\sigma}{1+\kappa\phi_\pi/\sigma} \end{bmatrix} \ , \quad \mathcal{R}(\phi) = \begin{bmatrix} \frac{-\kappa/\sigma}{1+\kappa\phi_\pi/\sigma} \\ \frac{-1/\sigma}{1+\kappa\phi_\pi/\sigma} \end{bmatrix} \ . \tag{5}$$

The vectors $\Gamma(\phi)$ and $\mathcal{R}(\phi)$ capture the impulse responses of the policy objectives to the structural shocks ξ_t and ϵ_t as a function of the reaction function ϕ.

An optimal reaction function is defined as any $\phi = (\phi_\pi, \phi_\xi)$ that minimizes the expected loss subject to the equations that describe the economy. Formally, let Φ denote the subset of reaction functions that lead to a unique equilibrium. The set of optimal reaction functions is given by

$$\Phi^{\text{opt}} = \left\{ \phi : \phi \in \underset{\phi\in\Phi}{\text{argmin}}\, \mathbb{E}\mathcal{L}_t \quad \text{s.t.} \quad (2)-(4) \right\} \ ,$$

which is non-empty irrespective of the structural parameter values due to the inclusion of ξ_t in (4) (e.g. Galí, 2015, page 133).

Reaction function evaluation

We will now illustrate how the impulse responses $\mathcal{R}(\phi)$ and $\Gamma(\phi)$ are sufficient statistics to evaluate a policy maker's reaction function.

Let $\phi^0 \in \Phi$ denote a central bank's reaction function which we assume is unknown to the researcher. To evaluate ϕ^0, we consider a thought experiment where ϕ^0_ξ —the reaction coefficient to the cost-push shock— is adjusted by some amount τ. Specifically, the adjusted policy rule becomes

$$i_t = \phi^0_\pi \pi_t + (\phi^0_\xi + \tau)\xi_t + \epsilon_t \ . \tag{6}$$

Following the same steps that led to (5), we can solve the model under that modified policy rule and express the endogenous variables as a function of exogenous shocks to get

$$Y_t = (\Gamma + \mathcal{R}\tau)\xi_t + \mathcal{R}\epsilon_t \ , \tag{7}$$

where $\Gamma \equiv \Gamma(\phi^0)$ and $\mathcal{R} \equiv \mathcal{R}(\phi^0)$ denote the responses to the structural shocks under the rule ϕ^0 and are defined as in (5).

Notice how the impulse response $\mathcal{R}$ contains all the information needed to compute the effect of changing the reaction coefficient ϕ^0_ξ by an amount τ. In other words, it is possible to learn the effect of a change to the rule coefficient ϕ_ξ from the impulse responses to policy

shock. This insight, which holds more generally in a large class of dynamic models (see section 4), is at the heart of our approach to evaluating reaction function from structural impulse responses.

To evaluate the reaction function, the idea is then to compute whether it is possible to adjust ϕ_ξ^0 and lower the loss function. Mathematically, we will look for a τ^* that can best lower the loss function, that is

$$\begin{aligned}\tau^* &= \underset{\tau}{\text{argmin}}\ \mathbb{E}\mathcal{L}_t \qquad \text{s.t.} \qquad Y_t = (\Gamma + \mathcal{R}\tau)\xi_t + \mathcal{R}\epsilon_t \\ &= \underset{\tau}{\text{argmin}}\ \sigma_\xi^2(\Gamma + \mathcal{R}\tau)'(\Gamma + \mathcal{R}\tau)\ , \end{aligned} \tag{8}$$

where the second equality uses that the structural shocks have mean zero and are uncorrelated. A closed form solution for τ^* is given by

$$\tau^* = -(\mathcal{R}'\mathcal{R})^{-1}\mathcal{R}'\Gamma \qquad \text{and} \qquad \phi_\xi^* = \phi_\xi + \tau^*. \tag{9}$$

We refer to the statistic τ^* as the *Optimal Reaction Adjustment*, or ORA, as it measures how much more (or less) the policy maker should have responded to the cost-push shock in order to minimize the loss function. In other words, τ^* captures the distance to the optimal reaction coefficient ϕ_ξ^*.[7]

Note how τ^* is given by the coefficient of the projection of Γ on $-\mathcal{R}$. A number of points are worth noting.

First, if the reaction function ϕ^0 was optimal, there should not exist *any* alternative reaction to ξ_t that can reduce loss, and the optimal adjustment τ^* should be zero. Geometrically, this means that Γ —the impulse responses to the cost-push shock— should be orthogonal to $\mathcal{R}$ —the impulse responses to the policy shock—.

Second, impulse responses to policy and non-policy shocks are sufficient to compute the optimal reaction to the non-policy shock, and a regression in impulse response space —regressing one impulse response on another— can determine that optimal coefficient. Intuitively, the goal of the ORA is to use the impulse responses to a monetary shock (which capture how a different reaction coefficient ϕ_ξ would affect the policy objectives) in order to better stabilize the impulse response to the non-policy shock.

Third, it may seem surprising to be able to assess a reaction function without specifying or estimating any policy rule. This reason this is possible, and the key insight underlying our approach, is that the effects of any reaction function are *encoded* in the impulse responses

[7]To prove the result $\phi^*_{\xi=\phi_\xi+\tau^*}$, use the expressions (5) for $\mathcal{R}$ and Γ under some starting rule $\phi = (\phi_\pi, \phi_\xi$. After some simple algebra, we obtain $\phi_\xi + \tau = (\kappa\sigma - \phi_\pi)/(1+\kappa^2) = \phi_\xi^{\text{opt}}$ for any $\phi_\pi > 1$, see Galí (2015, section 5).

Γ and $\mathcal{R}$, see (5) with Γ and $\mathcal{R}$ depending on ϕ. Even if we do not know the specific form of some past policy rule, that reaction function left a footprint on the effects of policy and non-policy shocks, and that footprint is sufficient to evaluate the reaction function. This is the essence of our sufficient statistics approach.

Comparing reaction functions

The ORA statistic can be used to compare the reaction functions of different policy makers. To avoid excessive notation at this stage, consider comparing two policy makers that used reaction functions ϕ^1 and ϕ^2, respectively, and let the economic environment that they faced by captured by the parameter vectors θ^1 and θ^2, respectively, which include all coefficients in the Phillips and IS curves.

For each policy maker we compute the distance between ϕ_ξ and ϕ^*_ξ. We have

$$\tau^*_j = -(\mathcal{R}^{j'}\mathcal{R}^j)^{-1}\mathcal{R}^{j'}\Gamma^j \qquad \text{for} \qquad j = 1, 2 \ ,$$

where $\mathcal{R}^j \equiv \mathcal{R}(\phi^j)$ and $\Gamma^j \equiv \Gamma(\phi^j)$. We rank the policy maker with ϕ^1 above policy maker with ϕ^2 if $|\tau^*_1| < |\tau^*_2|$.

The key insight is that while the environments are different (and thus the ϕ^*_ξ can be different across policy makers), the ORA statistics τ^*_1 and τ^*_2 measure the same quantity: the distance to the optimal reaction to the same non-policy shock ξ_t.

By evaluating the performance of a policy maker in a specific direction —the reaction to a specific non-policy shock—, we are able to evaluate and compare policy makers (or more generally policy institutions) operating in different economic environments and facing different histories of shocks.

In sum, this example illustrates how we can evaluate and compare policy makers based on their reaction function without specifying an explicit reaction function nor a specific structural macro model. Instead, the only requirement is to estimate the impulse responses Γ and $\mathcal{R}$ over a policy maker's term. The next sections show that these findings continue to hold for a general linear macro model and discuss the econometric implementation.

3 Environment

We consider a researcher who is interested in evaluating a policy maker based on the $M_y \times 1$ vector of gaps y_{t+h} —deviations of some variables from targets— for the current period $h = 0$ and possibly future periods $h = 1, 2, \ldots$. We denote by $\mathbf{Y}_t = (y'_t, y'_{t+1}, \ldots)'$ the path of the gaps that starts at a given time period t and we refer to this path as the objectives. For instance, in the illustrative example we had $y_t = (\pi_t, x_t)$ and did not include any future

objectives. We aggregate the objectives in the loss function

$$\mathcal{L}_t = \frac{1}{2}\mathbb{E}_t \mathbf{Y}_t' \mathcal{W} \mathbf{Y}_t \ , \tag{10}$$

where $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_t)$, with $\mathcal{F}_t$ the time t information set. The weighting matrix $\mathcal{W}$ is assumed to be diagonal allowing the researcher to place more or less weight on certain objectives and horizons.

Suppose that the policy maker has M_p instruments denoted by $p_t = (p_{1,t}, \ldots, p_{M_{p,t}})'$. At time t the policy maker can set the time t value of the instruments as well as their expected future path. We denote the path by $\mathbf{P}_t = (p_t', p_{t+1}', \ldots)'$ and the time t expected path is given by $\mathbb{E}_t \mathbf{P}_t$. A generic model for $\mathbf{Y}_t$ and $\mathbf{P}_t$ is given by

$$\left\{ \begin{array}{rcl} \mathcal{A}_{yy}\mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{yp}\mathbb{E}_t \mathbf{P}_t - A_{yw}\mathbb{E}_t \mathbf{W}_t &=& \mathcal{B}_{yx}\mathbf{X}_{-t} + \mathcal{B}_{y\xi}\mathbf{\Xi}_t \\ \mathcal{A}_{ww}\mathbb{E}_t \mathbf{W}_t - \mathcal{A}_{wp}\mathbb{E}_t \mathbf{P}_t - \mathcal{A}_{wy}\mathbb{E}_t \mathbf{Y}_t &=& \mathcal{B}_{wx}\mathbf{X}_{-t} + \mathcal{B}_{w\xi}\mathbf{\Xi}_t \end{array} \right. , \tag{11}$$

where $\mathbf{W}_t = (w_t', w_{t+1}', \ldots)'$ denotes the path of any additional endogenous variables, $\mathbf{X}_{-t} = (y_{t-1}', w_{t-1}', p_{t-1}', y_{t-2}', \ldots)'$ captures the initial conditions and $\mathbf{\Xi}_t = (\xi_t', \xi_{t,t+1}', \xi_{t,t+2}', \cdots)'$ is the vector of non-policy shocks. Specifically, ξ_t is the time t vector of structural non-policy shocks, while the shocks $\xi_{t,t+h}$, for $h = 1, 2, \ldots$, are news shocks: information revealed at time t about shocks that realize at time $t + h$. The maps $\mathcal{A}_{\cdot\cdot}$ and $\mathcal{B}_{\cdot\cdot}$ capture the economic environment, and we collect all these parameters of the non-policy block in $\theta = \{\mathcal{A}_{yy}, \mathcal{A}_{yp}, \mathcal{A}_{yw}, \mathcal{A}_{ww}, \mathcal{A}_{wp}, \mathcal{A}_{wy}, \mathcal{B}_{yx}, \mathcal{B}_{wx}, \mathcal{B}_{y\xi}, \mathcal{B}_{w\xi}\}$. The vector θ can be thought of as describing the economic environment that the policy maker faces.

This model is general and allows expected future policy decisions to affect current and expected future outcomes. Many models found the literature can be written in this form; prominent examples include New Keynesian models and more modern heterogeneous agents models. We normalize the non-policy news shocks $\mathbf{\Xi}_t$ to have unconditional mean zero, unit variance and to be uncorrelated with the initial conditions.

Turning to the policy block; we postulate that policy decisions can be written as

$$\mathcal{A}_{pp}\mathbb{E}_t \mathbf{P}_t - \mathcal{A}_{py}\mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{pw}\mathbb{E}_t \mathbf{W}_t = \mathcal{B}_{px}\mathbf{X}_{-t} + \mathcal{B}_{p\xi}\mathbf{\Xi}_t + \boldsymbol{\epsilon}_t \ , \tag{12}$$

where $\boldsymbol{\epsilon}_t = (\epsilon_t', \epsilon_{t,t+1}', \epsilon_{t,t+2}', \ldots)'$ a sequence of policy news shocks, i.e. exogenous shocks to the future path of $\mathbf{P}_t$ that are released at time t. The policy rule (12) imposes no restrictions as it allows the policy maker to respond to all available variables and shocks in the economy. We normalize all policy news shocks $\boldsymbol{\epsilon}_t$ to have unconditional mean zero with unit variance, and to be uncorrelated with all non-policy news shocks and initial conditions. We collect the coefficients of the policy rule in $\phi = \{\mathcal{A}_{pp}, \mathcal{A}_{py}, \mathcal{A}_{pw}, \mathcal{B}_{px}, \mathcal{B}_{p\xi}\}$ and refer to these as the

reaction function.

We denote by Φ the set of all reaction functions ϕ for which the model (11)-(12) implies a unique equilibrium. A reaction function is optimal if it minimizes the unconditional loss function. Specifically, we define the set of optimal reaction functions (there may exist more than one optimal reaction function) as follows

$$\Phi^{\text{opt}} = \left\{ \phi : \phi \in \underset{\phi \in \Phi}{\text{argmin}}\, \mathbb{E}\mathcal{L}_t \quad \text{s.t} \quad (11) - (12) \right\} . \tag{13}$$

The definition implies that we only consider optimal reaction functions that lie in Φ, i.e. the set of reaction functions which imply a unique equilibrium.

For any $\phi \in \Phi$ we can write the expected path of the objectives of the researcher as a linear function of the contemporaneous and expected future shocks as well as the initial conditions.

$$\mathbb{E}_t \mathbf{Y}_t = \Gamma(\phi)\mathbf{S}_t + \mathcal{R}(\phi)\boldsymbol{\epsilon}_t \ , \tag{14}$$

where $\mathbf{S}_t = (\mathbf{X}'_{-t}, \boldsymbol{\Xi}'_t)'$ captures the state of the economy, and the maps $\Gamma(\phi)$ and $\mathcal{R}(\phi)$ capture the causal effects of the structural shocks $\boldsymbol{\Xi}_t$ and $\boldsymbol{\epsilon}_t$ and the initial conditions $\mathbf{X}_{-t}$ on the objectives. It is useful to realize the similarity between (14) and (5) which was obtained for the illustrative New Keynesian model. Effectively, by augmenting the policy rule (12) with policy news shocks $\boldsymbol{\epsilon}_t$ we ensure that an equilibrium representation like (14) exists for a large class of dynamic models. The NK example is a special case where only the contemporaneous policy shock is needed and no initial conditions exist.

Clearly, the maps $\Gamma(\phi)$ and $\mathcal{R}(\phi)$ in (14) also depend on the environment as summarized by θ, but since θ is not under the control of the policy maker we omit this from the notation. The precise mapping from the model coefficients $\mathcal{A}_{..}$ and $\mathcal{B}_{..}$ to $\Gamma(\phi)$ and $\mathcal{R}(\phi)$ is provided in Appendix B, but we will not require knowledge of this mapping in the main text.

4 Measuring reaction function optimality

We take the position of a researcher who is interested in evaluating and ranking one or more policy makers or institutions based on their reaction function, i.e. how they responded on average to macro (non policy induced) shocks or initial conditions. We postulate that each policy maker faces an economy that can be represented by the generic model (11)-(12) where the parameters θ and ϕ may vary across policy makers. We first develop the methodology for evaluating the reaction function of a single policy maker in population. We then formalize how the methodology can be used to rank the performance of multiple policy makers.

4.1 Optimal reaction adjustments

Consider a policy maker with reaction function $\phi^0 \in \Phi$ and work in a setting where ϕ^0 is unknown to the researcher. Our evaluation of ϕ^0 is based on measuring how far ϕ^0 is from the set of optimal reaction functions Φ^{opt} as defined in (13). We compute this distance in the direction of the response to the state $\mathbf{S}_t$ which includes the structural shocks $\mathbf{\Xi}_t$ and the initial conditions $\mathbf{X}_{-t}$.

Consider the augmented policy rule

$$\mathcal{A}^0_{pp}\mathbb{E}_t\mathbf{P}_t - \mathcal{A}^0_{py}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}^0_{pw}\mathbb{E}_t\mathbf{W}_t = (\mathcal{B}^0_{px} + \mathcal{T}_X)\mathbf{X}_{-t} + (\mathcal{B}^0_{p\xi} + \mathcal{T}_\Xi)\mathbf{\Xi}_t + \boldsymbol{\epsilon}_t \ ,$$

where $\mathcal{T} = [\mathcal{T}_\Xi, \mathcal{T}_X]$ adjusts the ϕ^0 response to the state. Given that $\phi^0 \in \Phi$ we can combine this policy rule with the general model (11) to obtain the equilibrium representation

$$\mathbb{E}_t\mathbf{Y}_t = (\Gamma + \mathcal{R}\mathcal{T})\mathbf{S}_t + \mathcal{R}\boldsymbol{\epsilon}_t \ , \tag{15}$$

where $\mathcal{R} \equiv \mathcal{R}(\phi^0)$ and $\Gamma \equiv \Gamma(\phi^0)$. The equilibrium effect of $\mathcal{T}$ is found to be equal to $\mathcal{R}\mathcal{T}\mathbf{S}_t$ on the objectives and is proportional to the effect of the policy news shocks as captured by $\mathcal{R}$.

In other words, as in the simple example, it is possible to compute the effect of a different policy rule —changing the reaction to the non-policy shocks $\mathbf{\Xi}_t$ or the initial conditions $\mathbf{X}_t$— from the impulse responses to policy shocks. Such counter-factual analysis is fully robust to the Lucas critique provided that in the (unspecified) underlying economic model, the coefficients of the macro block in (11) are invariant to changes in the coefficients of the policy rule (see Barnichon and Mesters, 2022; McKay and Wolf, 2022).

Similar to the simple New Keynesian example, the *Optimal Reaction Adjustment* (ORA) is defined as the $\mathcal{T}$ that minimizes the expected loss function. Formally,

$$\mathcal{T}^* = \underset{\mathcal{T}}{\text{argmin}} \ \mathbb{E}\mathcal{L}_t \qquad \text{s.t.} \quad \mathbb{E}_t\mathbf{Y}_t = (\Gamma + \mathcal{R}\mathcal{T})\mathbf{S}_t + \mathcal{R}\boldsymbol{\epsilon}_t \ . \tag{16}$$

The ORA determines how the reaction coefficients in front of the state $\mathbf{S}_t$ should have been adjusted to minimize the unconditional loss. A given (i, j) entry of $\mathcal{T}$ measures how the policy maker should have responded differently to the jth non-policy input when setting the ith policy instrument.

An explicit expression is given by

$$\mathcal{T}^* = -(\mathcal{R}'\mathcal{W}\mathcal{R})^{-1}\mathcal{R}'\mathcal{W}\Gamma \ , \tag{17}$$

which exists provided that the inverse exists, and highlights that the ORA can be viewed

as the projection of the impulse responses to non-policy shocks on the impulse responses to policy shocks. The weighting matrix $\mathcal{W}$ incorporates the preferences of the researcher who aims to determine how the reaction function should have been adjusted to achieve the objective of minimizing the loss function (10).

The following proposition establishes the key property of the ORA.

Proposition 1. *Given the generic model* (11)-(12)*, with* Φ *non-empty, we have that*

$$\phi^* \in \Phi^{\text{opt}} \ , \quad \textit{where} \quad \phi^* = \{\mathcal{A}^0_{pp}, \mathcal{A}^0_{py}, \mathcal{A}^0_{pw}, \mathcal{B}^0_{px} + \mathcal{T}^*_X, \mathcal{B}^0_{p\xi} + \mathcal{T}^*_\Xi\} \ ,$$

where $\mathcal{T}^* = [\mathcal{T}^*_\Xi, \mathcal{T}^*_X]$ *is defined in* (16).

The vector $\mathbf{S}_t$ captures the inputs that enter the policy maker's problem at time t: the time t non-policy shocks and the initial conditions at t. As these shocks and initial conditions affect the current and expected future path of the objectives through the reaction function (among other channels), the ORA computes the optimal adjustment to the reaction coefficients $\mathcal{B}^0_{px}$ and $\mathcal{B}^0_{p\xi}$, i.e. $\mathcal{B}^0_{px} \to \mathcal{B}^0_{px} + \mathcal{T}^*_X$ and $\mathcal{B}^0_{p\xi} \to \mathcal{B}^0_{p\xi} + \mathcal{T}^*_\Xi$, in order to minimize the unconditional loss function.

Subset optimal reaction adjustments

One difficulty in practice is that there often exists insufficient exogenous variation to identify all impulse responses Γ and $\mathcal{R}$. For instance, estimating $\mathcal{R}$ would require being able to identify the contemporaneous shocks to all policy instruments as well as all the policy news shocks that create exogenous variation in the expected policy paths.

Fortunately, in practice we can still compute the distance to a subset of all reaction function coefficients. This allows to evaluate policy makers based on the specific structural shocks that can be identified or have already been identified by the literature.

To set this up, we denote by $\mathcal{R}_a$ the effects of $\boldsymbol{\epsilon}_{a,t}$ on $\mathbb{E}_t\mathbf{Y}_t$ under ϕ^0, where $\boldsymbol{\epsilon}_{a,t}$ is a vector formed from any subset or linear combination of the policy news shocks $\boldsymbol{\epsilon}_t$. Similarly, denote by Γ_b the impulse responses of subsets or linear combinations $\mathbf{S}_{b,t}$ on $\mathbb{E}_t\mathbf{Y}_t$.

The subset ORA is defined as

$$\mathcal{T}^*_{ab} = \underset{\mathcal{T}_{ab}}{\text{argmin}}\, \mathbb{E}\mathcal{L}_t \qquad \text{s.t.} \quad \mathbb{E}_t\mathbf{Y}_t = (\Gamma_b + \mathcal{R}_a\mathcal{T}_{ab})\mathbf{S}_{b,t} + \Gamma_{-b}\mathbf{S}_{-b,t} + \mathcal{R}\boldsymbol{\epsilon}_t \ , \tag{18}$$

where $\mathcal{R}$ separates as $[\mathcal{R}_a : \mathcal{R}_{-a}]$ and Γ as $[\Gamma_b : \Gamma_{-b}]$. Using that the structural shocks are mean zero and uncorrelated we can derive the closed form solution

$$\mathcal{T}^*_{ab} = -(\mathcal{R}_a'\mathcal{W}\mathcal{R}_a)^{-1}\mathcal{R}_a'\mathcal{W}\Gamma_b \ , \tag{19}$$

which shows that the subset ORA is equal to the projection of the identifiable non-policy impulse responses Γ_b on the identifiable policy impulse responses $\mathcal{R}_a$.

We have the following result.

Corollary 1. *Given the generic model* (11)-(12)*, with* Φ *non-empty, we have that the adjusted reaction function*

$$\phi^*_{ab} = \{\mathcal{A}^0_{pp}, \mathcal{A}^0_{py}, \mathcal{A}^0_{pw}, \mathcal{B}^0_{ab} + \mathcal{T}^*_{ab}, \mathcal{B}^0_{-a-b}\}$$

satisfies

$$\mathbb{E}\mathcal{L}_t(\phi^*_{ab}) \leq \mathbb{E}\mathcal{L}_t(\phi^0) \ ,$$

where $\mathcal{B}^0_{ab}$ *captures the entries of* $[\mathcal{B}^0_{px}, \mathcal{B}^0_{p\xi}]$ *corresponding to the responses of* $\mathbb{E}_t\mathbf{P}_{a,t}$ *to* $\mathbf{S}_{b,t}$ *and* $\mathcal{B}^0_{-a-b}$ *denotes the remaining entries.*

The corollary ensures that adjusting the given reaction function for the policy instruments in the direction of the identifiable non-policy shocks lowers the expected loss function. The intuition here is the same as in the previous section: adjusting $\mathcal{B}^0_{p_a\xi_b}$ does not change the coefficients in the economy and hence for the class of models (11) such change is robust to the Lucas critique.

The result is of great practically relevant as it shows that researchers never have to recover the entire causal maps Γ^0 and $\mathcal{R}^0$ to evaluate the reaction function. For instance, suppose a researcher is interested in testing the central bank's reaction to an oil price shock when setting the short term interest rate, then only the impulse responses of the researcher's objectives, say inflation and unemployment, to interest rate shocks and oil price shocks are needed. Ideally, in this scenario we would include all interest rate news shocks, but in practice we can conduct a subset test after only identifying, for instance, the contemporaneous policy shock.

Some further intuition

Next, we briefly discuss a few intuitive ways of interpreting and deriving the subset ORA statistic. For simplicity we consider the setting where $\epsilon_{a,t}$ is a scalar policy shock and $\xi_{b,t}$ is a scalar non policy shock.

First, we can interpret the ORA statistic as resulting from a two step procedure. We start from the definition of the $\mathcal{T}_{ab}$ adjusted equilibrium in (18), i.e. recall

$$\mathbb{E}_t\mathbf{Y}_t = (\Gamma_b + \mathcal{R}_a\mathcal{T}_{ab})\xi_{b,t} + \Gamma_{-b}\mathbf{S}_{-b,t} + \mathcal{R}\boldsymbol{\epsilon}_t \ .$$

In a first step we project $\mathbb{E}_t\mathbf{Y}_t$ on the non-policy shock, recalling that all policy shocks have

unit variance we have

$$\mathbb{E}(\xi_{b,t}\mathbb{E}_t\mathbf{Y}_t) = \Gamma_b + \mathcal{R}_a\mathcal{T}_{ab} \ .$$

This step effectively isolates our object of interest —the response to the *specific* non policy shock $\xi_{b,t}$—, and removes the confounding effects of the initial conditions and other shocks, i.e. $\mathbf{S}_{-b,t}$ and $\boldsymbol{\epsilon}_t$.

In the second step we then solve the policy problem in the projected space, i.e. we solve

$$\mathcal{T}^*_{ab} = \underset{\mathcal{T}_{ab}}{\operatorname{argmin}} \ (\Gamma_b + \mathcal{R}_a\mathcal{T}_{ab})'\mathcal{W}(\Gamma_b + \mathcal{R}_a\mathcal{T}_{ab}) \ .$$

Second, the ORA statistic has an intuitive and interesting direct instrumental variable interpretation. Specifically, the subset ORA statistic can be expressed as a population generalized moment estimator for $\mathcal{T}_{ab}$ in the expression

$$\xi_{b,t} = -\mathcal{T}_{ab}\epsilon_{a,t} + \eta_{ab,t} \ , \tag{20}$$

where the instruments used are $\mathbf{Y}_t$ and the weighting matrix is $\mathcal{W}$. Indeed, we have

$$\begin{aligned}\mathcal{T}^*_{ab} &= -\left(\mathbb{E}(\epsilon_{a,t}\mathbf{Y}_t')\mathcal{W}\mathbb{E}(\mathbf{Y}_t\epsilon_{a,t})\right)^{-1}\mathbb{E}(\epsilon_{a,t}\mathbf{Y}_t')\mathcal{W}\mathbb{E}(\mathbf{Y}_t\xi_{b,t}) \\ &= -(\mathcal{R}_a'\mathcal{W}\mathcal{R}_a)^{-1}\mathcal{R}_a'\mathcal{W}\Gamma_b \ .\end{aligned}$$

Note that the linear projection of $\xi_{b,t}$ on $\epsilon_{a,t}$ is zero by definition, as structural shocks are uncorrelated by construction. In contrast, both shocks are are correlated with the policy objectives (i.e. equation (14)) and the instrumental variable estimator is exactly the ORA statistic: the first stage projects the structural shocks on the policy objectives and the second stage projects the resulting impulse responses on each other.

This interpretation implies that if the structural shocks $\epsilon_{a,t}$ and $\xi_{b,t}$ are directly observed we can compute the ORA using a simply instrumental variable regression. In practice, we do not recommend this approach as structural shocks are usually not observable, or at least subject to measurement error (Stock and Watson, 2018), and replacing the shocks in (20) by proxies subject to measurement error is not inconsequential.

Third and finally, the ORA statistic is related to the Optimal Policy Perturbation (OPP) statistic that was introduced in Barnichon and Mesters (2022) for the evaluation of time-t policy paths $\mathbb{E}_t\mathbf{P}_t$. Specifically, they considered the problem: how to set the policy path $\mathbb{E}_t\mathbf{P}_t$ in order to minimize the *conditional* loss function $\mathcal{L}_t$. A key result was that the subset OPP statistic

$$\delta_{a,t} = -(\mathcal{R}_a'\mathcal{W}\mathcal{R}_a)^{-1}\mathcal{R}_a'\mathcal{W}\mathbb{E}_t\mathbf{Y}_t \ ,$$

corrects the policy path in order to lower the conditional loss, i.e. $\mathbb{E}_t\mathbf{P}_t + \mathcal{R}_{a,p}\boldsymbol{\delta}_{a,t}$ lowers the

loss when compared to the initial path $\mathbb{E}_t \mathbf{P}_t$. Here $\mathcal{R}_{a,p}$ is the impulse response of $\mathbf{P}_t$ to the policy shock $\epsilon_{a,t}$.

Hence, we can view the sequence of $\{\delta_{a,t}\}$ as corrections to the policy path over different time periods in the direction of the identified policy shock $\epsilon_{a,t}$. A limitation of $\delta_{a,t}$ is that it does not allow the researcher to distinguish between the different sources of policy mistakes, i.e. which shocks should the policy maker have responded to differently. The ORA statistic solves this limitation by asking; how much of the average policy mistakes can be attributed to any specific non-policy shock. Specifically, we have that

$$\begin{aligned} \mathcal{T}^*_{ab} &= \text{Cov}(\delta_{a,t}, \xi_{b,t}) \\ &= -(\mathcal{R}_a' \mathcal{W} \mathcal{R}_a)^{-1} \mathcal{R}_a' \mathcal{W} \Gamma_b \ . \end{aligned}$$

This derivation highlights that the ORA is an unconditional statistic, it captures the average adjustment to the reaction function, that would have lowered the *unconditional* loss $\mathbb{E}\mathcal{L}_t$.

4.2 Comparing policy institutions with ORAs

With the ORA and its properties established we now discuss how the ORA can be used to compare policy institutions or policy makers. As examples we can think of evaluating different central banks chairs based on their ability to control inflation and output gaps, or different presidents of a country based on their ability to keep output close to potential. Our comparisons are based on evaluating policy makers on their use of the same policy instruments for offsetting the same non-policy shocks. As such we may generally compare policy makers from the same institution across different time periods or policy maker from different but comparable institutions from different countries.

Suppose that there are p policy makers that the researcher aims to compare. Each policy maker faces an economy that can be described by the general model (11), but the parameters θ that govern the model may vary across policy makers, say θ_j for $j = 1, \ldots, p$. Similarly, each policy maker is assumed to set policy according to the generic rule (12), but may use a different reaction function ϕ_j. Following the notation defined above, let ϕ_j^0 denote the chosen reaction function of policy maker j. We note that while here we treat the parameters as fixed within the term of each policy maker, in our econometric implementation section below we show how we can extend the approach to allow the parameters to change within terms.

Using the methodology established above we can compute for any given policy maker the distance to the optimal reaction function in any identifiable direction using the (subset)

ORA statistic. Specifically, let

$$\mathcal{T}_{ab}^{j*} = -(\mathcal{R}_a^{j'}\mathcal{W}\mathcal{R}_a^j)^{-1}\mathcal{R}_a^{j'}\mathcal{W}\Gamma_b^j$$

denote the subset ORA statistic for policy maker j in the direction of responding to non-policy shocks $\mathbf{S}_{b,t}$ using the policy instruments moved by the policy shocks $\boldsymbol{\epsilon}_{a,t}$. We recall that $\mathcal{R}_a^j$ and Γ_b^j are the impulse responses of the objectives with respect to the policy and non-policy shocks computed under the reaction function ϕ_j^0 and given the economic environment θ_j.

The ORA statistics take into account the preferences of the researcher over the different objectives or ranking criteria. As such if the researcher has no further preferences over the types of shocks we may simply aggregate the entries of $\mathcal{T}_{ab}^{j*}$, i.e.

$$t_{ab}^{j*} = \|\text{vec}\mathcal{T}_{ab}^{j*}\| \ , \tag{21}$$

where any desired norm $\|\cdot\|$ can be used. We rank policy makers based on t_{ab}^{j*}, for $j = 1, \ldots, p$, where the smallest value corresponds to the best performing policy maker. For interpretation purposes it is generally useful to present the ranking separately for each combination of instrument and non-policy input as each ranking is informative about a specific dimension of policy.

In practice there could be cases where the researcher has preferences over the types of non-policy inputs. For instance, when evaluating a central bank one may find it more important that the central bank offset oil price shocks as opposed to TFP shocks. In such cases a weighted norm can be used in (21) to aggregate the entries of $\mathcal{T}_{ab}^{j*}$.

4.3 Computing ORA statistics

We discuss the computation of the optimal reaction function adjustments using observation data. Without loss of generality we will consider the subset ORA statistic defined in (19). Corollary 1 shows that if $\mathcal{T}_{ab}^* \neq 0$ then ϕ^0 is not optimal and adjusting by $\mathcal{T}_{ab}^*$ brings the reaction function closer to the set of optimal reaction functions.

The starting point is the equilibrium representation (14), restated here for convenience

$$\mathbb{E}_t\mathbf{Y}_t = \Gamma\mathbf{S}_t + \mathcal{R}\boldsymbol{\epsilon}_t \ .$$

Following empirical practice, we truncate the relevant horizon at H and define $\mathbf{Y}_{t:t+H} = (y_t', \ldots, y_{t+H}')'$ as the evaluation criteria of interest at time t. Further, suppose that the policy maker under consideration was active for periods $t = 1, \ldots, n$ during which the reaction

function ϕ^0 was used.

The causal effects $\mathcal{R}_a$ and Γ_b can be estimated by considering

$$\mathbf{Y}_{t:t+H} = \Gamma_b \mathbf{S}_{b,t} + \mathcal{R}_a \boldsymbol{\epsilon}_{a,t} + \mathbf{U}_{t:t+H} \ , \qquad t = 1, \ldots, n, \tag{22}$$

where $\mathbf{U}_{t:t+H}$ includes all other structural shocks, both policy and non-policy inputs that are not included in the selections a and b, respectively, as well as the forecast errors $\mathbf{Y}_{t:t+H} - \mathbb{E}_t \mathbf{Y}_{t:t+H}$.

We can recognize (22) as a system of stacked local projections (Jordà, 2005). This implies that given (i) an appropriate identification strategy and (ii) an accompanying estimation method, we can estimate the impulse responses $\mathcal{R}_a$ and Γ_b using standard local projection methods. Any identification strategy — short run, long run, sign, external instruments, etc — can be used, based on which an appropriate estimation method — OLS or IV, with or without shrinkage, etc — can be selected, see Ramey (2016) and Stock and Watson (2018) for different options. Moreover, we recall from Plagborg-Møller and Wolf (2021) that in population local projections and structural VARs estimate the same impulse responses; therefore all SVAR methods discussed in Kilian and Lütkepohl (2017), for instance, can also be adopted for estimating the impulse responses Γ_b and $\mathcal{R}_a$.

In sum, given (22), observational data that covers the term of the policy maker, and by making use of the existing literature we can recover estimates for $\mathcal{R}_a$ and Γ_b. Here we will not discuss any specific approach but instead directly postulate that the researcher is able to obtain estimates, say $\widehat{\mathcal{R}}_a$ and $\widehat{\Gamma}_b$, which can be approximated by

$$\operatorname{vec}\left(\begin{bmatrix} \widehat{\mathcal{R}}_a \\ \widehat{\Gamma}_b \end{bmatrix} - \begin{bmatrix} \mathcal{R}_a \\ \Gamma_b \end{bmatrix} \right) \overset{a}{\sim} F \ ,$$

where F is some known distribution function that can be estimated consistently by $\widehat{F}$. Such approximation can be obtained for many impulse response estimators using both frequentist (asymptotic and bootstrap) and Bayesian estimators.

Using the approximating distribution $\widehat{F}$ can simulate draws for $\mathcal{R}_a$ and Γ_b, and compute $\mathcal{T}^*_{ab} = -(\mathcal{R}_a' \mathcal{W} \mathcal{R}_a')^{-1} \mathcal{R}_a' \mathcal{W} \Gamma_b$ for each draw. Given the sequence of draws we can construct a confidence set for $\mathcal{T}^*_{ab}$, or any of its individual entries at any desired level of confidence. We note that if the distribution F is normal we can use the delta method to analytically compute the distribution of $\mathcal{T}^*_{ab}$.

5 Evaluating US monetary policy, 1879-2019

In this section we use our methodology to evaluate the conduct of monetary policy in the US over the 1876-2020 period using quarterly data. We consider four distinct periods: (i) the Gold Standard period 1879-1912 before the creation of the Federal Reserve, (ii) the early Fed years over 1913-1941, (iii) the post World War II period 1954-1984 and (iv) the post-Volcker period over 1990-2020.

During the Gold Standard period, there was no active monetary policy (the Federal Reserve did not exist yet), and we use this period as a benchmark to see what a fictional policy institution could have done in this period. A Gold Standard monetary regime is now generally considered a sub-optimal regime with excessive fluctuations in inflation and unemployment (e.g. Friedman and Schwartz, 1963). In that context, this passive monetary policy period is instructive as a benchmark against which we can compare later Fed performances. The early Fed period starts with the founding of the Fed in 1913 and ends with the US entering the second world war. The post-war period starts in 1951 with the Fed regaining some independence after the Treasury-Fed accord (e.g. Romer and Romer, 2004*a*).[8] The post Volcker period starts in 1984 —the beginning of the so-called Great Moderation period— and ends right before the pandemic.

We evaluate the Fed as a policy institution based on the loss function

$$\mathcal{L}_t = \frac{1}{2}\mathbb{E}_t \sum_{h=1}^{H} \beta^h (\pi_{t+h}^2 + \lambda u_{t+h}^2) \ , \tag{23}$$

where π_t denotes inflation, u_t the unemployment rate, β the discount factor and λ the preference parameter. Our baseline choice for the loss function sets $\beta = \lambda = 1$ and considers $H = 20$ quarters.[9] Inflation is measured as year-on-year inflation based on the output deflator from Balke and Gordon (1986). The unemployment rate before 1948 is taken from the NBER Macrohistory database over 1929-1948 and extended back to 1876 by interpolating the annual series from Weir (1992) and Vernon (1994). The top panel in Figure 1 shows the time series for inflation and unemployment.

[8] We exclude the period covering World War II until the Treasury-Fed accord of 1951, as the Fed was financing the war effort and had no independence. In April 1942, at the request of the Department of the Treasury, the Fed formally committed to maintaining a low interest-rate peg on short-term Treasury bills and also implicitly capped the rate on long-term Treasury bonds. The goal of the peg was to stabilize the securities market and allow the federal government to engage in cheaper debt financing of World War II. This system lasted until the 1951 Treasury-Fed accord separated government debt management from monetary policy (Romero, 2013).

[9] The robustness of our findings with respect to these choices is assessed in the web-appendix. We stress that the choice for the loss function is merely an evaluation criteria in our context what the true loss function of the different Fed chairs was is irrelevant from our perspective.

5.1 Naive approach

To provide a benchmark for our results, we first evaluate the Fed based on realized outcomes for inflation and unemployment, as shown in Figure 1. The first two rows of 1 display average inflation and unemployment over each period. While the Early Fed period stands out with a much higher average unemployment (driven by the Great Depression), it is the post WWII period that stands out with a much higher average inflation (driven by the Great Inflation of the 1970s).

While the pre-Fed period looks good in terms of averages, the conclusion changes dramatically when one looks at variances, which corresponds more closely to typical loss functions as posited in (10). Table 1 also shows the sample variances of inflation and unemployment.[10] We confirm a common wisdom about the early Fed period with very poor performances in terms of volatility in inflation and unemployment. But now the pre-Fed period also stands out with very high inflation volatility, at least much higher than after World War II. The post-Volcker period shows the most stable inflation by far, though the volatility of unemployment is comparable to the pre-Fed and post WWII periods.

As we discussed in the introduction however, these unconditional realized outcomes cannot be used to assess monetary policy performance. While they could be due to poor monetary policy —an inadequate reaction function—, many co-founding factors outside the Fed control could a priori also explain those results. For instance, the poor realizations in terms of inflation and unemployment over 1913-1941 could have been caused by bad luck (an unfortunate sequence of shocks), adverse initial conditions or by a difficult economic environment. To answer these questions, we will thus used the ORA methodology that we propose in this paper.

5.2 Econometric implementation for ORA

We first discuss the implementation details for computing the ORA statistics based on which we will compare the Fed across periods.

As we discussed in the general treatment, a complete evaluation of the reaction function requires evaluating the reaction of the entire policy path to all non-policy shocks. While this is not feasible in finite sample, we can evaluate performance from the subset ORA statistic; by evaluating the optimality of some coefficients of the policy rule in response to a subset of all possible non-policy shocks.

In our application, we will measure how well the contemporaneous policy rate reacted to five distinct disturbances: financial shocks, government spending shocks, energy price

[10] This exercise can be seen as computing the loss (23) using realized outcomes, i.e., using the unconditional variance of inflation and unemployment over each regime, if we assume that the sample means were the objectives of inflation and unemployment.

shocks, inflation expectation shocks and TFP shocks. A concrete example is useful to fix ideas. From (12) the generic policy rule for the contemporaneous policy rate is of the form

$$p_{t|t} = \rho_0 p_{t-1} + ... + a_0 \mathbb{E}_t \pi_t + ... + \phi_{\xi,0} \xi_{t,t} ... + \epsilon_{t,t}$$

with ρ_0, a_0 some coefficients, and we will measure the optimality of the reaction coefficient $\phi_{\xi,0}$ (where the 0 subscript denotes a coefficient of the contemporaneous policy rule) from the distance $\tau_t^* = \phi_{\xi,0} - \phi_{\xi,0}$. To estimate τ_t^*, we need to estimate the impulse responses to the non-policy shock $\xi_{t,t}$ and the contemporaneous policy shock $\epsilon_{t,t}$.

To estimate the impulse responses to policy and non-policy shocks we rely on a Bayesian structural vector autoregressive model (SVAR). We include a proxy for the policy shock, the non-policy shock, the outcome variables π_t and u_t, the growth rate of the monetary base, the policy rate, as well as possibly additional control variables w_t.

To capture the policy stance during the pre WWII periods, we include in the VAR the year-to-year growth rate of the monetary base (M0) and a short term interest rate. During the 1879-1912 Gold Standard period where there is no policy institution, we take the 3-months treasury rate as the "policy rate" that a fictitious central bank could have controlled. For the 1913-1941 early Fed period, we use the fed discount rate as the policy rate. To capture the policy stance during the post WWII periods, we use the fed funds rate as the policy rate. The specific additional variables w_t and instruments z_t are discussed in detail below. The historical monetary data are taken from (Balke and Gordon, 1986)

The SVAR is specified for $y_t = (z_t, \pi_t, u_t, p_t', w_t')'$, where the ordering may vary depending on the identification scheme considered. We have

$$A_0 y_t = A_1 y_{t-1} + \ldots + A_p y_{t-p} + e_t \ , \tag{24}$$

where $A_0, \ldots, A_p$ are the coefficient matrices and e_t captures the structural shocks. In each case our objective is to choose y_t such that e_t includes the subset policy shock $\boldsymbol{\epsilon}_{a,t}$ and the subset non-policy shock $\boldsymbol{\Xi}_{b,t}$. In our study these will always be scalars, $\boldsymbol{\epsilon}_{a,t}$ is chosen as the conventional contemporaneous monetary policy shock, and $\boldsymbol{\Xi}_{b,t}$ can correspond to shock to energy prices, financial intermediation, productivity, government spending and inflation expectations, see the discussion below.

We estimate the reduced form of the SVAR model using standard Bayesian methods, which shrink the reduced form VAR coefficients using a Minnesota style prior. The prior variance hyper-parameters follow the recommendations in Canova (2007). When no external instruments are included we generally rely on short, long or sign restrictions to identify the impulse responses. When instruments are included, we order them first and compute the impulse responses of the objectives to these shocks after normalizing them with respect to

the contemporaneous response of the corresponding policy rate or endogenous variable (e.g. Plagborg-Møller and Wolf, 2021). Alternatively, in some settings the series z_t can be of direct interest and we do not rescale the impulse responses.

We normalize all shocks such that they have unit variance which can be implemented in practice by computing the conventional one standard deviation impulse responses. This scaling is important to ensure comparability of the shocks across periods. With the draws of the parameters from the posterior density we compute the impulse responses $\mathcal{R}_a$ and Γ_b, and the subset ORA statistic $\mathcal{T}^*_{ab}$ using (18). Besides reporting Γ_b we also report the adjusted $\Gamma_b + \mathcal{R}_a\mathcal{T}^*_{ab}$ to assess how the impulse response to the non-policy shocks could have been adjusted.

Importantly, while we rely on a structural VAR for estimating the impulse responses, other estimation methods could have been used, for instance local projections (Jordà, 2005; Stock and Watson, 2018). Moreover, we insist that the ORA based assessment is very different from traditional VAR based assessments of policy, such as shutting down the interest rate channel (e.g. Sims and Zha, 2006*b*; Bernanke et al., 1997). A limitation of that approach is that it is subject to the Lucas critique as the counter-factual policy paths are imposed by repeatedly surprising agents with contemporaneous policy shocks.

In contrast, our ORA-based reaction function assessment is not subject to the Lucas critique. Instead of imposing a counter-factual policy path, our approach consists in evaluating whether different policy rule coefficients could have lowered the loss function, and the effects of such counter-factuals rule coefficients can be explored from the impulse responses to policy shocks in a manner that is fully robust to the Lucas critique (see Barnichon and Mesters, 2022; McKay and Wolf, 2022).[11]

5.3 Shock identification

For each of our for periods, we aim to identify a contemporaneous monetary policy shock and five non-policy shocks: financial shocks, government spending shocks, energy price shocks, inflation expectation shocks and TFP shocks.

5.3.1 Monetary policy shocks

We consider two main approaches for identifying monetary policy shocks. As our baseline identification scheme for monetary shocks, we use the state of the art in the literature, relying

[11]The key requirement to ensure robustness to the Lucas critique is that the general economy admits a representation like (11)-(12) where the coefficients of the non-policy block are independent of the coefficients of the policy block. As shown by McKay and Wolf (2022), this requirement holds in a very large class of macro models.

as much as possible on a narrative approach. As robustness, we also use a sign restriction identification.

Pre Fed regime For the Pre Fed Gold Standard period, there is no clear baseline identification approach to identify monetary shocks, and we propose a new approach that exploits the unique feature of the Gold Standard. Under a Gold Standard, the monetary base depends on the amount of gold in circulation, which can itself vary for exogenous reasons related to the random nature of gold discoveries or development of new extraction techniques (e.g., Barsky and De Long, 1991). As such, we use unanticipated large gold mine discoveries (discoveries that led to gold rushes) as an instrument for movements in the monetary base. To the extent that the timing of the discovery is unrelated to the state of the business cycle, gold mine discovery will be a valid instrument. Mirroring Gold discovery, we will also use peak mine extraction —the moment where the mine output reached its peak production—.

Figure 2 plots gold production along with our identified dates: the peak of the Comstock Lode mine in Nevada in 1877-Q1, the discovery of Gold in Alaska in 1896-Q3 (which led to the Alaska gold rush), the discovery of large Gold mines in Nevada in 1902-Q1 (which led to the Nevada gold rush), the Nevada gold mine maximum in 1910-Q1.[12] Given the unpredictability of the amount of gold available in any given region (either at the onset of a gold rush or at its zenith), we can consider these events as unanticipated. We code the Gold shocks as one when a new mine was discovered and minus one when the peak was reached.

Early Fed regime During the Early Fed period we use the Friedman and Schwartz (1963) dates extended by Romer and Romer (1989) as instruments to identify monetary policy shocks. We include five episodes —1920Q1, 1931Q3, 1933Q1, 1937Q1 and 1941Q3— where movements in money were "unusual given economic developments" (Romer and Romer, 1989). In the words of Romer and Romer (1989), these "unusual movements arose, in Friedman and Schwartz's view, from a conjunction of economic events, monetary institutions and the doctrines and beliefs of the time and of particular individuals determining policy". The Friedman and Schwartz (1963) dates are also modeled using dummies.

Post World War II regime For the Post World War II period we use the Romer and Romer (2004*b*) narrative identified monetary policy shocks as instruments.

Post Volcker regime For the Post Volcker period we use the high-frequency identification (HFI) approach, pioneered by Kuttner (2001) and Gürkaynak, Sack and Swanson (2005),

[12] Another important date is the Alaska gold mine maximum in 1915, though it is not useful in this context since the strict gold standard era stops in 1913 with the founding of the Fed.

and use surprises in fed funds futures prices (FF4) around FOMC announcement as proxies for monetary shocks.

Alternative identification scheme One limitation of using different shock proxies over the different regimes is that each proxy may have different characteristics that could affect the results and the ORA comparison across periods.[13] To guard ourselves against this possibility, we will also use an identification of monetary shock that is consistent across regimes, which will ensure that the monetary shocks are defined and identified in the exact same way across regimes. Specifically, we use sign restrictions, another popular method to identify monetary shocks (e.g., Uhlig, 2005). We restrict the impulse responses of inflation to be non-positive and the responses of unemployment and the interest rate to be non-negative for two quarters. All other responses are left unrestricted, see also Amir-Ahmadi, Matthes and Wang (2016) who implement the same restrictions. This approach has the benefit that it can be implemented over the entire sampling period. With the VAR including inflation, unemployment, the 3 months treasury rate, and we impose the following sign restrictions: a positive monetary shock raises the short-term rate in impact, lowers inflation after four quarters and raises unemployment after four quarter. Beyond quarter 4, the responses of inflation and unemployment are unconstrained. As we will see the results from the sign restricted identification scheme generally correspond to the results obtained with our baseline identification of monetary shocks.

5.3.2 Non-policy shocks

Financial shocks As our main financial shocks we use narrative identified bank panics as directly as an observable shock series. Each included panic was triggered by either a run on a particular trust fund or by foreign developments. The dates for the banking panics are taken from Reinhart and Rogoff (2009) and Schularick and Taylor (2012) and Romer and Romer (2017). To capture the severity of the bank run, each non-zero entry is rescaled by the change in the BAA yield at the time of the run, similar to the re-scaling of Bernanke et al. (1997) and in the spirit of Romer and Romer (2017)'s scaling of their financial distress index.[14] The use of banking panics comes with two concerns. First, a potential worry is that these banking panics are themselves due to the inflation and unemployment outlook, making the banking panic dummy endogenous. Second, there are no major bank runs over our post WWII period. Therefore, as an alternative strategy, we use innovations to the corporate

[13]For instance, exogeneity and relevance may differ across instrumental variables, see e.g., Barnichon and Mesters (2020) for a discussion of the different strength and limits of the Romer and Romer (2004*b*) and the Gürkaynak, Sack and Swanson (2005) shock proxies.

[14]Uing bank runs as 0-1 dummies does not change conclusions drastically though it makes the estimates a bit less precise.

credit spread (BAA-AAA) available over 1919-2019, effectively regressing the spread on current and past values of inflation, unemployment, real GDP growth, as well as lags of the 3-month and 10-year treasury rates. This identification amounts to a recursive ordering where the identifying assumption is that inflation and unemployment react with a quarter lag to innovations to the corporate credit spread.

Government spending shocks For government spending shocks we consider news shocks to defense spending as constructed in Ramey and Zubairy (2018).

Productivity shocks To identify productivity shocks we use the identification scheme of Gali (1999) and Barnichon (2010): we estimate bi-variate VARs with log output per hour and unemployment over each policy regime, and we impose long-run identifying restrictions, specifically that only productivity shocks can have permanent effects on productivity. The quarterly time series for output per hour is taken from Petrosky-Nadeau and Zhang (2021), though it only starts in 1890.

Energy shocks To identify energy shocks, we adapt the approach of Hamilton (1989) for monthly oil prices to quarterly energy prices. The modification from monthly to quarterly is common (e.g. Stock and Watson, 2012), but the change to energy prices is novel. The reason is that oil only became the pre-dominant source of energy after World War II. Before World War I, coal was the primary US energy source. To avoid the ad hoc connection of such series we measure energy price prices from the wholesale price index for fuel and lighting, available over 1890-2020. We follow Hamilton (1996) and Hamilton (2003) and define energy shocks as the value by which energy price rises above its 3-year maximum or falls below its 3-year minimum.

Inflation expectation shocks An important feature of a successful central bank is the anchoring of inflation expectations. In this context, we aim to measure how well the Fed has been responding to innovations to inflation expectations, with the clear example being the de-anchoring of inflation expectations in the 1970s. To do so, we aim to identify inflation expectation shocks, meant to capture threats to the anchoring of inflation expectations.

As measure of inflation expectations, we rely on the Livingston survey that has been continuously run over 1946-2019.[15] Of interest for us, the Livingston survey includes a question about 8-months ahead inflation expectations. Prior to World War II, there are no

[15]The Livingston survey is conducted with a pool of professional forecasters from non-financial businesses, investment banking firms, commercial banks, academic institutions, government, and insurance companies, see (Leduc, Sill and Stark, 2007).

systematic inflation expectation survey , so we instead rely on Cecchetti (1992)'s measure of inflation expectations for the Early Fed period.[16]

To identify innovations to inflation expectations, we proceed similarly to Leduc, Sill and Stark (2007) and project inflation expectations on a set of controls that include past values of inflation expectation, inflation, unemployment, lags of the 3-month and 10-year treasury rates. In addition, we also project on current and past values of the other identified non-policy shocks: financial, government spending, energy price and TFP. The idea of this exercise is to capture movements in inflation expectations that cannot be explained by the other shocks, i.e., that go above and beyond the typical effect of the non-policy shocks on inflation expectations.

5.4 Results

We split our results into two parts. First, we discuss the ORA statistics and compare the Fed over time. Second, we zoom in on the specific sub-periods and assess the economic magnitudes of the improvements in the reaction function.

ORA-based assessments of monetary policy over 1879-2019

Table 2 shows the baseline ORA statistics computed using our baseline monetary policy shocks and baseline non-policy shocks.

A central finding is that the performance of the Fed has substantially and continuously improved over time, with about threefold improvements in the ORA between the early Fed years and the post Volcker regime. For the pre and early Fed periods we find ORA statistics above 0.5, meaning that in response to a 1 standard deviation non-policy shock, the reaction coefficients should 0.5 point larger. In contrast, the post Volcker ORAs are much smaller in magnitude.

We now turn to more specific results.

A first thing to note is that the pre Fed and early Fed periods are broadly on par: being too passive in both regimes. While this passivity is normal in the absence of a central bank, it is more surprising after the creation of the Fed. Moreover, while the weak response of the early Fed in the face of bank runs and ensuing financial stress has been argued before (e.g., Friedman and Schwartz, 1963; Hamilton, 1987), our results show that this excessive passivity extends also to government spending shocks —letting government spending shocks

[16]Cecchetti (1992)'s measure of inflation expectations relies on Mishkin (1981)'s insight that the ex-ante real interest rate can be recovered from a projection of the ex-post real interest rate on the time t information set. The difference between the ex-ante and ex-post real interest rate provides a measure of inflation expectations.

excessively affect unemployment and inflation— or to inflation expectation shocks, for instance not reacting enough to the negative inflation expectation shocks of 1931-1932—.

US monetary policy during the 1970s has generally been considered poor (e.g., Romer and Romer, 2004*a*), in particular not responding more than one-to-one with changes in inflation (Clarida, Galí and Gertler, 2000) and violating the so-called Taylor principle. Table 2 confirms and generalizes this assessment to most aggregate shocks: the reaction function is too weak in the face of increasing government spending shocks, energy prices and especially inflation expectations. This latter result extends the findings of Leduc, Sill and Stark (2007) that the largest deviations from the Taylor principle —the failure to raise the real rate in the face of rising inflation— were in terms of the response to inflation expectation shocks. Importantly however, the ORA allows to go much further by measuring not just whether the Taylor principle holds or not but also by how much the policy rule should have been adjusted in order to minimize the loss function. We will come back to these important points in the next section.

Last, the post Volcker shows improvements in monetary policy across all dimensions, with ORA statistics much smaller than in the earlier regimes, and not statistically significant. The exception is the response to financial shocks, where the deviation from optimality is significant, though much smaller than during the Great Depression. The presence of the zero lower bound does seem to have limited somewhat the Fed's ability to best react to the 2007-2008 financial crisis, echoing earlier results that fed funds rate policy was constrained by the zero lower bound over 2009-2013 (Barnichon and Mesters, 2022).

Last, as robustness check Table 3 shows the same ORA statistics computed using sign-identified monetary policy shocks. The results are remarkably consistent with our baseline estimates, with ORAs generally of similar magnitudes and same levels of statistical significance.[17]

ORA-based adjustments to monetary policy over 1879-2019

The ORAs reported in Table 2 capture how much the Fed should have adjusted the reaction coefficient its contemporaneous policy rate to a one standard-deviation non-policy shock in order to minimize the unconditional loss function (23). While the ORAs are useful to compare reaction functions across periods, the magnitudes can be hard to interpret. The impulse responses below will translate these reaction function adjustments into policy path adjustments, which are easier to interpret. Another benefit of studying the impulse responses is that it allows us to better understand the mechanisms underlying the ORA statistics.

[17] As a third alternative identification of monetary shocks, we identified monetary shocks using short run restrictions (e.g. Sims, 1980). The results were also similar, so that overall our results are remarkably similar across the three different identification approaches.

The appendix display all the ORA-adjusted impulse responses, but we will discuss the most interesting ones in the main text.

Responding to financial shocks For the early Fed period, we estimated a particularly large ORA statistic in response to financial shocks, and Figure 3 displays the corresponding impulse responses underlying this ORA estimated over 1913-1940. For comparison, Figure 4 displays the same impulse responses but estimatd over 1990-2019. In both figures, the top row shows the impulse responses of inflation, unemployment and the interest rate to a monetary policy shock. There are no surprises in terms of directions, but it is worth noting that the magnitude of the response of inflation is very different across the periods; the responses being much small in the post Volcker period, consistent with the anchoring of inflation expectations in the recent period (e.g., Gürkaynak, Levin and Swanson, 2010). The bottom rows show the responses of the same variables to a financial shock: for both periods, inflation contracts whereas unemployment increases, though the inflation response is more muted for the post-Volcker period, again consistent with the anchoring of inflation expectations.

The ORA adjusted impulse responses $\Gamma_b + \mathcal{R}_a \mathcal{T}^*_{ab}$ in the bottom rows of Figures 3 and 4 (dashed green lines) show how adjusting the reaction coefficient to financial shocks would have changed the impulse responses of inflation, unemployment and the policy rate. For the early Fed period, we find that in response to an adverse shock, the Fed *raised* the discount rate. In other words, the Fed was not only too passive, but in fact was following a contractionary policy. Combined with the decline in inflation caused by the financial shock, this means that the real policy rate increased substantially following financial shocks. This finding echoes an earlier literature on the monetary factors behind the Great Depression (e.g., Friedman and Schwartz, 1963; Hamilton, 1987). Instead, the ORA calls for lowering the discount rate. Since this lower policy rate also mutes the inflation decline, this means that the real policy rate now declines in response to a financial shock: an expansionary monetary policy. In contrast, in the post Volcker period, the ORA calls for little adjustment to the path of the fed funds rate.

Responding to energy price and inflation expectation shocks Figures 5 and 6 plot impulse responses estimated over the post WWII period. The impulse response to monetary shocks are in line with earlier evidence (e.g., Coibion, 2012). In response to an energy price shock or an inflation expectation shock, inflation rises progressively, while the policy rate response is relatively mild, especially in response to energy price shocks. To see that, we can again consider the impulse response of the real policy rate: in both cases, the real interest rate declines following an energy price shock or inflation expectation schock. In other words,

the Taylor principle is not satisfied, a finding echoing an earlier literature on the performance of the Fed during the 1970s (Clarida, Galí and Gertler, 2000; Leduc, Sill and Stark, 2007). However, the ORA goes further and allows us to compute how the policy rate should have responded to these shocks, as displayed by the dashed green lines. Most striking is the large response of the policy rate to an inflation expectation shock (lower-right panel, Figure 6), which contrasts with the absence of response before ORA adjustment. And for both types of shocks, the ORA-adjusted policy rate responses do satisfy the Taylor principle —the ORA-adjusted real interest rate increases in response to the cost-push shocks—.

Responding to government spending shocks While earlier studies of Fed performance have generally focused on the (weak) response of the Fed to financial distress in the early 1930s or to the (weak) response of the Fed to inflationary pressures during the 1970s, our approach allows to study and compare performance in response to many other shocks. Figures 7 and 8 plot the impulse responses to negative news shocks to defense spending for the early Fed period and for the post WWII period.[18] In both periods, the Fed response is too weak, and in fact contractionary: as the decline in government spending raises unemployment (and also lowers inflation), the ORA calls for lowering the policy rate, again in order lower the real rate instead of raising it.

In other words, the Fed's too passive response extends beyond a too timid response to inflationary shocks: in response to all non-policy shocks we could identify —government spending shocks, energy price shocks, inflation expecation shocks and TFP shocks—, the Fed's response was too timid in the post WWII period.

Counterfactual historical policy scenarios

With the ORA in hand, we can also create counter-factual historical policy scenarios, in which the coefficients of the policy rule are adjusted according to our median ORA estimates. Figure 9 plots the effects of the ORA adjustments on the paths of the policy rate, the inflation rate and the unemployment rate.

Specifically, given the identified non-policy shocks $\mathbf{\Xi}_{b,t}$, we can compute Δy_t, the effects of the ORA adjustments on the macro objectives, and Δp_t, the effects of the ORA adjustments on the policy rate from

$$\Delta y_t = \mathcal{R}_{y,a} * \mathcal{T}^*_{ab}\mathbf{\Xi}_{b,t} \qquad \text{and} \qquad \Delta p_t = \mathcal{R}_{p,a} * \mathcal{T}^*_{ab}\mathbf{\Xi}_{b,t} \ , \tag{25}$$

where $*$ is the convolution operator, $\mathcal{R}_{y,a}$ and $\mathcal{R}_{p,a}$ are the impulse responses of the macro

[18] In the post WWII period for instance, the shocks to government spending are mostly positive and capture two large government programs related to US space program in the early 60s and the Vietnam war in the second half of the 60s.

objectives and the policy rate to a contemporaneous policy shock, and $\Xi_{b,t}$ denotes the five types of non-policy shocks that we considered: financial, government spending, energy price, inflation expectation and TFP. In addition, Figure 10 decomposes the total adjustments to the policy rate (Δp_t) over each period by isolating the contribution of each one of the five non-policy shocks. This decomposition allows us to pinpoint why the policy rate should have been different at any particular point in time.

We stress that the magnitudes of these counterfactuals Δy_t and Δp_t cannot be interpreted as magnitudes of policy "failures" across periods, as these magnitudes are not comparable across periods. The reason is that if the economic environments are different across periods (as is most likely the case), a given Optimal Rule Adjustment can have different effects on the endogenous variables (inflation, unemployment and the policy rate): the effects can can be amplified or reduced depending on the economic environment and on the other parameters of the policy rule. In other words, while the ORAs are comparable across periods —depending *only* on how well the policy maker reacted to a specific non-policy shock—, the counterfactuals Δy_t and Δp_t are affected by other factors outside the policy maker's control.

In the pre Fed period, there were substantial deviations from an optimal reaction coefficients, calling for lower interest rates in the aftermaths of the 1893 and 1907 bank runs and for higher interest rates in response to higher military spending following the war against Spain in 1898, and a strong build-up of the navy over 1902-1904 (Figure 10, left column). That said, over the per Fed period, our identified non-policy shocks explain only a small share of the total variance of inflation and unemployment over 1879-1912, so that the ORA corrections only have a moderate effect on the behavior of inflation and unemployment over that period.

The largest corrections to the policy instrument occur in the early Fed and the Post WWII periods. In the early Fed period, there are two noteworthy adjustments. First, the ORA calls for lowering the discount rate to close to (but still above) zero during 1931-1932 instead of the roughly 2ppt level observed at the time, and also for canceling the seesaw movements in the discount rate during 1930-1932. Such movements have often been blamed for turning the initial recession caused by the 1929 stock-market crash into a full blown depression (e.g., Hamilton, 1987). With this much lower discount rate, both the 1931-1932 deflation and the dramatic rise in unemployment could have been avoided to a large extent. The cost would have been temporarily extra inflation in 1933-1934. Second, another interesting finding is that the large increase in military spending over 1918-1919 is responsible for part of the inflation outburst in 1919-1920, and the ORA-implied higher discount rate would have tamed that increase.

Turning to the post world war II period, the ORA calls for substantially higher fed funds rate throughout the 1968-1979 period, initially only 1 ppt in the late 1960s but as

much as 6 ppt higher in response to the oil price shocks of 1974 and 1979 and the inflation expectation shocks occurring during that period. With such strong response, the trend inflation of the 1970s could have been altogether avoided, though at the cost of about 2ppt higher unemployment over the 1970s.

In the post Volcker period, policy rate adjustments are generally small except in the early phase of the Great Recession with the ORA calling for an additional 1 ppt drop in the fed funds rate in 2009. However, in contrast to our Great Depression counter-factual where the ZLB would *not* have been binding, the ZLB was restricting policy during the Great Recession. Given that constraint, fed funds rate policy was optimal.

6 Conclusion

In this paper, we showed that it is possible to evaluate and compare policy makers based on the distance-to-optimality of their reaction function coefficients to well-chosen non-policy shocks. We introduced ORA statistics to measure the distance and showed that these could be computed from two sufficient statistics: (i) the impulse responses of the macro objectives to non-policy shocks, and (ii) the same impulse responses to policy shocks. Importantly, explicit knowledge of the policy maker's reaction function is not necessary, because the effect of an (unspecified) reaction function is already encoded in the impulse responses to shocks, which are estimable.

Intuitively, our approach consists in projecting the policy objectives on the space spanned by non-policy shocks and then studying the optimal policy problem in that sub-space. Thanks to the projection, it is possible to evaluate the optimality of a *specific* reaction function coefficient: how "well" the policy maker or institution reacted to a specific non-policy shock over a given period. The idea is then to compare policy makers by comparing their distance to the optimal reaction coefficient to the same non-policy shock. Since computing the distance to the optimal reaction coefficient only requires impulse response estimates, it becomes possible to evaluate and compare policy makers or policy institutions across time or even space.

While this paper studied the performance of US monetary policy over the past 150 years, the methodology could be applied to many other important questions, such as comparing the performance of different central banks over the same period, e.g., the Fed vs the ECB during the Great Recession, or comparing the performances of policy makers over time, e.g., democrats vs republicans (Blinder and Watson, 2016).

References

Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang. 2016. "Drifts and volatilities under measurement error: Assessing monetary policy shocks over the last century." *Quantitative Economics*, 7(2): 591–611.

Balke, N., and R. J. Gordon. 1986. "Appendix B Historical Data." *The American Business Cycle: Continuity and Change*. National Bureau of Economic Research.

Barnichon, Regis. 2010. "Productivity and unemployment over the business cycle." *Journal of Monetary Economics*, 57(8): 1013–1025.

Barnichon, Regis, and Geert Mesters. 2020. "Identifying Modern Macro Equations using Old Shocks." *Quarterly Journal of Economics*, 135: 22552298.

Barnichon, Regis, and Geert Mesters. 2022. "A Sufficient Statistics Approach to Macro Policy Evaluation." working paper.

Barsky, Robert B, and J Bradford De Long. 1991. "Forecasting pre-World War I inflation: the Fisher effect and the gold standard." *The Quarterly Journal of Economics*, 106(3): 815–836.

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.

Blanchard, Olivier, and Jordi Galí. 2007. "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s so Different from the 1970s?" In *International Dimensions of Monetary Policy*. 373–421. National Bureau of Economic Research, Inc.

Blinder, Alan S., and Mark W. Watson. 2016. "Presidents and the US Economy: An Econometric Exploration." *American Economic Review*, 106(4): 1015–45.

Boivin, Jean. 2005. "Has US monetary policy changed? Evidence from drifting coefficients and real-time data."

Bordo, Michael D, and Finn E Kydland. 1995. "The gold standard as a rule: An essay in exploration." *Explorations in Economic History*, 32(4): 423–464.

Canova, Fabio. 2007. *Methods for Applied Macroeconomic Research.* Princeton University Press.

Cecchetti, Stephen G. 1992. "Prices During the Great Depression: Was the Deflation of 1930-1932 Really Unanticipated?" *The American Economic Review*, 141–156.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *The Quarterly Journal of Economics*, 115(1): 147–180.

Coibion, Olivier. 2012. "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics*, 4: 1–32.

Coibion, Olivier, and Yuriy Gorodnichenko. 2011. "Monetary policy, trend inflation, and the great moderation: An alternative interpretation." *American Economic Review*, 101(1): 341–370.

Fair, Ray C. 1978. "The Use of Optimal Control Techniques to Measure Economic Performance." *International Economic Review*, 19(2): 289–309.

Friedman, Milton, and Anna Jacobson Schwartz. 1963. *A monetary history of the United States, 1867-1960.* Princeton University Press.

Gali, Jordi. 1999. "Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations?" *American economic review*, 89(1): 249–271.

Galí, Jordi. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications.* Princeton University Press.

Gali, Jordi, and Mark Gertler. 2007. "Macroeconomic Modeling for Monetary Policy Evaluation." *Journal of Economic Perspectives*, 21(4): 25–46.

Galı, Jordi, J David López-Salido, and Javier Vallés. 2003. "Technology shocks and monetary policy: assessing the Fed's performance." *Journal of Monetary Economics*, 50(4): 723–743.

Gürkaynak, Refet S, Andrew Levin, and Eric Swanson. 2010. "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the US, UK and Sweden." *Journal of the European Economic Association*, 8(6): 1208–1242.

Gürkaynak, Refet, S., Brian Sack, and Eric Swanson. 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review*, 95: 425–436.

Hamilton, James D. 1987. "Monetary factors in the great depression." *Journal of Monetary Economics*, 19(2): 145–169.

Hamilton, James D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica*, 57(2): 357–384.

Hamilton, James D. 1996. "This is what happened to the oil price-macroeconomy relationship." *Journal of Monetary Economics*, 38(2): 215–220.

Hamilton, James D. 2003. "What is an oil shock?" *Journal of Econometrics*, 113(2): 363–398.

Jordà, Oscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *The American Economic Review*, 95: 161–182.

Judd, John P, and Glenn D Rudebusch. 1998. "Taylor's Rule and the Fed: 1970-1997." *Economic Review-Federal Reserve Bank of San Francisco*, 3–16.

Kilian, Lutz, and Helmut Lütkepohl. 2017. *Structural Vector Autoregressive Analysis.* Cambridge University Press.

Kuttner, Kenneth N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics*, 47(3): 523–544.

Leduc, Sylvain, Keith Sill, and Tom Stark. 2007. "Self-fulfilling expectations and the inflation of the 1970s: Evidence from the Livingston Survey." *Journal of Monetary economics*, 54(2): 433–459.

Leeper, Eric M., and Tao Zha. 2003. "Modest policy interventions." *Journal of Monetary Economics*, 50(8): 1673–1700.

Li, Dake, Mikkel Plagborg-Mller, and Christian K. Wolf. 2022. "Local Projections vs. VARs: Lessons From Thousands of DGPs." Working paper.

McKay, Alisdair, and Christian Wolf. 2022. "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" Working Paper.

Mishkin, Frederic S. 1981. "The real interest rate: An empirical investigation." Vol. 15, 151–200, Elsevier.

Mishkin, Frederic S. 2010. "Will Monetary Policy Become More of a Science?" In *The Science and Practice of Monetary Policy Today.* , ed. Wieland V., 81–103. Springer, Berlin.

Montiel Olea, Jos Luis, and Mikkel Plagborg-Møller. 2021. "Local Projection Inference Is Simpler and More Robust Than You Think." *Econometrica*, 89(4): 1789–1823.

Nakamura, Emi, and Jón Steinsson. 2018. "Identification in Macroeconomics." *Journal of Economic Perspectives*, 32(3): 59–86.

Orphanides, Athanasios. 2003. "Historical monetary policy analysis and the Taylor rule." *Journal of monetary economics*, 50(5): 983–1022.

Petrosky-Nadeau, Nicolas, and Lu Zhang. 2021. "Unemployment crises." *Journal of Monetary Economics*, 117: 335–353.

Plagborg-Møller, Mikkel, and Christian K. Wolf. 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica*, 89(2): 955–980.

Ramey, Valerie. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics.* , ed. J. B. Taylor and H. Uhlig. Amsterdam, North Holland:Elsevier.

Ramey, Valerie A., and Sarah Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data." *Journal of Political Economy*, 126.

Reinhart, Carmen M, and Kenneth S Rogoff. 2009. "This time is different." In *This Time Is Different.* princeton university press.

Romer, Christina D, and David H Romer. 1989. "Does monetary policy matter? A new test in the spirit of Friedman and Schwartz." *NBER macroeconomics annual*, 4: 121–170.

Romer, Christina D, and David H Romer. 2004*a*. "Choosing the Federal Reserve chair: lessons from history." *Journal of Economic Perspectives*, 18(1): 129–162.

Romer, Christina D., and David H. Romer. 2004*b*. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94: 1055–1084.

Romer, Christina D., and David H. Romer. 2017. "New Evidence on the Aftermath of Financial Crises in Advanced Countries." *American Economic Review*, 107(10): 3072–3118.

Romero, Jessie. 2013. "Treasury-Fed Accord." *Federal Reserve History*.

Schularick, Moritz, and Alan M Taylor. 2012. "Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008." *American Economic Review*, 102(2): 1029–61.

Sims, Christopher A. 1980. "Macroeconomics and reality." *Econometrica*, 1–48.

Sims, Christopher A, and Tao Zha. 2006*a*. "Does monetary policy generate recessions?" *Macroeconomic Dynamics*, 10(2): 231–272.

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

Stock, James H., and Mark W. Watson. 2012. "Disentangling the Channels of the 2007-09 Recession." *Brookings Papers on Economic Activity*, 81.

Stock, James H., and Mark W. Watson. 2016. "Chapter 8 - Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics." In . Vol. 2 of *Handbook of Macroeconomics*, , ed. John B. Taylor and Harald Uhlig, 415 – 525. Elsevier.

Stock, James H., and Mark W. Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal*, 128(610): 917–948.

Svensson, Lars EO. 2003. "What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules." *Journal of Economic Literature*, 41(2): 426–477.

Taylor, John B. 1999. "A historical analysis of monetary policy rules." In *Monetary policy rules*. 319–348. University of Chicago Press.

Uhlig, Harald. 2005. "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics*, 52(2): 381–419.

Vernon, J.R. 1994. "Unemployment rates in postbellum America: 18691899." *Journal of Macroeconomics*, 16(4): 701–714.

Weir, David R. 1992. "A century of US unemployment, 1890-1990: revised estimates and evidence for stabilization." *Research in Economic History*, 14(1): 301–46.

Wheelock, David C, et al. 2010. "Lessons learned? Comparing the Federal Reserve's responses to the crises of 1929-1933 and 2007-2009." *Federal Reserve Bank of St. Louis Review*, 92(Mar): 89–108.

Appendix A: Equilibrium relationships

We briefly discuss how the general model (11)-(12) can be written as (14). Define

$$\mathcal{A} = \begin{bmatrix} \mathcal{A}_{yy} & \mathcal{A}_{yw} & \mathcal{A}_{yp} \\ \mathcal{A}_{wy} & \mathcal{A}_{ww} & \mathcal{A}_{wp} \\ \mathcal{A}_{py} & \mathcal{A}_{pw} & \mathcal{A}_{pp} \end{bmatrix} , \quad \mathcal{B} = \begin{bmatrix} \mathcal{B}_{y\xi} \\ \mathcal{B}_{w\xi} \\ \mathcal{B}_{p\xi} \end{bmatrix} , \quad \mathbf{J} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{I} \end{bmatrix} \quad \text{and} \quad \mathbf{Z}_t = \begin{bmatrix} \mathbf{Y}_t \\ \mathbf{W}_t \\ \mathbf{P}_t \end{bmatrix} . \quad (26)$$

The model (11)-(12) is equivalent to

$$\mathcal{A}\mathbb{E}_t\mathbf{Z}_t = \mathcal{B}\mathbb{E}_t\mathbf{\Xi}_t + \mathbf{J}\mathbb{E}_t\boldsymbol{\epsilon}_t \ .$$

For any $\phi \in \Phi$ we have that there exists unique equilibrium representation. This implies that $\mathcal{A}$ is invertible and we obtain

$$\mathbb{E}_t\mathbf{Z}_t = \underbrace{\mathcal{A}^{-1}\mathcal{B}}_{=\mathcal{D}_1}\mathbb{E}_t\mathbf{\Xi}_t + \underbrace{\mathcal{A}^{-1}\mathbf{J}}_{=\mathcal{D}_2}\mathbb{E}_t\boldsymbol{\epsilon}_t \ .$$

The block structure of $\mathcal{D}_1$ and $\mathcal{D}_2$ is given by

$$\mathcal{D}_1 = \begin{bmatrix} \Gamma \\ \mathcal{D}_{1w} \\ \Theta_\xi \end{bmatrix} \quad \text{and} \quad \mathcal{D}_1 = \begin{bmatrix} \mathcal{R} \\ \mathcal{D}_{2w} \\ \Theta_\epsilon \end{bmatrix} ,$$

where the maps Γ and $\mathcal{R}$ appear in the first position as they capture the effects of the shocks on $\mathbb{E}_t\mathbf{Y}_t$. The other maps capture the effects of the shocks on the endogenous variables $\mathbb{E}_t\mathbf{W}_t$ and $\mathbb{E}_t\mathbf{P}_t$.

B: Proofs

Proof of Proposition 1. Let $\mathcal{T}$ be a linear map, sufficiently small such that $\phi = \{\mathcal{A}^0_{pp}, \mathcal{A}^0_{py}, \mathcal{A}^0_{pw}, \mathcal{B}^0_{p\xi} + \mathcal{T}\} \in \Phi$. If $\phi^0 \in \Phi^{\text{opt}}$, $\mathbb{E}\mathcal{L}_t$ cannot be lowered by any $\mathcal{T} \neq 0$. Similar as in (14) we obtain the equilibrium representation

$$\mathbf{Y}_t = (\Gamma^0 + \mathcal{R}^0\mathcal{T})\mathbb{E}_t\mathbf{\Xi}_t + \mathbf{V}_t \ ,$$

where $\mathbf{V}_t = \mathcal{R}^0\mathbb{E}_t\boldsymbol{\epsilon}_t + \mathbf{Y}_t - \mathbb{E}_t\mathbf{Y}_t$ and note that $\mathbb{E}[\mathbb{E}_t\mathbf{\Xi}_t\mathbf{V}_t'] = 0$. The expected loss $\mathbb{E}\mathcal{L}_t$ becomes

$$\begin{aligned} \mathbb{E}\mathcal{L}_t &= \frac{1}{2}\mathbb{E}\left([\Gamma^0 + \mathcal{R}^0\mathcal{T}]\mathbb{E}_t\mathbf{\Xi}_t + \mathbf{V}_t\right)' \mathcal{W}\left([\Gamma^0 + \mathcal{R}^0\mathcal{T}]\mathbb{E}_t\mathbf{\Xi}_t + \mathbf{V}_t\right) \\ &= \frac{1}{2}\text{Tr}\left\{[\Gamma^0 + \mathcal{R}^0\mathcal{T}]'\mathcal{W}[\Gamma^0 + \mathcal{R}^0\mathcal{T}]\mathbf{\Sigma}_\Xi\right\} + \frac{1}{2}\mathbb{E}\left(\mathbf{V}_t'\mathcal{W}\mathbf{V}_t\right) \end{aligned}$$

The derivative of the map $\mathcal{T} \to \mathbb{E}\mathcal{L}_t$ is given by

$$\mathcal{R}^{0\prime}\mathcal{W}(\Gamma^0 + \mathcal{R}^0\mathcal{T})\mathbf{\Sigma}_\Xi \ . \quad (27)$$

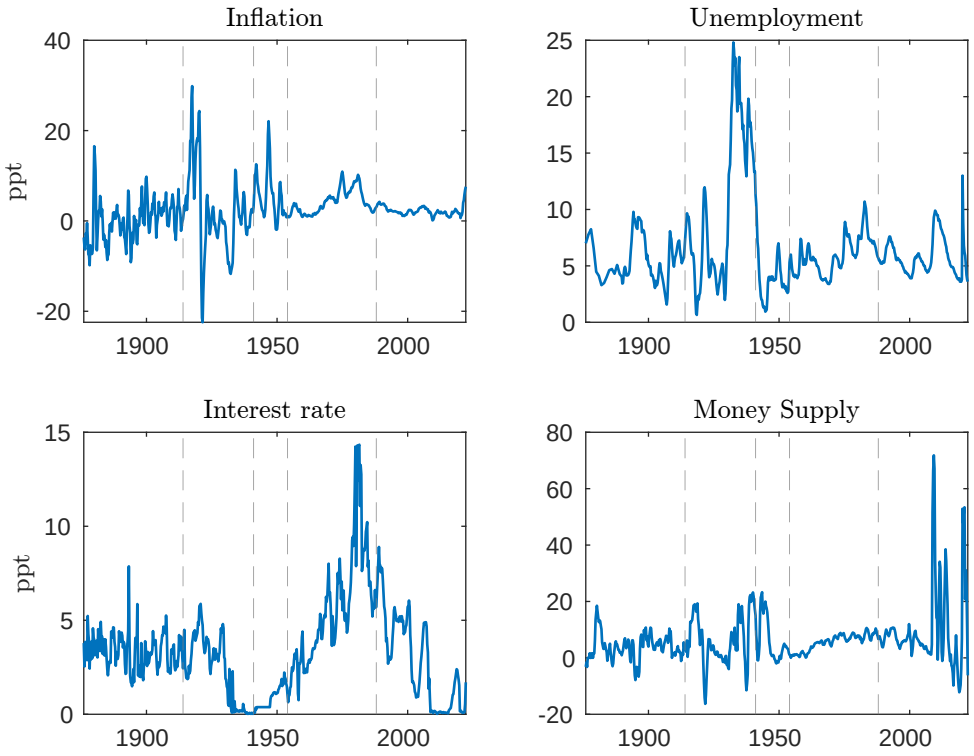
Evaluating at $\mathcal{T} = 0$ and setting the derivative to zero implies

$$\mathcal{R}^{0'}\mathcal{W}\Gamma^0 = 0 \ ,$$

is a necessary condition for optimality. Noting that $\mathcal{T} \to \mathbb{E}\mathcal{L}_t$ is a convex map, it follows that $\mathcal{R}^0\mathcal{W}\Gamma^0 = 0$ is also sufficient for $\mathcal{T} = 0$ being a global minimizer, and thus $\phi^0 \in \Phi^{\text{opt}}$. □

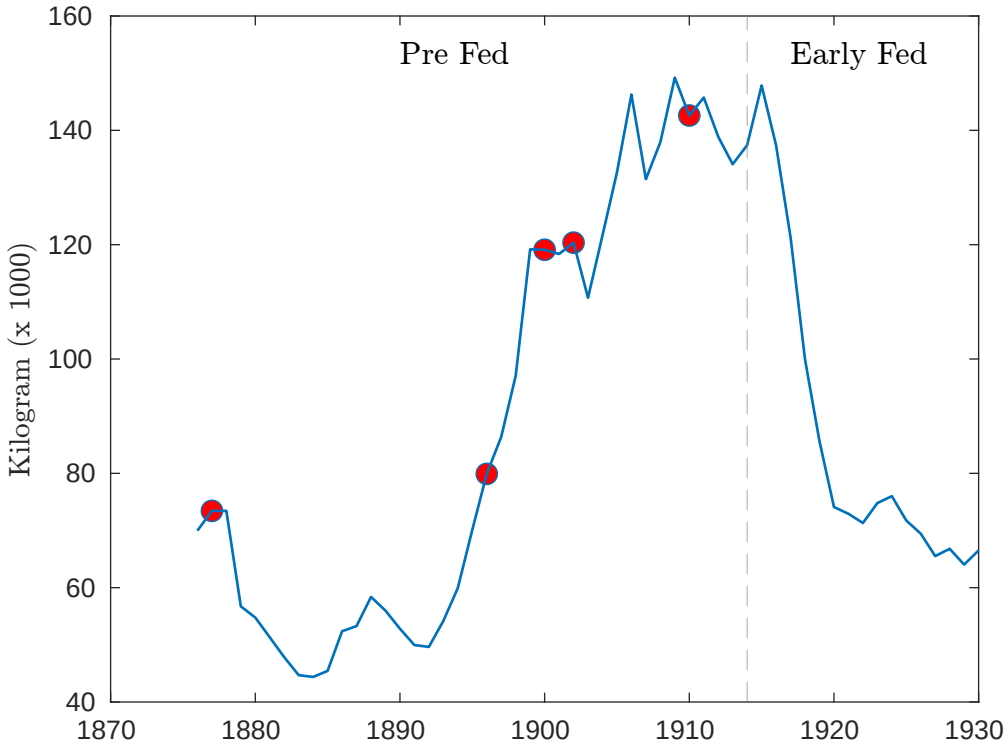
Proof of Corollary 1. The left hand side of (**??**) is equivalent to the subset of (27) corresponding to a, b and after post-multiplying (**??**) by $\mathbf{\Sigma}_\Xi^{-1}$. The right hand side follows by direct calculation. The second part follows directly by noting $\mathcal{T}_{a,b}^*$ sets the local gradient to zero, hence it lowers the loss convex loss function. □

Figure 1: Policy Objectives and Instruments, 1876–2020



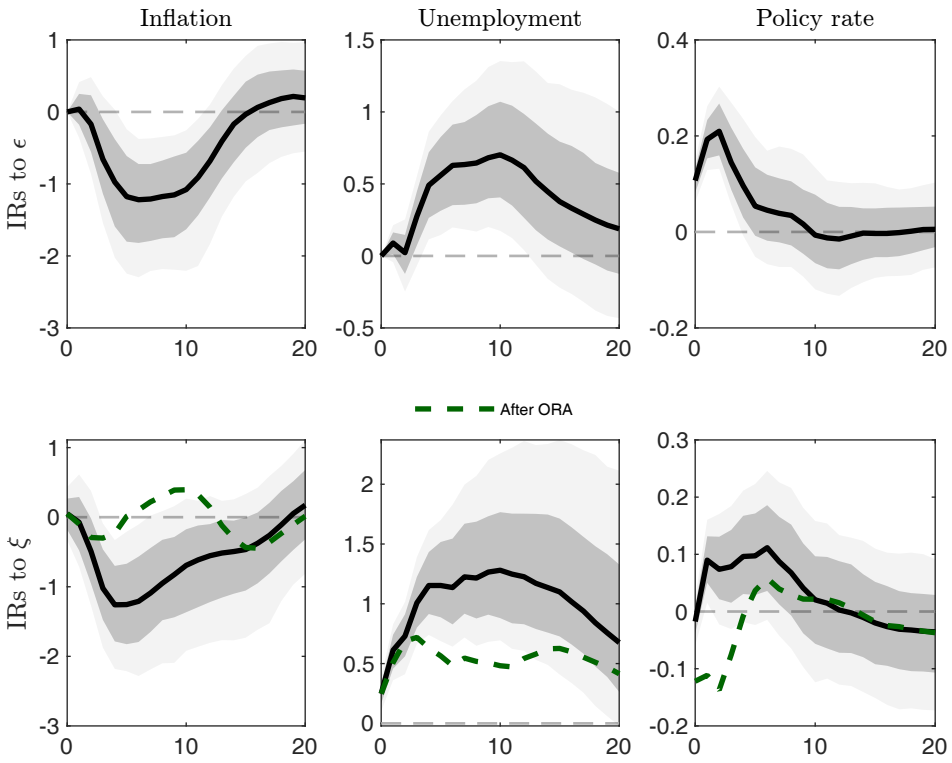
Notes: The top panel shows year-on-year inflation based on the output deflator and the unemployment rate extended backwards following Vernon (1994). The bottom panel shows the 3-month interest rate and the changes in the monetary base. The vertical lines highlight the different periods: pre-Fed 1879-1912, early-Fed 1913-1941, post-WWII 1951-1984 and post-Volcker 1991-2020.

Figure 2: US Gold Production, 1876–1930



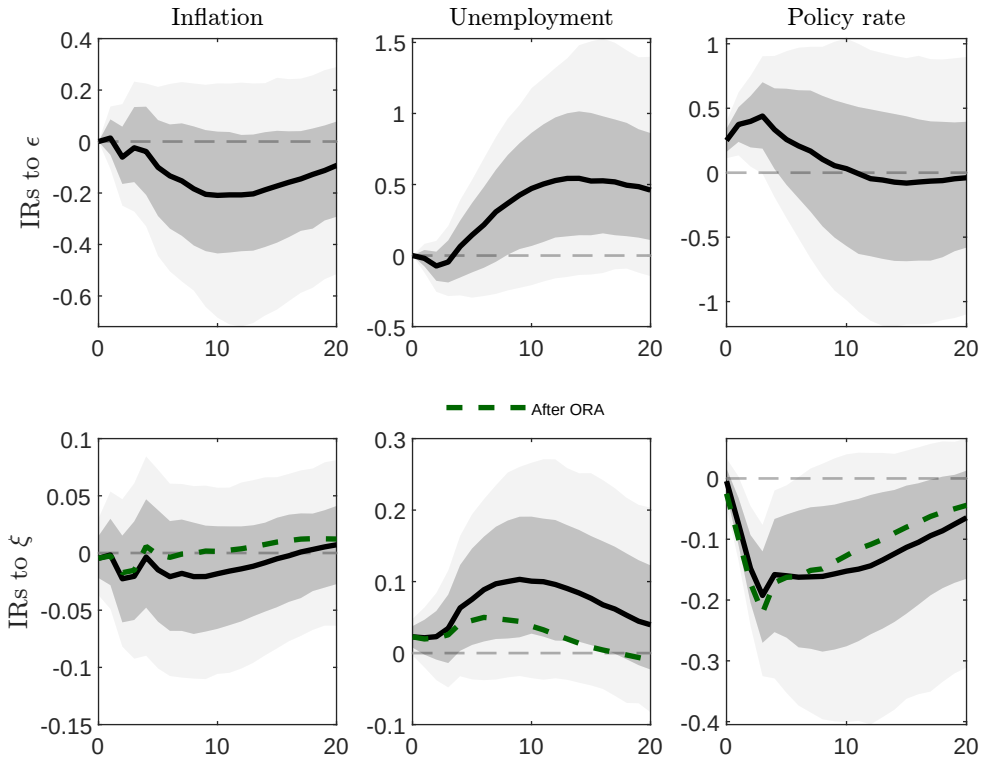
Notes: We show US gold production in kilograms (x1000). The red dots correspond to the peak in the Comstock lode mine in Nevada (1877-Q1), the discovery of Gold in Alaska (1896-Q3), the discovery of Gold mines in Nevada in (1902-Q1), the Nevada gold mine maximum (1910-Q1).

Figure 3: Early Fed, 1913-1941, Reaction to Financial shocks



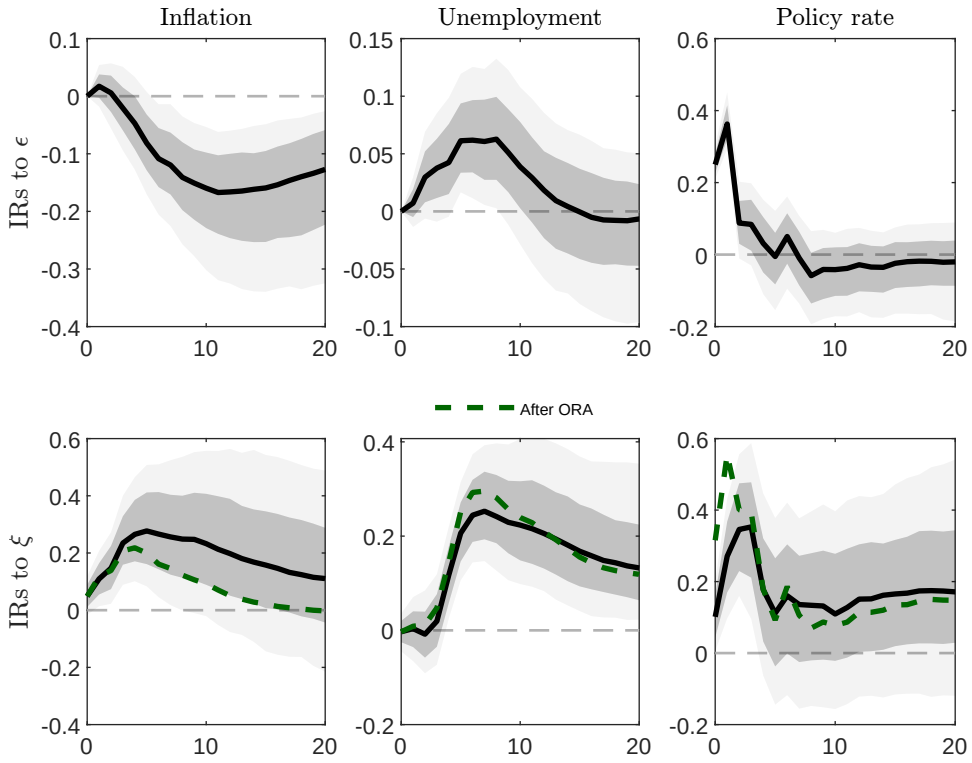
Notes: The top (resp.) row shows the median responses (thick line) of inflation, unemployment and the fed funds rate to a monetary policy shock (resp. financial shock). The dotted green lines show the ORA adjusted impulse responses $\Gamma_b^0 + \mathcal{R}_a^0 \mathcal{T}_{ab}^*$. The 95% and 67% credible sets are plotted as dark and light shaded areas, respectively.

Figure 4: Post Volcker Fed, 1990-2019, Reaction to Financial shocks



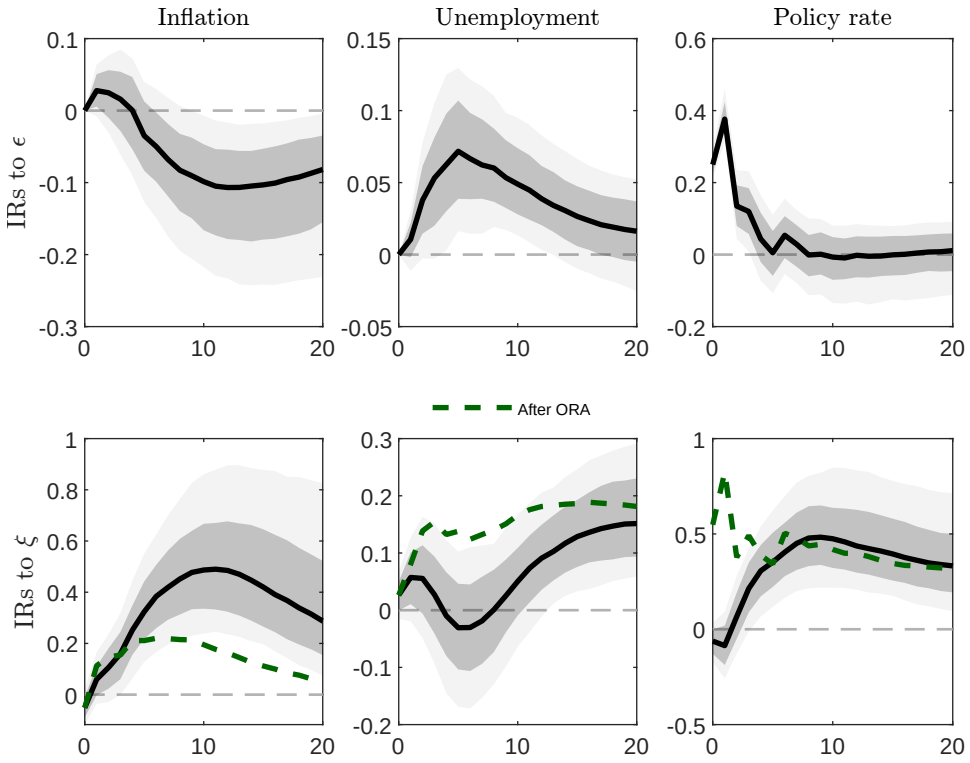
Notes: The top (resp.) row shows the median responses (thick line) of inflation, unemployment and the fed funds rate to a monetary policy shock (resp. financial shock). The dotted green lines show the ORA adjusted impulse responses $\Gamma_b^0 + \mathcal{R}_a^0 \mathcal{T}_{ab}^*$. The 95% and 67% credible sets are plotted as dark and light shaded areas, respectively.

Figure 5: Post WWII Fed, 1951-1985, Reaction to Energy shocks



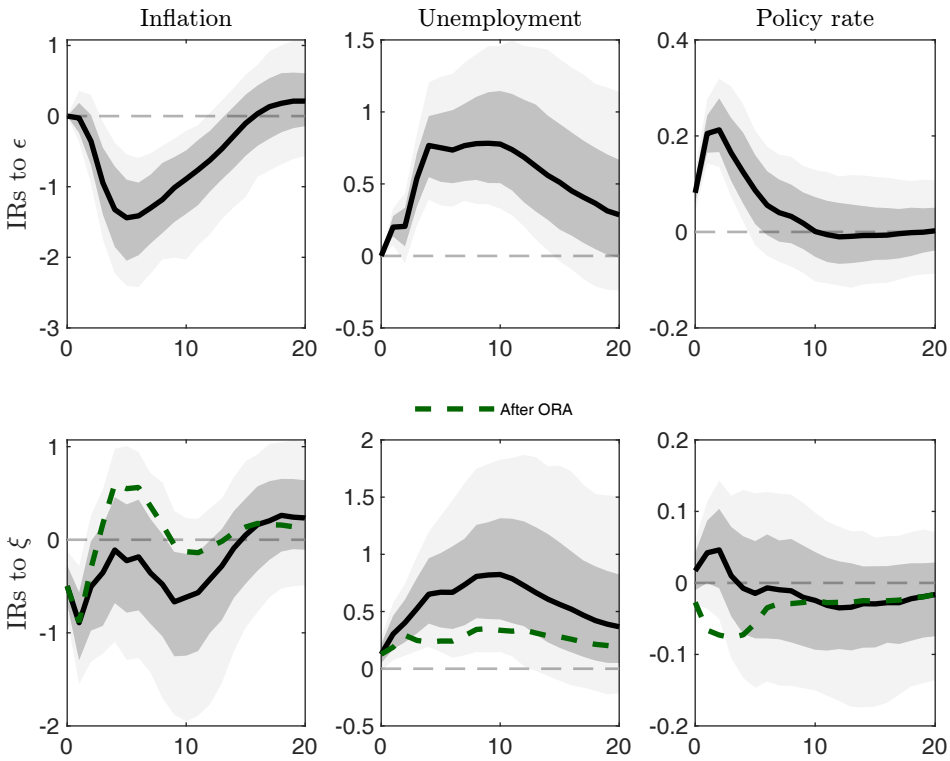
Notes: The top (resp. bottom) row shows the median responses (thick line) of inflation, unemployment and the fed funds rate to a monetary policy shock (resp. energy price shock). The dotted green lines show the ORA adjusted impulse responses $\Gamma_b^0 + \mathcal{R}_a^0 \mathcal{T}_{ab}^*$. The 95% and 67% credible sets are plotted as dark and light shaded areas, respectively.

Figure 6: Post WWII Fed, 1951-1985, Reaction to π^e shocks



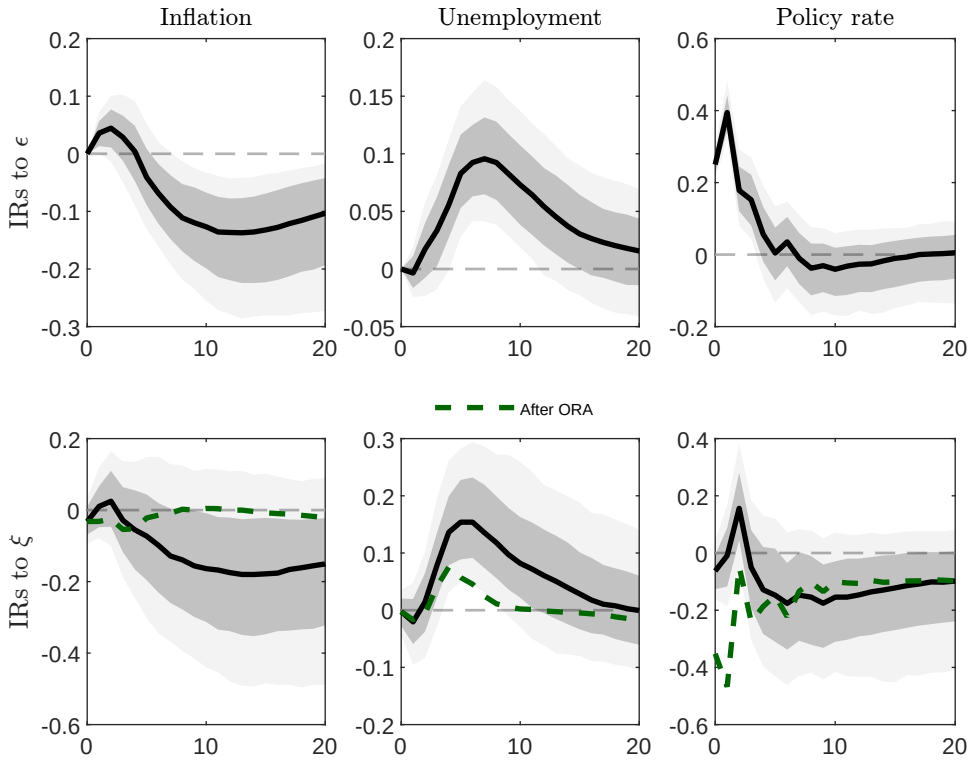
Notes: The top (resp. bottom) row shows the median responses (thick line) of inflation, unemployment and the fed funds rate to a monetary policy shock (resp. inflation expectations shock price shock). The dotted green lines show the ORA adjusted impulse responses $\Gamma_b^0 + \mathcal{R}_a^0 \mathcal{T}_{ab}^*$. The 95% and 67% credible sets are plotted as dark and light shaded areas, respectively.

Figure 7: Early Fed, 1913-1941, Reaction to G shocks



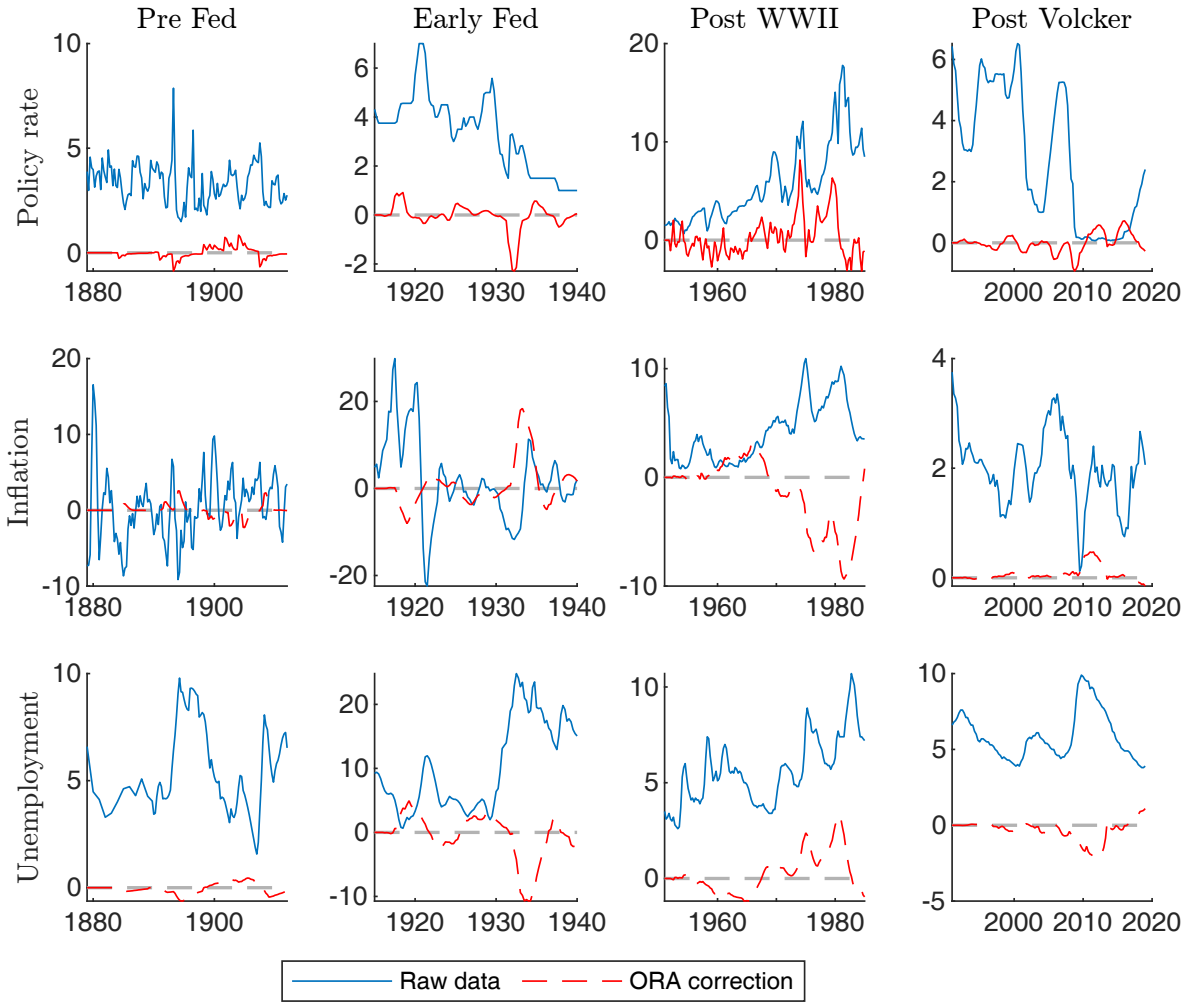
Notes: The top (resp. bottom) row shows the median responses (thick line) of inflation, unemployment and the fed funds rate to a monetary policy shock (resp. government spending shock). The dotted green lines show the ORA adjusted impulse responses $\Gamma_b^0 + \mathcal{R}_a^0 \mathcal{T}_{ab}^*$. The 95% and 67% credible sets are plotted as dark and light shaded areas, respectively.

Figure 8: Post WWII Fed, 1951-1985, Reaction to G shocks



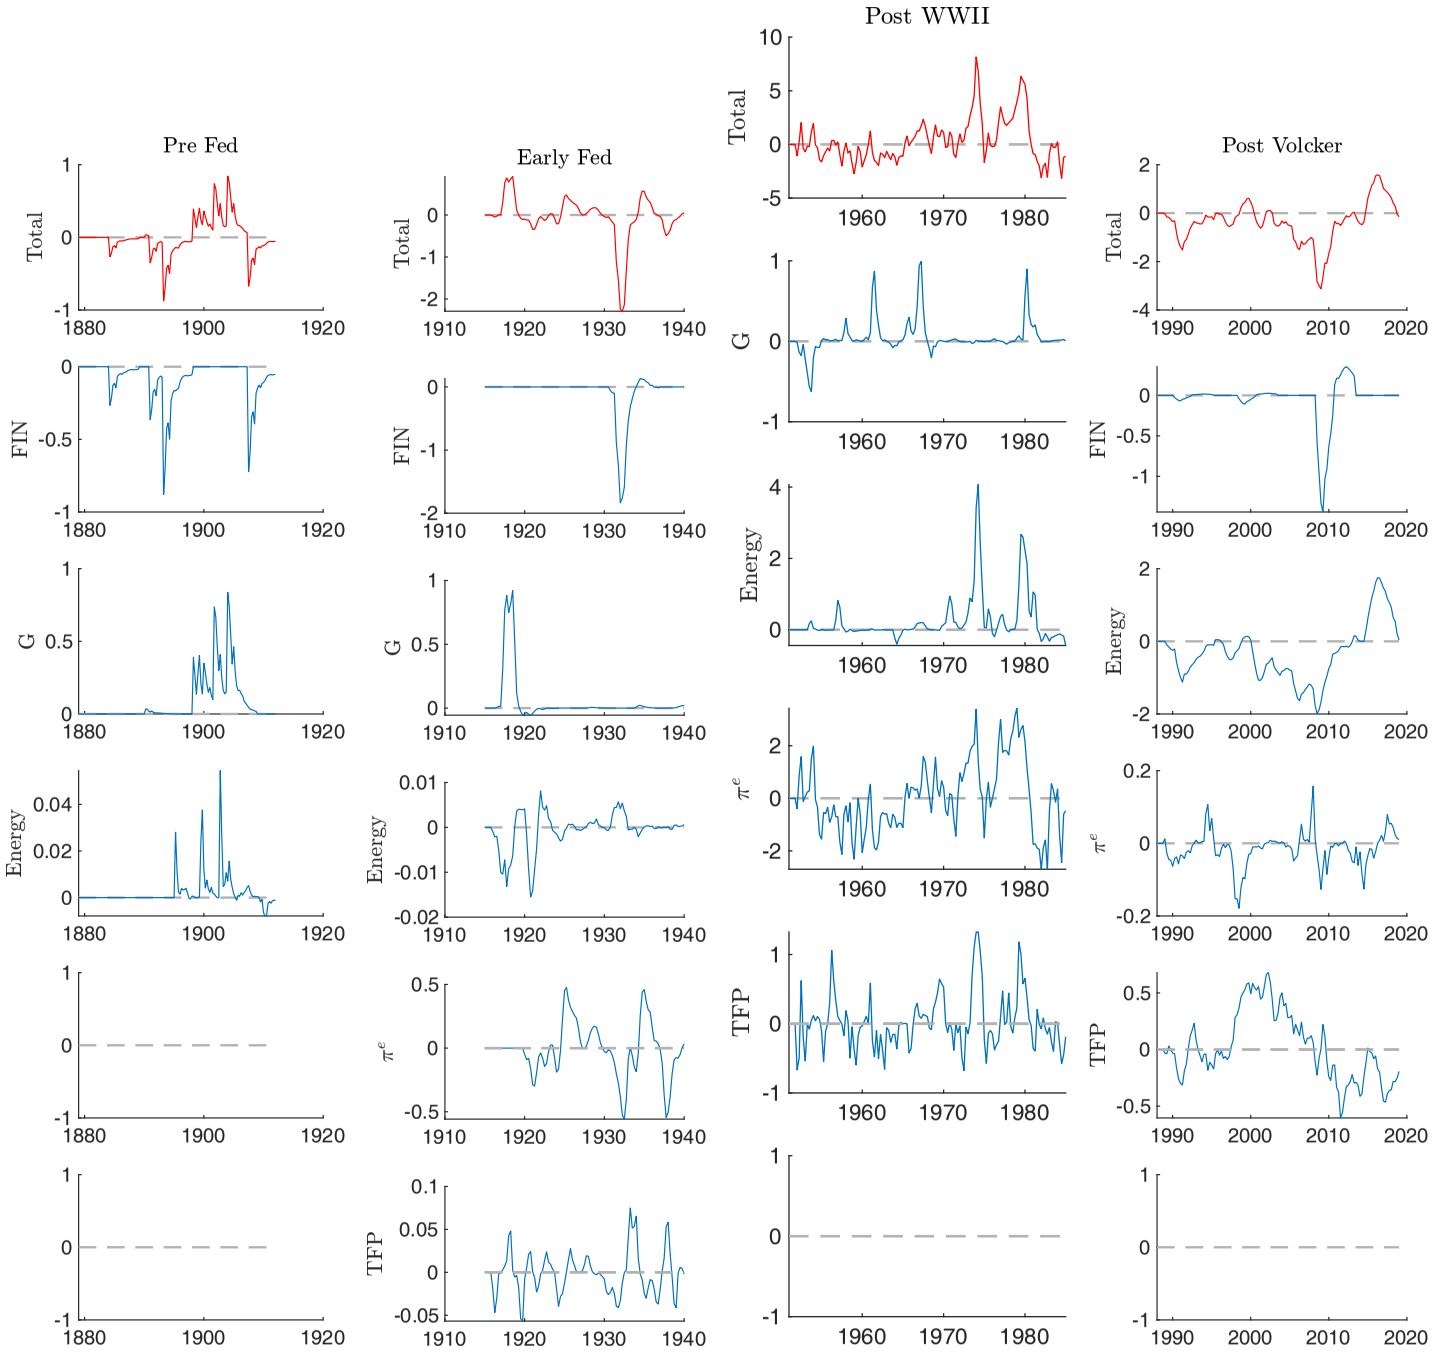
Notes: The top (resp. bottom) row shows the median responses (thick line) of inflation, unemployment and the fed funds rate to a monetary policy shock (resp. government spending shock). The dotted green lines show the ORA adjusted impulse responses $\Gamma_b^0 + \mathcal{R}_a^0 \mathcal{T}_{ab}^*$. The 95% and 67% credible sets are plotted as dark and light shaded areas, respectively.

Figure 9: ORA corrections over 1879-2019



Notes: The top row shows the policy rate ("raw data", blue plain line) along with the adjustment to the contemporaneous policy rate implied by the median ORA correction ("ORA correction", dashed red line) over each period, calculated following (25). The middle and bottom rows show the same information but for the inflation rate and the unemployment rate.

Figure 10: ORA DECOMPOSITION



Notes: Each column decomposes median ORA adjustment to the policy rate in the contribution of each non-policy shock: financial (FIN), government spending (G), Energy price (Energy), inflation expectations (π^e) and TFP for the four periods: Pre Fed, EArly Fed, Post WWII and Post Volcker. The total ORA corrections to the contemporaneous policy rate (and given by (25)) are depicted in the top row.

Table 1: Realized Outcomes

	Pre Fed 1879-1912	Early Fed 1913-1941	Post WWII 1951-1984	post-Volcker 1991-2019
$\overline{\pi}$	0.4	1.9	4.3	2.0
$\overline{u}$	5.3	10.2	5.6	5.9
$\text{Var}(\pi)$	19.4	90.1	7.7	0.5
$\text{Var}(u)$	3.5	48.6	3.2	2.6

Notes: $\overline{\pi}$ and $\overline{u}$ denote the sample means of inflation and unemployment, and $\text{Var}(\pi)$ and $\text{Var}(u)$ denote the sample variances of inflation and unemployment, as computed over the different periods.

Table 2: ORA statistics for US monetary policy

Non-policy shock Shock sign convention	Bank panics $u \uparrow$	BAA-AAA $u \uparrow$	G $u \uparrow$	Energy $\pi \uparrow$	π^e $\pi \uparrow$	TFP $\pi \uparrow$
Pre Fed	**−0.6*** (−1.2,−0.1)	—	**−0.7*** (−1.3,−0.2)	**−0.3*** (−0.6,0)	—	**0.7*** (0.3,1)
Early Fed	**−1.0*** (−1.7,−0.5)	**−1.0*** (−1.5,−0.6)	**−0.5*** (−1,−0.2)	**−0.1** (−0.4,0.3)	**0.5*** (0.2,0.8)	**0.2** (−0.2,0.6)
Post WWII	—	**−0.7*** (−1.3,−0.2)	**−0.6*** (−1.2,−0.2)	**0.8*** (0.2,1.4)	**1.3*** (0.7,2)	**0.4*** (−0.0,1)
Post Volcker	**−0.3** (−0.8,0.1)	**−0.2*** (−0.5,0)	**0.3** (−0.3,0.9)	**−0.5** (−1.4,0.7)	**0.2** (−0.1,0.6)	**−0.3** (−0.6,0.1)

Notes: Median ORA statistics together with 68% credible sets. The monetary policy shocks are identified as described in main text: using gold rush discovery in the pre-Fed period, Romer and Romer (1989)'s Friedman-Schwartz dates in the early Fed period, Romer and Romer (2004) monetary shocks for the post WWII period and high-frequency surprises in the post Volcker period. The financial shocks are bank panics or innovations to the BAA-AAA spread, the government spending shocks are from (Ramey and Zubairy, 2018, G), TFP shocks from (Gali, 1999, TFP), energy shocks are computed using the peak-over-threshold approach of Hamilton (1996), and inflation expectation shocks (π^e) are innovations to inflation expectations as measured from Cecchetti (1992) for Early Fed period and from the Livingston survey (Post WWII and Post Volcker periods). For the Pre Fed period the TFP, G and Energy ORAs are computed over the 1890-1913 period.

Table 3: ORA statistics for US monetary policy, Robustness

Non-policy shock Shock sign convention	Bank panics $u \uparrow$	BAA-AAA $u \uparrow$	G $u \uparrow$	Energy $\pi \uparrow$	π^e $\pi \uparrow$	TFP $\pi \uparrow$
Pre Fed	**−0.4*** (−0.6,−0.1)	—	**−0.4*** (−0.7,−0.2)	**0.0** (−0.3,0.3)	—	**0.3*** (0,0.6)
Early Fed	**−0.8*** (−1.1,−0.3)	**−0.9*** (−1.1,−0.6)	**−0.4*** (−0.6,0)	**0.1** (−0.3,0.3)	**0.5*** (0.2,0.8)	**0** (−0.4,0.3)
Post WWII	—	**−0.5*** (−0.9,−0.2)	**−0.1** (−0.4,0.3)	**0.6*** (0,1)	**0.7*** (0.1,1)	**0.4** (−0.2,0.9)
Post Volcker	**0.0** (−0.4,0.5)	**−0.3** (−0.7,0.1)	**0.3** (−0.5,0.8)	**−0.3** (−1,0.8)	**0.2** (−0.1,0.5)	**0** (−0.3,0.4)

Notes: Median ORA statistics together with 68% credible sets. The monetary policy shocks are identified using sign restrictions as described in main text. The financial shocks are bank panics or innovations to the BAA-AAA spread, the government spending shocks are from (Ramey and Zubairy, 2018, G), TFP shocks from (Gali, 1999, TFP), energy shocks are computed using the peak-over-threshold approach of Hamilton (1996), and inflation expectation shocks (π^e) are innovations to inflation expectations as measured from Cecchetti (1992) for Early Fed period and from the Livingston survey (Post WWII and Post Volcker periods). For the Pre Fed period the TFP, G and Energy ORAs are computed over the 1890-1913 period.