FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

# A Sufficient Statistics Approach for Macro Policy Evaluation

Regis Barnichon
Federal Reserve Bank of San Francisco
CEPR

Geert Mesters
Universitat Pompeu Fabra
BSE

April 2022

Working Paper 2022-15

https://www.frbsf.org/economic-research/publications/working-papers/2022/15/

**Suggested citation:**

Barnichon, Regis, Geert Mesters. 2022 "A Sufficient Statistics Approach for Macro Policy Evaluation," Federal Reserve Bank of San Francisco Working Paper 2022-15. https://doi.org/10.24148/wp2022-15

# A Sufficient Statistics Approach for Macro Policy Evaluation*

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

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

[(b)] Universitat Pompeu Fabra and BSE

April 27, 2022

**Abstract**

The evaluation of macroeconomic policy decisions has traditionally relied on the formulation of a specific economic model. In this work, we show that two statistics are sufficient to detect, often even correct, non-optimal policies, i.e., policies that do not minimize the loss function. The two sufficient statistics are (i) the effects of policy shocks on the policy objectives, and (ii) forecasts for the policy objectives conditional on the policy decision. Both statistics can be estimated without relying on a specific model. We illustrate the method by studying US monetary policy decisions.

*JEL classification*: C14, C32, E32, E52.

*Keywords*: Optimal policy, Impulse responses, Forecasting.

*First draft: March 2020. We thank Adrien Auclert, Davide Debortoli, Martin Eichenbaum, Stefano Eusepi, Jordi Galí, Alain Guay, Ethan Ilzetzki, Sylvain Leduc, Huiyu Li, Pascal Michaillat, Emi Nakamura, Robert Ulbricht, Nicolas Petrosky-Nadeau, Adam Shapiro, Jón Steinsson and seminar participants at Bank of England, Berkeley, Boston College, Brown, CREI, McGill, 2020 NBER Summer Institute (ME), Ohio State, Ottawa, UQAM, UT Austin, Cleveland Fed, St Louis Fed and San Francisco Fed for helpful comments. We graciously thank Refet Gürkaynak for sharing his data on high-frequency monetary surprises. Mesters acknowledge support from the Spanish Ministry of Economy and Competitiveness through the Ramon y Cajal fellowship (RYC2019-028287-I), the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), and the Netherlands Organization for Scientific Research (NWO) through the VENI research grant (016.Veni.195.036). 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

Despite impressive recent progress in structural macro modeling, policy makers often resort to heuristics to decide on policy; combining insights from different models and relying heavily on judgment calls and instincts.[1] This practical approach has benefits in terms of robustness to model mis-specification, but a major downside is that it can be difficult to identify the most appropriate course of policy. Without a specific economic model, how can a policy maker be confident that a policy decision is appropriate? For instance, how to determine that the magnitude and timing of a fiscal package is well calibrated, or that the monetary stance is appropriate, e.g., in a "Goldilocks zone" that best balances inflation and unemployment.

In this paper, we show that it is not necessary to know the full structure of the economy to evaluate a macro policy decision. Instead, only two statistics are sufficient to detect, and often correct, non-optimal policies, i.e., policies that do not minimize the loss function.

Our approach is based on the gradient of the loss function with respect to policy shocks. Although little studied in the literature, this gradient has two key attractive properties: (i) it is informative about the optimality of a policy decision, and (ii) it is relatively easy to compute, depending only on two well-known statistics.

First, for a large class of models —linear models, models with state dependent effects of policy, and even models with multiple policy regimes—, the gradient with respect to policy shocks must be zero under an optimal policy. From this necessary condition it follows that this gradient is sufficient to *evaluate* a policy decision, i.e., to detect a non-optimal policy. Moreover, for a smaller yet still very large class of models —linear models, models with state dependent effects of policy, but not models with multiple policy regimes—, setting the gradient to zero yields the optimal policy, such that the gradient is necessary and sufficient to characterize the optimal allocation. In other words, for that class of models, the gradient is also sufficient to *compute* the optimal policy.

Second, for a large class of loss functions the gradient with respect to the policy shocks is entirely determined by two simple statistics (i) forecasts for the policy objectives conditional on the policy decision, and (ii) the effects of policy shocks on the policy objectives. These two statistics are already central and well understood concepts for policy makers (e.g., Orphanides, 2019). Our contribution is to show that these two statistics alone can be used to rigorously evaluate and even set policy.

Importantly, the two sufficient statistics can be estimated without relying on a specific structural model. First, a large forecasting literature has shown how one can construct superior forecasts by combining large and disparate information sources, multiple (imperfect) models and possibly judgment (e.g. Stock and Watson, 2002*b*; Lawrence et al., 2006; Man-

[1] See e.g., Svensson (2003), Mishkin (2010) and Blinder (2020).

ganelli, 2009; Geweke and Amisano, 2012; Giacomini and Ragusa, 2014; Cheng and Hansen, 2015). In fact, policy makers already rely on that literature to construct such conditional forecasts as part of their decision making procedure. Second, a large macro-econometric literature has shown how to estimate the effects of policy shocks with minimal modeling assumptions, notably by using zero-, long-run, or inequality restrictions (e.g. Sims, 1980; Blanchard and Quah, 1989; Faust, 1998; Uhlig, 2005), or by using past exogenous variations as instrumental variables (e.g. Mertens and Ravn, 2013; Stock and Watson, 2018), see Ramey (2016) for a detailed review.

To formally implement our policy evaluation framework, we do not work directly with the gradient but instead with its rescaled version: the *Optimal Policy Perturbation* (OPP). Like the gradient, the OPP is entirely determined by the two sufficient statistics, but it also has a direct economic interpretation. The OPP is the adjustment to the policy instruments that exactly corrects an optimization failure when the loss function is quadratic and the (unspecified) underlying model is linear.[2]

Uncertainty in the estimates of the sufficient statistics has two notable effects. First, our evaluation of a policy choice will resemble a hypothesis test: a statement that the policy is not optimal at some confidence level. Increasing uncertainty will not invalidate our approach, but it will reduce the power of the test. Second, with uncertainty in the causal effect of policy, the OPP estimate faces an attenuation bias that echoes the Brainard (1967) conservatism principle familiar to policy makers: faced with uncertainty in the policy multipliers, policy makers should refrain from fully utilizing their instrument, and the OPP estimate will suggest a smaller policy adjustment than would be the case under certainty.

To illustrate our sufficient statistics approach to macro policy evaluation, we study US monetary policy decisions. We start from the Fed's dual inflation–unemployment mandate, and we estimate/recover the sufficient statistics underlying the OPP. We estimate causal effects using high-frequency monetary surprises as instrumental variables (e.g. Eberly, Stock and Wright, 2019), and we use as conditional forecasts the FOMC Survey of Economic Projections —the policy makers' own forecasts—.

While the contemporaneous fed funds rate has not been set exactly at its optimal level since 1990, the optimal adjustment (in absolute value) is overall small, averaging only 25 basis points over the full sample. There are however some noteworthy instances of non-optimal policies. Most notably, during the Great Recession the zero-lower bound (ZLB) kept the contemporaneous policy rate too high by about 1 percentage point. Moreover, given the information available in real time, we conclude that the Fed should have lowered the fed funds rate faster in early 2008, when the ZLB was not yet binding. Finally, we find that a

[2]This includes linearized New-Keynesian (NK) models, notably Heterogeneous Agents NK models (e.g., Auclert et al., 2021).

more active manipulation of the expected policy path could have improved the conduct of policy. During the Great Recession, the OPP calls for a more active use of forward guidance to lower the slope of the expected path of the fed funds rate, a conclusion echoing that of Eberly, Stock and Wright (2019).

The remainder of this paper is organized as follows. We continue the introduction by relating the OPP approach to existing approaches in the literature. In the next section we provide a simple example that informally explains how we can evaluate macro policy using sufficient statistics. Section 3 formally introduces the general environment and Section 4 presents the OPP statistic and its properties. Section 5 discusses the econometric details for computing the OPP statistic and its distribution. In Section 6 we apply our methodology to empirically study monetary policy decisions in the US. Section 7 extends the properties of the OPP for non-linear models, and section 8 concludes and provides potential avenues for further research.

### Relation to literature

In the wake of the Lucas (1976) critique, the macroeconomic literature has built increasingly elaborate micro-founded models in order to study, inter alia, the design of optimal policy rules; rules that specify once-and-for-all how the policy instruments should be set at all dates and states.[3] In this paper, we show that it is often not necessary to know the full structure of the economy to evaluate a policy decision and even compute the optimal policy.

Our sufficient statistics approach to macro policy evaluation naturally shares important similarities with the sufficient statistic approach that originated in public finance (e.g. Chetty, 2009). Both methods exploit the fact that the welfare consequences of a policy can be derived from high-level elasticities, allowing for policy evaluation without making parametric assumptions or estimating the structural primitives of fully specified models. One feature specific to our macro focus is that we postulate a loss function at the macro level, consistent with the fact the loss function is often determined by political factors or by statutory requirement. For instance, it is the US Congress that mandates the Federal Reserve to seek stable inflation and full employment. That said, our approach can equally be applied to problems with micro-founded loss functions.

Our treatment of uncertainty around the OPP shares similarities with the robust-control approach. In particular, OPP inference can be seen as developing a robust framework for handling parameter uncertainty and model mis-specification, similarly to the approach followed in the context of structural models, see Hansen and Sargent (2001), Onatski and Stock (2002), Onatski and Williams (2003) and Hansen and Sargent (2008), among others.

[3]See e.g., Chari, Christiano and Kehoe (1994); Woodford (2010); Michaillat and Saez (2019).

From a general macro perspective, our paper uncovers an important but so far overlooked link between the structural impulse response literature (e.g. Ramey, 2016) and the optimal policy literature. While impulse response estimates have primarily been used as a guide for model building (e.g., Ramey, 2016), our paper provides a novel and important role for impulse response estimates: as a testbed for the optimality of policy.[4] In fact, our paper formalizes how two core econometric tasks —economic forecasting and dynamic causal effects estimation— are crucial ingredients for macro policy making, as they deal exactly with the estimation of the two sufficient statistics underlying our approach to policy evaluation. Progress in estimation precision and in forecasting performance will directly improve the ability to detect and improve policy decisions with minimal structural assumptions on the underlying model.

Finally, our sufficient statistics approach to policy evaluation can be seen as a key input in the context of forecast-targeting (e.g., Svensson, 2019). Forecast targeting is a general approach to policy making that consists in selecting a policy path so that "the forecasts of the target variables look good, meaning appears to best fulfill the mandates and return to their target at an appropriate pace" (Svensson, 1999). However, unless the forecast targeting rule is tied to a specific model (Woodford, 2010; Giannoni and Woodford, 2017), a "looking good" criterion is imprecise and leaves the policy maker uncertain about the optimality of the policy choice. Our sufficient statistics approach to policy evaluation precisely fills this gap.

# 2 A simple example

Before formally describing our general framework, we first present a simple example to illustrate how two key statistics are sufficient to evaluate and improve macro policies. The example is based on Galí (2015, Section 5.1.1), which discusses the optimal policy problem under discretion in the baseline New Keynesian model.[5]

Consider a central bank with loss function

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

[4]In recent work subsequent to ours, McKay and Wolf (2022) add to this insight that structural impulse responses can also be used to construct policy rule counter-factuals provided that the model coefficients do not change with the parameters of the policy rules. An important contribution of their work is the realization that our approach also holds in modern macro models, most notably HANK models.

[5]In the web-appendix we show that the an alternative example can be formulated for the optimal policy problem under commitment in the same baseline New Keynesian model (e.g. Galí, 2015, Section 5.1.2). We focus on the discretionary case as it simplifies the exposition, but our sufficient statistics approach applies independently of the nature of the optimization problem, i.e., whether under discretion or under commitment.

with $\pi_t$ the inflation gap and $x_t$ the output gap. The central bank has only one instrument: the nominal interest rate $i_t$.

The log-linearized baseline New-Keynesian model is defined by a Phillips curve and an intertemporal (IS) curve 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 $\xi_t$ an iid cost-push shock.

### The optimal targeting rule

The optimal allocation can be characterized by minimizing the loss function with respect to $\pi_t$, $x_t$ and $i_t$ subject to the Phillips curve and (IS) curve constraints. This gives the well-known optimal targeting rule

$$x_t = -\kappa \pi_t \ . \tag{4}$$

Since the optimization problem is convex, the optimal targeting rule is necessary and sufficient to characterize the optimal policy which we denote by $i_t^{\text{opt}}$.

A limitation of this approach to characterize the optimal policy is that it requires the full underlying model, that is the exact specification and coefficients of the Phillips and (IS) curves. As we discussed in the introduction, this information requirement is unlikely to be met for policy makers in practice.

### An alternative characterization of the optimal targeting rule

We will now see that there is an alternative approach to characterize the optimal targeting rule and the optimal policy; an approach that does not require knowing the details of the model.

Consider a simple instrument policy rule augmented with a policy shock $\epsilon_t$, i.e.,

$$i_t = \phi \pi_t + \epsilon_t \ , \tag{5}$$

with $\phi > 1$ to guarantee a unique equilibrium. As shown in e.g., Galí (2015), the optimal allocation can be implemented with that rule by setting $\phi^{\text{opt}} = \kappa\sigma$ and $\epsilon_t = 0$.[6]

[6] We posit $\kappa\sigma > 1$ to ensure a unique equilibrium. Our argument does not rely on this parameter restriction, and we only make it for clarity of exposition. Provided that the functional form of the rule is unrestricted, there exists an optimal instrument rule that can deliver a unique equilibrium; for instance, a rule of the form $i_t = \phi_\pi \pi_t + \phi_\xi \xi_t$ (Galí, 2015, page 133).

Under the assumption that the rule (5) leads to a unique equilibrium we can combine (2), (3) and (5) to solve the model, i.e., express the endogenous variable $Y_t = (\pi_t, x_t)'$ as a function of the structural shocks. We have

$$Y_t = \mathcal{R}\epsilon_t + \mathcal{C}\xi_t \tag{6}$$

with

$$\mathcal{R} = \begin{bmatrix} \frac{-\kappa/\sigma}{1+\kappa\phi/\sigma} \\ \frac{-1/\sigma}{1+\kappa\phi/\sigma} \end{bmatrix} \qquad \text{and} \qquad \mathcal{C} = \begin{bmatrix} \frac{1}{1+\kappa\phi/\sigma} \\ \frac{-\phi/\sigma}{1+\kappa\phi/\sigma} \end{bmatrix} ,$$

where $\mathcal{R}$ captures the effect of a policy shock on the policy objectives.

We can now establish our main result: the optimal targeting rule can be derived by setting the gradient of the loss function with respect to policy shocks to zero. Formally, we have the equivalence:

$$i_t = i_t^{\text{opt}} \qquad \Longleftrightarrow \qquad \left.\frac{\partial \mathcal{L}_t}{\partial \epsilon_t}\right|_{i_t} = \mathcal{R}'Y_t = 0 \ , \tag{7}$$

where $\mathcal{R}'Y_t = 0$ *is* the optimal targeting rule.

To prove the result, use (6) to compute the effect of a policy shock on the gradient of the loss function:

$$\begin{aligned} \left.\frac{\partial \mathcal{L}_t}{\partial \epsilon_t}\right|_{i_t} &= \mathcal{R}'Y_t \\ &= \frac{1}{\sigma + \kappa\phi}\,(-\kappa, -1)\,(\pi_t, x_t)' \\ &= \frac{-1}{\sigma + \kappa\phi}\,(\kappa\pi_t + x_t) \ . \end{aligned} \tag{8}$$

The term in parenthesis is zero under the optimal targeting rule (4): $x_t = -\kappa\pi_t$, which establishes the equivalence.

The equivalence (7) states that setting the gradient of the loss function with respect to a policy shock to zero is necessary and sufficient to characterize the optimal policy and optimal targeting rule. In other words, a policy is optimal if and only if the policy maker has no incentive to deviate from it with a policy surprise, i.e., a shock.

Intuitively, for a linear system the optimal allocation can be attained by a linear policy rule. As a result, at the optimal policy, there should not exist *any* deviation from that linear rule —including surprise deviations— that can improve the allocation: the gradient of the loss function with respect to policy shocks must be zero. Moreover, since the problem is convex, that gradient condition is also sufficient to characterize the optimal solution.

The equivalence relation has two specific implications: First, the $\Longrightarrow$ relation implies that

$\mathcal{R}'Y_t = 0$ forms a testable condition to evaluate whether a given policy decision is optimal. Second, the $\Longleftarrow$ relation can be exploited to correct a non-optimal policy using *only* $\mathcal{R}$ and $Y_t$.

### Policy evaluation with sufficient statistics

We first illustrate how the two statistics $\mathcal{R}$ and $Y_t$ can be used to evaluate a policy decision.

Consider a policy maker following the rule (5) and proposing the policy $i_t^0$ given by the pair $(\phi^0, \epsilon_t^0)$. The policy can be non-optimal for two reasons: $\phi^0 \neq \phi^{\text{opt}}$ or $\epsilon_t^0 \neq 0$. For $i_t^0$ to be optimal, we just saw that the gradient of the loss function evaluated at $i_t^0$ must be zero, i.e., that

$$\left.\frac{\partial \mathcal{L}_t}{\partial \epsilon_t}\right|_{i_t^0} = \mathcal{R}^{0\prime} Y_t^0 = 0 \ , \tag{9}$$

where $\mathcal{R}^0$ is the effect of policy shocks under $\phi^0$, and $Y_t^0$ is the allocation under $i_t^0$. Equation (9) forms a testable condition to evaluate policy decisions: if $\mathcal{R}^{0\prime} Y_t^0 \neq 0$, we can conclude that $i_t^0$ is not optimal.

A key benefit of this approach to policy evaluation is that it is possible to construct estimates of $\mathcal{R}^0$ and $Y_t^0$, *even* if the full model structure is unknown. For that purpose, we can rely on recent advances in the econometrics literature. First, a large macro-econometric literature has focused on estimating causal effects $\mathcal{R}^0$ with minimal modeling assumptions, notably using past exogenous variations in policy as instrumental variables (e.g. Ramey, 2016; Stock and Watson, 2018). Second, $Y_t^0$, the allocation at $i_t^0$, is simply a conditional oracle forecast, i.e., the expectation of $Y_t^0$ conditional on the policy $i_t^0$.[7] To estimate $Y_t^0$, a large forecasting literature has shown how one can construct superior forecasts by combining multiple (imperfect) models, shrinkage methods, judgment and large and disparate information sources (e.g. Stock and Watson, 2002*b*; Lawrence et al., 2006; Manganelli, 2009; Geweke and Amisano, 2012; Giacomini and Ragusa, 2014; Cheng and Hansen, 2015). In fact, policy makers heavily rely on that literature to construct conditional forecasts as part of their decision making procedure.

With both $\mathcal{R}^0$ and $Y_t^0$ known with uncertainty, the gradient can only be computed with uncertainty, and our evaluation of the optimality of a policy choice will resemble a hypothesis test: a statement that the policy is not optimal for some confidence level.

### Policy improvement with sufficient statistics

Should the gradient be non-zero, we can go further and use $\mathcal{R}^0$ and $Y_t^0$ to find the magnitude of the adjustment to $i_t^0$ that restores optimality, that is we can combine $\mathcal{R}^0$, $Y_t^0$ and $i_t^0$ to

[7] In this simple static example, we slightly abuse the term "forecasting" in anticipation of our general treatment where dynamics will figure prominently.

compute the optimal policy $i_t^{\text{opt}}$.

Let $\delta_t$ denote a fixed deterministic adjustment to the policy choice $i_t^0$, that is consider changing the policy to $i_t^0 + \delta_t$. Proceeding similarly to our derivation of (6) and solving the model, the effect of the adjustment $\delta_t$ on the allocation is given by

$$\begin{aligned} Y_t &= \mathcal{R}^0(\epsilon_t^0 + \delta_t) + \mathcal{C}^0 \xi_t \\ &= Y_t^0 + \mathcal{R}^0 \delta_t \ . \end{aligned}$$

Thus the effect of a policy adjustment on $Y_t$ is the same irrespective of the nature of the adjustment, i.e., whether it is random (a shock) or deterministic: in both cases, the effect is given by $\mathcal{R}^0$.

To compute the optimal policy, the idea is to find the policy adjustment $\delta_t^*$ that ensures that $\mathcal{R}' Y_t = 0$ holds, as the $\Longleftarrow$ relation in (7) would then imply that $i_t^0 + \delta_t^*$ *is* the optimal policy. Specifically, we have

$$\mathcal{R}^{0\prime} Y_t = \mathcal{R}^{0\prime}(Y_t^0 + \mathcal{R}^0 \delta_t^*) = 0 \qquad \Longleftrightarrow \qquad \delta_t^* = -\left(\mathcal{R}^{0\prime}\mathcal{R}^0\right)^{-1} \mathcal{R}^{0\prime} Y_t^0 \ . \tag{10}$$

The statistic $\delta_t^*$ is what we call the Optimal Policy Perturbation (OPP). Clearly, the OPP has the same property as the gradient — $\delta_t^* = 0$ implies $i_t^0 = i_t^{\text{opt}}$—. In addition, starting from the initial policy choice $i_t^0$, we just saw that $i_t^0 + \delta_t^* = i_t^{\text{opt}}$: the OPP $\delta_t^*$ allows us to *compute* the optimal policy with sufficient statistics alone.

Importantly, we emphasize that $\delta_t^*$ is *not a shock* and does not introduce any *new* exogenous variation into the policy decision. To see that, let us express the policies $i_t^{\text{opt}}$ and $i_t^0$ as functions of the structural shocks. Proceeding as with the derivation of (6), we have

$$i_t^{\text{opt}} = \Theta_\xi^{\text{opt}} \xi_t \qquad \text{and} \qquad i_t^0 = \Theta_\xi^0 \xi_t + \Theta_\epsilon^0 \epsilon_t^0 \tag{11}$$

with

$$\Theta_\epsilon = \frac{1}{1 + \kappa\phi/\sigma} \quad \text{and} \quad \Theta_\xi = \frac{\phi}{1 + \kappa\phi/\sigma} \ , \qquad \text{for } \phi = \left\{\phi^0, \phi^{\text{opt}}\right\} \tag{12}$$

where $\Theta_\xi^{\text{opt}}$ is the effect of the cost-push shock under the optimal rule $\phi^{\text{opt}}$ and $\Theta_\xi^0$ is the effect under the rule $\phi^0$.

The OPP $\delta_t^* = i_t^{\text{opt}} - i_t^0$ can then be written as

$$\delta_t^* = (\Theta_\xi^{\text{opt}} - \Theta_\xi^0)\xi_t - \Theta_\epsilon \epsilon_t^0 \ . \tag{13}$$

This shows that the OPP corrects the two possibles sources of optimization failure: (i) an exogenous policy mistake ($\epsilon_t^0 \neq 0$), or (ii) a non-optimal policy rule ($\Theta_\xi^0 \neq \Theta_\xi^{\text{opt}}$), and we can see the OPP adjustment as consisting in (i) removing the effect of the policy mistake $\epsilon_t^0$, and (ii) changing the policy rule from $\phi^0$ to $\phi^{\text{opt}}$.[8]

In sum, this simple example shows that it is possible to assess and improve a policy proposal without having complete knowledge of the underlying structure of the economy. The underlying property is the equivalence (7): in a linear model, setting the gradient of the loss function with respect to the policy shock to zero is necessary and sufficient to characterize the optimal allocation. The next sections generalize this insight for a general dynamic macro environment that includes the majority of macro models encountered in the literature but without committing to a particular one.

# 3 General framework and objectives

In this section, we describe the policy problem, the generic economic environment and we formalize our objectives.

## 3.1 The policy problem

Consider a policy maker at time $t$ who aims to stabilize the expected path of the policy objectives $\mathbf{Y}_t = (y_t', y_{t+1}', \ldots)'$ where $y_t = (y_{1,t}, \ldots, y_{M_y,t})'$ is a vector of stationary variables. The policy maker can form expectations about the future paths of $\mathbf{Y}_t$, based on the time $t$ information set $\mathcal{F}_t$. We denote the expectation operator by $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_t)$.

The objective of the policy maker is to minimize the expected loss function

$$\mathcal{L}_t = \frac{1}{2}\mathbb{E}_t \mathbf{Y}_t' \mathcal{W} \mathbf{Y}_t \ , \tag{14}$$

where $\mathcal{W} = \text{diag}(\beta \otimes \lambda)$ denotes a diagonal matrix of preferences with $\lambda = (\lambda_1, \ldots, \lambda_{M_y})'$ capturing the weights on the different variables and $\beta = (\beta_0, \beta_1, \ldots)'$ the discount factors for the different horizons. While we consider a quadratic loss function in the baseline treatment, the web-appendix shows that our approach can be easily modified to accommodate any convex loss function.

To minimize the loss function the policy maker can set $M_p$ policy instruments at time $t$, denoted by $p_t = (p_{1,t}, \ldots, p_{M_p,t})'$. In addition, the policy maker can set the time-$t$ expected values for $p_{t+1}$, $p_{t+2}, \ldots$, etc... We denote by $\mathbf{P}_t^e = \mathbb{E}_t(p_t', p_{t+1}', \ldots)'$ the corresponding expected future policy path as a function of the time-$t$ information set.

[8] We have $i_t + \delta_t^* = \Theta_\xi^{\text{opt}} \xi_t = \phi^{\text{opt}} \pi_t^{\text{opt}}$ from (6), (12) and using $\phi^{\text{opt}} = \kappa\sigma$.

## 3.2 Environment

We consider a linear environment which can be justified for small fluctuations around a steady-state, but we will later consider more general treatments that relax this assumption. A generic model for the non-policy block of the economy at time $t$ is given by

$$\begin{cases} \mathcal{A}_{yy}\mathbb{E}_t\mathbf{Y}_t - A_{yw}\mathbb{E}_t\mathbf{W}_t - \mathcal{A}_{yp}\mathbf{P}^e_t &= \mathcal{B}_{yx}\mathbf{X}_{-t} + \mathcal{B}_{y\xi}\mathbb{E}_t\Xi_t \\ \mathcal{A}_{ww}\mathbb{E}_t\mathbf{W}_t - \mathcal{A}_{wy}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{wp}\mathbf{P}^e_t &= \mathcal{B}_{wx}\mathbf{X}_{-t} + \mathcal{B}_{w\xi}\mathbb{E}_t\Xi_t \end{cases} , \tag{15}$$

where $\mathbf{W}_t = (w'_t, w'_{t+1}, \ldots)'$ is a path for additional endogenous variables, the vector $\mathbf{X}_{-t} = (y'_{t-1}, w'_{t-1}, p'_{t-1}, y'_{t-2}, \ldots)'$ captures the initial conditions as summarized by the history of the variables $y_t, w_t$ and $p_t$, and $\Xi_t = (\xi'_t, \xi'_{t+1}, \cdots)'$ denotes the path of the structural shocks $\xi_t$. The linear maps $\mathcal{A}_{..}$ and $\mathcal{B}_{..}$ are conformable. After taking expectations we can interpret $\mathbb{E}_t\Xi_t$ as shocks (including news shocks) to the fundamentals of the economy that are released at time $t$ (e.g. Chahrour and Jurado, 2018).

This model is general and it accommodates a large class models found in the literature, not only standard New-Keynesian (NK) models (e.g., Smets and Wouters, 2007), but also modern heterogeneous agents NK models (Auclert et al., 2021). Numerous specific examples can be found in Woodford (2003) and Walsh (2017).

## 3.3 Optimal policy

The optimal policy can be characterized by considering a planner who chooses the paths $\mathbf{Y}_t, \mathbf{W}_t$ and $\mathbf{P}_t$ in order to minimize the loss function, i.e.,

$$\min_{\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t} \mathcal{L}_t \qquad \text{s.t.} \qquad (15) \ . \tag{16}$$

Denote by $\mathbf{P}^{e\,\text{opt}}_t$ the corresponding optimal policy. Note that the problem defines the entire optimal policy path as a function of the information available at time $t$.

## 3.4 Objectives

Our aim is to evaluate policy decisions, that is to detect and possibly correct policy choices that deviate from the optimal policy $\mathbf{P}^{e\,\text{opt}}_t$.

Without loss of generality, a policy decision can be written as the sum of two terms, a component determined in response to the state of the economy captured by all time-$t$ measurable variables —the policy rule— and an exogenous component. Specifically, we write a generic model for the policy block with

$$\mathcal{A}_{pp}\mathbf{P}_t^e - \mathcal{A}_{py}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{pw}\mathbb{E}_t\mathbf{W}_t = \mathcal{B}_{p\xi}\mathbb{E}_t\mathbf{\Xi}_t + \mathcal{B}_{px}\mathbf{X}_{-t} + \boldsymbol{\epsilon}_t^e \ , \tag{17}$$

where $\boldsymbol{\epsilon}_t^e = \mathbb{E}_t\boldsymbol{\epsilon}_t$ are shocks to the expected policy paths, where $\boldsymbol{\epsilon}_t = (\epsilon_t', \epsilon_{t+1}', \ldots)$ denotes shocks to the policy path with $\epsilon_t = (\epsilon_{1,t}, \ldots, \epsilon_{M_p,t}')'$. In other words, by taking the expected value of $\boldsymbol{\epsilon}_t$, we transform these policy shocks into "policy news shocks" revealed at time $t$, i.e. time-$t$ shocks to $\mathbf{P}_t^e$. These policy news shocks $\boldsymbol{\epsilon}_t^e$ are assumed to be uncorrelated with the initial conditions and all other structural shocks.

We collect all the elements of the policy rule in $\phi = \{\mathcal{A}_{pp}, \mathcal{A}_{py}, \mathcal{A}_{pw}, \mathcal{B}_{px}, \mathcal{B}_{p\xi}\}$. A policy choice $\mathbf{P}_t^e$ is then determined by a pair $(\phi, \boldsymbol{\epsilon}_t^e)$.

Consider a policy maker who proposes the expected policy path $\mathbf{P}_t^{e0}$ which is determined by the pair $(\phi^0, \boldsymbol{\epsilon}_t^{e0})$.[9] Our objective is to evaluate whether the path $\mathbf{P}_t^{e0}$ corresponds to the optimal path $\mathbf{P}_t^{e\text{opt}}$ as defined in (16). Formally, we are interested in testing

$$H_0 : \mathbf{P}_t^{e0} = \mathbf{P}_t^{e\text{opt}} \qquad \text{vs} \qquad H_1 : \mathbf{P}_t^{e0} \neq \mathbf{P}_t^{e\text{opt}} \ . \tag{18}$$

In addition, if $\mathbf{P}_t^{e0}$ is not optimal, we would like to improve the policy choice such that it becomes closer to $\mathbf{P}_t^{e\text{opt}}$.

# 4 Policy evaluation with sufficient statistics

In this section we show how, for the generic class of model captured by (15), we can test $H_0$ using only sufficient statistics. To do so we make two assumptions:

**Assumption 1.** *The optimal policy* $\mathbf{P}_t^{e\text{opt}}$ *is unique.*

**Assumption 2.** *The rule* $\phi^0$ *underlying the proposed policy path* $\mathbf{P}_t^{e0}$ *leads to a unique and determinate equilibrium.*

The first assumption simplifies the exposition, but is not essential. In fact, our main results continue to hold when replacing $\mathbf{P}_t^{e\text{opt}}$ with a set of optimal policies for which each element of the set solves (16)

The second assumption imposes that the proposed policy rule $\phi^0$ leads to a unique and determinate equilibrium. This is a necessary condition for the existence of the gradient with respect to policy shocks and thus the OPP statistic. Note that explicit knowledge of $\phi^0$

[9]We do not take a stance on the specific formulation that led to $\mathbf{P}_t^{e0}$, clearly multiple reaction functions (i.e., multiple $\phi^0$) can lead to the same path $\mathbf{P}_t^{e0}$.

is not required. Given that our underlying policy rule (17) is unrestricted, we view this assumptions as very mild.

Under Assumptions 1 and 2 we have the following characterization of the optimal policy.

**Proposition 1.** *Given the generic model* (15)-(17) *and Assumptions 1-2, we have the equivalence*

$$\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\,\mathrm{opt}} \qquad \Longleftrightarrow \qquad \nabla_{\boldsymbol{\epsilon}_t} \mathcal{L}_t|_{\mathbf{P}_t^{e0}} = \mathcal{R}^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0 = 0 \tag{19}$$

*where* $\mathbb{E}_t \mathbf{Y}_t^0$ *is the allocation under* $\mathbf{P}_t^{e0}$ *and* $\mathcal{R}^0$ *captures the causal effects of policy shocks under the rule* $\phi^0$ *with*

$$\mathbf{Y}_t^0 = \mathcal{R}^0 \boldsymbol{\epsilon}_t^{e0} + \boldsymbol{\Upsilon}_t^0 \ , \qquad \mathbb{E}(\boldsymbol{\epsilon}_t^{e0} \boldsymbol{\Upsilon}_t^{0'}) = 0 \ . \tag{20}$$

All proofs are stated in the appendix.

The proposition shows that the equality $\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\,\mathrm{opt}}$ holds if and only if the gradient of the loss function with respect to policy shocks is equal to zero. The term $\boldsymbol{\Upsilon}_t^0$ is a linear combination of the structural shocks, initial conditions and future errors $\mathbf{Y}_t^0 - \mathbb{E}_t \mathbf{Y}_t^0$. The exact expression is shown in the appendix, but for our purposes it is sufficient to note that $\boldsymbol{\Upsilon}_t^0$ is orthogonal to the policy shocks which allows to identify $\mathcal{R}^0$ from (20).

The proposition generalizes our equivalence (7) for the simple NK model. The same mechanism is at play —relying on the linearity of the underlying model—, and we do not repeat the intuition.

## 4.1 The OPP statistic

Building on Proposition 1, we can now establish how two statistics —$\mathcal{R}^0$ and $\mathbb{E}_t \mathbf{Y}_t^0$—are sufficient to detect, even correct, a non-optimal policy decision $\mathbf{P}_t^{e0}$, i.e., a decision where $\mathbf{P}_t^{e0} \neq \mathbf{P}_t^{e\,\mathrm{opt}}$. To formally implement this, we will not work directly with the gradient of Proposition 1 but instead with its rescaled version: the *Optimal Policy Perturbation* (OPP). Like the gradient, the OPP is entirely determined by the two statistics, $\mathcal{R}^0$ and $\mathbb{E}_t \mathbf{Y}_t^0$ but it also has a direct economic interpretation.

We define the *optimal policy perturbation* (OPP) statistic as the gradient of the loss function at $\mathbf{P}_t^{e0}$ rescaled by the inverse Hessian or

$$\begin{aligned} \boldsymbol{\delta}_t^* &= (-\nabla^2_{\boldsymbol{\epsilon}_t} \mathcal{L}_t|_{\mathbf{P}_t^{e0}})^{-1} \nabla_{\boldsymbol{\epsilon}_t} \mathcal{L}_t|_{\mathbf{P}_t^{e0}} \\ &= -(\mathcal{R}^{0'} \mathcal{W} \mathcal{R}^0)^{-1} \mathcal{R}^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0 \ . \end{aligned} \tag{21}$$

The OPP statistic is analog to the version introduced for the simple example, see equation (10). The only differences are the dimensions and the weighting of the policy objectives. It is clear that the OPP depends only on the two statistics $\mathcal{R}^0$ and $\mathbb{E}_t \mathbf{Y}_t^0$. The OPP has the following properties

**Proposition 2.** *Given the generic model* (15)-(17) *and Assumptions 1-2, we have that*

1. $$\boldsymbol{\delta}_t^* = 0 \qquad \Longleftrightarrow \qquad \mathbf{P}_t^{e0} = \mathbf{P}_t^{e\text{opt}}$$

2. $$\mathbf{P}_t^{e0} + \boldsymbol{\delta}_t^* = \mathbf{P}_t^{e\text{opt}} \quad .$$

The formal proof is in the Appendix, but the properties of the OPP essentially derive from the equivalence results stated in Proposition 1: to characterize the optimal policy, it is necessary and sufficient to ensure that the gradient with respect to shocks is zero. The same intuition presented in Section 2 is at work, and we do not repeat it here. Instead, we propose three perspectives —an optimization perspective, an econometrics perspective, and an economics perspective— on how the OPP (instead of the gradient alone) can allow to *compute* the optimal policy starting from some initial (non-optimal) policy.

**Intuition: an optimization perspective** First, notice that the OPP is the first-step of a gradient descent algorithm, specifically Newton's method. Since the gradient depends only on the two sufficient statistics $\mathcal{R}^0$ and $\mathbb{E}_t \mathbf{Y}_t^0$, we can use gradient-descent to improve the policy choice. When the problem is linear-quadratic as in generic model (15)-(17), Newton's algorithm can be used to exactly get to the optimum in one step.[10] The OPP is precisely that first-step, starting from $\mathbf{P}_t^{e0}$.

**Intuition: an econometrics perspective** Relatedly, the OPP formula looks like the formula of a weighted least squares regression: $\boldsymbol{\delta}_t^*$ is minus the coefficient estimate of a regression of $\mathbb{E}_t \mathbf{Y}_t^0$ on $\mathcal{R}$, weighted by $\mathcal{W}$. To see why, consider adjusting the policy choice from $\mathbf{P}_t^{e0}$ to $\mathbf{P}_t^{e1} \equiv \mathbf{P}_t^{e0} + \boldsymbol{\delta}_t^*$ : an OPP policy adjustment. Such adjustment can also be viewed as an adjustment to $\boldsymbol{\epsilon}_t^{e0}$. This implies that we can use Proposition 1 to compute the effect of this policy adjustment on the policy objectives and get $\mathbb{E}_t \mathbf{Y}_t^1 = \mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}^0 \boldsymbol{\delta}_t^*$ with $\mathbb{E}_t \mathbf{Y}_t^1$ the allocation under the new policy $\mathbf{P}_t^{e1}$. Rewriting, this gives

$$\mathbb{E}_t \mathbf{Y}_t^0 = -\mathcal{R}^0 \boldsymbol{\delta}_t^* + \mathbb{E}_t \mathbf{Y}_t^1 \ . \tag{22}$$

The goal of the OPP adjustment $\boldsymbol{\delta}_t^*$ can then be seen as follows: use $\mathcal{R}^0$ —the causal effects of an OPP adjustment— in order to minimize the (weighted) sum-of-squares of $\mathbb{E}_t \mathbf{Y}_t^1$, the

[10]Newton's algorithm is a search algorithm designed to solve optimization problems, and it works by approximating the optimization problem with its linear-quadratic approximation (which has a closed-form solution).

expected paths for the policy objectives after the OPP adjustment. This is nothing but a regression of $\mathbb{E}_t \mathbf{Y}_t^0$ on $-\mathcal{R}^0$.[11]

**Intuition: an economics perspective** From an economics perspective, we will now see that the OPP adjustment corrects a non-optimal policy by changing the policy maker's reaction function (and making it optimal).

The mechanism is similar to that of the New-Keynesian model described in Section 2 and in equation (13). To see that more formally, first assume that there exists a rule that underlies $\mathbf{P}_t^{e\text{opt}}$, say $\phi^{\text{opt}}$, that leads to a unique equilibrium. Then, together with Assumption 2, we can write both $\mathbf{P}_t^{e0}$ and $\mathbf{P}_t^{e\text{opt}}$ as functions of (i) the initial conditions and (ii) the current and expected future shocks, that is

$$\mathbf{P}_t^{e\text{opt}} = \mathbf{\Theta}_\xi^{\text{opt}} \mathbb{E}_t \mathbf{S}_t \qquad \text{and} \qquad \mathbf{P}_t^{e0} = \mathbf{\Theta}_\xi^0 \mathbb{E}_t \mathbf{S}_t + \mathbf{\Theta}_\epsilon^0 \boldsymbol{\epsilon}_t^{e0} , \tag{23}$$

where $\mathbf{S}_t = (\mathbf{\Xi}_t', \mathbf{X}_{-t}')'$ and $\mathbf{\Theta}_\xi^{\text{opt}}$, $\mathbf{\Theta}_\xi^0$ and $\mathbf{\Theta}_\epsilon^0$ are conformable matrices for which specific expressions are given in the appendix.

For $\mathbf{P}_t^{e\text{opt}}$, the map $\mathbf{\Theta}_\xi^{\text{opt}}$ captures the policy maker's optimal response to the state of the economy $\mathbb{E}_t \mathbf{S}_t$. For $\mathbf{P}_t^{e0}$, expression (23) captures the two possible reasons for a non-optimal policy: (i) the policy maker does not respond appropriately to the state of the economy ($\mathbf{\Theta}_\xi^0 \neq \mathbf{\Theta}_\xi^{\text{opt}}$), or (ii) the policy maker makes exogenous policy mistakes ($\boldsymbol{\epsilon}_t^{e0} \neq 0$).

We can then establish the following corollary

**Corollary 1.** *Given the generic model* (15)-(17) *and Assumptions 1-2, if there exists a rule* $\phi^{\text{opt}}$ *underlying* $\mathbf{P}_t^{e\text{opt}}$ *that leads to a unique equilibrium, we have that*

$$\boldsymbol{\delta}_t^* = \left(\mathbf{\Theta}_\xi^{\text{opt}} - \mathbf{\Theta}_\xi^0\right) \mathbb{E}_t \mathbf{S}_t - \mathbf{\Theta}_\epsilon^0 \boldsymbol{\epsilon}_t^{e0}. \tag{24}$$

Expression (24) shows that the OPP statistic corrects the two possible sources of optimization failures: (i) a non-optimal reaction function ($\mathbf{\Theta}_\xi^0 \neq \mathbf{\Theta}_\xi^{\text{opt}}$), or (ii) policy mistakes, i.e., non-zero exogenous policy shocks ($\boldsymbol{\epsilon}_t^{e0} \neq 0$). The OPP adjustment then corrects the non-optimal policy by canceling out the policy shocks and by changing how the policy maker responds to the state of the economy, i.e., by changing $\mathbf{\Theta}_\xi^0$ to $\mathbf{\Theta}_\xi^{\text{opt}}$.

## 4.2 Subset OPP

In practice, computing the entire matrix of causal effects $\mathcal{R}^0$ can be infeasible as identifying shocks to any element of the expected policy path can be hard. To avoid stringent iden-

[11]In other words, $\boldsymbol{\delta}_t^*$ is *minus* the coefficient estimate of a regression of $\mathbb{E}_t \mathbf{Y}_t^0$ on $\mathcal{R}$, because the goal of the OPP adjustment is not to best fit the path for $\mathbb{E}_t \mathbf{Y}_t^0$ with $\mathcal{R}^0$, but instead to best "undo" it.

tification assumption we modify our objective and outline an approach for evaluating and improving any subset of the proposed policy path. Specifically, let $\mathbf{P}^{e0}_{a,t}$ denote the subset or linear combination of the proposed policy path for which the corresponding policy shocks $\boldsymbol{\epsilon}^{e0}_{a,t}$ can be identified. The subset of the causal effects $\mathcal{R}^0_a$ measures the effect of $\boldsymbol{\epsilon}^{e0}_{a,t}$ on the policy objectives.[12]

The optimality of the proposed subset path $\mathbf{P}^{e0}_{a,t}$ can then be evaluated using the following subset-OPP statistic

$$\boldsymbol{\delta}^*_{a,t} = -(\mathcal{R}^{0\prime}_a \mathcal{W} \mathcal{R}^0_a)^{-1} \mathcal{R}^{0\prime}_a \mathcal{W} \mathbb{E}_t \mathbf{Y}^0_t \ . \tag{25}$$

The main benefit of the subset OPP statistic over the OPP statistic is its smaller information requirement —$\mathcal{R}^0_a$ is a subset of the full matrix of causal effects $\mathcal{R}^0$—, which will prove useful in empirical applications.

The subset-OPP statistic is able to detect non-optimal policies and adjusting by the subset statistic can bring the policy closer to optimality.

**Corollary 2.** *Given the generic model* (15)-(17) *and Assumptions 1-2, we have that*

1. $\boldsymbol{\delta}^*_{a,t} \neq 0 \quad \Longrightarrow \quad \mathbf{P}^{e0}_t \neq \mathbf{P}^{e0}_t$
2. $\mathcal{L}_t(\mathbf{P}^{e1}_t) \leq \mathcal{L}_t(\mathbf{P}^{e0}_t)$, *where* $\mathbf{P}^{e1}_t$ *replaces* $\mathbf{P}^{e0}_{a,t}$ *with* $\mathbf{P}^{e0}_{a,t} + \boldsymbol{\delta}^*_{a,t}$ *in* $\mathbf{P}^{e0}_t$.

Similar as in proposition 2, if the subset OPP statistic is non-zero the policy $\mathbf{P}^{e0}_t$ is non-optimal. Indeed, for an optimal policy the gradient with respect to the policy news shocks should be zero in *all* directions, i.e., for any element of $\mathbf{P}_t$ there should be no adjustment possible. Moreover, adjusting the subset of the proposed policy by the subset OPP lowers the loss function, but will generally not give the optimal allocation obtained with $\mathbf{P}^{e\,\mathrm{opt}}_t$.

# 5 Econometrics of the OPP

Section 4 derived attractive properties of the *population* (subset) OPP statistic. Thanks to the OPP, it is in theory not necessary to know the full structure of the economy to evaluate a macro policy decision: only sufficient statistics are required to detect and correct non-optimal policies. In practice of course, such statistics need to be estimated and section will formalize inference for the OPP based on estimated sufficient statistics. We will focus on the subset-OPP statistic (25) as it is the more empirically relevant statistic —requiring only the causal effects of shocks to a subset of the policy path—. For convenience we restate the subset-OPP statistic

$$\boldsymbol{\delta}^*_{a,t} = -(\mathcal{R}^{0\prime}_a \mathcal{W} \mathcal{R}^0_a)^{-1} \mathcal{R}^{0\prime}_a \mathcal{W} \mathbb{E}_t \mathbf{Y}^0_t \ ,$$

[12]Formally, starting from (20) we have that $\mathbf{Y}^0_t = \mathcal{R}^0 \boldsymbol{\epsilon}^{e0}_t + \boldsymbol{\Upsilon}^0_t$ which we decompose as $\mathbf{Y}^0_t = \mathcal{R}^0_a \boldsymbol{\epsilon}^{e0}_{a,t} + \mathcal{R}^0_{a^\perp} \boldsymbol{\epsilon}^{e0}_{a^\perp,t} + \boldsymbol{\Upsilon}^0_t$, where $\boldsymbol{\epsilon}^{e0}_{a,t}$ denotes the subset or linear combination of policy shocks that can be identified.

which relies on the causal effects $\mathcal{R}_a^0$ and the oracle forecasts $\mathbb{E}_t\mathbf{Y}_t^0$. We impose the following high-level assumption.

**Assumption 3.** *The underlying economic model is unknown, but one can construct estimates for $\mathcal{R}_a^0$ and $\mathbb{E}_t\mathbf{Y}_t^0$, denoted by $\widehat{\mathcal{R}}_a^0$ and $\widehat{\mathbf{Y}}_t^0$, that satisfy*

- $\mathbb{E}_t\mathbf{Y}_t^0 = \widehat{\mathbf{Y}}_t^0 + \mathbf{U}_t^y$ *and* $\mathbf{U}_t^y \overset{a}{\sim} F_{Y_t^0}$
- $\mathcal{R}_a^0 = \widehat{\mathcal{R}}_a^0 + \mathbf{U}_t^{\mathcal{R}_a^0}$ *and* $\mathbf{U}_t^{\mathcal{R}_a^0} \overset{a}{\sim} F_{\mathcal{R}_a^0}$

*with $F_{Y_t^0}$ and $F_{\mathcal{R}_a^0}$ some known or estimable distribution functions.*

Note that Assumption 3 is weaker than typically imposed in the literature. Unlike the rest of the literature, we do not assume that the model is explicitly known (in the case of model (15), that the matrices $\mathcal{A}_{..}$ and $\mathcal{B}_{..}$ are known). Instead, we only impose that it is possible to estimate the dynamics causal effect $\mathcal{R}_a^0$ and to produce a forecast $\widehat{\mathbf{Y}}_t^0$ that approximates $\mathbb{E}_t\mathbf{Y}_t^0$.

We think Assumption 3 is reasonable for two reasons. First, the assumption is in line with the reality of policy making. Policy makers do not rely on one specific model, but they do center the decision making process around the two statistics $\mathbb{E}_t\mathbf{Y}_t^0$ and $\mathcal{R}_a^0$. In the language of policy makers, the estimate of $\mathbb{E}_t\mathbf{Y}_t^0$ is often referred to as the *economic outlook*, while the estimate of $\mathcal{R}_a^0$, the effect of the policy instruments on the objectives, is often referred to as the *policy multiplier*. The statistics are routinely estimated (or at least extensively discussed) as part of their decision making process (e.g. Orphanides, 2019).[13] Second, a large econometrics literature has shown that it is possible to estimate these sufficient statistics without relying on a specific economic model, i.e., without assuming that the underlying model is fully specified. We now discuss these points in more detail.

**Causal effects.** To estimate the causal effects ($\mathcal{R}_a^0$), one can rely on a large macroeconometric literature on the estimation of impulse responses to policy shocks. Different approaches can be considered, and we do not take a specific stand on which method should be used. The review of Ramey (2016) provides a wealth of options ranging from local projection (LP) methods to structural VAR methods and outlines a number of possible identification strategies. Generally, the estimation of impulse responses requires three types of assumptions: (i) an identification assumption, (ii) an assumption on the class of reduced form econometric models (e.g., a linear model) and (iii) some regularity conditions. Calling a suitable estimator $\widehat{\mathcal{R}}_a^0$ and defining $\hat{r}_a = \text{vec}(\widehat{\mathcal{R}}_a^0)$ and $r_a^0 = \text{vec}(\mathcal{R}_a^0)$, we can construct

$$r_a^0 \overset{a}{\sim} N(\hat{r}_a, \widehat{\Sigma}_a) \ , \tag{26}$$

[13] That said, the two statistics have not been explicitly used in the context of policy assessment, and thus not fully exploited, as we show in this paper.

where $\widehat{\Sigma}_a$ is the estimate for the variance of $\hat{r}_a$.[14]

**Forecasts.** Turning to $\mathbb{E}_t\mathbf{Y}_t^0$, a large forecasting literature has studied how one can construct superior forecasts by combining multiple (imperfect) models, judgment and large and disparate information sources (e.g. Stock and Watson, 2002*a*,*b*; Lawrence et al., 2006; Manganelli, 2009; Geweke and Amisano, 2012; Giacomini and Ragusa, 2014; Cheng and Hansen, 2015). Denoting such point forecast by $\widehat{\mathbf{Y}}_t$, we also need to approximate the distribution of model uncertainty: the distribution of $\mathbb{E}_t\mathbf{Y}_t^0 - \widehat{\mathbf{Y}}_t$, the difference between the oracle forecast $\mathbb{E}_t\mathbf{Y}_t^0$ and the observed point forecast $\widehat{\mathbf{Y}}_t$. To estimate the distribution of model uncertainty, a conservative approach consists in approximating the distribution of $\mathbb{E}_t\mathbf{Y}_t^0 - \widehat{\mathbf{Y}}_t$ by the distribution of the historical forecast errors $\{\mathbf{Y}_s - \widehat{\mathbf{Y}}_s\}_{s=t^0}^t$. Using this sequence one can estimate the historical bias and variance and use these to *upper-bound* the distribution of model misspecification error (i.e., model uncertainty) using a normality assumption.[15] This approach yields an approximation for the distribution of $\mathbb{E}_t\mathbf{Y}_t^0 - \widehat{\mathbf{Y}}_t$ which we denote by

$$\mathbb{E}_t\mathbf{Y}_t^0 - \widehat{\mathbf{Y}}_t \overset{a}{\sim} F_{Y_t^0} \; . \tag{27}$$

**Computing the subset OPP.** After obtaining the approximating distributions of the dynamic causal effects and the oracle forecasts, we can compute the distribution of the OPP using simulation methods for a given preference matrix $\mathcal{W}$. Specifically, we simulate dynamic causal effects from (26) and forecasts misspecification errors from (27) and compute $\boldsymbol{\delta}_{a,t}^j$ the simulated OPP statistic. We repeat this for a large number of draws $j = 1, \ldots, S_d$, and report the average OPP and the confidence interval for some $\alpha \in (0,1)$:

$$\widehat{\boldsymbol{\delta}}_{a,t} = \frac{1}{S_d}\sum_{j=1}^{S_d}\boldsymbol{\delta}_{a,t}^j \qquad \text{and} \qquad \left[\boldsymbol{\delta}_{a,t}^{(\alpha S_d)}, \boldsymbol{\delta}_{a,t}^{((1-\alpha)S_d)}\right] \; , \tag{28}$$

where $\boldsymbol{\delta}_{a,t}^{(k)}$ denotes the (element wise) $k$th largest draw of $\{\boldsymbol{\delta}_{a,t}^j, j = 1, \ldots, S_d\}$.

**Policy evaluation and improvement.** Based on Corollary 2 we will conclude that a policy $\mathbf{P}_t^{e0}$ is not optimal whenever the confidence bands of $\boldsymbol{\delta}_{a,t}^*$ exclude zero at any desired level of confidence. Moreover, we can improve the policy choice by adjusting $\mathbf{P}_t^{e0}$ with the

[14]This approximation can follow from both frequentist and Bayesian arguments. For instance, in a frequentist setting many estimators, under suitable assumptions, are asymptotically normal implying that $\sqrt{n}(\hat{r}_a - r_a^0) \overset{d}{\to} N(0, \Sigma_a)$. Given that such a result applies and that the asymptotic variance can be consistently estimated, we can obtain the approximation (26).

[15]Forecast errors mix two sources of uncertainty: (i) misspecification, i.e., model uncertainty, *and* (ii) future uncertainty. As a result, the variance of forecast errors will upper-bound the variance of mis-specification uncertainty.

subset OPP mean $\widehat{\boldsymbol{\delta}}_{a,t}$. In fact, we have

$$\widehat{\boldsymbol{\delta}}_{a,t} = (\widehat{\mathcal{R}}_a^{0\prime}\mathcal{W}\widehat{\mathcal{R}}_a^0 + \Omega_a)^{-1}\widehat{\mathcal{R}}_a^{0\prime}\mathcal{W}(\widehat{\mathbf{Y}}_t^0 + \bar{\mathbf{U}}_t^y) \ , \tag{29}$$

where $\Omega_a$ is a function of $\mathcal{W}$ and the variance-covariance matrix of the distribution $F_{\mathcal{R}_a^0}$, and $\bar{\mathbf{U}}_t^y$ is the bias in the forecast.

Note how $\widehat{\boldsymbol{\delta}}_{a,t}$ does not correspond to $(\widehat{\mathcal{R}}^{0\prime}\mathcal{W}\widehat{\mathcal{R}}^0)^{-1}\widehat{\mathcal{R}}^{0\prime}\mathcal{W}\widehat{\mathbf{Y}}_t^0$, a naive plug-in estimator of $\boldsymbol{\delta}_{a,t}^*$, where we would use the point-estimates for $\mathcal{R}_a^0$ and $\mathbb{E}_t\mathbf{Y}_t^0$. There are two reasons for that. First, the term $\Omega_a$ can be thought of as capturing an attenuation bias coming from uncertainty in our knowledge of the effect of policy (uncertainty in $\mathcal{R}_a^0$). This result goes back to the seminal Brainard (1967) conservatism principle. Brainard's principle states that in the face of parameter uncertainty, a policy maker should be more conservative in its use of the policy instruments and refrain from fulling minimizing the loss function. Exactly the same happens when the policy maker uses the OPP mean estimate (29) to update the proposed policy, i.e. $\mathbf{P}_{a,t}^{e0} + \widehat{\boldsymbol{\delta}}_{a,t}$ is a conservative adjustment to policy.

Second, the possibly of a non-zero bias in the forecast $\bar{\mathbf{U}}_t^y$ affects the policy adjustment. If the bias is estimable (as we assumed in Assumption 3), say from historical forecast errors, one simply corrects the point forecast $\widehat{\mathbf{Y}}_t^0$ by adding the bias term $\bar{\mathbf{U}}_t^y$. Alternatively, if the bias $\bar{\mathbf{U}}_t^y$ cannot be estimated but one can impose a bound on the magnitude of the bias, it is possible to consider a worst case scenario where $\bar{\mathbf{U}}_t^y$ is chosen to minimize the OPP statistic. This would correspond to a minimax rule in the spirit of Hansen and Sargent (2008).

# 6 Illustration: US monetary policy

In this section we illustrate how the OPP statistic can be used in practice to evaluate and improve monetary policy decisions. As loss function we posit the dual inflation-full employment mandate imposed by the US Congress on the Fed:

$$\mathcal{L}_t = \|\Pi_t\|^2 + \lambda \|U_t\|^2 \ , \tag{30}$$

with $\Pi_t = (\pi_t - \pi_t^*, \ldots, \pi_{t+H} - \pi_t^*)'$ the vector of inflation gaps and $U_t = (u_t - u_t^*, \ldots, u_{t+H} - u_t^*)'$ the vector of unemployment gaps. We truncate the paths at a horizon of $H = 5$ years, and the discount rate is implicitly set to $\beta_h = 1$ for all $h$. Since the Fed publicly announced that it follows a balanced approach to its dual mandate, i.e., $\lambda = 1$ (Bernanke, 2015), our baseline results are based on $\lambda = 1$. In the web-appendix, we discuss the choice of $\lambda$ and show results for a range of $\lambda$ over $[0.2, 2]$.[16]

[16]Arguably, the main beneficiaries of a sufficient statistic approach to policy evaluation are the policy makers themselves, as we discussed in the introduction. In that case, $\lambda$ (or $\beta$) is a preference parameter, i.e.,

We evaluate Fed decisions over 1990-2018. After discussing the construction of the OPP statistic (and associated distribution), we present and discuss the sequence of OPP statistics over 1990-2018 —a systematic study of all policy decisions over that sample period—, and we then focus on three specific decisions: 1990M6, 2008M4 on the eve of the Great Recession, 2010M4 in the middle of the Great recession—.

## 6.1 Computing the subset OPP

We first describe how we construct the subset OPP statistics (and associated distributions).

**The subset OPP** To assess policy decisions, we will compute a subset OPP vector with two elements, each corresponding to a specific policy "experiment": (i) a *short-rate OPP* corresponding to an innovation to the contemporaneous fed funds rate and (ii) a *slope OPP* corresponding to an innovation to the slope of the fed funds rate path. In other words, the short-rate OPP will capture the optimal adjustment to the contemporaneous fed funds rate and the slope OPP will capture the optimal adjustment to the slope of the fed funds rate path.

**Estimating dynamic causal effects** To estimate the impulse responses to these policy shocks, we follow Kuttner (2001) and Eberly, Stock and Wright (2019) and assume that the Fed's reaction function was stable over 1990-2018 (such that the matrix $\mathcal{R}_a$ is constant over that period),[17] and we use as instrumental variables the monetary policy surprises measured around the FOMC announcements within a 30 minute window. First, we use surprises to the fed funds rate —the difference between the expected fed funds rate (as implied by current-month federal funds futures contracts) and the actual fed funds rate— to identify the effects of a shock to the contemporaneous fed funds rate. Second, we use surprises to the ten-year on-the-run Treasury yield (orthogonalized with respect to surprises to the current fed funds rate) to identify the effects of shocks to the slope of the expected fed funds rate path. We then use these surprises as instrumental variables in local projections to compute the dynamic causal effects $\widehat{\mathcal{R}}_a^0$.[18]

a choice for the policy maker. But for a retrospective and external analysis of policy decisions, presenting results for a range of values for $\lambda$ is useful to understand which alternative values for $\lambda$ *could* explain some decisions. In the web-appendix, we also propose a conservative (i.e., robust) approach to elicit $\lambda$. Specifically, the procedure consists in picking the $\lambda$ that is least favorable to rejecting that a policy was optimal.

[17]In the web-appendix, we test this assumption by testing the stability of our $\mathcal{R}_a$ estimates using the structural change tests proposed in Hall, Han and Boldea (2012). We find no evidence of parameter instability.

[18]Using SVAR-IV (e.g. Montiel Olea, Stock and Watson, 2020) gives similar results, see the online Appendix. As in Eberly, Stock and Wright (2019), the effect of shocks to the fed funds rate is estimated over 1990-2007, which avoids the zero-lower bound period, and the effect of shocks to the slope of the yield curve is estimated over 2008-2018, the period during which the Fed was actively trying to affect that slope.

**Measuring $\mathbb{E}_t\mathbf{Y}_t^0$** For the conditional forecasts $\widehat{\mathbf{Y}}_t$, we use the median FOMC forecasts reported in the Survey of Economic Projections (SEP).[19] Since the SEP forecasts only extend up to three years out, we complement them with the median FOMC estimates for long-run inflation and unemployment, and we posit that these long-run values are reached after 5 years. Since SEP projections are annual, we linearly interpolate them in order to project them on the estimated effects of the policy instruments (available at a quarterly frequency). As long-run targets $u_t^*$ and $\pi_t^*$, we again use the FOMC estimates for long-run inflation and unemployment. To capture the uncertainty around these point forecasts, we use the Board staff assessment of forecast uncertainty, as reported in the Tealbook.

**The OPP statistics** Based on $\widehat{\mathcal{R}}_a^0$, and $\widehat{\mathbf{Y}}_t$, we compute the mean subset OPP statistics $\widehat{\boldsymbol{\delta}}_{a,t}$ and construct confidence bands as described in Section 5. The subset OPP $\widehat{\boldsymbol{\delta}}_{a,t}$ is a vector with two elements: (i) the short-rate OPP denoted by $\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}$, and (ii) the slope OPP denoted by $\widehat{\boldsymbol{\delta}}_{\Delta,t}$.

In terms of confidence interval, we will report both the 68% confidence bands and the 95%.[20] Importantly, we note that the objective of a policy optimality test is different from the traditional objective of statistic learning. Specifically, from the perspective of the policy maker it is not clear that high significance is the most interesting/appropriate criteria. Consider the main outcome of the OPP test: "With $X$% confidence, the proposed policy choice is not optimal". A policy maker particularly averse to making a non-optimal decision may want to change the proposed policy choice at a relatively low $X$ level, say 68% instead of the usual 95%, as she may want to discard a policy that is non-optimal with a 68% probability. A trade-off however is that too low a threshold may lead a policy maker to change policy course too often. While such a decision problem is outside the scope of this paper, it highlights that for the OPP test the classical dichotomy of hypothesis testing (i.e. preference for type 1 vs type 2 errors) really depends on the preference of the policy maker, e.g., making non-optimal decisions vs. changing course frequently. In the empirical application, we will thus show the rejection probability for different significance levels.

## 6.2 A retrospective analysis of US monetary policy

Figure 4 displays the time series for the two elements of the subset OPP —the short-rate OPP and the slope OPP— along with their confidence intervals, as implied by both impulse response and model uncertainty.

[19] As an alternative to the SEP, we explored using the Board staff Greenbook forecasts. The results were very similar.

[20] Macroeconomic forecasts are typically noisy, and many policy makers report 68% confidence bands (instead of say 95%), see for instance the discussion on page 549 in Stock and Watson (2011).

While the contemporaneous fed funds rate has not been set exactly at its optimal level since 1990, the optimal adjustment (in absolute value) is overall relatively small averaging only 25 basis points over the full sample. There is however a few interesting cases of non-optimal policies.

We start with the short-rate OPP, which evaluates the optimality of the contemporaneous fed funds rate, the traditional tool of monetary policy.

First, notice the large negative short-rate OPP during the Great Recession: the contemporaneous fed funds rate was about 1 percentage point too high over 2009-2013. This is not an optimization failure per se since the fed funds rate was stuck at the zero lower bound, but it conveys the magnitude of the ZLB constraint, i.e., by how much the ZLB was restraining the Fed's ability to stabilize the economy with its traditional tool.

Second, on the eve of the Great Recession (when the ZLB was not yet binding), the short-rate OPP indicates (based on solely on real time information) that the FOMC should have lowered rates faster in early 2008.

Third, another case of a suboptimal fed funds rate occurs in the late 1990s, when the short-rate OPP indicates that the fed funds rate was too low by about .25ppt. This finding echoes earlier arguments that the Fed may have found itself falling behind the curve in the late 1990s tightening cycle (e.g., Blinder and Reis, 2005).

The slope OPP, which assesses the optimality of the slope of the policy path, confirms the conclusion of the short-rate OPP. During the Great Recession for instance, the slope OPP indicates that forward-guidance could have been used more aggressively, a conclusion echoing that of Eberly, Stock and Wright (2019). In 2009, the slope OPP drops rapidly to about -1ppt and only slowly revert back to zero. In fact, the slope OPP remains significantly different from zero over the whole 2009-2013 period. Overall, these results indicate that a more active manipulation of the expected policy path through forward guidance holds considerable promise for improvements in the conduct of policy.

To better understand these results and also to illustrate the workings of the OPP, we consider three instructive case studies: June 1990 where the OPP is essentially zero with the Fed successfully balancing conflicting objectives, (ii) April 2008 where a non-optimal policy is detected by the short-rate OPP at the onset of the Great Recession, (iii) April 2010 where a non-optimal policy path is detected by the slope OPP.

## 6.3 Three case studies

**Fed funds rate policy as of June 1990** In the first case study, we evaluate the optimality of the contemporaneous fed funds rate as of June 1990. In June 1990, the FOMC was confronted with a classic inflation-unemployment trade-off: while it would have liked to lower the fed funds rate to fight excess unemployment, it was prevented to do so by the high and on-going inflation (Bluebook, June 2006). The question for the policy maker at the time (and thus for the OPP) is whether the level of the fed funds rate optimally balanced that trade-off.

Figure 1 depicts graphically all the information needed to construct the short-rate OPP $\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}$. The top-left panel reports the FOMC conditional expected path for the inflation gap, that is it reports $\widehat{\Pi}_t$. The bottom-left panel reports $\widehat{\mathcal{R}}_{\mathrm{i}}^{0\pi}$, the estimated effect of a 1ppt shock to the contemporaneous fed funds rate on inflation. The right column reports the same information for unemployment: $\hat{U}_t$ and $\widehat{\mathcal{R}}_{\mathrm{i}}^{0u}$. For illustration purposes, in this first case study we omit confidence bands and treat the causal effect estimates and forecasts as fixed.

To illustrate the workings of the OPP, we first consider the case of a strict inflation targeter with $\lambda = 0$. In the top-left panel, the red empty-circles display how the expected path for inflation would change if we adjusted the fed funds rate by $\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}^{\pi}$, the OPP for a strict inflation targeter that did not care about the path of unemployment ($\lambda = 0$). With $\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}^{\pi} \approx 0.7 > 0$, the short-rate OPP calls for a more contractionary policy in order to bring down inflation faster (empty red circles). To understand the workings of the OPP, recall that the effect of a policy adjustment $\delta_t$ can be obtained by simply adding (or subtracting) to the baseline inflation forecast the impulse response of inflation to a fed funds rate shock (scaled by the magnitude of the policy adjustment).[21] The goal of the OPP —the optimal adjustment— when $\lambda = 0$ is then to best "use" the impulse response of inflation in order to best stabilize the forecast for the inflation gap. Because of the lag in the effect of policy, the Fed can do little about the contemporaneous burst of inflation but it can bring down inflation in two years by raising rates today.

In the top-right panel, the blue empty-circles plot a similar counter-factual exercise but for a strict unemployment targeter that ignores the path of inflation ($\lambda = \infty$). This time, the OPP statistic calls for a more expansionary policy ($\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}^{u} \approx -0.2 < 0$) in order to use the (negative of) impulse response of unemployment to bring down unemployment faster.

With a dual inflation–unemployment mandate ($\lambda = 1$), the FOMC is facing an inflation-unemployment trade-off, as the two OPP statistics $\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}^{\pi}$ and $\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}^{u}$ call for opposite policies.[22] Intuitively, when computing the dual-mandate OPP, the goal is to best use *both* impulse responses of inflation and unemployment in order to best stabilize *both* the inflation and the unemployment forecasts. Given that the impulse responses of inflation and unemployment

[21] As we saw in Section 4, we have $\mathbb{E}_t \mathbf{Y}_t^1 = \mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}^0 \boldsymbol{\delta}_t$ with $\mathbb{E}_t \mathbf{Y}_t^1$ the allocation after the $\boldsymbol{\delta}_t$ adjustment.

[22] We can re-write $\boldsymbol{\delta}_{\mathrm{i},t}^{*}$ as a weighted-average of the OPP for each mandate with

$$\boldsymbol{\delta}_{\mathrm{i},t}^{*} = (1-\omega)\boldsymbol{\delta}_{\mathrm{i},t}^{\pi *} + \omega \boldsymbol{\delta}_{\mathrm{i},t}^{u *} \ , \tag{31}$$

with $\boldsymbol{\delta}_{\mathrm{i},t}^{v*} = -(\mathcal{R}_{\mathrm{i}}^{0^{v'}} \mathcal{R}_{\mathrm{i}}^{0^{v}})^{-1} \mathcal{R}_{\mathrm{i}}^{0^{v'}} \mathbb{E}_t V_t^0$ the OPP for a single mandate with $V_t = (v_t - v^*, \ldots, v_{t+H} - v^*)'$ for $v = \pi$ or $u$, and $\omega = \frac{1}{1+\kappa^2/\lambda}$.

move in opposite direction, it is not possible to simultaneously lower both inflation and unemployment: the Fed is facing a trade-off, and an optimization failure —a non-zero OPP $\widehat{\boldsymbol{\delta}}_t$— can come from a failure to appropriately balance conflicting objectives, in this case $\widehat{\boldsymbol{\delta}}^{\pi}{}_{\mathsf{i},t} > 0$ but $\widehat{\boldsymbol{\delta}}^{u}_{\mathsf{i},t} < 0$.

We find $\widehat{\boldsymbol{\delta}}_{\mathsf{i},t} = -0.11$ (see Table 1), meaning that the two mandates were roughly balanced and the departure from optimality is economically small. Moreover, taking estimation and model uncertainty into account, the 68 percent confidence interval includes zero, and we cannot discard that the contemporaneous fed funds rate was set optimally.

**Fed funds rate policy as of April 2008** In the second case study, we evaluate the fed funds rate policy as of April 2008, in the early stage of the financial crisis: Lehman Brothers was still 6 months away from failing, unemployment was only at 5 percent, and few anticipated the magnitude of the recession that was going to ensue. In fact, the fed funds rate was still at 2.25ppt so the Fed still had room to use conventional policies to stimulate activity.[23] At that meeting, the fed funds rate was lowered by .25ppt to 2 percent, but it remained at that level until October 2008, i.e., until the collapse of Lehman brothers.

As is clear from the April Tealbook and forecast narratives reported by the FOMC, the central bank was facing two conflicting issues in April 2008: (i) a marked deterioration in the growth outlook due declining housing prices and tensions in the financial market, and (ii) upside risks to inflation coming from "persistent surprises to energy and commodity prices" (Kohn, 2008).

An interesting question in hindsight is thus whether the 2008-M4 decision was optimal. Figure 2 has the same structure as Figure 1 except that we now report the confidence intervals for the impulse response estimates, as well as the confidence intervals capturing the model uncertainty surrounding the Fed's forecast, as judged by the Board staff in the April 2008 Tealbook.

The two issues of the time —poor economic outlook and inflationary pressures from high energy prices— are visible in the FOMC forecasts in the first row of Figure 2.

The short-rate OPP comes out at $\widehat{\boldsymbol{\delta}}_{\mathsf{i},t} = -0.37$ (Table 1), calling for an additional 25 or 50 basis points cuts (depending on rounding). The 68% confidence interval excludes zero, indicating that there is a less than 32 percent chance that the policy $\mathsf{i}^0_t$ was optimal, i.e., balanced the expected paths of the inflation and unemployment gaps. The unfilled dots plot the counter-factual expected paths for the policy objectives after adjusting the policy path with the short-rate OPP. We can see that the FOMC could have brought down expected unemployment faster at a small inflationary cost. Indeed, the effect of monetary policy on

[23]By the end of 2008 however, unemployment had reached 7.3 percent, and the Fed had dropped the fed funds rate by almost 2ppt (to the zero lower bound) in the span of only three months (September-December) following the failure of Lehman Brothers in September 2008.

inflation is so delayed that the extra expected inflation caused by a lower fed funds rate will only materialize two years later, i.e., after the commodities-driven burst in inflation has died down.

**Slope (QE) policy as of April 2010** It is interesting to contrast the 2008-M4 situation with that of two years later; in 2010-M4. There, the Fed funds rate was stuck at zero but the Fed could have further used forward-guidance to better stabilize the economy. To test this possibility, we can use the slope OPP statistic.

Figure 3 displays the situation in 2010-M4 where the bottom panels show the effects on inflation and unemployment of a 1ppt shock to the slope of the yield curve. In Table 1 we find that the mean estimate of the slope OPP is given by $\widehat{\boldsymbol{\delta}}_{\Delta,t} = -0.90$, calling for a substantial decline in the slope of the expected policy path (Figure 4, top-right panel). The deviations from targets are so large that we can easily discard optimality even at the 95 percent confidence level.

The unfilled dots plot the counter-factual expected paths for the policy objectives after adjusting the policy path with the slope OPP. The FOMC could have brought down expected unemployment faster in exchange for a small overshoot in expected inflation in 2011.

# 7 Generalizing the OPP approach

So far we have developed our sufficient statistics approach for policy evaluation in the context of linear models that can be written as in (15). In this section we explore for which other classes of models the statistics $\mathcal{R}$ and $\mathbb{E}_t\mathbf{Y}_t^0$ are sufficient to evaluate, and possibly improve, policy decisions. Key examples that we consider are models with state dependence (e.g. Auerbach and Gorodnichenko, 2013) and models with multiple policy regimes (e.g. Sims and Zha, 2006).

Recall that the properties of the OPP derive from the equivalence

$$\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\,\text{opt}} \qquad \Longleftrightarrow \qquad \nabla_{\boldsymbol{\epsilon}_t}\mathcal{L}_t|_{\mathbf{P}_t^{e0}} = 0 \ , \tag{32}$$

which we have shown to hold for linear models of the generic form (15). In fact, the equivalence holds for all models that can be viewed as conditionally linear. In the web-appendix we provide a high level framework that exactly spells out the necessary conditions on the underlying economy that ensure that our sufficient statistics approach applies.

Here we keep the discussion concrete, and we will make two specific points. First, as long as the model is linear conditional on time-$t$ predetermined variables, the equivalence continues to hold and all our previous results hold. This case notably includes models of state dependent policy effects that are often considered in the empirical literature (e.g. Auerbach and Gorodnichenko, 2013). Second, in a model with multiple regimes conditioned by the policy rule (e.g. Sims and Zha, 2006) as long as the economy can only be in a finite number of regimes, the equivalence no longer holds, but a non-zero gradient still implies that the proposed policy choice is non-optimal. In other words, the two statistics $\mathcal{R}^0$ and $\mathbb{E}_t\mathbf{Y}_t^0$ are still sufficient to evaluate a policy decision, but adjusting the policy choice with the OPP is no longer guaranteed to yield a superior policy decision.

## 7.1 State dependence

Numerous works have documented evidence for various forms of state dependence in the effects of fiscal and monetary policy, where the state dependence is governed by some time-$t$ pre-determined variable that is independent of the policy decision.[24] Our main results continue to hold in this setting, and the two statistics $\mathcal{R}^0$ and $\mathbb{E}_t\mathbf{Y}_t^0$ are sufficient to detect and correct non-optimal policy decisions. The only difference is that the statistic $\mathcal{R}^0$ needs to be conditioned on the state of the economy. As an illustration, consider the specification of Auerbach and Gorodnichenko (2013) where the economy can be in two states, depending on the value of some state variable $z_t$. In that case, the generic model is modified in that the maps $\mathcal{A}_{..}$ and $\mathcal{B}_{..}$ become functions of $z_t$, i.e., $\mathcal{A}_{..}(z_t)$ and $\mathcal{B}_{..}(z_t)$.

With two states, each map can be written as $\mathcal{A}_{..}(z_t) = F(z_t)\mathcal{A}_{..(1)} + (1-F(z_t))\mathcal{A}_{..(2)}$ where $F(z_t)$ can be interpreted as a measure of probability of being in state 1 at time $t$ based on some time $t$ predetermined variable $z_t$.[25] The effects of policy shocks will then be given by on the state and can take two values $\mathcal{R}^0_{(1)}$ or $\mathcal{R}^0_{(2)}$.

The corresponding state dependent OPP statistic can be written as

$$\boldsymbol{\delta}_t^*(z_t) = -(\mathcal{R}^0(z_t)'\mathcal{W}\mathcal{R}^0(z_t))^{-1}\mathcal{R}^0(z_t)'\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 \ ,$$

where $\mathcal{R}^0(z_t) = F(z_t)\mathcal{R}^0_{(1)} + (1-F(z_t))\mathcal{R}^0_{(2)}$. As shown formally in the web-appendix, the state dependent OPP inherits all properties of the baseline OPP: $\boldsymbol{\delta}_t^*(z_t) \neq 0$ implies that $\mathbf{P}_t^{e0}$ is non-optimal, and $\mathbf{P}_t^{e0} + \boldsymbol{\delta}_t^*(z_t)$ corrects the optimization failure.

[24] See e.g., Auerbach and Gorodnichenko (2012, 2013); Ramey and Zubairy (2018); Barnichon, Debortoli and Matthes (2021) for studies on whether fiscal policy is more or less effective when the economy is in a high unemployment state, and Tenreyro and Thwaites (2016); Ascari and Haber (2021); Eichenbaum, Rebelo and Wong (2022) for studies on whether monetary policy is more or less effective when unemployment is high.

[25] A popular functional form for $F(.)$ is $F(z_t) = \exp(-\gamma z_t)/[1+\exp(-\gamma z_t)]$ with $\gamma$ a tuning parameter.

## 7.2 Multiple policy regimes

Next, we consider an economy with a finite number of policy regimes, where the model coefficients (equations (15)) can depend on the policy regime. To give a concrete and relevant example in monetary policy, inflation expectations can be "anchored" —fixed at some value— or "unanchored", in that inflation expectations depend the state of the economy (expectation formation could be e.g., rational or adaptive), and the anchoring of inflation expectations likely depends on the central bank's policy rule or objective function (e.g., Bernanke, 2007).

In this type of model, the policy rule can affect the coefficients of the non-policy block (15). In that case, we can still use the OPP to detect optimization failures, but there is no longer any guarantee (without additional structural assumption) that an OPP adjustment would improve policy. In other words, the two statistics $\mathcal{R}^0$ and $\mathbb{E}_t\mathbf{Y}_t^0$ are still sufficient to evaluate a policy decision, but they may not be sufficient to correct a non-optimal policy decision.

To see that, consider an economy described by $N$ regimes indicated by $\vartheta \in \{\vartheta_1, \ldots, \vartheta_N\}$. The regime is determined by the policy rule $\phi$ chosen by the policy maker. In each regime the effects of policy shocks can be different, and $\mathcal{R}^0(\vartheta_i)$ captures the causal effect of policy shocks under regime $\vartheta_i$. Given a policy proposal $\mathbf{P}_t^{e0}$, implied by a rule $\phi^0$, which correspond to the regime $\vartheta^0 \in \{\vartheta_1, \ldots, \vartheta_N\}$, the regime specific OPP statistic is

$$\boldsymbol{\delta}_t^*(\vartheta^0) = -(\mathcal{R}^0(\vartheta^0)'\mathcal{W}\mathcal{R}^0(\vartheta^0))^{-1}\mathcal{R}^0(\vartheta^0)'\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 \ .$$

The regime specific OPP retains the property that $\boldsymbol{\delta}_t^*(\vartheta^0) \neq 0$ implies that $\mathbf{P}_t^{e0}$ is non-optimal: $\mathcal{R}^0(\vartheta^0)$ and $\mathbb{E}_t\mathbf{Y}_t^0$ are still sufficient to evaluate a policy decision.

Intuitively, our detection of a non-optimal policy relies on the gradient being non-zero, and the gradient captures the effect of an infinitesimally small change in policy. Such infinitesimally small change in the reaction function will not trigger a regime change, as in our our example of anchored vs. unanchored inflation expectations. In that case, the two statistics are still sufficient to detect a non-optimal policy.

However, adjusting the policy by the OPP is no longer guaranteed to give the optimal policy. Indeed, recall from Corollary 1 that an OPP adjustment amounts to a change in the reaction function. Thus adjusting $\mathbf{P}_t^{e0}$ by $\boldsymbol{\delta}_t^*(\vartheta^0)$ could lead to a new regime, call it $\vartheta^1$, where the effect of policy shocks $\mathcal{R}^0(\vartheta^1)$ is different. In that case, the two statistics $\mathcal{R}^0(\vartheta^0)$ and $\mathbb{E}_t\mathbf{Y}_t^0$ are no longer sufficient to improve policy.

Intuitively, with multiple policy regimes the optimization problem is no longer linear-quadratic, as the model is no longer linear with respect to the coefficients of non-policy block. As a result, the problem may not be convex, and the first-order condition of optimality may not be sufficient: the gradient could be zero, because the policy choice is only a *local* minimum.

# 8 Conclusion

In this work, we show that it is possible to evaluate macro policy decisions from two sufficient statistics that determine the gradient of the loss function with respect to policy shocks. The two sufficient statistics are (i) the dynamic causal effects of the policy instruments on the policy objectives —the policy multiplier—, and (ii) forecasts for the policy objectives conditional on the policy decision —the economic outlook—. These statistics can be estimated without a specific economic model. The *Optimal Policy Perturbation* (OPP) is the gradient of the loss function (rescaled by the Hessian) and we show how the OPP is sufficient to detect, and often correct, non-optimal policies.

Importantly, the monetary policy setting considered in this paper is only one of many potential applications of a sufficient statistics approach to evaluating macro policy problems. Other fruitful uses include the many areas where macro policy makers must balance difficult trade-offs in complex settings: fiscal policy (e.g., balancing growth considerations with risks to debt sustainability), exchange rate management (balancing monetary independence with exchange rate stability), foreign-reserve management (e.g., balancing the cost of holding reserves with the insurance against sudden stops in capital flows), or even climate change policy (e.g., balancing the costs of climate change with the costs of preventive actions), among others.

# References

**Ascari, Guido, and Timo Haber.** 2021. "Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy." *The Economic Journal*, 132(641): 37–57.

**Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub.** 2021. "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models." *Econometrica*, 89(5): 2375–2408.

**Auerbach, Alan J., and Yuriy Gorodnichenko.** 2012. "Measuring the Output Responses to Fiscal Policy." *American Economic Journal: Economic Policy*, 4: 1–27.

**Auerbach, Alan J., and Yuriy Gorodnichenko.** 2013. "Output Spillovers from Fiscal Policy." *American Economic Review*, 103(3): 141–46.

**Barnichon, Regis, Davide Debortoli, and Christian Matthes.** 2021. "Understanding the Size of the Government Spending Multiplier: It's in the Sign." *The Review of Economic Studies*, 89(1): 87–117.

**Bernanke, Ben.** 2007. "Inflation expectations and inflation forecasting." Board of Governors of the Federal Reserve System (US).

**Bernanke, Ben S.** 2015. "The Taylor Rule: A benchmark for monetary policy?" *Ben Bernanke's Blog*, 28.

**Blanchard, Olivier Jean, and Danny Quah.** 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review*, 79(4): 655–673.

**Blinder, Alan, and Ricardo Reis.** 2005. *Understanding the Greenspan standard.* Citeseer.

**Blinder, Alan S.** 2020. "What does Jerome Powell know that William McChesney Martin didn't—and what role did academic research play in that?" *The Manchester School.*

**Brainard, William C.** 1967. "Uncertainty and the Effectiveness of Policy." *The American Economic Review*, 57(2): 411–425.

**Chahrour, Ryan, and Kyle Jurado.** 2018. "News or Noise? The Missing Link." *American Economic Review*, 108(7): 1702–36.

**Chari, Varadarajan V, Lawrence J Christiano, and Patrick J Kehoe.** 1994. "Optimal fiscal policy in a business cycle model." *Journal of polítical Economy*, 102(4): 617–652.

**Cheng, Xu, and Bruce E. Hansen.** 2015. "Forecasting with factor-augmented regression: A frequentist model averaging approach." *Journal of Econometrics*, 186(2): 280–293.

**Chetty, Raj.** 2009. "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods." *Annu. Rev. Econ.*, 1(1): 451–488.

**Eberly, Janice C., James H. Stock, and Jonathan H. Wright.** 2019. "The Federal Reserve's Current Framework for Monetary Policy: A Review and Assessment." Working paper.

**Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong.** 2022. "State Dependent Effects of Monetary Policy: The Refinancing Channel." *American Economic Review*. forthcoming.

**Faust, Jon.** 1998. "The robustness of identified VAR conclusions about money." *Carnegie-Rochester Conference Series on Public Policy*, 49: 207 – 244.

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

**Geweke, John, and Gianni Amisano.** 2012. "Prediction with Misspecified Models." *American Economic Review*, 102(3): 482–86.

**Giacomini, Raffaella, and Giuseppe Ragusa.** 2014. "Theory-coherent forecasting." *Journal of Econometrics*, 182(1): 145–155.

**Giannoni, Marc P, and Michael Woodford.** 2017. "Optimal target criteria for stabilization policy." *Journal of Economic Theory*, 168: 55–106.

**Hansen, Lars Peter, and Thomas J Sargent.** 2001. "Robust control and model uncertainty." *American Economic Review*, 91(2): 60–66.

**Hansen, Lars Peter, and Thomas J. Sargent.** 2008. *Robustness.* Princeton University Press.

**Kohn, Donald.** 2008. "Survey of Economic Projections." Report to the FOMC.

**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.

**Lawrence, Michael, Paul Goodwin, Marcus O'Connor, and Dilek Önkal.** 2006. "Judgmental forecasting: A review of progress over the last 25 years." *International Journal of forecasting*, 22(3): 493–518.

**Lucas, Robert Jr.** 1976. "Econometric policy evaluation: A critique." *Carnegie-Rochester Conference Series on Public Policy*, 1(1): 19–46.

**Manganelli, Simone.** 2009. "Forecasting with judgment." *Journal of Business & Economic Statistics*, 27(4): 553–563.

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

**Mertens, Karel, and Morten O. Ravn.** 2013. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review*, 103(4): 1212–1247.

**Michaillat, Pascal, and Emmanuel Saez.** 2019. "Optimal public expenditure with inefficient unemployment." *The Review of Economic Studies*, 86(3): 1301–1331.

**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, José L., James H. Stock, and Mark W. Watson.** 2020. "Inference in Structural Vector Autoregressions identified with an external instrument." *Journal of Econometrics.* forthcoming.

**Onatski, Alexei, and James H. Stock.** 2002. "Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy." *Macroeconomic Dynamics*, 6(1): 85?110.

**Onatski, Alexei, and Noah Williams.** 2003. "Modeling Model Uncertainty." *Journal of the European Economic Association*, 1(5): 1087–1122.

**Orphanides, Athanasios.** 2019. "Monetary Policy Strategy and its Communication." Working paper, presented at Jackson Hole Conference August 22, 2019.

**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.

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

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

**Smets, Frank, and Rafael Wouters.** 2007. "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American economic review*, 97(3): 586–606.

**Stock, James H., and Mark W. Watson.** 2002*a*. "Forecasting Using Principal Components from a Large Number of Predictors." *Journal of the American Statistical Association*, 97(460): 1167–1179.

**Stock, James H, and Mark W Watson.** 2002*b*. "Macroeconomic Forecasting Using Diffusion Indexes." *Journal of Business & Economic Statistics*, 20(2): 147–162.

**Stock, James H., and Mark W. Watson.** 2011. *Introduction to Econometrics (3rd edition).* Addison Wesley Longman.

**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 E.O.** 1999. "Inflation targeting as a monetary policy rule." *Journal of Monetary Economics*, 43(3): 607 – 654.

**Svensson, Lars E.O.** 2003. "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules." *Journal of Economic Literature*, 41(2): 426–477.

**Svensson, Lars E.O.** 2019. "Monetary policy strategies for the Federal Reserve." 4–5.

**Tenreyro, Silvana, and Gregory Thwaites.** 2016. "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions." *American Economic Journal: Macroeconomics*, 8(4): 43–74.

**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.

**Walsh, Carl E.** 2017. *Monetary theory and policy.* MIT press.

**Woodford, Michael.** 2003. *Interest and prices: Foundations of a theory of monetary policy.* princeton university press.

**Woodford, Michael.** 2010. "Optimal monetary stabilization policy." In *Handbook of Monetary Economics*. Vol. 3, 723–828. Elsevier.

# Appendix

*Proof of Proposition 1.* We first characterize the optimal policy that is defined as the solution to the planners problem (16), that is

$$\min_{\mathbf{Y}_t,\mathbf{W}_t,\mathbf{P}_t} \mathcal{L}_t \qquad \text{s.t.} \qquad (15)\ . \tag{33}$$

The Lagrange function for this problem is given by

$$\begin{aligned}\mathsf{L}_t =& \mathbb{E}_t \Big\{ \frac{1}{2}\mathbf{Y}_t'\mathcal{W}\mathbf{Y}_t + \boldsymbol{\mu}_1'(\mathcal{A}_{yy}\mathbf{Y}_t - A_{yw}\mathbf{W}_t - \mathcal{A}_{yp}\mathbf{P}_t - \mathcal{B}_{yx}\mathbf{X}_{-t} - \mathcal{B}_{y\xi}\boldsymbol{\Xi}_t) \\ &+\boldsymbol{\mu}_2'(\mathcal{A}_{ww}\mathbf{W}_t - \mathcal{A}_{wy}\mathbf{Y}_t - \mathcal{A}_{wp}\mathbf{P}_t - \mathcal{B}_{wx}\mathbf{X}_{-t} - \mathcal{B}_{w\xi}\boldsymbol{\Xi}_t)\Big\}\ ,\end{aligned}$$

where $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ denote the Lagrange multipliers. The first order conditions for $\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t$ are given by

$$\begin{aligned}\mathbf{0} &= \mathcal{W}\mathbb{E}_t\mathbf{Y}_t + \mathcal{A}_{yy}'\boldsymbol{\mu}_1 - \mathcal{A}_{wy}'\boldsymbol{\mu}_2 \\ \mathbf{0} &= -\mathcal{A}_{yw}'\boldsymbol{\mu}_1 + \mathcal{A}_{ww}'\boldsymbol{\mu}_2 \\ \mathbf{0} &= -\mathcal{A}_{yp}'\boldsymbol{\mu}_1 - \mathcal{A}_{wp}'\boldsymbol{\mu}_2\ ,\end{aligned}$$

and from Assumption 1 it follows that this system of equations implies a unique solution $\mathbf{P}_t^{e\,\text{opt}}$.

Next, we consider the fictitious policy problem of a policy maker considering deviating from the fixed rule $\phi^0$ with some sequence of policy shocks $\boldsymbol{\epsilon}_t$.

$$\min_{\mathbf{Y}_t,\mathbf{W}_t,\mathbf{P}_t,\boldsymbol{\epsilon}_t} \mathcal{L}_t \qquad \text{s.t.} \qquad (15) \quad \text{and} \quad (17)\ . \tag{34}$$

The Lagrange function for this problem is given by

$$\begin{aligned}\mathsf{L}_t^f =& \mathbb{E}_t \Big\{ \frac{1}{2}\mathbf{Y}_t'\mathcal{W}\mathbf{Y}_t + \boldsymbol{\mu}_1'(\mathcal{A}_{yy}\mathbf{Y}_t - A_{yw}\mathbf{W}_t - \mathcal{A}_{yp}\mathbf{P}_t - \mathcal{B}_{yx}\mathbf{X}_{-t} - \mathcal{B}_{y\xi}\boldsymbol{\Xi}_t) \\ &+\boldsymbol{\mu}_2'(\mathcal{A}_{ww}\mathbf{W}_t - \mathcal{A}_{wy}\mathbf{Y}_t - \mathcal{A}_{wp}\mathbf{P}_t - \mathcal{B}_{wx}\mathbf{X}_{-t} - \mathcal{B}_{w\xi}\boldsymbol{\Xi}_t) \\ &+\boldsymbol{\mu}_3'(\mathcal{A}_{pp}^0\mathbf{P}_t - \mathcal{A}_{py}^0\mathbf{Y}_t - \mathcal{A}_{pw}^0\mathbf{W}_t - \mathcal{B}_{px}^0\mathbf{X}_{-t} - \mathcal{B}_{p\xi}^0\boldsymbol{\Xi}_t - \boldsymbol{\epsilon}_t)\Big\}\ ,\end{aligned}$$

which leads to the first order conditions for $\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t, \boldsymbol{\epsilon}_t$ given by

$$\begin{aligned}\mathbf{0} &= \mathcal{W}\mathbb{E}_t\mathbf{Y}_t + \mathcal{A}_{yy}'\boldsymbol{\mu}_1 - \mathcal{A}_{wy}'\boldsymbol{\mu}_2 - \mathcal{A}_{py}^{0'}\boldsymbol{\mu}_3 \\ \mathbf{0} &= -\mathcal{A}_{yw}'\boldsymbol{\mu}_1 + \mathcal{A}_{ww}'\boldsymbol{\mu}_2 - \mathcal{A}_{pw}^{0'}\boldsymbol{\mu}_3 \\ \mathbf{0} &= -\mathcal{A}_{yp}'\boldsymbol{\mu}_1 - \mathcal{A}_{wp}'\boldsymbol{\mu}_2 + \mathcal{A}_{pp}^{0'}\boldsymbol{\mu}_3 \\ \mathbf{0} &= \boldsymbol{\mu}_3\ .\end{aligned}$$

Since $\boldsymbol{\mu}_3 = \mathbf{0}$, it is easy to verify that the first order conditions of the fictitious policy problem (34) are identical to the first order conditions of the planner's policy problem (33). Since both optimization problems are convex, this means that the two optimization problems have the same solution $\mathbf{P}_t^{e\,\text{opt}}$.

Next, note that Assumption 2 imposes that the rule $\phi^0$ underlying the proposed policy leads to a unique equilibrium. Hence under this rule we can write the solution as

$$\begin{aligned}\begin{bmatrix} \mathbb{E}_t\mathbf{Y}_t \\ \mathbb{E}_t\mathbf{W}_t \\ \mathbf{P}_t^e \end{bmatrix} &= \begin{bmatrix} \mathcal{A}_{yy} & -\mathcal{A}_{yw} & -\mathcal{A}_{yp} \\ -\mathcal{A}_{wy} & \mathcal{A}_{ww} & -\mathcal{A}_{wp} \\ -\mathcal{A}^0_{py} & -\mathcal{A}^0_{pw} & \mathcal{A}^0_{pp} \end{bmatrix}^{-1} \begin{bmatrix} \mathcal{B}_{yx} & \mathcal{B}_{y\xi} & \mathbf{0} \\ \mathcal{B}_{wx} & \mathcal{B}_{w\xi} & \mathbf{0} \\ \mathcal{B}^0_{px} & \mathcal{B}^0_{p\xi} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{X}_{-t} \\ \mathbb{E}_t\mathbf{\Xi}_t \\ \boldsymbol{\epsilon}_t^e \end{bmatrix} \\ &= \begin{bmatrix} \mathcal{C}^0_x & \mathcal{C}^0_\xi & \mathcal{R}^0 \\ \mathcal{C}^0_{wx} & \mathcal{C}^0_{w\xi} & \mathcal{C}^0_{w\epsilon} \\ \mathcal{C}^0_{px} & \mathcal{C}^0_{p\xi} & \mathcal{C}^0_{p\epsilon} \end{bmatrix} \begin{bmatrix} \mathbf{X}_{-t} \\ \mathbb{E}_t\mathbf{\Xi}_t \\ \boldsymbol{\epsilon}_t^e \end{bmatrix} , \end{aligned} \tag{35}$$

where the expressions for the coefficients, e.g. $\mathcal{C}^0_x, \mathcal{C}^0_\xi, \mathcal{R}^0$, can be derived explicitly from inverting the map, but we will not require such expressions: existence as imposed by Assumption 2 is sufficient for our purposes. This representation shows that the fictitious policy problem can be alternatively stated (by substituting out the variables $\mathbf{W}_t, \mathbf{P}_t$ in terms of the shocks and initial conditions) as

$$\min_{\boldsymbol{\epsilon}_t} \mathcal{L}_t \qquad \text{s.t.} \qquad \mathbb{E}_t\mathbf{Y}_t = \mathcal{R}^0\boldsymbol{\epsilon}_t^e + \mathcal{C}^0_x\mathbf{X}_{-t} + \mathcal{C}^0_\xi\mathbb{E}_t\mathbf{\Xi}_t \ .$$

This leads to the first order condition

$$\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = \mathbf{0} \ .$$

This shows that the optimal policy allocation $\mathbf{P}_t^{e\,\text{opt}}$ can be characterized by the condition $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = 0$. Since, $\mathbf{P}_t^{e\,\text{opt}}$ is unique we have that $\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\,\text{opt}}$ if and only if $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 = \mathbf{0}$. The final observation in the theorem follows directly as under $\mathbf{P}_t^{e0}$ we have

$$\mathbf{Y}_t^0 = \mathcal{R}^0\boldsymbol{\epsilon}_t^{e0} + \underbrace{\mathcal{C}^0_x\mathbf{X}_{-t} + \mathcal{C}^0_\xi\mathbb{E}_t\mathbf{\Xi}_t + \mathbf{Y}_t - \mathbb{E}_t\mathbf{Y}_t^0}_{\boldsymbol{\Upsilon}_t^0} ,$$

and since the time $t$ news shocks $\boldsymbol{\epsilon}_t^{e0}$ are orthogonal to the initial conditions and all other shocks we have that $\mathbb{E}(\boldsymbol{\epsilon}_t^{e0}\boldsymbol{\Upsilon}_t^{0'}) = 0$. □

*Proof of Proposition 2.* Part 1: From Proposition 1 it follows that $\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\,\text{opt}}$ if and only if $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 = \mathbf{0}$. Since, $\boldsymbol{\delta}_t^* = -(\mathcal{R}^{0'}\mathcal{W}\mathcal{R}^0)^{-1}\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0$ we have that $\boldsymbol{\delta}_t^* = 0$ if and only if $\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\,\text{opt}}$. Part 2: First, if $\boldsymbol{\delta}_t^* = \mathbf{0}$ the claim follows immediately from Part 1. Hence, suppose that $\boldsymbol{\delta}_t^* \neq \mathbf{0}$ which implies that $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 \neq \mathbf{0}$. From Proposition 1 it follows that $\mathbf{Y}_t^0 = \mathcal{R}^0\boldsymbol{\epsilon}_t^{e0} + \boldsymbol{\Upsilon}_t^0$. Now consider adjusting $\boldsymbol{\epsilon}_t^{e0}$ by $\boldsymbol{\delta}_t$ such to ensure that the first order condition holds. Specifically we solve

$$\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\left(\mathcal{R}^0(\boldsymbol{\epsilon}_t^{e0} + \boldsymbol{\delta}_t) + \boldsymbol{\Upsilon}_t^0\right) = \mathbf{0} ,$$

for $\boldsymbol{\delta}_t$. This gives

$$\boldsymbol{\delta}_t^* = -(\mathcal{R}^{0'}\mathcal{W}\mathcal{R}^0)^{-1}\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 ,$$

i.e. the OPP statistic sets the gradient condition to zero and thus imposes that the policy choice $\boldsymbol{\epsilon}_t^{e0} + \boldsymbol{\delta}_t^*$ is optimal. Finally, note that adding $\boldsymbol{\delta}_t^*$ to $\boldsymbol{\epsilon}_t^*$ is equivalent to adding $\boldsymbol{\delta}_t^*$ to $\mathbf{P}_t^{e0}$ as the model is linear. □

*Proof of Corollary 1.* First from Assumption (2), (35) implies that we can write $\mathbf{P}_t^{e0}$ as

$$
\begin{aligned}
\mathbf{P}_t^{e0} &= \mathcal{C}_{px}^0 \mathbf{X}_{-t} + \mathcal{C}_{p\xi}^0 \mathbb{E}_t \boldsymbol{\Xi}_t + \mathcal{C}_{p\epsilon} \epsilon_t^{e0} \\
&= \boldsymbol{\Theta}_\xi^0 \mathbb{E}_t \mathbf{S}_t + \boldsymbol{\Theta}_\epsilon^0 \boldsymbol{\epsilon}_t^{e0}
\end{aligned}
$$

where $\boldsymbol{\Theta}_\xi^0 = [\mathcal{C}_{p\xi}^0, \mathcal{C}_{px}^0]$, $\boldsymbol{\Theta}_\epsilon^0 = \mathcal{C}_{p\xi}^0$ and $\mathbf{S}_t = (\boldsymbol{\Xi}_t', \mathbf{X}_{-t}')'$. Second, as —$\phi^{\text{opt}}$ leads to a unique equilibrium— we can use the same reasoning and write

$$
\mathbf{P}_t^{e\text{opt}} = \boldsymbol{\Theta}_\xi^{\text{opt}} \mathbb{E}_t \mathbf{S}_t \ ,
$$

where we note that the definition of the optimal policy $\mathbf{P}_t^{e\text{opt}}$ in (16) implies that $\mathbb{E}_t \boldsymbol{\epsilon}_t = 0$ under the optimal rule. The corollary now follows directly from proposition (2) part 2:

$$
\begin{aligned}
\boldsymbol{\delta}_t^* &= \mathbf{P}_t^{e\text{opt}} - \mathbf{P}_t^{e0} \\
&= \boldsymbol{\Theta}_\xi^{\text{opt}} \mathbb{E}_t \mathbf{S}_t - \boldsymbol{\Theta}_\xi^0 \mathbb{E}_t \mathbf{S}_t - \boldsymbol{\Theta}_\epsilon^0 \boldsymbol{\epsilon}_t^{e0} \\
&= (\boldsymbol{\Theta}_\xi^{\text{opt}} - \boldsymbol{\Theta}_\xi^0) \mathbb{E}_t \mathbf{S}_t - \boldsymbol{\Theta}_\epsilon^0 \boldsymbol{\epsilon}_t^{e0} \ .
\end{aligned}
$$

□

*Proof of Corollary 2.* Part 1: From Proposition 1 it follows that $\mathcal{R}^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0 = 0$ if and only if $\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\text{opt}}$. Since, $\mathcal{R}_a^0$ is a subset (or linear combination) of the columns of $\mathcal{R}^0$ it follows that $\mathcal{R}_a^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0 \neq 0$ implies that $\mathbf{P}_t^{e0} \neq \mathbf{P}_t^{e\text{opt}}$. Part 2: Using Proposition 1 we can partition

$$
\begin{aligned}
\mathbf{Y}_t^0 &= \mathcal{R}^0 \boldsymbol{\epsilon}_t^{e0} + \boldsymbol{\Upsilon}_t^0 \\
&= \mathcal{R}_a^0 \boldsymbol{\epsilon}_{a,t}^{e0} + \mathcal{R}_{a^\perp}^0 \boldsymbol{\epsilon}_{a^\perp,t}^{e0} + \boldsymbol{\Upsilon}_t^0 \ ,
\end{aligned}
$$

where $\mathcal{R}_{a^\perp}^0$ denotes the causal effects that cannot be identified. Using this notation we can characterize $\boldsymbol{\delta}_{a,t}^*$ as follows

$$
\boldsymbol{\delta}_{a,t}^* = \underset{\boldsymbol{\delta}_{a,t}}{\operatorname{argmin}} \, \mathcal{L}_t \qquad \text{s.t} \qquad \mathbf{Y}_t = \mathbf{Y}_t^0 + \mathcal{R}_a^0 \boldsymbol{\delta}_{a,t} \ .
$$

Indeed solving this problem gives $\boldsymbol{\delta}_{a,t}^* = -(\mathcal{R}_a^{0'} \mathcal{W} \mathcal{R}_a^0)^{-1} \mathcal{R}_a^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0$. This implies that

$$
\begin{aligned}
\mathcal{L}_t(\mathbf{P}_t^{e0}) &= \frac{1}{2} \mathbb{E}_t \mathbf{Y}_t^{0'} \mathcal{W} \mathbf{Y}_t^0 \\
&\geq \frac{1}{2} \mathbb{E}_t (\mathbf{Y}_t^0 + \mathcal{R}_a^0 \boldsymbol{\delta}_{a,t}^*)' \mathcal{W} (\mathbf{Y}_t^0 + \mathcal{R}_a^0 \boldsymbol{\delta}_{a,t}^*) \\
&= \mathcal{L}_t(\mathbf{P}_t^{e1})
\end{aligned}
$$

where $\mathbf{P}_t^{e1}$ replaces $\mathbf{P}_{a,t}^{e0}$ with $\mathbf{P}_{a,t}^{e0} + \boldsymbol{\delta}_{a,t}^*$. □

Figure 1: Fed funds rate policy in June 1990

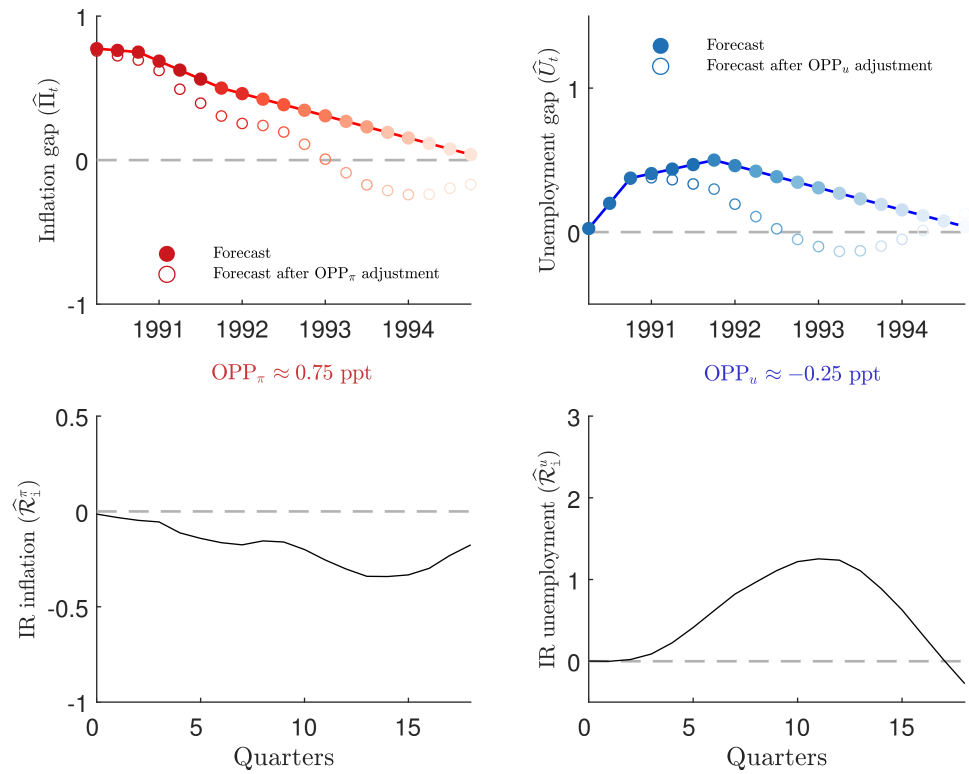


*Notes:* Top panel: median FOMC forecasts for the inflation and unemployment gaps as of 1990-M6. Filled circles denote the forecasts conditional of the policy choice $\mathbf{P}_t^{e\,0}$, and empty circles denote the forecasts after an OPP adjustment. In the left panel, the counter-factual path (red empty circles) is for a policy maker aiming to stabilize only inflation ($\lambda = 0$) and adjusting its policy with $\widehat{\boldsymbol{\delta}}^{\pi}_{\mathsf{i},t}$ (labeled OPP$_\pi$). In the right panel, the counter-factual path (blue empty circles) is for a policy maker aiming to stabilize only unemployment ($\lambda \to \infty$) and adjusting its policy with $\widehat{\boldsymbol{\delta}}^{u}_{\mathsf{i},t}$ (labeled OPP$_u$). Bottom panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock.

Figure 2: Fed funds rate policy in April 2008

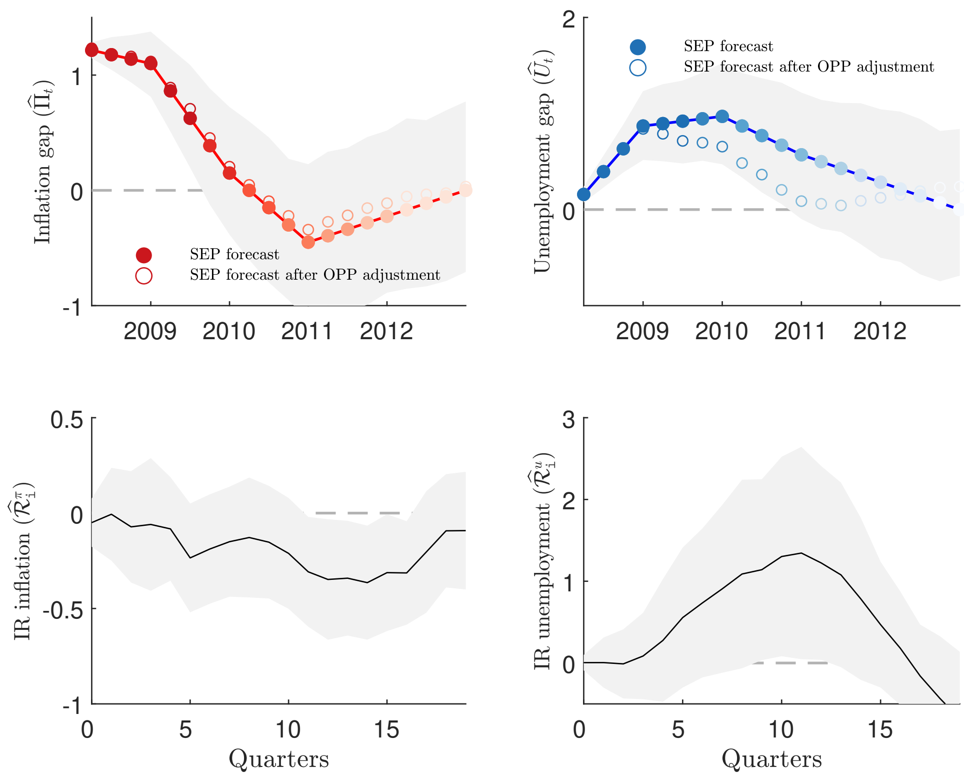


*Notes:* Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2008-M4 (in red and blue) along with the 68 percent confidence bands. Filled circles denote the forecasts conditional of the policy choice $\mathbf{P}_t^{e\,0}$, and empty circles denote the forecasts after the short-rate OPP adjustment $\widehat{\boldsymbol{\delta}}_{\mathrm{i},t}$, the optimal adjustment for a policy maker with a dual inflation–unemployment mandate ($\lambda = 1$). Bottom panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock with 95 percent confidence intervals.

Figure 3: Slope policy in April 2010

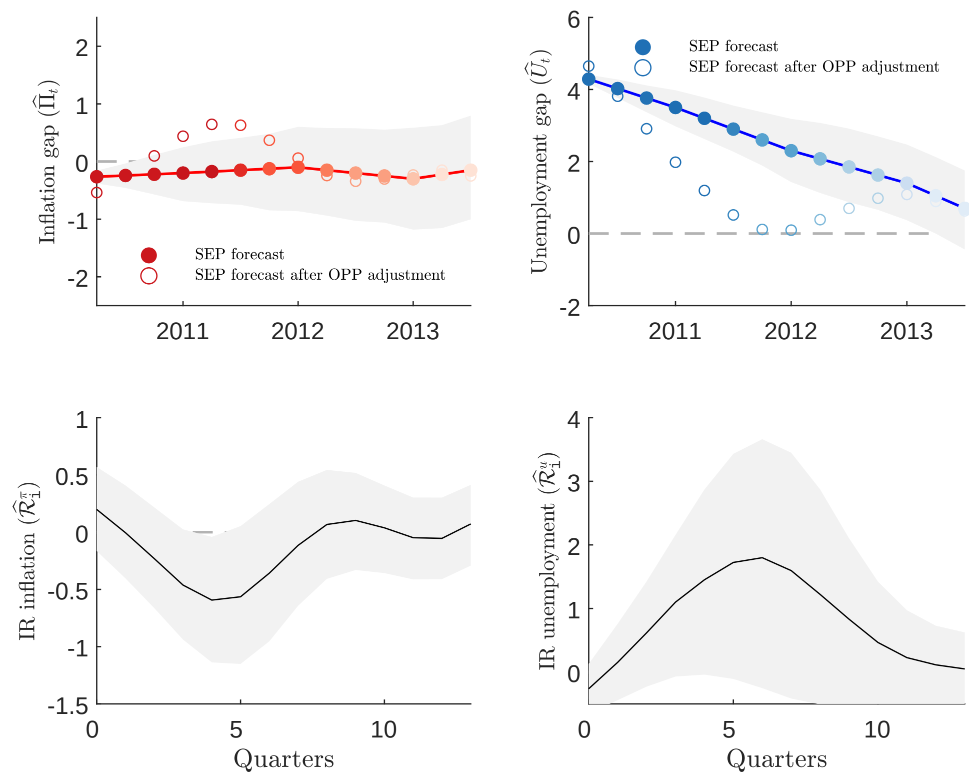


*Notes:* Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2010-M4 (in red and blue) along with the 68 percent confidence bands uncertainty. Filled circles denote the forecasts conditional of the policy choice $\mathbf{P}_t^{e\,0}$, and empty circles denote the forecasts after the slope OPP adjustment $\widehat{\boldsymbol{\delta}}_{\boldsymbol{\Delta},t}$, the optimal adjustment for a policy maker with a dual inflation–unemployment mandate ($\lambda = 1$). Bottom panel: impulse responses of the inflation and unemployment gaps to a slope policy shock with the 95 percent confidence intervals.

Figure 4: A sequence of OPP for Fed monetary policy (1990-2019)

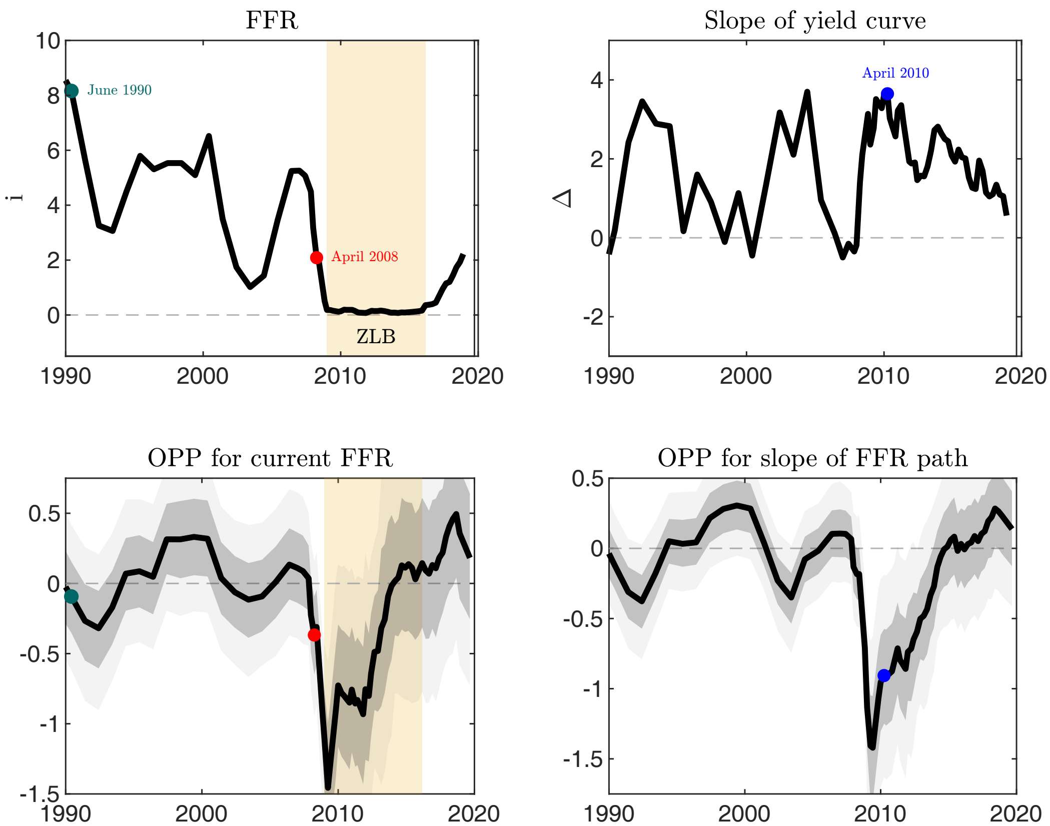


*Notes:* Top panels: the fed funds rate ("FFR", left-panel) and the difference between the 10-year bond yield and the fed funds rate ("Slope of yield curve", right panel). The yellow shaded area denotes the zero-lower bound (ZLB) period. Bottom panels: time series for the two elements of the subset OPP: the short-rate OPP (labeled "OPP for current FFR", left panel) and the slope OPP (labeled "OPP for slope of FFR path", right panel) over 1990-2019 for a policy maker with a dual inflation–unemployment mandate ($\lambda = 1$). The grey areas capture impulse response and model uncertainty at 68% (darker shade) and 95% (lighter shade) confidence. The three case studies are marked as 3 points: June 1990 (green), April 2008 (red) and April 2010 (blue).

Table 1: OPP ESTIMATES FOR CASE STUDIES

| FFR | 1990M6 | 2008M4 | Slope | 2010M4 |
|---|---|---|---|---|
| $\widehat{\delta}_{\mathsf{i},t}$ | -0.11 | -0.37 | $\widehat{\delta}_{\mathbf{\Delta},t}$ | -0.90 |
| | [-0.3$_{-0.5}$ ,0.2 $_{0.5}$] | [-0.6$_{-0,9}$,-0.1$_{0.1}$] | | [-1.2$_{-1.7}$,-0.6$_{-0.1}$] |
| $\widehat{\delta}^{\pi}_{\mathsf{i},t}$ | 0.7 | 0.1 | $\widehat{\delta}^{\pi}_{\mathbf{\Delta},t}$ | -0.1 |
| $\widehat{\delta}^{u}_{\mathsf{i},t}$ | -0.2 | -0.4 | $\widehat{\delta}^{u}_{\mathbf{\Delta},t}$ | -1.3 |

*Notes:* Top row: $\widehat{\delta}_{\mathsf{i},t}$ and $\widehat{\delta}_{\mathbf{\Delta},t}$ denote the mean estimates for the fed funds OPP and the slope OPP for a policy maker with a dual inflation–unemployment mandate ($\lambda = 1$). In brackets, the 68 percent confidence intervals (95 percent in lower case) from impulse response and model uncertainty. Bottom rows: $\widehat{\delta}^{\pi}_{\mathsf{i},t}$ and $\widehat{\delta}^{\pi}_{\mathbf{\Delta},t}$ denote resp. the mean short-rate OPP estimate and the slope OPP estimate for a strict inflation targeter ($\lambda = 0$). Similarly, $\widehat{\delta}^{u}_{\mathsf{i},t}$ and $\widehat{\delta}^{u}_{\mathbf{\Delta},t}$ denote resp. the mean short-rate OPP estimate and the mean slope OPP estimate for a strict unemployment targeter ($\lambda \to \infty$).