Testing Macroeconomic Policies with Sufficient Statistics*

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

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

[(b)] Universitat Pompeu Fabra, Barcelona GSE and VU Amsterdam

February 5, 2021

Abstract

The evaluation of macroeconomic policy decisions typically requires the formulation of a specific economic model. In this work, we present a framework to assess policy decisions with minimal assumptions on the underlying structure of the economy. Given a policy maker's loss function, we propose a statistic —the *Optimal Policy Perturbation* (OPP)— to test whether a policy decision is optimal, i.e., whether it minimizes the loss function. The computation of the OPP does not rely on specifying an underlying model and it can be computed from two sufficient statistics: (i) forecasts for the policy objectives conditional on the policy choice, and (ii) the causal effects of the policy instruments on the policy objectives. We illustrate the OPP by studying past US monetary policy decisions.

JEL classification: C14, C32, E32, E52.

Keywords: Macroeconomic stabilization, Optimal policy, Dynamic causal effects, Forecast targeting.

*We thank 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, 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. 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

Macroeconomic policy decisions have to be made in complex settings, and policy makers often rely on a combination of models, judgment calls, and instinct to decide on policy. Traditionally, the evaluation of such policy decisions has relied on the careful analysis of a *specific* economic model.[1] While this approach can produce great insights, a worry is that assumed model structure may be too stylized, relative to the complexity of the economy, to evaluate policy decisions made in practice (e.g., Blanchard, 2018).

In this work, we propose a framework to evaluate macroeconomic policy decisions with minimal assumptions on the underlying economic model. Given a policy maker's loss function, we construct a statistic —the *Optimal Policy Perturbation*, OPP— to detect "optimization failures" in the policy decision process, i.e., instances when the policy decision does not minimize the loss function. The framework can be applied to a broad range of macro policy problems encountered in practice, such as a central bank interested in stabilizing both inflation and unemployment, or a government interested in smoothing business cycle fluctuations but concerned about excessive deficits. The OPP can then be used as a tool to help policy makers in real time with their decision making process or as a tool to retrospectively study past policy decisions.

Our starting point is a high-level quadratic loss function, as specified by a policy maker, for instance a central banker interested in minimizing the squared deviations of inflation and unemployment from some target levels. The idea underlying our approach is to explore whether deviating from the current policy choice is desirable, i.e., whether a perturbation to the policy instruments can lower the loss function. At the optimum, a perturbation should have no first-order effects on the loss function: the gradient of the loss function should be zero. If this is not the case, we will conclude that the policy is not set optimally.

The Optimal Policy Perturbation (OPP) is the gradient of the loss function, evaluated at the proposed policy choice and rescaled appropriately. If the policy choice is optimal, the gradient is zero and so is the OPP statistic. This property will form the basis of our approach to assessing the optimality of a given policy. In addition, thanks to an appropriate rescaling the OPP statistic has an economic meaning, and it can be interpreted as the magnitude of the deviation from optimality. Specifically, the OPP statistic is the discretionary adjustment to the policy instruments that would correct the optimization failure.

A key benefit of our framework is that the OPP statistic can be computed even if the *specific* underlying economic model is unknown. The reason is that the OPP only depends on two typically known or estimable sufficient statistics: (i) the forecasts for the policy objectives conditional on the policy choice, and (ii) the dynamic causal effects of the policy instruments

[1] See Chari, Christiano and Kehoe (1994) and Woodford (2003) for prominent examples in the context of fiscal and monetary policy.

on the policy objectives. Conditional forecasts are routinely constructed by policy makers as part of the policy decision process. The causal effects of the policy instruments can be estimated – under appropriate assumptions – using methods from the treatment-evaluation literature, most notably instrumental variable methods.

In practice, the OPP cannot be measured exactly, because causal effects estimates and forecasts are uncertain. In particular, causal effects estimates face the usual estimation uncertainty, and the policy makers' conditional forecasts can be mis-specified and thus face mis-specification uncertainty. Because of these two sources of error, our evaluation of a policy choice will resemble a hypothesis test: a statement about whether we can reject the null of optimality at some level of confidence. In economic terms, the test allows to make claims such as "With $X\%$ confidence, the proposed policy choice is not appropriate". Another practical consideration is that the preferences of policy makers —how they weigh different policy objectives in the loss function— may not be explicitly available. To avoid that the evaluation of policy decisions depends on ad hoc choices for such parameters, we propose a robust approach to testing macro policies. Specifically, when preference parameters are unknown, we propose to use the preference parameters that are least favorable to rejecting optimality (on average) over a sequence of past macro decisions.

The non-optimal policies that we detect are those policies that do not minimize the loss function. Clearly, if the world was described by a specific macro model like a New-Keynesian model, such failures should not occur, because the policy maker could simply solve the optimization problem. In practice however, the underlying model is highly complex, forcing policy makers to rely on a combination of models, judgment calls, and instinct to decide on policy. This heuristic approach is not guaranteed to reach an optimum, and the goal of the OPP test is to identify instances where the policy choice could be improved, all the while making minimal modeling assumptions and thus preserving the ability of the policy maker to incorporate a large amount of information, both quantitative and qualitative, into the decision making process.

An optimization failure could be due to a one-time optimization failure, but it could also be the sign of a systematic optimization failure, whereby the policy maker systematically under- or over-reacts to some variables. A *single* OPP statistic can detect both types of optimization failures, but cannot distinguish between them. However, a *sequence* of OPP statistics can separate the two sources of optimization failures, i.e. inform whether systematic mistakes were made. Intuitively, if the systematic conduct of policy is optimal, an OPP sequence should not display any systematic, i.e., predictable, movements. This provides a testable moment condition for detecting systematic optimization failures. We thus propose a second test to make claims such as "A systematically stronger/weaker policy response to movements in X would be more appropriate to achieve the policy maker's objectives".

To clarify the working of the OPP statistic and illustrate its usefulness for policy makers we conduct two exercises in the context of monetary policy decisions where the policy maker is the central bank.

First, we illustrate the properties of the OPP in the standard New Keynesian model (e.g. Galí, 2015). This is a theoretical exercise that shows that *if* the economy can be described by the equations of the New Keynesian model the OPP statistic can (i) detect optimization failures and (ii) determine whether they are due to a non-optimal systematic component of monetary policy, i.e., to a non-optimal reaction function.

Second, in an empirical study we assess US monetary policy over the 1990-2018 period. We summarize the Fed's monetary policy instruments into two groups: a first one captures conventional monetary policy and operates through the fed funds rate; and a second one, available since 2007, captures a broad class of unconventional monetary policies that operate through the slope of the yield curve, as in Eberly, Stock and Wright (2019). We estimate the dynamic causal effects of interest using external instruments derived from changes in asset prices around FOMC announcements (Kuttner, 2001; Gürkaynak, Sack and Swanson, 2005).

We find several instances in which the Fed's monetary policy decisions were non-optimal. Most notably, *given* the information available in early 2008 on the eve of the Great recession, we can reject that fed funds rate policy was set optimally at the time, with the OPP calling for more aggressive interest rate cuts. We can also reject the optimality of unconventional monetary policy operations in the middle of the Great recession, with the OPP calling for a more aggressive use of unconventional policy measures LSAP or QE to lower the slope of the yield curve.

As last exercise, we then use sequences of OPP statistics over 1990-2018 to study the optimality of the systematic component of Fed policy. While we do not find any systematic relationship between the OPP and inflation, we do find that unemployment systematically affects the OPP statistic. This points to a non-optimal systematic monetary policy function in that a more aggressive response to unemployment fluctuations could have been more appropriate to achieve the Fed's objectives.

The remainder of this paper is organized as follows. We continue the introduction by carefully 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 detect optimization failures with minimal assumptions. Section 3 formally introduces the environment in which the policy maker and the researcher operate. Section 4 presents the OPP statistic and discusses its theoretical properties. These properties are further illustrated for a New Keynesian model in Section 5. Inference for the OPP approach is discussed in Section 6. In Section 7 we apply our methodology to empirically study monetary policy decisions in the US. Section 8 concludes and provides some potential avenues for further research.

Relation to literature

Since Lucas (1976), the literature on macroeconomic policy evaluation has largely focused on the "ex-ante" analysis of the optimal allocation in the context of fully-specified forward-looking economic models. This often involves solving the Ramsey policy problem and finding simple policy rules that can approximate the Ramsey allocation (e.g., Chari, Christiano and Kehoe, 1994; Woodford, 2010; Michaillat and Saez, 2019). In this context, an important agenda has been to derive, from first principles, the appropriate loss function that the policy maker *should* be considering. In this context popular questions include the desirability of price level targeting versus inflation targeting, the optimal rate of inflation, or the desirability of a single price stability mandate versus a dual inflation-unemployment mandate (e.g. Woodford, 2003; Schmitt-Grohé and Uribe, 2010; Coibion, Gorodnichenko and Wieland, 2012; Debortoli et al., 2019).

This paper takes a different starting point. We take the policy maker's objectives as given, and we propose a methodology for assessing whether a proposed policy choice is the most appropriate to achieve the policy maker's objectives. Hence our methodology is complementary to the existing literature that focuses on defining the appropriate objectives of the policy maker.[2] Because we rely on minimal assumptions on the underlying model, the method can be used in the context of macroeconomic policy decisions that do not rely on one specific macroeconomic model and are based on heuristics; combining multiple models, judgment calls and instincts. As far as we know, there are no alternatives in the literature (e.g. Bénassy-Quéré et al., 2018; Kocherlakota, 2019).

Importantly, and different from the earlier reduced-form macroeconometric literature, the OPP approach takes into account the limits of non-structural approaches such as ours (e.g. Lucas, 1976; Sargent, 1981) and relies on a proper identification of the causal effects of changes in policy (e.g. Sims, 1980). In a nutshell, the OPP framework is valid *even* under the Lucas critique, because the OPP is only used to test the optimality of a policy choice. As such, the OPP framework does not require "counter-factuals", which could be subject to the Lucas critique (absent a micro-founded structural model). Instead, the OPP only requires the gradient of the loss function evaluated at the current policy choice, that is the effect of an infinitesimally small change in policy, which is not subject to the Lucas (1976) critique.

The OPP framework shares important similarities with the sufficient statistic approach (e.g. Chetty, 2009; Kleven, 2020), in that both methods exploit the fact that the consequences of a policy can be derived from high-level elasticities: the causal effects of moving the policy

[2]An alternative, equally valid, starting point for the OPP is to consider a researcher with her own view of what the policy maker's objectives *should be*. In that context, the OPP framework can be used to assess whether a policy maker's decision is most appropriate to minimize that loss function.

instruments. Our approach thus rests on the "estimability" of these elasticities, i.e., on the possibility to use quasi-experimental variations to infer the causal effects of the policy instruments, just as in the sufficient-statistic literature. In a macro context, this requires being in a stable environment, a stable macro environment and a stable policy regime, for some period prior to the policy decision. Different from the sufficient-statistic literature however, the OPP framework exploits another statistic —the policy maker's forecasts— to bypass the need for a fully specified model in order to find the equilibrium allocation under the desired policy choice. This additional information allows us to evaluate macro policy decisions without committing to one specific model and the method can be used without any change to current operating procedures. Using policy makers' forecasts as an input is unusual in the optimal policy literature, but in practice it merely amounts to replacing the structural model's forecast with the policy maker's forecast (which can involve structural and reduced-form models, model combination, judgment, etc..). While the jury is still out on determining the best forecasting method, we note that policy makers' forecasts often do perform well (e.g., Romer and Romer, 2000; Sims, 2002) and are thus natural alternatives to evaluate the expected equilibrium allocation under the desired policy choice.

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.

Finally, the OPP framework can be useful in the context of forecast-targeting rules often used by policy makers, notably for monetary policy (e.g., Svensson, 2019). A forecast targeting rule is a general approach to policy making that consists in selecting a policy rate and policy-rate 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, 2017, 2019).[3] 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. The OPP can precisely quantify that "looking good" criterion while imposing minimal assumptions on the underlying economic model.[4]

[3]As argued by Svensson, the forecast targeting approach is attractive for its flexibility and capacity to incorporate all relevant information and to accommodate judgment adjustments. This is in contrast to Taylor-type rules that can be "too restrictive and mechanical, not taking into account all relevant information, and the ability to handle the complex and changing situations faced by policy makers" (Svensson, 2017).

[4]For instance, the OPP is immediately applicable to Bernanke (2015)'s interpretation of the Fed's rule of conduct: "The Fed has a rule. The Fed's rule is that we will go for a 2 percent inflation rate; we will go for the natural rate of unemployment; we put equal weight on those two things; we will give you information about our projections, our interest rate. That is a rule and that is a framework that should clarify exactly

2 A simple example

Before formally describing our general framework, we first informally present a simple example to illustrate how we can evaluate macroeconomic policies with minimal assumptions on the underlying economic model.

Consider a central bank with an equally weighted dual mandate

$$\mathcal{L} = \frac{1}{2}\left(\pi^2 + x^2\right) \ ,$$

where π is inflation and x some real activity mandate. The central bank has one instrument p, e.g. the short-term interest rate, that affects each mandate according to

$$\underbrace{\begin{bmatrix} \pi \\ x \end{bmatrix}}_{Y} = \underbrace{\begin{bmatrix} f_\pi(p) \\ f_x(p) \end{bmatrix}}_{f(p)} , \tag{1}$$

where the function $f(p)$ captures how policy affects the vector of mandates $Y = (\pi, x)'$. The central bank proposes implementing the policy p^0, which implies the equilibrium (Y^0, p^0) with $Y^0 = f(p^0)$.

Consider a researcher interested in testing whether (Y^0, p^0) minimizes the loss function. Crucially, the function $f(.)$ is not available, either because it is unknown to the researcher, or because it is too complex to write down.

Our approach rests on the idea that, at the optimum, the gradient of the loss function should be zero —$\nabla_p \mathcal{L}|_{p=p^0} = 0$— and we will assess the optimality of p^0 by directly computing the gradient

$$\nabla_p \mathcal{L}|_{p=p^0} = \mathcal{R}^{0\prime} Y^0 \ , \qquad \text{where} \qquad \mathcal{R}^0 = \left.\frac{\partial f(p)}{\partial p}\right|_{p=p^0} . \tag{2}$$

At the optimum, $\mathcal{R}^0$ —the causal effect of a marginal change in policy— should be orthogonal to Y^0 —the allocation at p^0—, i.e., $\mathcal{R}^{0\prime} Y^0 = 0$ as illustrated in Figure 1(a).[5]

The simple insight underlying this paper is that, even if the researcher does not have complete knowledge of the model $f(\cdot)$, computing the gradient at p^0 is possible because (i) causal effects $\mathcal{R}^0$ are estimable, and (ii) Y^0 is typically available from the policy maker. Specifically, methods from the treatment-effect literature, most notably instrumental variable

---

what the Fed is doing." In that context, the OPP can be used to assess whether the Fed is indeed doing the best it can to satisfy its two objectives.

[5] At the optimum $\mathcal{R}^0_\pi \pi^0 + \mathcal{R}^0_u u^0 = 0$. If a marginally higher p^0 lowers the loss function by stabilizing inflation ($\mathcal{R}^0_\pi \pi^0 < 0$), the effect is exactly compensated by a destabilizing effect on unemployment ($\mathcal{R}^0_u u^0 = -\mathcal{R}^0_\pi \pi^0 > 0$), leaving welfare unchanged.

Figure 1: Illustration of the gradient test

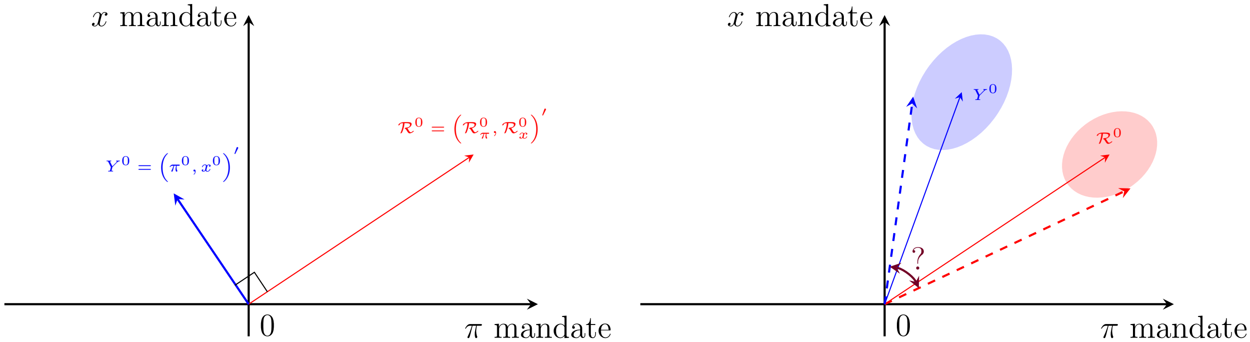

(a) The orthogonality condition

(b) An inference problem

Notes: The vector Y^0 collects the values of the π and x mandates at the policy choice p^0. The vector $\mathcal{R}^0$ collects the marginal effects of a policy change on π and x. Left panel: At the optimum, $\mathcal{R}^0$ should be orthogonal to Y^0. Right panel: The filled ellipses denote uncertainty (at some confidence level X) in $\mathcal{R}^0$ and Y^0. With uncertainty, we can only infer the presence of an optimization failure (at confidence level X), when the largest angle (modulo 180°) between $\mathcal{R}^0$ and Y^0 does not include 90°.

methods, can be adopted to estimate $\mathcal{R}^0$ with minimal assumptions. Further, the allocation Y^0 implied by the policy choice p^0 is typically published by policy makers, who routinely construct and publish their conditional forecasts.[6]

Thanks to these two pieces of information, it is possible to compute the relevant gradient, and thus assess the optimality of the policy choice with minimal assumptions on the underlying model. In contrast, in the standard approach to macro policy evaluation the researcher would have to rely on a complete specification of $f(\cdot)$ in order to determine the optimal allocation and compare such allocation to (p^0, Y^0).

In practice, both $\mathcal{R}^0$ and Y^0 are not known exactly, and the gradient in (2) can only be computed with uncertainty: (i) $\mathcal{R}^0$ must be estimated and thus faces *estimation uncertainty*, and (ii) the policy maker faces *model uncertainty* and thus may not report the exact Y^0. As a result, 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.

Figure 1(b) illustrates the effect of uncertainty on our ability to assess optimality from the orthogonality condition between $\mathcal{R}^0$ and Y^0. The filled ellipses represent the uncertainty (at some confidence level X) in our estimates of the vectors $\mathcal{R}^0$ and Y^0. With uncertainty, we can only infer the presence of an optimization failure (at confidence level X), when the *largest* angle (the dotted arrows) between $\mathcal{R}^0$ and Y^0 does not include 90°. In the example displayed in Figure 1(b), the largest angle does not include 90°, so we can conclude that there is an optimization failure at confidence level X. Increasing uncertainty in the estimate of $\mathcal{R}^0$ and in the policy maker's ability to report the accurate Y^0 reduces the power of our

[6] We do not have expectations in this toy model, but they will figure prominently in our general framework.

test. To see this, just imagine increasing the size of the ellipses in Figure 1(b); at some point the dotted arrows will become orthogonal to each other making it impossible to reject that the policy is optimal.

The remainder of this paper will now develop these simple ideas for a very generic class of dynamic models that encompasses most macro models encountered in the literature, but without committing to a particular one.

3 Environment

In this section we describe the policy maker's objectives and instruments, as well as the generic structure of the economy. There are three players in our setting: the policy maker that makes an initial policy choice, the researcher that aims to verify whether the policy maker's choice is optimal, and nature that determines the distribution of the variables.

The loss function

A policy maker aims to stabilize an $M \times 1$ vector of policy objectives y_t, for instance inflation, unemployment or GDP growth. We posit that the policy maker's loss function takes the form:[7]

$$\mathcal{L}_t = \mathbb{E}_t \sum_{h=0}^{H} \sum_{m=1}^{M} \lambda_m \beta_h \left(y_{m,t+h} - y^*_{m,t+h} \right)^2 \ , \tag{3}$$

where $y_{m,t+h}$ denotes the value of policy objectives $m = 1, \ldots, M$ at horizon $h = 0, \ldots, H$. The horizon H is arbitrary and can be considered infinite.

The target value for $y_{m,t+h}$ is denoted by $y^*_{m,t+h}$ and preferences across variables and horizons are captured by λ_m —the weight on variable m— and β_h —the discount factor for horizon h. The expectation is taken with respect to the time t information set $\mathcal{F}_t$, e.g. $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_t)$, which we define below. The preference parameters λ_m and β_h are taken as given for now, but in Section 6 we discuss how they can be estimated.

The policy instruments

The policy maker has J instruments to minimize the loss function.[8] For each instrument $j \in 1, \ldots, J$, the policy maker determines a "policy plan" which sets the instrument's value at time t as well as its expected path $p_{j,t|t}, \ldots, p_{j,t+H|t}$, $j = 1, \ldots, J$, where $p_{j,t+h|t}$ denotes

[7] A quadratic specification can be seen as an approximation of a more general loss function for small deviations of the policy objectives from their target.

[8] For instance, a fiscal policy maker sets taxes, government spending and transfers, or a monetary policy maker sets the short-term interest rate and chooses how much longer maturity securities to buy/sell.

the level of instrument j for period $t+h$ that is proposed at time t. We stack the different policy plans in the $K \times 1$ policy vector

$$p_t = \left(p_{1,t|t}, \ldots, p_{1,t+H|t}, \ldots, p_{J,t|t}, \ldots, p_{J,t+H|t}\right)' \ ,$$

which implies that $K = J(H+1)$.

Without loss of generality, the policy vector p_t can be written as the sum of two components: (i) a *predictable* (or systematic) component based on time t observables, and (ii) an *unpredictable* (or discretionary) component:[9]

$$p_t = g(y_t, X_t) + \epsilon_t \ , \tag{4}$$

where g is a function that depends on the targets $y_t = (y_{1,t}, \ldots, y_{M,t})'$ and possibly additional observables X_t. The $(K \times 1)$ vector ϵ_t —the discretionary component— is by construction uncorrelated with the variables y_t and X_t. In this paper we do not impose a priori restrictions on the reaction function g, which can be an arbitrarily complex function of time t observables.[10]

A generic economic model

We postulate a generic model to describe the behavior of the target deviations $y_{m,t+h} - y^*_{m,t+h}$. For convenience, we stack all policy objectives at all horizons in the $M(H+1) \times 1$ vector $Y_t = [y_{m,t+h} - y^*_{m,t+h}]_{m=1,\ldots,M,h=0,\ldots,H}$ and simply refer to this vector as the targets. We emphasize that this vector depends on future observations, but to keep the notation minimal we denote it simply by Y_t.

The generic model for the target is given by

$$Y_t = \mathcal{R}(g)p_t + f(X_t; g) + \xi_t \ , \tag{5}$$

which depends on three components: (i) the $M(H+1) \times K$ matrix $\mathcal{R}(g)$ that captures the

[9] In this paper, we do not take a stand on whether policy makers have an interest or not in making their policy decisions predictable. We simply note that a large literature has emphasized the benefits of communication and transparency for policy makers (e.g., Blinder et al., 2008). In that context, a policy maker could want to make policy as predictable as possible (i.e., minimize the contribution of ϵ_t to the variance of p_t). For instance, monetary policy makers do try to convey their reaction function, but do not consider themselves bound by it (e.g., Bernanke, 2015; Kocherlakota, 2016, in the context of monetary policy)

[10] To give a concrete example, Taylor (1993) showed that the central bank's interest rate policy could be well described with a simple "Taylor rule" with $g(y_t, X_t) = (\phi_x, \phi_\pi)(x_t, \pi_t)' + X_t$ where $X_t = r^*_t$ the equilibrium real interest rate and x_t and π_t the deviation of economic activity and inflation from their targets. The same function $g(.)$ could describe the expected interest at horizon h with $i_{t+h|t} = \mathbb{E}_t g(y_{t+h}, X_{t+h}) + \epsilon_{t,t+h}$ and where $\epsilon_{t,t+h}$ denote expected future discretionary adjustments, like forward guidance for instance (e.g. McKay, Nakamura and Steinsson, 2016).

causal effects of the K policy variables in p_t on the $M(H+1)$ target variables in Y_t, (ii) an arbitrary function $f(\cdot)$ which depends on additional time t measurable variables X_t, and (iii) current and future shocks ξ_t. The structure of the economy, as captured by $\mathcal{R}(g)$ and $f(X_t; g)$, may depend on the reaction function g of the policy maker.

We stress that the linearity assumption – with respect to p_t – is made for convenience only and the main results of this paper continue to apply for more general nonlinear models.[11] Moreover, we note that many commonly used recursive macroeconomic models, such as vector autoregressive models and dynamic stochastic general equilibrium models, can be expressed as special cases of model (5) simply by iterating forward.[12]

Based on the generic model we can now properly define the information set for (3) by setting $\mathcal{F}_t = \{y_s, p_s, X_s, s \leq t\}$ where X_s includes any, time-s measurable, variables other than y_s and p_s. The current variables (y_t, p_t) can still be adjusted by the time t policy choice, but the past variables (y_s, p_s), with $s < t$, are fixed as their outcomes have been realized in the past.

The policy maker's choice

The policy maker's loss function can be conveniently expressed as

$$\mathcal{L}_t = \mathbb{E}_t \|\mathcal{W}^{1/2} Y_t\|^2 \ , \qquad Y_t = \mathcal{R}(g) p_t + f(X_t; g) + \xi_t \ , \tag{6}$$

where $\|\mathcal{W}^{1/2} Y_t\|^2 = Y_t' \mathcal{W} Y_t$ and $\mathcal{W} = \text{diag}(\beta \otimes \lambda)$ denotes a diagonal matrix of preferences, with $\beta = (\beta_0, \dots, \beta_H)'$ and $\lambda = (\lambda_1, \dots, \lambda_M)'$.

The policy maker's *perceived* solution for minimizing the loss function $\mathcal{L}_t$ in (6) is the pair (g^0, ϵ_t^0) comprising g^0 —the policy maker's reaction function— and ϵ_t^0 —a vector of discretionary adjustments to the reaction function. Based on these policy choices the expected equilibrium allocations $(\mathbb{E}_t Y_t^0, p_t^0)$ are given by

$$\begin{cases} \mathbb{E}_t Y_t^0 = \mathcal{R}^0 p_t^0 + \mathbb{E}_t f(X_t; g^0) + \mathbb{E}_t \xi_t \\ p_t^0 = g^0(y_t^0, X_t) + \epsilon_t^0 \end{cases} , \tag{7}$$

[11] Specifically, in Appendix A we show that for $Y_t = f(p_t, X_t; g) + \xi_t$, where $f(\cdot, \cdot; \cdot)$ is some nonlinear function of p_t and possibly additional variables X_t, we can derive a similar test statistic that allows to detect optimization failures. This extension can be important for certain applications, but conceptually nothing is lost by considering the partially linear model (5).

[12] An important example is the simultaneous equations model

$$\mathcal{A}(g) Y_t = \mathcal{B}(g) p_t + \tilde{\xi}_t \ ,$$

which can be expressed as

$$Y_t = \underbrace{\mathcal{A}(g)^{-1} \mathcal{B}(g)}_{\mathcal{R}(g)} p_t + \underbrace{\mathcal{A}(g)^{-1} \tilde{\xi}_t}_{\xi_t} \ .$$

where $\mathbb{E}_t \epsilon_t^0 = 0$ and $\mathcal{R}^0 \equiv \mathcal{R}(g^0)$ are the dynamic causal effects implied by the reaction function g^0. We will generally impose that $\mathcal{R}^0$ has full column rank.

Defining an optimization failure

The policy choice (g^0, ϵ_t^0) that led to the policy vector p_t^0 may not be optimal in that the expected equilibrium $(\mathbb{E}_t Y_t^0, p_t^0)$ given by equation (7) may not minimize the policy maker's loss function (6). This can happen for a variety of reasons: (i) the policy maker used a sub-optimal reaction function, (ii) the policy maker made sub-optimal discretionary adjustments, (iii) both the reaction function and the discretionary adjustments are sub-optimal.

To formally define optimization failures we need to assume the existence of at least one policy choice that minimizes the loss function.[13] To do so let $\mathcal{G}$ be an arbitrary class of reaction functions with $g^0 \in \mathcal{G}$.

Assumption 1. (Existence of optimum)
There exists a non-empty set $\mathcal{G}^{\text{opt}}$ such that

$$\mathbb{E}_t \|\mathcal{W}^{1/2} Y_t(g, 0)\|^2 \leq \mathbb{E}_t \|\mathcal{W}^{1/2} Y_t(\tilde{g}, \tilde{\epsilon}_t)\|^2, \qquad \forall\ g \in \mathcal{G}^{\text{opt}},\ \tilde{g} \in \mathcal{G} \backslash \mathcal{G}^{\text{opt}}, \tilde{\epsilon}_t \neq 0\ ,$$

where $Y_t(g, \epsilon_t) = \mathcal{R}(g) p_t + f(X_t; g) + \xi_t$ and $p_t = g(y_t, X_t) + \epsilon_t$.

The assumption simply ensures that the optimization problem is well-posed: there exists a non-empty set of reaction functions $\mathcal{G}^{\text{opt}}$, which minimize the loss function (3). For ease of exposition, we normalized the optimal policy choice to be of the form $(g, 0)$, i.e., the discretionary adjustments corresponding to $g \in \mathcal{G}^{\text{opt}}$ are set equal to zero. Importantly, this is simply a normalization of the optimal solution, since we do not restrict the class $\mathcal{G}$ the function g can be arbitrarily complex. In fact, the only requirement is that there exists at least one pair (g, ϵ_t) for which the gradient of the loss function is zero.

Based on Assumption 1, we can now define an optimization failure.

Definition 1. (Optimization failure)
An optimization failure is a policy choice (g, ϵ_t) such that either $g \notin \mathcal{G}^{\text{opt}}$ or $\epsilon_t^0 \neq 0$.

Detecting such optimization failures is relevant for any policy maker who is not constrained by commitments from before time t, such as following a policy rule or honoring past promises. We view this as the baseline case, and in Appendix B we show that our methodology can be generalized to detect optimization failures in the presence of commitments from the past (e.g. Kydland and Prescott, 1977; Barro and Gordon, 1983; McKay, Nakamura and Steinsson, 2016).

[13] Equations (5) and (4) merely provide a generic description of the economy and hence existence is not guaranteed by the model.

4 Detecting optimization failures

In this section we take the perspective of a researcher interested in assessing whether the policy maker's choice (g^0, ϵ_t^0) is optimal, i.e. in detecting a possible *optimization failure*, and we derive necessary conditions under which the policy maker's choice is (sub-)optimal.

We assume that the researcher does not have access to the functions $f(\cdot; g^0)$ and $g^0(\cdot, \cdot)$, either because they are unknown to the researcher or because they are too complex to write down explicitly. Of course, if the functions f and g^0 were known and could be written down explicitly, one could first minimize the loss function with respect to Y_t, subject to the constraints imposed by the model, and then characterize this solution in terms of the policy variables to obtain an optimal reaction function. That solution could then be compared to p_t^0, the policy choice of the policy maker, in order to assess the optimality of p_t^0. This approach is the traditional route followed by the literature in the context of fully specified macro models (e.g. Chari, Christiano and Kehoe, 1994; Woodford, 2003).

In practice however the economy is a complex system, and the assumption that one can explicitly specify the underlying economic model may be too strong. Assessing whether the policy choices are appropriate or optimal in this context requires a different approach.

In this section, we introduce the OPP statistic and derive its theoretical properties, notably its ability to detect an optimization failure. Second, we show that a sequence of OPP statistics can be used to detect *systematic* optimization failures, that is optimization failures coming from a non-optimal reaction function.

4.1 The OPP statistic

The idea underlying our approach is to explore whether deviating from the current policy vector is desirable, i.e., whether a perturbation $\delta_t = (\delta_{1,t}, \ldots, \delta_{K,t})'$ to the policy choice p_t^0 can lower the loss function. At the optimum, a perturbation should have no first-order effects on the loss function: the gradient of the loss function should be zero. If this is not the case, we will conclude that the policy is not set optimally.

The policy perturbation that we use is the (rescaled) gradient of the loss function evaluated at p_t^0, i.e.,

$$\delta_t^* = -(\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 \ , \tag{8}$$

where $Y_t^0 = \mathcal{R}^0 p_t^0 + f(X_t; g^0) + \xi_t$. For reasons that will become clear shortly, we refer to δ_t^* as the *optimal policy perturbation*, or OPP. Under model (5), it is easy to verify that the OPP is indeed proportional to the gradient at p_t^0 with the scaling factor equal to the inverse of the Hessian:[14]

$$\delta_t^* = -(\nabla_{p_t}^2 \mathcal{L}_t^0)^{-1} \nabla_{p_t} \mathcal{L}_t^0 \ . \tag{9}$$

The next proposition formalizes two attractive properties of the OPP.

Proposition 1. *Given an economy defined by equations* (4) *and* (5)*, we have that under Assumption 1:*

1. $\delta_t^* \neq 0$ *implies that* $g^0 \notin \mathcal{G}^{\text{opt}}$ *or* $\epsilon_t^0 \neq 0$*;*
2. $p_t^0 + \delta_t^* = p_t^*$*, where* $p_t^* = \arg\min_{p_t \in \mathbb{R}^K} \mathbb{E}_t \|\mathcal{W}^{1/2}(\mathcal{R}^0 p_t + f(X_t; g^0) + \xi_t)\|^2$.

All proofs are provided in Appendix C. We now discuss the intuition behind these two properties of the OPP.

Intuition: Detecting optimization failures from the OPP

By being proportional to the gradient of the loss function, the intuition that was discussed in Section 2 carries over to the OPP: If δ_t^* is not equal to zero, it means that the gradient is non-zero and there is an optimization failure, as stated in the first part of Proposition 1. The optimization failure could come from a non-optimal reaction function, a non-optimal discretionary adjustment, or both.

The OPP is based on the same idea as a standard score test, or Lagrange multiplier test. Under the null, the gradient (or score) should be zero. The important benefit of a score test is that it does not require the estimation of the parameters that are fixed under the null. In our case, the null is that the policy choice is optimal with $p_t^0 = g^0(y_t^0, X_t)$. Thus, p_t^0 is "fixed under the null", and there is no need to estimate the optimal policy. This feature is the reason why we can avoid strong modeling assumptions, and why we do not need to specify the underlying economic model or to estimate the reaction function g.

The same feature is the reason why the OPP approach is not subject to the Lucas critique. Detecting a non-optimal policy only requires knowing the gradient of the loss function, that is the effect of an infinitesimally small change in policy on the loss. But infinitesimally small changes in policy do not shift agents beliefs and thus does not generate expectations-formation effects of the kind Lucas (1976) emphasizes (e.g. Leeper and Zha, 2003).

Figure 2: The OPP policy thought experiment

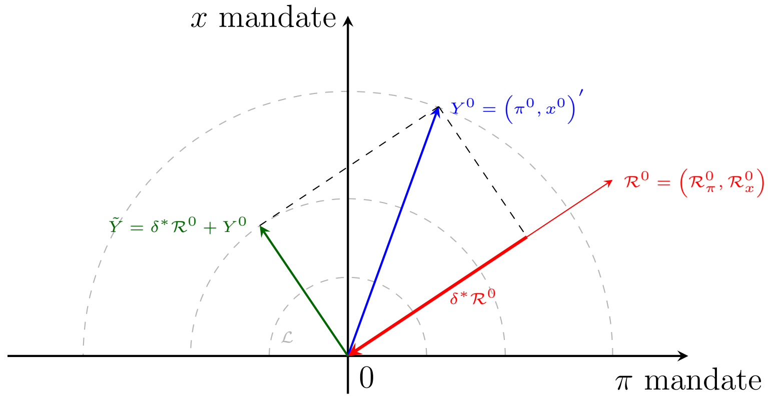

Notes: Starting from a policy choice p^0 with allocation Y^0 and loss $\mathcal{L}^0 = \|Y^0\|^2$, the OPP thought experiment consists in using $\mathcal{R}^0$ —the marginal effect of a policy adjustment— to find the change in policy that would most "shrink" the vector Y^0 (and thereby the loss $\mathcal{L}^0$). Geometrically, this amounts to projecting *out* Y^0 on $\mathcal{R}^0$. The "residual" is the vector $\tilde{Y} = \delta^* \mathcal{R}^0 + Y^0$, the allocation after adjustment by the OPP. The dashed grey curves depict isoloss curves $\mathcal{L} = \frac{1}{2}\|Y\|^2$ corresponding to different $Y = (\pi, x)'$ pairs, with curves further away from the origin implying higher loss.

Intuition: The OPP as a policy thought experiment

A second attractive feature of the OPP is that its level can be *interpreted* as the magnitude of the deviation from optimality.

While the Lucas critique prevents us from using the OPP as a *prescription* for optimal policy making, part 2 of Proposition 1 establishes that the OPP can be interpreted as the discretionary adjustment to the policy instruments that would correct the optimization failure (under the reaction function g^0).[15] This is the reason why we refer to δ_t^* as the *optimal* policy perturbation.

In other words, we can conceive the OPP as a policy thought experiment where the goal is to try to improve on the policy maker's choice by adjusting p_t^0 by a discretionary amount. To see that most clearly, we can go back to the simple static example of Section 2 with two

[14] More formally, the OPP is the first-step of a Gauss-Newton algorithm that starts at p_t^0. The Gauss-Newton algorithm is based on the Newton line search algorithm and is designed to solve non-linear least squares problems such as ours. Compared to a Newton line search algorithm, the Gauss-Newton algorithm approximates the Hessian with first-derivatives only (e.g., Nocedal and Wright, 2006). Under model (5) where policy has a linear effect on the target variables, the Hessian approximation is exact and $\nabla^2_{p_t} \mathcal{L}_t^0 = \mathcal{R}^{0'} \mathcal{W} \mathcal{R}^0$.

[15] In other words, δ_t^* is the discretionary adjustment to p_t^0 that minimizes the policy maker's problem:

$$\min_{\delta_t} \mathbb{E}_t \|\mathcal{W}^{1/2} \tilde{Y}_t\|^2 \ , \qquad \text{where} \qquad \tilde{Y}_t = \mathcal{R}^0 \delta_t + Y_t^0 \ .$$

mandates. In that case, the OPP becomes

$$\delta^* = -(\mathcal{R}^{0\prime}\mathcal{R}^0)^{-1}\mathcal{R}^{0\prime}Y^0.$$

Starting from a policy choice p^0 with allocation Y^0 and loss $\mathcal{L}^0 = \|Y^0\|^2$, the OPP thought experiment (depicted in Figure 2) consists in using $\mathcal{R}^0$ —the marginal effect of a policy adjustment— to find the change in policy that would most "shrink" the vector Y^0 (and thereby the loss $\mathcal{L}^0$). Geometrically, this amounts to projecting *out* Y^0 on $\mathcal{R}^0$. The "residual" is the vector $\tilde{Y} = \delta^*\mathcal{R}^0 + Y^0$, the allocation after adjustment by the OPP. [16]

4.2 Testing the optimality of the reaction function

The OPP statistic δ_t^* can detect an optimization failure, but it cannot disentangle systematic and discretionary failures: an optimization failure could come from a non-optimal reaction function, a non-optimal discretionary adjustment, or both.

The following proposition establishes a moment condition that must hold if the systematic component of policy is set optimally. If we find that this moment condition does not hold in the data, we will be able to conclude that some of the optimization failures have been systematic, i.e., that the policy maker has been using a non-optimal reaction function.

Proposition 2. *Given an economy described by equations* (4) *and* (5)*, we have that under Assumption 1 and* $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$

$$\mathbb{E}(\delta_t^*|\mathcal{F}_t) = 0 \ . \tag{10}$$

The proposition captures the idea that if the reaction function is optimal (i.e., under the null $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$), the optimization failures must be of a discretionary nature only, and the OPP statistic should be unrelated to the information set $\mathcal{F}_t$. Intuitively, if the policy maker's reaction function is optimal, an OPP sequence should not display any systematic, i.e., predictable, movements: the OPP statistic should resemble a policy shock (e.g., Ramey, 2016) and be orthogonal to any variable in $\mathcal{F}_t$.

4.3 Testing subsets of the policy plan

In practice, estimating the full matrix of causal effects $\mathcal{R}^0$ may be difficult because of sample size issues or because of difficulties in finding convincing identification strategies. Instead, the

[16] Another interpretation is to view the OPP statistic as the coefficient of an OLS regression, as δ^* is minus the coefficient estimate of a regression of Y^0 on $\mathcal{R}^0$: the goal of the optimal perturbation δ^* is to use the causal effects ($\mathcal{R}^0$) in order to minimize the squared deviations of Y^0. This is nothing but a regression of Y^0 on $\mathcal{R}^0$, except one with a minus sign in front of the coefficient estimate since the goal is not to best fit the path for Y^0, but instead to best "undo" Y^0.

researcher may only be able to estimate the causal effect of a subset, or a linear combination, of the different policy instruments. In this section, we show that the logic of the OPP test carries over to that setting as well. Intuitively, since the OPP framework only aims at testing a necessary condition —the gradient should be zero at the optimum—, testing the null of optimality for a subset of the policy instruments is a trivial extension of the baseline setting. Specifically, propositions 1 and 2 continue to apply when considering only a subset, or a linear combination of the policy plan p_t.

To set this up, let $S' = (S'_a, S'_{a^\perp})$ (with subscript a for *available*) denote a $K \times K$ orthogonal selection matrix that splits the matrix of causal effects $\mathcal{R}^0$ into columns (or linear combinations) of $\mathcal{R}^0$ that a researcher can estimate ($\mathcal{R}^0_a = \mathcal{R}^0 S'_a$) and cannot estimate ($\mathcal{R}^0_{a^\perp} = \mathcal{R}^0 S'_{a^\perp}$).[17] In other words, the $K_a \times K$ matrix S_a is the selection matrix that determines $p_{a,t}$ the elements (or linear combinations) of the policy plan p_t that the researcher can perturbate with the OPP thought experiment, i.e.,

$$p_{a,t} = S_a p_t .$$

To detect optimization failures for the subset of policies $p_{a,t}$ we proceed similarly as before except that we only perturbate the policy instruments $p^0_{a,t} = S_a p^0_t$ and rely on the subset-OPP statistic

$$\delta^*_{a,t} = -(\mathcal{R}^{0\prime}_a \mathcal{W} \mathcal{R}^0_a)^{-1} \mathcal{R}^{0\prime}_a \mathcal{W} \mathbb{E}_t Y^0_t \ . \tag{11}$$

The subset OPP $\delta^*_{a,t}$ has the same theoretical properties as the OPP δ^*_t, which we summarize in the following proposition.

Proposition 3. *Given an economy defined by equations* (4) *and* (5)*, we have that under Assumption 1:*

1. $\delta^*_{a,t} \neq 0$ *implies that* $g^0 \notin \mathcal{G}^{\text{ opt}}$ *or* $\epsilon^0_{a,t} \neq 0$*, where* $\epsilon^0_{a,t} = S_a \epsilon^0_t$*;*
2. $p^0_{a,t} + \delta^*_{a,t} = p^*_{a,t}$*, where* $p^*_{a,t} = \arg\min_{p_{a,t} \in \mathbb{R}} \mathbb{E}_t \|\mathcal{W}^{1/2}(\mathcal{R}^0_a p_{a,t} + \mathcal{R}^0_{a^\perp} p^0_{a^\perp,t} + f(X_t; g^0) + \xi_t)\|^2$
3. *we have that under* $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$

$$\mathbb{E}(\delta^*_{a,t} | \mathcal{F}_t) = 0 \ .$$

In other words, all the attractive properties of the OPP continue to apply when we only test a subset of the policy instruments.

[17] To give an example, in the context of monetary policy where the policy plan is an interest rate path $p_t = (i_{t|t}, \ldots, i_{t+H|t})'$, there has been a lot of research on identifying the causal effect of a change in the current policy rate $i_{t|t}$ (e.g., Ramey, 2016), but much less on identifying the effects of changes to the expected policy path. In that case, we can focus on the optimality of the current policy rate alone and take $S_a = \begin{bmatrix} 1 & 0 & \ldots & 0 \end{bmatrix}$ such that $p_{a,t} = S_a p_t = i_{t|t}$ and $\mathcal{R}^0_a = \mathcal{R}^0 S'_a$, which selects the first column of $\mathcal{R}^0$.

5 Illustration: the OPP in a structural macro model

To illustrate the theoretical properties of the OPP we consider a stylized New Keynesian model where the policy maker is the central bank that sets the short term interest rate. The goal of this section is to highlight how a researcher would use the OPP to test macro policies *if* the observed variables were generated by the New Keynesian model. The example also serves to contrast our approach with the standard approach used to verify the optimality of a policy choice. Importantly, we emphasize that this example is only for illustrative purposes as the premise of our paper is to test macro policies without postulating a specific underlying model.

5.1 Detecting an optimization failure

We first illustrate Proposition 1 and show how the OPP statistic can be used to detect an optimization failure in the policy decision process.

The log-linearized baseline New-Keynesian model (Galí, 2015) is defined by a Phillips curve and an intertemporal (IS) curve given by

$$\pi_t = \mathbb{E}_t \pi_{t+1} + \kappa x_t + e^s_t \ , \tag{12}$$

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}(i_t - \mathrm{E}_t \pi_{t+1} - r^*_t) \ , \tag{13}$$

taking a discount rate of 1 and with x_t the welfare-relevant output gap, i_t the nominal interest rate set by the central bank, r^*_t the equilibrium real interest rate and e^s_t an iid cost-push shock.

To illustrate the working of the OPP in an analytically tractable example we consider the case where the central bank has the loss function

$$\mathcal{L}_t = (\pi_t^2 + \lambda x_t^2) \ , \tag{14}$$

with λ the weight on output gap fluctuations. The central bank has only one instrument, $p_t = i_t$ the interest rate at time t.[18]

The standard approach

Denote by i^0_t the policy implemented by the policy maker. The traditional approach to evaluate i^0_t is to contrast this policy decision with that implied by that of a planner choosing

[18]In the online Appendix, Section S2, we show that a similar analysis can be conducted for a central bank who can set the entire expected interest rate plan $p_t = (i_t, i_{t+1|t}, \ldots)'$ in order to minimize the loss function $\mathbb{E}_t \sum_{h=0}^{\infty} \beta^h (\pi^2_{t+h} + \lambda x^2_{t+h})$.

directly π_t and x_t to minimize the loss function:

$$\min_{\pi_t, x_t} \ (\pi_t^2 + \lambda x_t^2) \qquad \text{s.t.} \qquad \pi_t = \mathbb{E}_t \pi_{t+1} + \kappa x_t + e_t^s. \tag{15}$$

The optimality condition $x_t = -\frac{\kappa}{\lambda}\pi_t$ combined with the (IS) curve then gives one possible optimal reaction function $g^{\text{opt}} \in \mathcal{G}^{\text{opt}}$ given by

$$\begin{aligned} i_t^{\text{opt}} &= g^{\text{opt}}(r_t^*, \pi_t) \\ &= r_t^* + \phi_\pi \pi_t \end{aligned} \tag{16}$$

with $\phi_\pi = \frac{\kappa\sigma}{\lambda}$ (e.g., Galí, 2015).[19] One can then assess the optimality of i_t^0 by comparing it to i_t^{opt}.

The OPP approach

The standard approach is not possible when the underlying Phillips curve (12) is not available to the researcher. Instead, our approach consists in directly computing the gradient of the loss function with respect to the policy instrument, in this case the interest rate.

To illustrate the workings of the OPP statistic in this baseline New-Keynesian model, consider a central bank following a non-optimal Taylor rule:

$$i_t^0 = r_t^* + \phi_\pi^0 \pi_t + \epsilon_t^0 \qquad \text{with} \qquad \phi_\pi^0 = \phi_\pi(1+\gamma^0)\ . \tag{17}$$

The policy choice i_t^0 is non-optimal for two reasons. First, with $\gamma^0 > 0$ the central bank is not following the optimal reaction function and is reacting too strongly to movements in the inflation gap. In our general notation this implies that g^0 is not optimal: the central bank is making a systematic optimization failure. Second, with $\epsilon_t^0 \neq 0$ the central bank is making a discretionary (i.e., non-systematic) optimization failure.

Calculating the gradient (and the OPP) requires the two statistics $\mathbb{E}_t Y_t^0$ and $\mathcal{R}^0$. Under i_t^0 and assuming determinacy, the vector $\mathbb{E}_t Y_t^0 = Y_t^0$ is given by

$$\mathbb{E}_t Y_t^0 = \left(\pi_t^0, x_t^0\right)' \ , \qquad \begin{cases} \pi_t^0 = \omega^0 \left(e_t^s - \frac{\kappa}{\sigma}\epsilon_t^0\right) \\ x_t^0 = -\frac{1}{\kappa}(1-\omega^0)e_t^s + \omega^0 \frac{1}{\sigma}\epsilon_t^0 \end{cases} .$$

with $\omega^0 = \frac{1}{1+\frac{\kappa^2}{\lambda}(1+\gamma^0)}$. Moreover, the causal effects of a unit change in i_t are given by

$$\mathcal{R}^0 = \left(\mathcal{R}_\pi^0, \mathcal{R}_x^0\right)' = \left(-\kappa\sigma^{-1}, -\sigma^{-1}\right)' \ ,$$

[19] We posit that the model parameters satisfy $\frac{\kappa\sigma}{\lambda} > 1$ to ensure determinacy, i.e., the existence of a unique equilibrium.

so that we can compute the OPP

$$\begin{aligned}\delta_t^* &= -(\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}\,\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 \\ &= -\underbrace{\gamma^0}_{\neq 0}\phi_\pi \omega \pi_t^0 - \omega^0\underbrace{\epsilon_t^0}_{\neq 0} \end{aligned} \tag{18}$$

where $\mathcal{W} = \text{diag}(1, \lambda)$ and $\omega = \frac{1}{1+\kappa^2/\lambda}$.

Expression (18) makes clear that we can discard optimality when that orthogonality condition is not satisfied (and thus when $\delta_t^* \neq 0$), which is the case when $\gamma^0 \neq 0$ or $\epsilon_t^0 \neq 0$. That being said, since $\delta_t^* \neq 0$ mixes the two sources of optimization failure —a systematic mistake and a discretionary mistake—, a single OPP statistic is not informative about the source of the optimization failure. For that, we will rely on a sequence of OPPs as we illustrate in the next section.

Going beyond the sole orthogonality condition, Proposition 1 also implies that the *magnitude* of the OPP has an economic meaning: it corresponds to the discretionary adjustment to the policy i_t^0 that brings the policy to $i_t^* = i_t^0 + \delta_t^*$; the interest rate that achieves the constrained optimal policy — the policy setting that minimizes the loss function $\mathcal{L}_t$ *under the constraint* that the reaction function is g^0.[20] To see that, note that minimizing $\mathcal{L}_t$ with respect to ϵ_t^0 and subject to eqs (12), (13) and (17) gives the constrained optimal allocation $(\pi_t^*, x_t^*) = \left(\omega e_t^s, -\frac{1-\omega}{\kappa} e_t^s\right)$, which is precisely the allocation achieved under $i_t^0 + \delta_t^*$.

5.2 Testing the optimality of the reaction function

We now illustrate Proposition 2 by showing how $\{\delta_t^*\}$, a sequence of OPP statistics, can be used to test the optimality of the reaction function $g^0(r_t^*, \pi_t) = r_t^* + \phi_\pi^0 \pi_t$. Specifically, under H_0: $g^0 = g^{\text{opt}}$, the moment condition $\mathbb{E}_t\left(\delta_t^* | \mathcal{F}_t\right) = 0$ of Proposition 2 can be tested by studying the effect of inflation on the OPP statistic.

Since the reaction function is linear we can do so using the regression

$$\delta_t^* = b\pi_t^0 + \eta_t \ ,$$

where b captures the effect of inflation on the OPP and η_t is the error term. Since δ_t^* and π_t^0 are simultaneously determined, the coefficient b needs to be identified using an instrumental variable approach.[21] In the New Keynesian model the cost-push shock e_t^s provides a valid

[20]In this simple example, the policy setting i_t^* achieves even more, in that the OPP brings the economy to the optimal allocation achieved under the optimal policy rule g_t^{opt}. This is not a general result however. It holds in this simple example, because the causal effects $\mathcal{R}^0$ are the same under g^0 and g^{opt}.

[21]To clarify, if the researcher could observe π_t right before the policy choice, there would be no simultaneity problem. In practice, however we only observe the equilibrium outcomes, hence the simultaneity issue arises.

instrument as equation (18) implies that e^s_t only affects δ^*_t via π^0_t.

Using e^s_t as instrumental variable, the coefficient b is identified by

$$b = \frac{\mathbb{E}(e^s_t \delta^*_t)}{\mathbb{E}(e^s_t \pi^0_t)} = -\gamma^0 \phi_\pi \omega \ . \tag{19}$$

Intuitively, it is only when the reaction function is non-optimal ($\gamma^0 \neq 0$) that inflation has a systematic (i.e., predictable) effect on the OPP and $b \neq 0$. In contrast, when the reaction function is optimal, inflation should have no effect on the OPP. In that case ($\gamma^0 = 0$), expression (18) implies $\delta^*_t \propto \epsilon^0_t$, and the OPP statistic resembles a monetary shock: the OPP is orthogonal to any time-t variable including e^s_t, and we get $\mathbb{E}(e^s_t \epsilon^0_t) = 0$ and $b = 0$.

6 Testing macro-economic policy: OPP inference

In this section we discuss inference for the OPP. We consider the setting where the researcher aims to test the optimality of a subset $p_{a,t}$ of the policy vector p_t, since inference for the full vector p_t constitutes a special case.

The computation of the subset OPP $\delta^*_{a,t}$ requires two key statistics: (i) the dynamic causal effects $\mathcal{R}^0_a$, and (ii) the conditional expectations $\mathbb{E}_t Y^0_t$. While the previous sections treated these statistics as given, in practice (i) the researcher needs to estimate the causal effects $\mathcal{R}^0_a$ and (ii) the policy maker only provides an estimate of the optimal forecasts $\mathbb{E}_t Y^0_t$.[22]

Taking these constraints into consideration, we now develop an approach to "testing" the optimality of macro policies.[23] Specifically, as illustrated by the simple example in Figure 1(b), the goal of our approach is to make statements of the type "With $X\%$ confidence, we conclude that the policy choice $p^0_{a,t}$ is not optimal", with the confidence level X taking into account the uncertainty surrounding the estimates of $\mathcal{R}^0_a$ and $\mathbb{E}_t Y^0_t$.

We also present the implementation of Proposition 2 to test for the optimality of the reaction function.

Finally, we consider the case where the policy maker's preference parameters are unknown to the researcher, and we present a robust approach to testing macro policies.

[22]The conditional expectations $\mathbb{E}_t Y^0_t$ can generally only be approximated for two reasons: (i) the model that the policy maker used may be incorrectly specified, e.g. a case of function mis-specification or (ii) the measure underlying $\mathbb{E}_t$ may be incorrectly specified, e.g. a case of distribution mis-specification.

[23]The analogy with hypothesis testing is useful for conceptualizing, but formally incorrect as the OPP is a function of the optimal forecast which is a random variable and not a parameter.

6.1 Inference for $\mathcal{R}_a^0$

In this section, we discuss the estimation of the dynamic causal effects $\mathcal{R}_a^0$ based on instrumental variables and state the key assumptions for obtaining a valid approximation to the limiting distribution of the estimator. Overall, the assumptions imposed are similar to those found in other treatment-effect or sufficient statistic type approaches and the approach is relatively standard. Therefore we keep the exposition brief and intuitive and we refer the reader to the online appendix, Section S1, for a rigorous treatment.

We assume that the researcher observes the sample $\{Y_s^0, p_{a,s}^0\}_{s=t_0}^t$, with sample size $n = t - t_0 + 1$, and aims to use this sample to estimate $\mathcal{R}_a^0$. To estimate the relevant causal effects we require that the economy was in a constant regime over the sampling period.

Assumption 2. (Constant regime)
For periods $s = t_0, \ldots, t$ the economy can be described by

$$\begin{array}{rcl} Y_s^0 & = & \mathcal{R}^0 p_s^0 + f(X_s; g^0) + \xi_s \\ p_s^0 & = & g^0(y_s^0, X_s) + \epsilon_s^0 \end{array} . \tag{20}$$

The idea is then to estimate $\mathcal{R}_a^0$ from the model

$$Y_s^0 = \mathcal{R}_a^0 p_{a,s}^0 + \zeta_s \ , \qquad \text{where} \qquad \zeta_s = \mathcal{R}_{a^\perp}^0 p_{a^\perp,s}^0 + f(X_s; g^0) + \xi_s \ .$$

To do so, two endogeneity problems must be handled. First, there is a simultaneity problem as the reaction function $g^0(y_s^0, X_s^p)$ implies that $p_{a,s}^0$ is simultaneously determined with Y_s^0. Second, there is an omitted variable problem as the researcher does not have access to all components of ζ_s, notably $f(X_s; g^0)$ for instance. Both problems imply that the policy paths $p_{a,s}^0$ are correlated with the error term ζ_s.

To solve the endogeneity problem we assume that there exists instrumental variables z_s that correlate only with the policy choice $p_{a,s}^0$.

Assumption 3. (Instrumental variables)
The instrumental variables z_s with dimension $L \times 1$ and $L \geq K_a$ satisfy

1. $\mathbb{E}(z_s \zeta_s') = 0$ *for all* s
2. $\frac{1}{n} \sum_{s=t_0}^t \mathbb{E}(z_s p_{a,s}^{0'})$ *has uniformly full column rank*

The first part of the assumption imposes that the instruments are exogenous whereas the second part imposes that they are relevant, i.e. correlated with $p_{a,s}^0$.

In practice, policy shocks are good instruments as the discretionary component $\epsilon_{a,s}^0$ is an exogenous component of $p_{a,s}^0$ which only affects Y_s^0 via its influence on $p_{a,s}^0$. While $\epsilon_{a,s}^0$ is typically not observable, the literature has produced a variety of proxies that can be used as instruments, see Ramey (2016) and Stock and Watson (2018) for a detailed discussion.

With assumptions 2 and 3 and some regularity conditions stated in the online appendix, we can construct a consistent moment estimator, say $\hat{r}_a$, for the stacked vector of causal effects $r_a^0 = \text{vec}(\mathcal{R}_a^0)$, and we can show that $\hat{r}_a$ is asymptotically normal. We use a consistent approximation to the asymptotic distribution, say $\sqrt{n}(\hat{r}_a - r_a^0) \overset{a}{\sim} N(0, \widehat{\Omega}_a)$, to characterize parameter uncertainty, (e.g. White, 2000).

In fact, in the simplest case of one policy instrument of interest ($K_a = 1$) and one mandate ($M = 1$), our estimator reduces to the popular LP-IV estimator discussed in Stock and Watson (2018) which is based on the local projection framework of Jordà (2005).[24] The only difference (in this partially linear setting) is that we simultaneously estimate all dynamic causal effects, because the confidence region around the *entire* vector $\hat{r}_a$ is required to conduct inference on the OPP.

6.2 Model misspecification uncertainty

Besides the estimation of the causal effects, the researcher must acknowledge that the policy maker will often not be able to provide the oracle forecasts $\mathbb{E}_t Y_t^0$ due to model misspecification. Instead, the policy maker typically provides a point forecast $\widehat{Y}_t$ which can be regarded as an approximation to $\mathbb{E}_t Y_t^0$. In this section, we outline two methods to approximate the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_t$ which we denote by $F_{Y_t^0}$.

First, a researcher can construct its own approximation to the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_t$. A difficulty that arises in practice is that historical misspecification errors $\{\mathbb{E}_s Y_s^0 - \widehat{Y}_s\}_{s=t_0}^t$ are not observable, and we cannot exploit such sequence to predict the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_t$. To see that, note that we have

$$\underbrace{Y_t - \widehat{Y}_{t|t}}_{\text{forecast error}} = \underbrace{Y_t - \mathbb{E}_t Y_t^0}_{\text{future error}} + \underbrace{\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}}_{\text{misspecification error}} .$$

There are two sources of forecast errors: (i) misspecification, i.e., model uncertainty, and (ii) future uncertainty. Model uncertainty takes the form of errors about the current state of the economy —the initial condition—, particularly the nature and the magnitude of the shocks that hit the economy, and about the transmission mechanisms of these putative shocks. Future uncertainty comes from the realization of future shocks that may hit the

[24] Under the partial linearity assumption embedded in model (5), the matrix $\mathcal{R}^0$ can be seen as a matrix of impulse responses. Each column captures the impulse response of a one-unit change in one of the policy instruments on one of the targets $y_{m,t}$. In this paper, we prefer to the terminology *dynamic causal effects* instead of *impulse responses*, as the terminology of causal effects generalizes more naturally to non-linear settings, see Appendix A.

economy. Unfortunately, the two sources of forecast error —misspecification and future uncertainty— are indistinguishable, since only forecast errors are observable.

One strategy is to assume that the misspecification error $\mathbb{E}_t Y_t^0 - \widehat{Y}_t$ follows a normal distribution,[25] and then to upper bound the variance of that distribution with the variance of the forecast errors, which are observable.[26] Naturally, this will be a conservative approach, that may limit the power of the OPP test. Going forward however, there might be more refined methods to sharpen our assessment of model uncertainty, for example by imposing restrictions on the time series properties of misspecification error. For instance, one could impose that misspecification error is a persistent process.

Another strategy is to rely on policy makers' self assessment of model uncertainty. Policy makers spend a lot of effort in assessing the risk around their baseline conditional forecasts, and often experiment with different parameter settings and model specifications in order to assess the effects of model uncertainty. Consistent with this practice, a number of policy institutions publish the distribution $F_{Y_t^0}$ that characterizes the uncertainty around the forecasts.[27] A simple solution is then to assume that this distribution $F_{Y_t^0}$ corresponds to the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_t$. Again, this can be a conservative approach however, because policy makers' assessment of uncertainty often conflate model uncertainty with future shock uncertainty.[28]

6.3 Confidence interval for the OPP

In the previous two sections we characterized parameter uncertainty, $r_a^0 = \text{vec}(\mathcal{R}_a^0) \overset{a}{\sim} N(\hat{r}_a, n^{-1}\widehat{\Omega}_a)$, and model misspecification uncertainty, $\mathbb{E}_t Y_t^0 - \widehat{Y}_t \sim F_{Y_t^0}$. Based on these approximations we can use simulation methods to approximate the distribution of $\delta_{a,t}^* = -(\mathcal{R}_a^{0\prime}\mathcal{W}\mathcal{R}_a^0)^{-1}\mathcal{R}_a^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0$. We denote by $\widehat{\delta}_{a,t}$ the mean of that distribution and whenever the confidence bounds around $\widehat{\delta}_{a,t}$ exclude zero, we will conclude that $p_{a,t}$ is not set optimally (with confidence level X).

[25] Since the true model is not known to the researcher, bootstrap methods, as in Wolf and Wunderli (2015), cannot be adopted, and we must resort to the classical construction of the prediction interval (e.g. Scheffe, 1953), which is based on a normality assumption. Note also that, as argued in Wolf and Wunderli (2015), asymptotic arguments cannot be used to justify the normal approximation. It is an assumption in our setting.

[26] This requires the assumption that the covariance between the future error and the misspecification error is zero.

[27] For instance, monetary policy documents often include an "Assessment of Forecast Uncertainty" section, see for instance the Fed Tealbook or the Bank of England fan-charts.

[28] Looking forward however, policy makers could be more specific about model uncertainty alone. This is the nature of the alternative simulations considered in the Tealbook: studying how variations in the initial conditions or in the transmission mechanism of shocks affects the baseline forecast. To date, however the Fed has not assigned probability weights to these different alternative scenarios.

6.4 Inference for systematic failures

In order to determine whether the optimization failures have been systematic over the sampling period $s = t_0, \ldots, t$, we rely on part 3 of Proposition 3. Specifically, under the null that the reaction function is optimal ($g^0 \in \mathcal{G}^{\text{opt}}$), we have that $\mathbb{E}(\delta^*_{a,s}|\mathcal{F}_s) = 0$.

To test this moment condition we restrict the class of reaction functions $\mathcal{G}$ to be linear, and we consider the model

$$\widehat{\delta}_{a,s} = Bw_s + \eta_s \ , \qquad s = t_0, \ldots, t \ , \tag{21}$$

where w_s includes a constant and any variable that the researcher suspects to be an important (but perhaps overlooked) argument of the reaction function. The error term η_s includes the estimation error of the OPP statistic $\widehat{\delta}_{a,s} - \delta^*_{a,s}$ and possibly other omitted variables.

The consistent estimation of B generally requires an identification strategy as w_s is likely to be correlated with η_s. Indeed, even under $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$ and Assumption 1, we have that $\delta^*_{a,s} \propto \epsilon^0_{a,s}$ which is simultaneously determined with y^0_s. Hence, if y^0_s is included in w_s there exists a simultaneity problem as we illustrated in our New Keynesian example above.

To solve the identification problem we suggest to use instrumental variables. Valid instruments in this context can be series of shocks or variables that are uncorrelated with η_s and correlated with w_s. In our New Keynesian example cost push shocks provided valid instruments, but in general any variable or shock that is not simultaneously determined with the discretionary adjustment, for instance lagged macro variables, can also be used.

Provided that the researcher has access to valid instrumental variables, we show in the online appendix (under standard regularity conditions) that the parameters in B can be consistently estimated and one can construct a standard Wald test of joint significance for the B elements. If we can reject that the B elements are jointly zero, we can reject the null hypothesis that the policy maker's reaction function g^0 is optimal: the policy maker has systematically over- or under-reacted to the variables in w_s.

6.5 Robust OPP inference

The computation of the OPP statistic $\delta^*_{a,s} = -(\mathcal{R}^{0\prime}_a \mathcal{W} \mathcal{R}^0_a)^{-1} \mathcal{R}^{0\prime}_a \mathcal{W} \mathbb{E}_t Y^0_t$ relies on knowing the preference matrix $\mathcal{W} = \text{diag}(\beta \otimes \lambda)$. In certain cases, λ and β may be unknown to the researcher, either because the policy maker does not want to communicate such preferences or because the policy maker has difficulties eliciting preferences precisely.

In this context, we outline an approach for conducting *preference robust* OPP inference, that is we propose a method to test the optimality of a policy decision when the researcher does not have access to the preference parameters. The idea is to exploit a sequence of past policy decisions to find the preferences that gives the smallest deviations from optimality on average. This approach can thus be seen as considering a worse-case scenario for rejecting optimality.

Specifically, we write $\omega = \lambda \otimes \beta$, the elements of the preference matrix $\mathcal{W}$, as a function of the $d_\theta \times 1$, parameter vector θ, i.e., $\omega = \omega(\theta)$, with $d_\theta \leq K$, and we estimate θ by numerically solving[29]

$$\hat{\theta} = \arg\min_{\theta \in \Theta} \widehat{S}(\theta) \ , \qquad \widehat{S}(\theta) = \left\| \frac{1}{n} \sum_{s=t_0}^{t} \widehat{\delta}_{a,s}(\theta) \right\|^2 \ , \tag{22}$$

where $\widehat{\delta}_{a,s}(\theta)$ denotes the mean OPP estimate (that depends on θ). In the online appendix section S1.5 we provide a more general econometric implementation and the details for inference.

With the estimated $\hat{\theta}$ in hand a researcher can base the optimality assessment on $\widehat{\delta}_{a,t}(\hat{\theta})$, which ensures that deviations from optimality are not due to potentially arbitrary choices for the preference parameters.

7 Illustration: Testing US monetary policy

In this section we illustrate how the OPP framework can be used to evaluate macroeconomic policy choices in practice. In general, the OPP statistic can be used to answer three types of questions:

1. Is there an optimization failure?

2. What are the economic reasons for this optimization failure, i.e., what trade-offs were overlooked when setting policy at a given point in time? For instance, did the policy maker stabilize one mandate too much relative to another one? Or did the policy maker put too much weight on stabilizing mandates in the short-run at the expense of the long-run?

3. Over the past n periods, would a systematically stronger/weaker policy response to movements in observable X have been more appropriate to achieve the policy maker's objectives?

We will answer these questions by computing the OPP for Fed policies over 1990-2018, focusing on three case studies: 1990M6, 2008M4 on the eve of the Great Recession, 2010M4 in the middle of the Great recession.

[29] To give a specific example, suppose that $M = 2$ and $K = 2$, then we can take $\theta = (\theta_1, \theta_2)'$ and set $\omega(\theta) = (1, \theta_1)' \otimes (\theta_2^0, \ldots \theta_2^H)'$, which implies that $\lambda = (1, \theta_1)'$ and $\beta = (\theta_2^0, \ldots \theta_2^H)'$.

7.1 Implementation details

We evaluate the optimality of the Fed monetary policy with respect to its interest rate policy. As loss function we posit the usual dual inflation-unemployment mandate

$$\mathcal{L}_t = \mathbb{E}_t \left\|\Pi_t\right\|^2 + \lambda \mathbb{E}_t \left\|U_t\right\|^2 \ , \tag{23}$$

with $\Pi_t = (\pi_t - \pi^*, \dots, \pi_{t+H} - \pi^*)'$ the vector of inflation gaps and $U_t = (u_t - u_t^*, \dots, u_{t+H} - u_{t+H}^*)'$ the vector of unemployment gaps. The discount rate is implicitly set to $\beta_h = 1$ for all h, and we take a horizon of $H = 5$ years. The choice for λ is discussed below.

In our analysis we assume that the economy was in a stable regime over the 1990-2018 period and Assumption 2 applies. This is consistent with the widely held belief that since 1985 the US economy has evolved in a stable monetary regime (e.g., Clarida, Galí and Gertler, 2000).

The policy instruments

We focus on the optimality of two elements of the set of Fed instruments: (i) the current level of the fed funds rate, or fed funds rate policy for short, and we take $p_{a,t} = i_{t|t}$, and (ii) the slope policy instrument, whereby the Fed aims to affect the slope of the yield curve —the spread between the 10-year treasury and the fed fund rate— through forward-guidance or asset purchases (QE), and we take $p_{a,t} = i_{t+10\text{yr}|t} - i_{t|t}$.

Similarly to Eberly, Stock and Wright (2019) we consider these policy instruments separately as they were used over different sampling periods. In particular, the fed funds rate policy is active over the entire 1990-2018 sampling period, with the exception of the zero lower bound period. In contrast, slope policies, such as forward-guidance or asset purchases, have only been considered after 2007.

Estimating $\mathcal{R}_a^0$

To estimate $\mathcal{R}_a^0$, we follow Kuttner (2001) and Eberly, Stock and Wright (2019), 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 causal effects changes in the current interest rate $i_{t|t}$. 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 capture the effect of changes in the slope of the yield curve. We then use these surprises as instrumental variables in equation (20) and compute the dynamic causal effects $\widehat{\mathcal{R}}_a$ as described in Section

6.

Measuring $\mathbb{E}_t Y_t^0$

For the conditional forecasts $\widehat{Y}_t$, we use the median FOMC forecast reported in the Survey of Economic Projections (SEP).[30] To capture the uncertainty around these point forecasts, we use the Board staff assessment of forecast uncertainty, as reported in the Tealbook.

The preference matrix $\mathcal{W}$

The last piece of information necessary to evaluate policy decisions is the policy maker's preference parameter λ that enters the preference matrix $\mathcal{W}$. There are two possible routes to set λ. A first strategy consists in using the policy maker's published value for λ, if available. In the case of the Fed for instance, the Fed has repeatedly stated its commitment to a "balanced approach" between its two mandates and consistent with this view the Board staff uses the value $\lambda = 1$ in its optimal policy simulations for the Tealbook. A second strategy consists in using the value for λ that is least favorable to reject optimality on average, as described in Section 6.5. That second strategy is more conservative by nature, but it is also more robust as it avoids the assumption that $\lambda = 1$ is the correct value.

In this empirical application, we will use that robust approach. Using the estimator (22) for policy decisions taken over 1990-2018, we estimate $\hat{\lambda} = 0.6$, which we will use from now on in our assessment of policy decisions.[31]

Computing the OPP

Based on $\widehat{\mathcal{R}}_a$, $\widehat{Y}_t$ and $\hat{\lambda}$, we compute the mean subset OPP statistics $\widehat{\delta}_{a,t}$ and construct confidence bands as described in section 6. We obtain two subset OPP statistics: (i) the short rate OPP denoted by $\widehat{\delta}_{\mathsf{i},t}$, and (ii) the slope OPP denoted by $\widehat{\delta}_{\boldsymbol{\Delta},t}$.

7.2 Fed funds rate policy as of June 1990

Narrative

In the first case study, we evaluate the fed funds rate policy as of June 1990. At the time, the FOMC was confronted with a classic inflation-unemployment trade-off: while it would

[30]Since 2006, SEP data include the median forecasts at a three-year ahead horizon. We complement these forecasts with the median FOMC estimate of the long-run projections for inflation and unemployment. We set the horizon for the long-run FOMC projections to equal 5 years. Since the 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).

[31]As it turns out, using instead $\lambda = 1$ in our assessment of policy decisions gives very similar results.

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 OPP is thus whether the level of the fed funds rate optimally balanced that trade-off.

Computing the OPP

Figure 3 depicts graphically all the information needed to assess the optimality of the fed funds rate at the time. The top-left panel reports the FOMC expected paths for inflation conditional on the current policy choice, that is it reports $\hat{\Pi}_t$. The bottom-left panel reports the causal effects on inflation of a 1ppt innovation to the current fed funds rate. We will refer to that causal effect as $\widehat{\mathcal{R}}_{\mathsf{i}}^{\pi}$. The right column reports the same information for unemployment: $\hat{U}_t$ and $\widehat{\mathcal{R}}_{\mathsf{i}}^{u}$. For illustration purposes, in this first case study we omit confidence bands and treat the causal effect estimates and forecasts as fixed.

A trade-off across mandates

With one mandate, an optimization failure is a failure to best stabilize that one particular mandate. However, with multiple mandates, an optimization failure can come from a failure to balance conflicting mandates, as was the case with inflation and unemployment in June 1990. To see that, we can re-write $\delta_{\mathsf{i},t}^{*}$ as a weighted-average of the OPP for each mandate with

$$\delta_{\mathsf{i},t}^{*} = (1-\omega)\delta_{\mathsf{i},t}^{\pi *} + \omega \delta_{\mathsf{i},t}^{u *} \tag{24}$$

with $\delta_{\mathsf{i},t}^{w*} = -(\mathcal{R}_{\mathsf{i}}^{w\prime}\mathcal{R}_{\mathsf{i}}^{w})^{-1}\mathcal{R}_{\mathsf{i}}^{w\prime}\mathbb{E}_t W_t^0$ the OPP for a single mandate with $W_t = (w_t - w^*, \ldots, w_{t+H} - w^*)'$ for $w = \pi$ or u, and with

$$\omega = \frac{1}{1+\kappa^2/\lambda} \tag{25}$$

a scalar weight that depends on the ratio of the policy maker's preference between the two mandates (λ), and the central bank's instrument "average" ability to transform unemployment into inflation $\kappa = \|\mathcal{R}_{\mathsf{i}}^{\pi}\|/\|\mathcal{R}_{\mathsf{i}}^{u}\|$.[32] From (24), we can see that an optimization failure —a non-zero OPP δ_t^{*}— can come from a failure to appropriately balance $\delta_{\mathsf{i},t}^{u*} > 0$ with $\delta_{\mathsf{i},t}^{u*} < 0$, or vice-versa.

The OPP thought experiment

The thought experiment underlying the OPP test is to consider whether a discretionary adjustment to the policy choice could better stabilize the policy maker's mandates. To

[32]In a New-Keynesian model like that of section 5, κ reduces to the slope of the Phillips curve. The weight ω is the parameter that controls how the policy maker should balance the two mandates, and it depends on κ the policy instrument's ability to influence one mandate versus another, as well as λ the policy maker's preference for stabilizing one mandate versus another.

illustrate this thought experiment and illustrate the inflation-unemployment trade-off at play, Figure 3 shows how such hypothetical adjustments to the fed funds rate would alter the conditional forecasts.

Specifically, 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 $\widehat{\delta}^{\pi}_{\mathrm{i},t}$, the OPP for a strict inflation targeter that did not care about the path of unemployment ($\lambda = 0$). With $\hat{\delta}^{\pi}_{\mathrm{i},t} \approx 0.7 > 0$, the OPP thought experiment calls for a more contractionary policy in order to close the positive inflation gap faster. In fact, the adjustment $\widehat{\delta}^{\pi}_{\mathrm{i},t}$ is the discretionary adjustment to the fed funds rate that best stabilizes the expected path for inflation.[33]

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 thought experiment calls for a more expansionary policy ($\widehat{\delta}^{u}_{\mathrm{i},t} \approx -0.2 < 0$) in order to close the unemployment gap faster. With $\widehat{\delta}^{\pi}_{\mathrm{i},t}$ and $\widehat{\delta}^{u}_{\mathrm{i},t}$ calling for opposite policies, the FOMC was indeed facing an inflation-unemployment trade-off.

Real-time policy assessment

The optimization failure is economically small with $\hat{\delta}_{\mathrm{i},t} \approx -0.1$ (Table 1), meaning that the two mandates were roughly balanced. Moreover, taking estimation and model uncertainty into account, the confidence interval includes zero, and we cannot discard that the fed funds rate was set optimally.

7.3 Fed funds rate policy as of April 2008

Narrative

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.[34] 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.

[33]Recall that policy is not set optimally whenever the causal effect of a change in policy is *not* orthogonal to the expected path of the inflation gap ($\widehat{\mathcal{R}}^{\pi\prime}_{\mathrm{i}}\widehat{\Pi}_t \neq 0$). When this is the case, it is possible to project *out* $\widehat{\Pi}_t$ on $\widehat{\mathcal{R}}^{\pi}_{\mathrm{i}}$ and obtain a "shorter" vector $\widetilde{\Pi}_t = \widehat{\Pi}_t + \widehat{\mathcal{R}}^{\pi}_{\mathrm{i}}\widehat{\delta}^{\pi}_{\mathrm{i},t}$ —a lower sum of squared inflation gaps—, i.e., a lower loss function as illustrated in Figure 2.

[34]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.

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. In other words, should the Fed have done more and lowered its fed funds rate by more *given* the FOMC forecasts of the time and *given* the uncertainty attached to its forecasts?[35]

Computing the OPP

Figure 4 has the same structure as Figure 3 except that we now report the 68 percent confidence intervals for the impulse response estimates, as well as the 68 percent confidence interval capturing the model uncertainty surrounding the Fed's forecast, as judged by the Board staff in the April 2008 Greenbook.

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 4. While this could suggest the existence of an inflation-unemployment trade-off as in June 1990, the OPP shows that there was no inflation-unemployment trade-off at the time. As shown in Table 1, *both* OPPs are negative with $\hat{\delta}^u_{\mathrm{i},t} \approx -.5$ and $\hat{\delta}^\pi_{\mathrm{i},t} \approx -.1$.

A trade-off across horizons

Before presenting our policy assessment, we briefly mention the reason for having a negative value for $\widehat{\delta}^\pi_{\mathrm{i},t}$. Indeed, $\widehat{\delta}^\pi_{\mathrm{i},t} < 0$ calls for a lower fed funds rate to stabilize inflation. This may seem surprising, since the FOMC was expecting a *positive* inflation gap in the near term, and that positive gap was the reason why the FOMC did not pursue a more aggressive cut to the fed funds rate.

The reason for this result is that the Fed must also balance a trade-off across horizons. In particular, Figure 4 shows that a discretionary adjustment to the fed funds rate has a very delayed effect of inflation —the well-known transmission lag of monetary policy—. As a result, the OPP for inflation is determined by the longer-term developments in inflation, and not by the short-term positive inflation gap.[36] Compared to June 1990, the positive inflation

[35]Following a more aggressive policy was a real possibility at the time. The three alternative monetary strategies prepared by the Board staff for the April Bluebook —the three strategies ultimately discussed by the FOMC— comprised a no-change option, a 25bp cut (ultimately chosen by the FOMC) *and* a 50bp cut.

[36]Stated differently, the causal effect $\mathcal{R}^\pi_{\mathrm{i}}$ is close to zero at short horizon, such that the orthogonality condition underlying the OPP test (here $\mathcal{R}^{\pi\prime}_{\mathrm{i}}\mathbb{E}_t\Pi^0_t = 0$) is always satisfied at short horizon. This captures a common wisdom of central banking: because of the delayed effects of monetary policy, central banks should "look through" transitory inflationary episodes.

gap in April 2008 was expected to be much more transitory and even to turn negative after three years. As a result, and in contrast to June 1990, we get a slightly negative OPP for inflation.

Real-time policy assessment

The dual-mandate OPP comes out at $\hat{\hat{\delta}}_{\mathsf{i},t} = -0.41$ (Table 1), calling for an additional 50 basis points cuts (after rounding at the nearest quarter percentage point). Moreover, the 68% confidence interval excludes zero, indicating that there is a less than 32 percent chance that the chosen policy i_t^0 actually balances the true conditional expectations for inflation and unemployment. In words, despite large uncertainty, we can discard optimality: the expected increase in unemployment was large enough to justify (in real time) a more aggressive monetary stimulus.

7.4 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 used its slope instrument to stabilize the economy.

Figure 5 displays the situation in 2010-M4 where the bottom panels show the causal effects of inflation and unemployment to a 1ppt increase in the slope of the yield curve. In Table 1 we find that the mean estimate of the slope OPP is given by $\widehat{\delta}_{\Delta,t} = -1.02$. The deviations from targets are so large that we can easily discard optimality at the 68 percent confidence level.[37]

7.5 Testing the Fed reaction function

The previous section illustrated the workings of the OPP for a few noteworthy policy decisions over the past 30 years. Naturally, we can proceed more systematically and evaluate all policy decisions made over 1990-2018. The bottom panel of Figure 6 shows the sequences of OPP estimates for the two policy instruments that we consider: the OPP for the fed funds rate and the OPP for slope policy.

Except during the Great recession, the deviations from optimality are economically small, averaging only a quarter ppt (in absolute value) for the fed funds rate. However, note how the OPP sequence for the fed funds rate is also displaying some cyclicality, which could suggest some systematic deviation from optimality, i.e., a non-optimal reaction function.

[37]That being said, this conclusion is reached with the benefit of hindsight, since our evidence on the effect of the slope instrument comes precisely from that time period.

To study the systematic component of the Fed's policy, we can rely on Proposition 2 and use the OPP sequence to test whether the Fed's reaction function has been optimal. Specifically, under the assumption of a constant policy regime over 1990-2018, the test consists in studying whether some information available at time t causes systematic movements in the OPP. Inspired by the empirical success of the basic Taylor rule specification for the fed funds rate (Taylor, 1993), whereby the contemporaneous policy rate follows a linear function of inflation and unemployment $g^0(\pi_t, u_t) = \phi^0_\pi \pi_t + \phi^0_u u_t$, we will test whether inflation or unemployment have a non-zero effect on the OPP statistic.[38] We follow the methodology outlined in Section 6.4 to conduct inference.

Specifically, we estimate the regression

$$\widehat{\delta}_{x,s} = c + b_\pi \pi^0_s + b_u u^0_s + \eta_s \ , \qquad x = \mathsf{i}, \Delta \ , \qquad s = t_0, \ldots, t \ ,$$

where non-zero coefficients for c, b_π or b_u indicate a non-optimal reaction function. As instrumental variables, we use lagged values of the right-hand side variables as instruments.

Table 2 presents the results, with the bottom row displaying the Wald test of joint significance. While inflation has no clear effect on either OPP, unemployment does have a statistically significant effect on the OPP, regardless of whether we use the OPP for the current fed funds rate or the OPP for the slope policy. In other words, we can reject the null that the Fed has been using an optimal reaction for the fed funds rate or for its slope instrument. This indicates that a more systematic reaction to the unemployment gap would have been more appropriate to achieve the policy maker's objectives.

8 Conclusion

In this work, we developed a framework to evaluate macroeconomic policy decisions with minimal assumptions on the underlying economic model. Given a policy maker's loss function, we proposed a statistic —the *Optimal Policy Perturbation*, OPP— to detect "optimization failures" in the policy decision process, those are instances when the policy decision does not minimize the loss function. Further, we showed that a sequence of OPP statistics can disentangle systematic and discretionary optimization failures.

Going forward, there are several ways in which the power of the OPP test can be improved. First, any improvement in the precision of the forecasts or in the precision of the causal effect estimates directly improves the ability to detect optimization failures by shrinking the confidence bands of the OPP. Relatedly, more accurate estimates of the causal effects

[38] Recall that to reject the optimality of the policy maker's reaction function, Propositions 2 and 3 states that we only need *one* time-t variable that causes movements in the OPP δ^*_t. Thus, considering a simpler Taylor rule is not restrictive.

of policy changes, possibly allowing for time-variation or state dependence could improve the detection of optimization failures. Second, policy makers should provide an explicit *quantitative* assessment of the model uncertainty around the forecasts. Currently, few policy makers provide a rigorous assessment of the model uncertainty attached to their forecasts. For instance, the Fed only provides a qualitative assessment of the uncertainty around its forecasts. In the context of fiscal policy, governments often only publish their point forecasts, with no sense of uncertainty around these forecasts. However, our OPP framework shows that model mis-specification measurements are crucial inputs into the assessment of policy choices, and thus crucial inputs into the decision making process.

The monetary policy setting considered in this paper is only one of the potential applications for the OPP approach. For instance, in the context of fiscal policy, a number of rules are being used to prevent excessive deficit, such as the European "Stability and Growth Pact" that limits budget deficits in EU member countries to 3 percent of GDP. These rules are rigid and do not take into account other important objectives of policy makers, such as avoiding large drops in GDP and excessive unemployment.[39] The OPP could be used in this context to modernize the deficit rule with a "forecast deficit targeting" approach to fiscal discipline, in the same way that forecast inflation targeting replaced strict monetary growth targets. The OPP then provides a quantitative criterion to formalize how a policy maker should balance the debt burden with growth and unemployment considerations.

There are many other possible applications of the OPP, for instance a government interested in exchange rate management, or a country interested in foreign-exchange reserve management.

[39] In practice such trade-offs are central to policy makers, as reflected by the number of EU countries, which chose to violate the rule during the 2007-2009 crisis.

References

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

Barro, Robert J., and David B. Gordon. 1983. "A Positive Theory of Monetary Policy in a Natural Rate Model." *Journal of Political Economy*, 91(4): 589–610.

Bénassy-Quéré, Agnès, Benot Coeuré, Pierre Jacquet, and Jean Pisani-Ferry. 2018. *Economic Policy: Theory and Practice. Second edition.* Oxford University Press.

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

Blanchard, Olivier. 2018. "On the future of macroeconomic models." *Oxford Review of Economic Policy*, 34(1-2): 43–54.

Blinder, Alan S, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. 2008. "Central bank communication and monetary policy: A survey of theory and evidence." *Journal of Economic Literature*, 46(4): 910–45.

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.

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

Clarida, Richard H, et al. 2020. "The Federal Reserve's New Framework: Context and Consequences: a speech at the" The Economy and Monetary Policy" event hosted by the Hutchins Center on Fiscal and Monetary Policy at the Brookings Institution, Washington, DC (via webcast), November 16, 2020."

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

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland. 2012. "The optimal inflation rate in New Keynesian models: should central banks raise their inflation targets in light of the zero lower bound?" *Review of Economic Studies*, 79(4): 1371–1406.

Debortoli, Davide, Jinill Kim, Jesper Lindé, and Ricardo Nunes. 2019. "Designing a simple loss function for central banks: Does a dual mandate make sense?" *The Economic Journal*, 129(621): 2010–2038.

Dennis, John E., and Robert B. Schnabel. 1996. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations.* Siam.

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.

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

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

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

Hansen, Bruce. 2020. *Econometrics.* Available at https://www.ssc.wisc.edu/ bhansen/econometrics.

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.

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

Kleven, Henrik J. 2020. "Sufficient Statistics Revisited." *Annual Review of Economics*. forthcoming.

Kocherlakota, Narayana. 2016. "Rules versus discretion: A reconsideration." *Brookings Papers on Economic Activity*, 2016(2): 1–55.

Kocherlakota, Narayana R. 2019. "Practical policy evaluation." *Journal of Monetary Economics*, 102: 29 – 45.

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.

Kydland, Finn E., and Edward C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, 85(3): 473–491.

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

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

McKay, Alisdair, Emi Nakamura, and Jón Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review*, 106(10): 3133–58.

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

Newey, Whitney K., James L. Powell, and Francis Vella. 1999. "Nonparametric Estimation of Triangular Simultaneous Equations Models." *Econometrica*, 67(3): 565–603.

Nocedal, Jorge, and Stephen Wright. 2006. *Numerical optimization.* Springer Science & Business Media.

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.

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

Romer, Christina D., and David H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review*, 90(3): 429–457.

Sargent, Thomas J. 1981. "Interpreting economic time series." *Journal of political Economy*, 89(2): 213–248.

Scheffe, Henry. 1953. "A Method for Judging all Contrasts in the Analysis of Variance." *Biometrika*, 40(1/2): 87–104.

Schmitt-Grohé, Stephanie, and Martin Uribe. 2010. "The optimal rate of inflation." In *Handbook of monetary economics*. Vol. 3, 653–722. Elsevier.

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

Sims, Christopher A. 2002. "The role of models and probabilities in the monetary policy process." *Brookings Papers on Economic Activity*, 2002(2): 1–40.

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.

Su, Liangjun, and Aman Ullah. 2008. "Local polynomial estimation of nonparametric simultaneous equations models." *Journal of Econometrics*, 144(1): 193 – 218.

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

Svensson, Lars E.O. 2017. "What Rule for the Federal Reserve? Forecast Targeting." National Bureau of Economic Research.

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

Taylor, John B. 1993. "Discretion versus policy rules in practice." Vol. 39, 195–214, Elsevier.

White, Halbert L. Jr. 2000. *Asymptotic Theory for Econometrians - 2nd edition.* San Diego, California:Academic Press.

Wolf, Michael, and Dan Wunderli. 2015. "Bootstrap Joint Prediction Regions." *Journal of Time Series Analysis*, 36(3): 352–376.

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 A: Relaxing linearity

For some applications assuming a linear relationship between Y_t and p_t may be too strong. In this appendix we show that it is quite easy to relax this assumption and obtain a generalized OPP statistic that remains able to detect optimization failures. However, inference for such statistic requires estimating the dynamic causal effects of interest using non-parametric instrumental variable methods, which is more data demanding and has therefore not been adopted often in macroeconomics. Nevertheless the methodology exists (e.g. Su and Ullah, 2008) and can be adopted for our purposes.

Generalized OPP

To derive the generalized OPP (GOPP) statistic for nonlinear models we consider a more general description of the economy:

$$Y_t = f(p_t, X_t; g) + \xi_t \ , \qquad p_t = g(y_t, X_t) + \epsilon_t \ , \tag{26}$$

where $f(p_t, X_t; g)$ now specifies a general nonlinear mapping between the policy instruments p_t and Y_t. The policy maker again aims to minimize the loss function $\mathcal{L}_t = \mathbb{E}_t \|\mathcal{W}^{1/2} Y_t\|^2$ and his proposed solution is given by g^0 and ϵ_t^0. We define

$$\mathfrak{R}_t^0 \equiv \left. \frac{\partial f(p_t + \delta_t, X_t; g^0)}{\partial \delta_t'} \right|_{\delta_t = 0} \ , \tag{27}$$

which are the relevant dynamic causal effects for the nonlinear model (26). The GOPP is given by

$$\delta_t^{g*} = -(\mathfrak{R}_t^{0\prime} \mathcal{W} \mathfrak{R}_t^0)^{-1} \mathfrak{R}_t^{0\prime} \mathcal{W} \mathbb{E}_t Y_t^0 \ ,$$

where $Y_t^0 = f(p_t^0, X_t; g^0) + \xi_t$. It is easy to see that under model (5) we have that $\mathfrak{R}_t^0 = \mathcal{R}^0$ and the GOPP reduces to the OPP. As we show below, this generalized OPP has retains the ability of the OPP to detect optimization failures, but adjusting p_t^0 by δ_t^{g*} does generally not lead to the optimal attainable policy anymore, i.e. part 2 of Proposition 1 does not hold anymore. Additionally, estimating $\mathfrak{R}_t^0$ is considerably more difficult when compared to $\mathcal{R}^0$. To see this, just recall that typically f is unknown and may depend on the – potentially high dimensional – vector of state variables X_t. Estimating the derivative is difficult in such settings and certainly requires further assumptions on f and the dimension of X_t.

We first discuss the theoretical properties of the generalized OPP after which we outline possible strategies for estimating the derivative function $\mathfrak{R}_t^0$.

Properties of the generalized OPP

In this section we formalize the properties of the generalized OPP δ_t^{g*}. In particular, we give the (increasingly stronger) conditions under which the GOPP can be used to: (i) reject that the policy choice p_t^0 is optimal, (ii) bring p_t^0 closer to the optimal policy and (iii) make policy optimal given the reaction function g^0 as the baseline OPP does.

(i) Discarding optimality

In order to use the GOPP to discard that p_t^0 is optimal, we essentially only require that the underlying model is well defined and that it is continuously differentiable with respect to the policy choice. Formally, we make the following assumption.

Assumption 4. *Let $X_t \in \mathcal{X}$ be a random vector and let $\mathcal{D}$ be an open convex subset of $\mathbb{R}^K$. We assume that*

1. *the function $\mathfrak{R}_t(p_t, X_t; g) \equiv \partial f(p_t, X_t; g)/\partial p_t$ exists for all $X_t \in \mathcal{X}$, $p_t \in \mathcal{D}$ and $g \in \mathcal{G}$, and $\operatorname{rank}(\mathfrak{R}_t^0) = K$*
2. *there exists a non-empty set $\mathcal{G}^{\text{opt}}$ such that*

$$\mathbb{E}_t \|\mathcal{W}^{1/2} Y_t(g, 0)\|^2 \leq \mathbb{E}_t \|\mathcal{W}^{1/2} Y_t(\tilde{g}, \tilde{\epsilon}_t)\|^2, \qquad \forall\ g \in \mathcal{G}^{\text{opt}},\ \tilde{g} \in \mathcal{G} \backslash \mathcal{G}^{\text{opt}}, \tilde{\epsilon}_t \neq 0\ ,$$

where $Y_t(g, \epsilon_t) = f(p_t, X_t; g) + \xi_t$ and $p_t = g(y_t, X_t) + \epsilon_t$.

Part 1 of this assumption imposes that the derivative function $\mathfrak{R}_t$ exists and has full column rank at p_t^0 and g^0. Part 2 is equivalent to Assumption 1 in the main text, and imposes the existence of a well defined optimum. Note that the GOPP formula (27) involves $\mathfrak{R}_t^0 = \mathfrak{R}_t(p_t^0, X_t; g^0)$, as the GOPP is computed using the derivatives evaluated at p_t^0. The following proposition formalizes the notion of discarding optimality using the GOPP.

Proposition 4. *Given model* (26) *and Assumption 4, we have that $\delta_t^{g*} \neq 0$ implies $g^0 \notin \mathcal{G}^{\text{opt}}$ or $\epsilon_t^0 \neq 0$.*

The proposition implies that if δ_t^{g*} is not equal to zero there exists an optimization failure, either due to the systematic part g^0, the discretionary part ϵ_t^0, or both. The proposition shows that the GOPP has the same ability to detect optimization failures as the OPP but now for nonlinear models. Perhaps surprisingly this result requires virtually no conditions on the functional form of f.

(ii) Improving policy

Part 2 of proposition 1 does not carry over to the generalized OPP statistic. In particular, we require additional conditions under which the perturbation δ_t^{g*} can bring p_t^0 closer to the optimal choice $p_t^* = \arg\min_{p_t \in \mathcal{D}} \|f(p_t, X_t; g^0) + \xi_t\|^2$. Recall that p_t^* is the minimal loss the policy maker can attain given the reaction function g^0. We make the following additional assumption.

Assumption 5. *We assume that*

1. *$\mu_{\min} > 0$, where $\mu_{\min}$ is the smallest eigenvalue of $\mathfrak{R}_t(p_t^*, X_t; g^0)' \mathcal{W} \mathfrak{R}_t(p_t^*, X_t; g^0)$ uniformly over $X_t \in \mathcal{X}$*
2. *$\|(\mathfrak{R}_t(p_t, X_t; g^0) - \mathfrak{R}_t(p_t^*, X_t; g^0))' \mathcal{W} \mathbb{E}_t[f(p_t^*, X_t; g^0) + \xi_t]\| \leq c\|p_t - p_t^*\|$, with constant $c < \mu_{\min}$ for all $(p_t, X_t) \in \mathcal{D} \times \mathcal{X}$.*
3. *$\mathfrak{R}_t$ is Lipschitz continuous with respect to p_t on $\mathcal{D}$ with parameter γ.*

The first part of the assumption assumes that the effects of the different policy instruments are not linearly dependent. The second part ensures that the loss function is not too nonlinear in the neighborhood of p_t^* and the third part imposes a smoothness condition on the causal effects.

The assumption allows us the formalize the following notion of a policy improvement.

Proposition 5. *Given Assumptions 4 and 5 we have there exists $e > 0$ such that for all $p_t^0 \in \mathcal{N}(p_t^*, e)$ we have*[40]

$$\|p_t^0 + \delta_t^{g*} - p_t^*\| \leq \|p_t^0 - p_t^*\|$$

where $p_t^* = \arg\min_{p_t \in \mathcal{D}} \|f(p_t, X_t; g^0) + \xi_t\|^2$.

The proposition states that, if the policy choice of the policy maker is in the neighborhood $\mathcal{N}(p_t^*, e)$ of the optimal policy, the OPP will bring p_t^0 closer to the optimum. Importantly, as we show in the proof of the proposition, the "size" e of the neighborhood $\mathcal{N}(p_t^*, e)$ depends on the degree of non-linearity in the effects of policy. The more non-linear the effect of policy —the more non-linear the function $\mathfrak{R}_t$—, the smaller the neighborhood has to be.[41] Proposition 5 is weaker then part 2 of Proposition 2, which states that under linearity the OPP adjustment brings the given policy choice p_t^0 directly to p_t^*. Proposition 5 show that for nonlinear models the GOPP can only bring the policy choice closer to the constrained minimum p_t^*.

While the Assumption 4 imposed virtually no restrictions on the functional form of $\mathfrak{R}_t(p_t, X_t; g^0)$, Assumption 5 restricts the effect of p_t on $\mathfrak{R}_t$. These types of smoothness conditions are typically required for the non-parametric estimation of $\mathfrak{R}_t^0$. In general, as the GOPP only depends on first order derivatives, getting closer to the optimal policy given g^0 is the best one can do with the GOPP.

(iii) Getting to the optimal policy

Finally, we impose the more stringent condition under which the GOPP brings us directly to the constrained optimal policy choice p_t^*. This happens when policy has a linear effect on the targets, as in model 5, but with the added generality that we also allow for state dependence in that $\mathfrak{R}_t$ can vary with X_t. For instance, the effect of a policy instrument could depends on the level of some macro variable (e.g., the unemployment rate Auerbach and Gorodnichenko, 2012, in the case of fiscal policy).

To facilitate a comparison with assumptions 4 and 5 we impose

Assumption 6. *$\mathfrak{R}_t$ is independent of p_t.*

When compared to assumptions 4 and 5 assumption 6 rules out any dependence of the derivatives on the policy choice. Given this assumption we obtain the following result.

Proposition 6. *Given Assumptions 4 and 6, we have $p_t^0 + \delta_t^{g*} = p_t^*$.*

This proposition is the same as stated in Proposition 1 part 2, it implies that under a linearity assumption the GOPP, can be used to determine the distance to the constrained optimal policy p_t^*. In this case the policy problem is strictly convex in p_t and there exists a unique minimizer p_t^* which can be reached in one-step regardless of the starting point p_t^0.

[40] The neighborhood $\mathcal{N}(p_t^*, e)$ is defined in the usual way: $\mathcal{N}(p_t^*, e) = \{p_t \in \mathcal{D} : \|p_t - p_t^*\| < e\}$.

[41] Note that for any specific model f the neighborhood can be determined exactly.

Inference for the generalized OPP

The previous part showed that it was relatively easy to define a generalized OPP statistic that retained the ability to detect optimization failures. Notably proposition 4 provided a strong result that illustrates the ability of the GOPP statistic to detect optimization failures.

In this section we discuss how inference can proceed for the generalized OPP statistic. To estimate the derivative function $\mathfrak{R}_t$ with minimal assumptions non-parametric IV methods need to be used.[42] In particular, applicable methods are described in Newey, Powell and Vella (1999) and more recently in Su and Ullah (2008). We omit the details, but we stress that invariably these methods require (i) a stable policy regime (similar to Assumption 2), (ii) exogenous variation in the form of instrumental variables (similar to Assumption 3) and (iii) a variety of regularity conditions. Importantly, the regularity conditions require the existence of higher order derivatives of f with respect to p_t, implying that in order to conduct inference using the GOPP statistic requires more assumptions on the model, e.g. besides Assumption 4.

Apart from the estimation of the derivative function $\mathfrak{R}_t^0$ inference for the GOPP proceeds exactly the same as discussed in Section 6. Moreover, similar as in Section 4.3, we can define a subset GOPP for which the properties of Proposition 3 carry over.

Appendix B: The constrained OPP

In this appendix we show that the OPP approach can be easily adjusted to take into account constraints on the policy choices. Such constraints could arise for instance from a priori commitments dating from before time t. For instance, in the context of forward guidance, a monetary policy maker could have promised interest rates that will be "lower for longer", as the Fed did in December 2012 (e.g., Clarida et al., 2020).

The simple idea is to perturbate the policy maker's loss function subjected to the additional constraints. In other words, instead of testing whether the *gradient of the loss function* is zero at the optimal policy choice, we will test whether the *gradient of the Lagrangian* is zero at the constrained-optimal choice. In other words, the least-square regression underlying the OPP thought experiment will be replace by a constrained least-squares regression, see Section 8.2 in Hansen (2020).

For illustrative purposes we postulate that the constraints can be formulated in terms of the linear system

$$Cp_t = c \ , \tag{28}$$

where C is a known $\#r \times K$ matrix with full row rank and c is a known $\#r \times 1$ vector. To incorporate these constraints, we re-write the policy maker's problem as

$$\min_{p_t} \mathcal{L}_t \qquad s.t. \qquad Cp_t = c \ .$$

Similar to Assumption 1 in the unconstrained case, the assumption posits the existence of at least one reaction function $\mathcal{G}^{\text{copt}}$ that minimizes the loss function.

Assumption 7. Existence of constrained optimum

[42]If the parametric form of f is known nonlinear GMM methods can typically be adopted to estimate $\mathfrak{R}_t$.

There exists a non-empty set $\mathcal{G}^{\text{copt}}$ such that

$$\mathbb{E}_t\|\mathcal{W}^{1/2}Y_t(g,0)\|^2 \leq \mathbb{E}_t\|\mathcal{W}^{1/2}Y_t(\tilde{g},\tilde{\epsilon}_t)\|^2, \qquad \forall\ g \in \mathcal{G}^{\text{opt}},\ \tilde{g} \in \mathcal{G}\backslash\mathcal{G}^{\text{copt}}, \tilde{\epsilon}_t \neq 0\ ,$$

such that $Cp_t = c$ and $C\tilde{p}_t = c$, for $p_t = g(y_t, X_t)$ and $\tilde{p}_t = \tilde{g}(y_t, X_t) + \tilde{\epsilon}_t$, where $Y_t(g, \epsilon_t) = \mathcal{R}(g)p_t + f(X_t; g) + \xi_t$.

To construct a policy perturbation that takes into account the constraint on policy, we simply mimic the constrained policy maker's problem and consider

$$\min_{\delta_t} \mathbb{E}_t\|\mathcal{W}^{1/2}\tilde{Y}_t\|^2 \quad s.t. \qquad C(p_t^0 + \delta_t) = c \qquad \text{where} \quad \tilde{Y}_t = \mathcal{R}^0(p_t^0 + \delta_t) + f(X_t; g^0) + \xi_t\ .$$

which gives the constrained OPP (cOPP) δ_t^{c*}

$$\delta_t^{c*} = \delta_t^* - (\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}C'\mu_t$$

where $\delta_t^* = (\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_tY_t^0$ is the unrestricted OPP and $\mu_t = \left(C(\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}C'\right)^{-1}C\delta_t^*$ is the Lagrange multiplier. Combining gives

$$\delta_t^{c*} = \delta_t^*\,(I - Q) \tag{29}$$

where $Q = \mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}C'\left(C(\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}C'\right)^{-1}C$.

Expression (29) is the standard constrained least-square estimator, where the unconstrained OPP is corrected to incorporate the shadow value of the constraint. The constrained OPP has the same properties as the unconstrained OPP but now with respect to Assumption 7.

Proposition 7. *Given an economy defined by equations* (4) *and* (5)*, and policy constraints* (28)*, we have that under Assumption 7:*

1. *$\delta_t^{c*} \neq 0$ implies that $g^0 \notin \mathcal{G}^{\text{copt}}$ or $\epsilon_t^0 \neq 0$;*
2. *$p_t^0 + \delta_t^{c*} = p_t^{c*}$, where $p_t^{c*} = \arg\min_{p_t\in\mathbb{R}} \mathbb{E}_t\|\mathcal{W}^{1/2}(\mathcal{R}^0p_t + f(X_t; g^0) + \xi_t)\|^2$ s.t. $Cp_t = c$.*

The proof is similar as for Proposition 1.

We conclude that if the policy problem is constrained in a known way, the OPP approach continues to apply. Going beyond linear restrictions works in exactly the same way, but the OPP can then generally not be derived in closed form. However, the problem subjected to nonlinear or inequality constraints can still be solved numerically if the inputs $\mathcal{R}^0$ and $\mathbb{E}_tY_t^0$ are known or estimable.

Example: the OPP in Barro-Gordon

To illustrate how the OPP framework can be extended to incorporate a constraint such as a commitment to a rule, we study what the OPP and the constrained OPP would compute in the seminal model of Barro and Gordon (1983).

We consider a static version of the Barro and Gordon (1983) model, where the loss function is given by

$$\mathcal{L} = \frac{1}{2}\left[\pi^2 + (y - ky^n)^2\right]$$

with y denoting output, y^n the efficient output level, which is independent of monetary policy. With $k > 1$, the policy maker wants to push outout above its efficient level. For simplicity we have set the target $\pi^* = 0$ and the preference parameter $\lambda = 1$. Output is determined by

$$y = y^n + \frac{1}{\kappa}(\pi - \pi^e) \tag{30}$$

where π^e are private sector inflation expectations. To remain within our general framework (5), we postulate without loss of generality that the central bank can use its policy instrument i to influence inflation according to

$$\pi = \mathcal{R}_\pi i + tip \tag{31}$$

where the coefficients $\mathcal{R}_\pi$ captures the response of inflation to the policy rate i and where tip denotes "terms independent of policy". Similarly, we can write $\pi^e = \mathcal{R}_{\pi^e} i + tip$. From the Phillips curve-type relation (30), we immediately get $y = \mathcal{R}_y i + tip$ with $\mathcal{R}_y = \frac{1}{\kappa}(\mathcal{R}_\pi - \mathcal{R}_{\pi^e})$.

In this setting, consider the case where the central bank follows the Barro-Gordon optimal rule and *committed* itself to set

$$\pi^0 = 0$$

by systematically setting the interest rate i^0 such that (31) implies $\pi = 0$ at all time. Under that rule, inflation expectations would be anchored at zero ($\pi^e = 0$), which implies $\mathcal{R}^0_{\pi^e} = 0$ and $\mathcal{R}^0_y = \frac{1}{\kappa}\mathcal{R}^0_\pi$.

First, if we computed the (unconstrained) OPP, we would get

$$\delta^* = \frac{\kappa \mathcal{R}^{0^{-1}}_\pi}{1 + \kappa^2}(1 - k)y^n$$

using $\mathcal{R}^0 = \left(\mathcal{R}^0_\pi, \mathcal{R}^0_y\right)' = \mathcal{R}^0_\pi \left(1, \frac{1}{\kappa}\right)'$ and $Y^0_t = (\pi^0, y^0 - ky^n)' = (0, (1-k)y^n)'$. Even though the policy maker is following the optimal Barro-Gordon rule $\pi^{\text{opt}} = 0$, the OPP is non-zero. In fact, the OPP is capturing the magnitude of the Barro-Gordon temptation to renege on a previous commitment: since inflation expectations are fixed at zero under the rule, the optimal discretionary policy is to "cheat" and create surprise inflation by lowering the policy rate below i^0.

Now, let us compute the constrained OPP, which takes into account the policy maker's commitment. By incorporating the constraint $\pi^0_t = 0$, the cOPP is given by

$$\delta^{c*} = \delta^* (I - Q) = 0$$

as we can verify that $Q = I$ when $C = \mathcal{R}^0_\pi$ in this setting with only one policy instrument. In other words, the constrained OPP correctly identifies that the policy is set optimally given the previous commitment.

To recap, the key point underlying our analysis is that a researcher who aims to detect optimization failures in the presence of commitment need to re-write the policy maker's loss function in order to incorporate all previous promises. Of course, this relies on these previous commitments being publicly available. We think this is reasonable for two reasons. First, assuming knowledge of previous commitments is in line with the starting point of our paper —positing knowledge of the policy maker's objectives—, and consistent with our aim to help policy makers in their decision making process. Second, to the extent that previous commitments are made in order to steer private agents' expectations, these commitments should be made public and thus should be available to the researcher.

Appendix C: Proofs

Proof of Proposition 1. Part 1. Given that $\mathcal{L}_t$ is a strictly convex function of p_t, a sufficient condition for the optimality of p_t^0 is

$$\left.\frac{\partial \mathcal{L}_t}{\partial p_t}\right|_{p_t=p_t^0} = 0 \ .$$

Using the definition of the loss function (6) we find that

$$\left.\frac{\partial \mathcal{L}_t}{\partial p_t}\right|_{p_t=p_t^0} = 2\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 = 0 \ .$$

Since, $\mathcal{R}^0$ has full column rank and the diagonal elements of $\mathcal{W}$ are non-zero, we have $\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0 \succ 0$ and thus

$$\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 = 0 \qquad \Rightarrow \qquad \delta_t^* = (\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 = 0 \ .$$

Since Assumption 1 imposes that $\mathcal{L}_t$ is minimized for any $g \in \mathcal{G}^{\text{opt}}$ and $\epsilon_t = 0$, which implies

$$\left.\frac{\partial \mathcal{L}_t}{\partial p_t}\right|_{p_t=p_t^{\text{opt}}} = 0 \ ,$$

for $p_t^{\text{opt}} = g^{\text{opt}}(y_t, X_t)$ where $g^{\text{opt}} \in \mathcal{G}^{\text{opt}}$,[43] we have that $\delta_t^* \neq 0$ implies that $p_t^0 \neq p_t^{\text{opt}}$ which can arise because $g^0 \notin \mathcal{G}^{\text{opt}}$ or $\epsilon_t^0 \neq 0$.
Part 2. Let $p_t = p_t^0 + \delta_t$, and plug this in the gradient to get

$$\frac{\partial \mathcal{L}_t}{\partial p_t} = 2\mathcal{R}^{0\prime}\mathcal{W}(\mathbb{E}_t Y_t^0 + \mathcal{R}^0\delta_t)$$

Setting the gradient to zero and solving for δ_t gives

$$\delta_t^* = -(\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 \ ,$$

which implies that in order to minimize the loss function with respect to p_t we must take $p_t^* = p_t^0 + \delta_t^*$. □

Proof of Proposition 2. First, if $\delta_t^* = 0$ the claim is trivially true. Now let $\delta_t^* \neq 0$. Since, under $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$, $\delta_t^* \neq 0$ can only be caused by $\epsilon_t^0 \neq 0$, we have that $p_t^{\text{opt}} - p_t^0 = \epsilon_t^0$, and

[43] To show this formally, suppose that $\left.\frac{\partial \mathcal{L}_t}{\partial p_t}\right|_{p_t=p_t^{\text{opt}}} \neq 0$, then one could find $\epsilon_t^* \neq 0$ such that $p_t^{\text{opt}} + \epsilon_t^*$ satisfies $\left.\frac{\partial \mathcal{L}_t}{\partial p_t}\right|_{p_t=p_t^{\text{opt}}+\epsilon_t^*} = 0$ which since $\mathcal{L}_t$ is strictly convex in p_t implies that $p_t = p_t^{\text{opt}} + \epsilon_t^*$ leads to a lower loss $\mathcal{L}_t$, thus contradicting Assumption 1.

thus by proposition 1 part 2., we have $\delta_t^* = \epsilon_t^0$ (as under H_0, $p_t^* = p_t^{\text{opt}}$). Since, $\mathbb{E}(\epsilon_t^0|\mathcal{F}_t) = 0$, we have that $\mathbb{E}(\delta_t^*|\mathcal{F}_t) = \mathbb{E}(\epsilon_t^0|\mathcal{F}_t) = 0$. □

Proof of Proposition 3. The proof is identical to the proof of Proposition 1 for parts 1 and 2, and identical to the proof of Proposition 2 for part 3. □

Proof of Proposition 4. By Assumption 4 part 1, the loss function $\mathbb{E}_t\|f(p_t, X_t; g) + \xi_t\|^2$ is continuously differentiable on $\mathcal{D}$, thus by Lemma 4.3.1 in Dennis and Schnabel (1996) and Assumption 4 part 2 the optimal policy $p_t^{\text{opt}} = g^{\text{opt}}(y_t, X_t)$ for any $g^{\text{opt}} \in \mathcal{G}^{\text{opt}}$ satisfies the gradient condition $\frac{\partial}{\partial p_t}\mathbb{E}_t\|f_t(p_t, X_t; g^{\text{opt}}) + \xi_t\|^2|_{p_t=p_t^{\text{opt}}} = 0$. Hence, if p_t^0 is optimal we must have that $\mathfrak{R}_t^{0\prime}\mathcal{W}\mathbb{E}_tY_t^0 = 0$, with $Y_t^0 = f(p_t^0, X_t; g^0) + \xi_t$, which since $\mathfrak{R}_t^0$ has full column rank (see Assumption 4) implies that $\delta_t^* = -(\mathfrak{R}_t^{0\prime}\mathcal{W}\mathfrak{R}_t^0)^{-1}\mathfrak{R}_t^{0\prime}\mathcal{W}\mathbb{E}_tY_t^0$ must satisfy $\delta_t^* = 0$ if p_t^0 is optimal. If $\delta_t^* \neq 0$ we have that either $g^0 \notin \mathcal{G}^{\text{opt}}$ or $\epsilon_t^0 \neq 0$ or both. □

Proof of Proposition 5. For convenience let $\mathbb{E}_tY_t^* = \mathbb{E}_tf(p_t^*, X_t; g^0) + \mathbb{E}_t\xi_t$. Let κ be a fixed constant in $(1, \mu_{\min}/c)$ and note that such constant exists as $c < \mu_{\min}$ by assumption 5 part 1. Note that $\mathfrak{R}_t(p_t^0, X_t; g^0)'\mathcal{W}\mathfrak{R}_t(p_t^0, X_t; g^0)$ is non-singular and thus there exists a constant $\epsilon_1 > 0$ such that

$$\|(\mathfrak{R}_t(p_t^0, X_t; g^0)'\mathcal{W}\mathfrak{R}_t(p_t^0, X_t; g^0))^{-1}\| \leq \frac{\kappa}{\mu_{\min}} \qquad \forall p_t^0 \in \mathcal{N}(p_t^{\text{opt}}, \epsilon_1) \ .$$

Let

$$e = \min\left\{\epsilon_1, \frac{\mu_{\min} - \kappa c}{\kappa\Delta\gamma}\right\} \ .$$

Now consider

$$\begin{aligned}
p_t^0 + \delta_t^{g*} - p_t^* &= p_t^0 - p_t^* - (\mathfrak{R}_t^{0\prime}\mathcal{W}\mathfrak{R}_t^0)^{-1}\mathfrak{R}_t^{0\prime}\mathcal{W}\mathbb{E}_tY_t^0 \\
&= -(\mathfrak{R}_t^{0\prime}\mathcal{W}\mathfrak{R}_t^0)^{-1}\left[\mathfrak{R}_t^{0\prime}\mathcal{W}\mathbb{E}_tY_t^0 - (\mathfrak{R}_t^{0\prime}\mathcal{W}\mathfrak{R}_t^0)(p_t^* - p_t^0)\right] \\
&= -(\mathfrak{R}_t^{0\prime}\mathcal{W}\mathfrak{R}_t^0)^{-1}\left[\mathfrak{R}_t^{0\prime}\mathcal{W}\mathbb{E}_tY_t^* - \mathfrak{R}_t^{0\prime}\mathcal{W}\left(\mathbb{E}_tY_t^* - \mathbb{E}_tY_t^0 - \mathfrak{R}_t^0(p_t^* - p_t^0)\right)\right] \ .
\end{aligned}$$

Now the Lipschitz assumption (e.g. Assumption 5 - part 3) implies that

$$\|\mathbb{E}_tY_t^* - \mathbb{E}_tY_t^0 - \mathfrak{R}_t(p_t^* - p_t^0)\| \leq \frac{\gamma}{2}\|p_t^* - p_t^0\|^2 \ ,$$

see Lemma 4.1.12 in Dennis and Schnabel (1996). Note that at the constrained optimum p_t^* we have $R_t(p_t^*, y_t, X_t; g^0)'\mathcal{W}\mathbb{E}_tY_t^* = 0$, and thus we have by Assumption 5 - part 2 that

$$\|\mathfrak{R}_t^{0\prime}\mathbb{E}_tY_t^*\| \leq c\|p_t^0 - p_t^*\| \ .$$

Combining the bounds gives

$$\begin{aligned}
\|p_t^0 + \delta_t^{g*} - p_t^*\| &\leq \|(\mathfrak{R}_t^{0\prime}\mathcal{W}\mathfrak{R}_t^0)^{-1}\|\left[\|\mathfrak{R}_t^{0\prime}\mathcal{W}\mathbb{E}_tY_t^*\| + \|\mathfrak{R}_t\|\|\mathcal{W}\|\|\mathbb{E}_tY_t^* - \mathbb{E}_tY_t^0 - \mathfrak{R}_t^0(p_t^* - p_t^0)\|\right] \\
&\leq \frac{\kappa}{\mu_{\min}}\left[c\|p_t^0 - p_t^*\| + \frac{\gamma\Delta}{2}\|p_t^* - p_t^0\|^2\right] \ ,
\end{aligned}$$

which can be simplified using the definition of e to obtain

$$
\begin{aligned}
\|p_t^0 + \delta_t^{g*} - p_t^*\| &\leq \|p_t^0 - p_t^*\| \left[\frac{\kappa c}{\mu_{\min}} + \frac{\kappa \gamma \Delta}{2\mu_{\min}} \|p_t^0 - p_t^*\| \right] \\
&\leq \|p_t^0 - p_t^*\| \left[\frac{\kappa c}{\mu_{\min}} + \frac{\mu_{\min} - \kappa c}{2\mu_{\min}} \right] \\
&= \frac{\kappa c + \mu_{\min}}{2\mu_{\min}} \|p_t^0 - p_t^{\text{opt}}\| \\
&\leq \|p_t^0 - p_t^*\| \ .
\end{aligned}
$$

This completes the proof. □

Proof of Proposition 6. Note that since $\mathfrak{R}_t$ is independent of p_t under assumption 6 the derivative of every element of $\mathfrak{R}_t$ with respect to p_t is equal to zero. Therefore when we expand the model f around $\delta_t = 0$ we have

$$
f(p_t^0 + \delta_t, X_t; g^0) = f_t(p_t^0, X_t; g^0) + \mathfrak{R}_t^0 \delta_t
$$

which implies that the policy problem is strictly convex in δ_t. Now note that the first order conditions at $p_t^0 + \delta_t$ are equal to $\mathfrak{R}_t^{0\prime} \mathcal{W} \mathbb{E}_t[f_t(p_t^0 + \delta_t, X_t; g^0) + \xi_t] = 0$, which using the expansion can be rewritten as

$$
\mathfrak{R}_t^{0\prime} \left(\mathbb{E}_t Y_t^0 + \mathfrak{R}_t^0 \delta_t \right) = 0
$$

solving for δ_t gives δ_t^{g*}, which implies that for $p_t^0 + \delta_t$ to be optimal we must take $\delta_t = \delta_t^{g*}$. □

Figure 3: Fed funds rate policy in June 1990

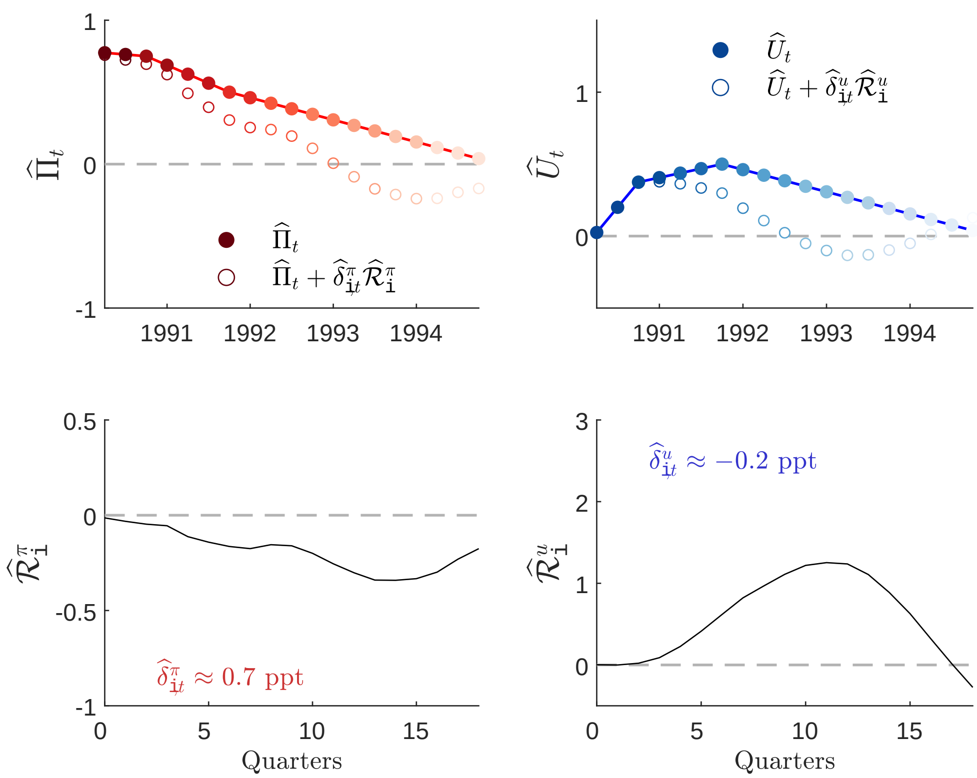

Notes: Top panel: median FOMC forecasts for the inflation and unemployment gaps as of 1990-M6. Bottom panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock. In red (blue) is the OPP $\delta^{*\pi}$ (δ^{*u}) for a strict inflation (unemployment) targeter.

Figure 4: Fed funds rate policy in April 2008

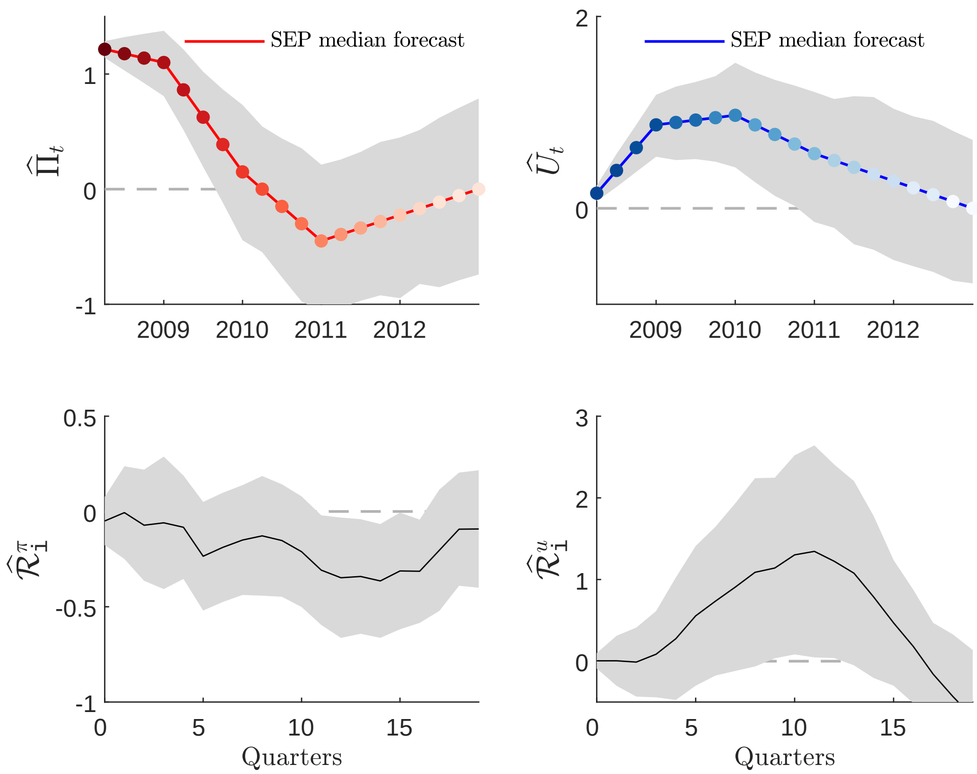

Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2008-M4 (in red and blue) along with the 68 percent confidence bands. Bottom panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock with the 95 percent confidence intervals.

Figure 5: Slope policy in April 2010

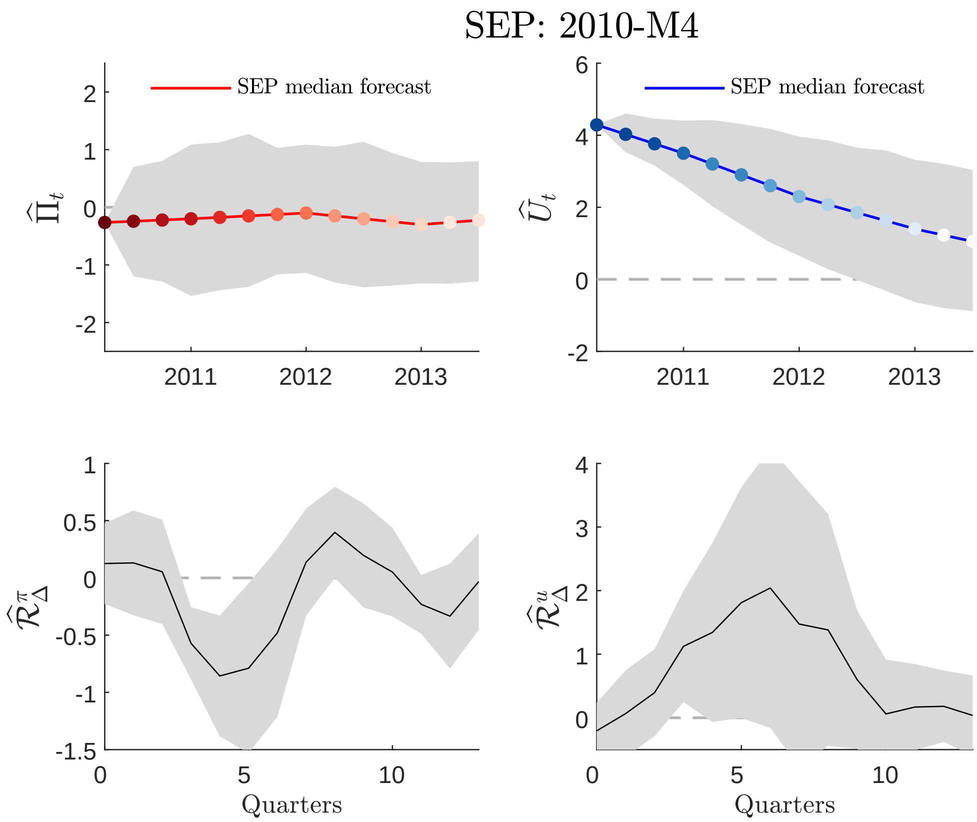

Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2010-M4 (in red and blue) along with the 68 percent confidence bands uncertainty. Bottom panel: impulse responses of the inflation and unemployment gaps to a slope policy shock with the 95 percent confidence intervals.

Figure 6: A sequence of OPP for Fed monetary policy (1990-2018)

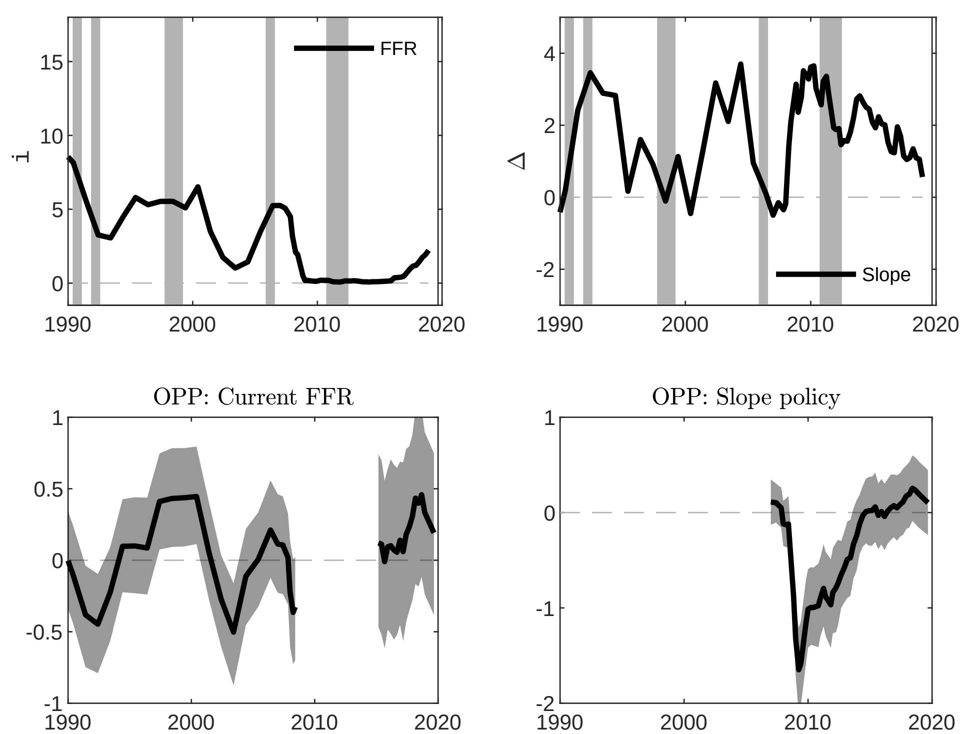

Notes: Top panels: the fed funds rate ("FFR", left-panel) and the difference between the 10-year bond yield and the fed funds rate ("slope", right panel). Grey bars denote NBER recessions. Bottom panels: OPP for the fed funds rate at time t (left-panel) and OPP for the slope instrument at time t (right-panel). The grey area captures both impulse response and mis-specification uncertainty: OPP values outside the shaded-areas can be excluded with a 68 percent probability.

Table 1: OPP ESTIMATES FOR CASE STUDIES

OPP: current FFR	1990M6	2008M4	OPP: slope policy	2010M4
$\widehat{\delta}_{\mathsf{i},t}$	-0.10 [-0.5,0.3]	-0.41 [-0.8,-0.1]	$\widehat{\delta}_{\Delta,t}$	-1.02 [-1.4,-0.6]
$\widehat{\delta}^{\pi}_{\mathsf{i},t}$	0.7	-0.1	$\widehat{\delta}^{\pi}_{\Delta,t}$	-0.1
$\widehat{\delta}^{u}_{\mathsf{i},t}$	-0.2	-0.5	$\widehat{\delta}^{u}_{\Delta,t}$	-1.4

Notes: $\widehat{\delta}_{\mathsf{i},t}$ and $\widehat{\delta}_{\Delta,t}$ denote the mean estimates for the dual-mandate OPP ($\hat{\lambda} = .6$) for the short rate policy and the slope policy, respectively. In brackets, the 68 percent confidence interval from estimation uncertainty *and* model mis-specification uncertainty. Similarly, $\widehat{\delta}^{\pi}_{s,t}$ and $\widehat{\delta}^{u}_{t}$ denote the mean OPP estimates for a strict inflation targeter ($\lambda = 0$) and a strict unemployment targeter ($\lambda = \infty$), for $s = \mathsf{i}, \Delta$.

Table 2: TESTING THE OPTIMALITY OF THE FED REACTION FUNCTION

OPP	Current FFR		Slope policy	
	OLS	IV	OLS	IV
c	$-.06$ [.03]	.01 [.10]	$-.00$ [.05]	$-.02$ [.05]
b_π	$-.07$ [.08]	.02 [.10]	$-.06$ [.03]	$-.09$ [.03]
b_u	$-.31$ [.02]	$-.28$ [.02]	$-.33$ [.02]	$-.33$ [.02]
Wald test (p-val. joint sig.)	[< .01]	[< .01]	[< .01]	[< .01]
Sample	1990-2007	1990-2007	2007-2018	2007-2018

Note: As instrumental variables for the IV regressions, we use two lags of inflation and unemployment. Newey-West standard-errors are reported in brackets.