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

November 30, 2020

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 economic model. 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 objectives. We illustrate the OPP in a stylized New Keynesian model as well as empirically for US monetary policy.

JEL classification: C14, C32, E32, E52.

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

*We thank Martin Eichenbaum, Ethan Ilzetzki, Jinil Kim, Sylvain Leduc, Huiyu Li, Emi Nakamura, Nicolas Petrosky-Nadeau, Adam Shapiro, Jón Steinsson and seminar participants at Berkeley, Bank of England, McGill, Ohio State, University of Ottawa, 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

The evaluation of macroeconomic policy decisions is typically based on the careful analysis of a specific economic model, see Chari, Christiano and Kehoe (1994) and Woodford (2003) for prominent examples in the context of fiscal and monetary policy, respectively. While this approach can produced great insights, a worry for policy makers is that the underlying model structure may be too stylized to inform policy decisions made in complex and data-rich real-life settings.

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, those are 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, a government interested in smoothing business cycle fluctuations but concerned about excessive deficits, or a government interested in promoting growth but concerned about income inequality.

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 insight is that the OPP 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 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".

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.

In general, an optimization failure could occur for two main reasons: (i) the policy maker is using a non-optimal reaction function —a systematic optimization failure—, or (ii) the policy maker made a one-time optimization failure —a discretionary mistake—. A *single* OPP statistic cannot distinguish between the different sources of failures, i.e. a systematic or a discretionary failure. However, a *sequence* of OPP statistics can separate the two sources of optimization failures. Intuitively, if the policy maker's reaction function 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 aimed at rejecting the null that the reaction function is optimal. This test allows 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 policy rule.

Second, in an empirical study we revisit US monetary policy over the 1990-2018 period. We use FOMC forecasts for inflation and unemployment from historical records of monetary reports to Congress which are conditional on the Fed following an optimal policy, as judged by the FOMC members. This allows us to assess the optimality of the Feds actions back to 1990. We summarize the Feds 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. In some instances the uncertainty in the causal effects and forecasts makes it impossible to formally reject optimality, but the optimality of unconventional monetary policy operations during the great recession is convincingly rejected. In particular, during the Great recession the optimality deviations are large, peaking at -2ppt at the onset of the crisis, and we can reject optimality over 2009-2012. This suggests that unconventional monetary policy measures LSAP or QE could have been used more aggressively to bring the slope of the yield curve down in line with optimality.

We then use the sequence of OPPs over 1990-2018 to study the optimality of the Fed's reaction function. 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 reaction function, in that the Fed could have achieved a lower loss over the 1990-2018 by responding more aggressively to unemployment.

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 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 from the US. Section 8 summarizes the framework and provides some potential avenues for further research.

Relation to literature

In the wake of 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 question has been to derive, from first principles, the appropriate loss function that the policy maker *should* be considering. To give a few examples in monetary policy, 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 preferences *as given*, and we propose a methodology for assessing whether the proposed policy choices are optimal.[1] This is a different objective, which does not require explicitly specifying the underlying structure of the economy. As a consequence, the OPP testing approach is attractive for evaluating *practical* macroeconomic policy decisions, which do not rely on one specific macroeconomic model and are often based on heuristics; combining multiple models, judgement calls and instincts. As far as we know, there are no alternatives in the literature (e.g. Bénassy-Quéré et al., 2018; Kocherlakota, 2019).

In a public finance context, the OPP framework shares important similarities with the sufficient statistic approach, 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 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 (e.g. Chetty, 2009; Kleven, 2020). 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 the straightjacket of committing to one specific model.[2]

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, judgement, etc..). While the jury is still out on determining the best forecasting method, we

[1] As such the OPP approach is not tied to the debate on rules versus discretion (e.g. Kydland and Prescott, 1977; Barro and Gordon, 1983). We take the loss function and the policy instruments as given, and hence the set-up can correspond to a policy problem under commitment or discretion.

[2] Another benefit of using policy makers' forecast as an input is that our approach can be readily used to inform policy decisions, without any change to current operating procedures.

note that policy makers' forecasts often do perform well, and are thus natural alternatives to evaluate the expected equilibrium allocation under the desired policy choice.[3]

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 seen as a key element of the forecast-targeting rules used in practice by policy makers, notably in 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).[4] However, unless the forecast targeting rule is tied to a specific model (Woodford, 2010; Giannoni and Woodford, 2017), the informal "looking good" criterion used in practice 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.[5]

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 inflation-unemployment mandate

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

The central bank has one instrument p, e.g., the short-term interest rate, that affects each

[3]For instance, Romer and Romer (2000), Sims (2002) and Gavin and Mandal (2003) found that the Board staff forecasts are often comparable or more accurate to various statistical benchmark models.

[4]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).

[5]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 2percent 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 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.

mandate according to

$$\underbrace{\begin{bmatrix} \pi \\ u \end{bmatrix}}_{Y} = \underbrace{\begin{bmatrix} f_\pi(p) \\ f_u(p) \end{bmatrix}}_{f(p)} \tag{1}$$

where the function $f(p)$ captures how policy affects the vector of mandates $Y = (\pi, u)'$. The central bank proposes implementing the policy p^0, which implies the equilibrium (Y^0, 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.[6]

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 can assess the optimality of p^0 by directly computing the gradient

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

At the optimum, the causal effect of policy – $\mathcal{R}^0$ – should be orthogonal to the allocation $Y^0 = f(p^0)$ at p^0, i.e., $\mathcal{R}^{0'} Y^0 = 0$, and it is not possible to better stabilize one mandate without destabilizing the other one by the exact same amount, leaving welfare unchanged.[7]

The simple insight idea 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) the causal effects $\mathcal{R}^0$ are estimable, and (ii) Y^0 is typically available from the policy maker.

First, methods from the treatment-effect literature, most notably instrumental variable methods, can be adopted to estimate $\mathcal{R}^0$ with minimal assumptions. Second, 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.[8]

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 rely on a complete specification for $f(\cdot)$ in order to determine the optimal allocation and compare such allocation to (p^0, Y^0). Since the allocation is, except in a few simple cases, not available in closed form, optimal policy analysis has traditionally relied

[6]For instance, if the computation of Y^0 requires the combination of multiple models, instinct and judgement calls, the function $f(.)$ will not be available in closed-form. Note that such heuristic approach are very common in practice.

[7]At the optimum $\mathcal{R}^0_\pi \pi^0 + \mathcal{R}^0_u u^0 = 0$. If a 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 ($\mathcal{R}^0_u u^0 = -\mathcal{R}^0_\pi \pi^0 > 0$).

[8]Note that we do not have expectations in this toy model, but they will figure prominently in our general framework.

on numerical methods, which can be computationally demanding. By exploiting the information published by policy makers, we can bypass this problem, and our setting remains computationally trivial.

In practice, both $\mathcal{R}^0$ and Y^0 are not known exactly, and the gradient 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, the gradient can only be measured with uncertainty, and our evaluation of a policy choice will resemble a hypothesis test: a statement that the policy is not optimal for some confidence level.[9] The larger the uncertainty around the effect of the policy instrument or the larger the uncertainty around the policy maker's estimate of Y^0 and the weaker will be our test.

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 underlying structure of the economy and the policy maker's objectives and choices. In general, 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.

We start by postulating a general set of equations that describe the economy. Based on this underlying structure we define the policy maker's objectives and the corresponding policy choices.

The economy

The economy is defined for an $M \times 1$ vector of policy objectives y_t that the policy maker aims to stabilize, for instance inflation, unemployment or GDP growth. We postulate that these state variables can be described at time t for any horizon $h = 1, \ldots, H$, by the model

$$y_{t+h} = \mathcal{R}_h(g)p_t + f_h(y_t, X_t; g) + \xi_{t+h} \ , \qquad (3)$$

which depends on three components: (i) the $M \times K$ matrix $\mathcal{R}_h(g)$ that captures the causal effects of the policy choice p_t, (ii) an arbitrary function $f_h()$ which depends on the contemporaneous variables y_t and additional time t measurable variables X_t, and (iii) the future

[9] In other words, incorporating these two sources of uncertainty allows our approach to be robust to parameter uncertainty and model mis-specification.

shocks ξ_{t+h} that satisfy $\mathbb{E}_t\xi_{t+h} = 0$ for $h > 0$. The structure of the economy, as captured by $\mathcal{R}_h(g)$ and $f_h(y_t, X_t; g)$, may depend on the reaction function g of the policy maker, which is explained below together with p_t. The horizon H is arbitrary and can be considered infinite.

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.[10] 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 (3) simply by iterating forward.

The policy variables p_t

The policy maker has a set of instruments to steer the state variables y_t. For instance, a fiscal policy maker sets taxes, government spending and transfers (e.g. Alesina, Favero and Giavazzi, 2019), alternatively a monetary policy maker sets the short-term interest rate, can choose to buy/sell longer-maturity securities and adjust the size of its balance sheet through Quantitative Easing (QE) type programs.

For each policy instrument $j \in 1, \ldots, J$, the policy maker typically determines a "policy plan" which sets the instrument's value at time t as well as its expected path $\left(p_{j,t|t}, \ldots, p_{j,t|t+H}\right)$, $j = 1, \ldots, J$, where $p_{j,t|t+h}$ denotes 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 choice vector

$$p_t = \left(p_{1,t|t}, \ldots, p_{1,t|t+H}, \ldots, p_{J,t|t}, \ldots, p_{J,t|t+H}\right) \ ,$$

which implies that $K = J(H+1)$. Without loss of generality, the policy choice p_t can be written as a the sum of two components: (i) a systematic component whereby policy is set in a systematic fashion as a function of time t observables, and (ii) a discretionary component:

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

where g is the reaction function that depends on the observables $y_t = (y_{1,t}, \ldots, y_{M,t})'$ and possibly additional observables X_t. The discretionary component ϵ_t is, by construction, uncorrelated with the variables y_t and X_t. The reaction function g enters the model (3) for y_{t+h} via $\mathcal{R}_h(g)$ and $f_h(y_t, X_t; g)$ and thus determines the equilibrium outcome. In this paper we do not impose restrictions on the reaction function and note that it can be an arbitrarily complex function of time t observables.

[10] Specifically, in Appendix A we show that for $y_{t+h} = f_h(p_t, y_t, X_t; g) + \xi_{t+h}$, where $f(\cdot, \cdot, \cdot; \cdot)$ is some non-linear function of p_t and possibly additional variables y_t, 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 linear case.

To give a concrete example, a central bank's plan for the short-term interest rate is often defined by means of a Taylor, for instance

$$\begin{cases} i_{t|t} = r^e_t + \phi_\pi \mathbb{E}_t \pi_{t+1} + \epsilon_{t,t} \\ i_{t+h|t} = \mathbb{E}_t r^e_{t+h} + \phi_\pi \mathbb{E}_t \pi_{t+h+1|t} + \epsilon_{t,t+h}, \qquad \forall\ h \in \{1,..,H\} \end{cases}$$

with $i_{t+h|t}$ the expected interest path at horizon h, r^e_t the efficient real interest rate and π_t the deviation of inflation from its target. In this example, the policy choice $p_t = (i_{t|t}, \ldots, i_{t|t+H})'$, for arbitrary H, depends on the policy rule $g()$ —a linear function with parameter(ϕ_π)—, as well as current and (expected) future discretionary adjustments, which could be zero.[11]

The policy maker

The policy maker aims to minimize a loss function of the form

$$\mathcal{L}_t = \mathbb{E}_t \sum_{h=0}^{H} \sum_{m=1}^{M} \lambda_m \beta_h \left(y_{m,t+h} - y^*_{m,t+h} \right)^2 \ , \tag{5}$$

where $y_{m,t+h}$ denotes the value of policy objective $m = 1, \ldots, M$ at horizon $h = 0, \ldots, H$. 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 conditional on the time t information set $\mathcal{F}_t$, e.g. $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_t)$, which is given by $\mathcal{F}_t = \{y_s, p_s, X_s, s \leq t\}$, where 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.

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

$$\mathcal{L}_t = \mathbb{E}_t Y_t^{'} \mathcal{W} Y_t, \qquad Y_t = \mathcal{R}(g) p_t + f(y_t, X_t; g) + \Xi_t \ , \tag{6}$$

where the weighting matrix $\mathcal{W}$ includes the non-zero preference parameters and discount factors and is defined as $\mathcal{W} = \text{diag}(\beta \otimes \lambda)$, with $\beta = (\beta_0, \ldots, \beta_H)'$ and $\lambda = (\lambda_1, \ldots, \lambda_M)'$. The generic model for Y_t is obtained by stacking model (3) across h. Formally, we have that $\mathcal{R}(g) = (\mathcal{R}_0'(g), \ldots, \mathcal{R}_H'(g))'$, $f(y_t, X_t; g) = (f_0'(y_t, X_t; g), \ldots, f_H'(y_t, X_t; g))'$ and $\Xi_t = (\xi_t', \ldots, \xi_{t+H}')'$.

[11] Expected future changes capture announcements of future deviations from the systematic conduct of policy, for instance "forward guidance" with promises to keep short-term rates lower-for-longer (e.g. McKay, Nakamura and Steinsson, 2016).

The policy maker's choice

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

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

with $\mathcal{R}^0 \equiv \mathcal{R}(g^0)$ the dynamic causal effects implied by the reaction function g^0. We will generally impose that $\mathcal{R}^0$ has full column rank.

Possible optimization failure

The choices (g^0, ϵ_t^0) that led to the policy paths 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 chose 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. To see this, note that equations (3) and (4) merely provide a generic description of the economy and hence existence is not guaranteed by the model.

To do so let $\mathcal{G}$ be an arbitrary class of reaction functions with $g^0 \in \mathcal{G}$.[12] We postulate that there exists, at least one, reaction function that minimizes the loss function.

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

The assumption imposes the existence of a non-empty set of reaction functions $\mathcal{G}^{\text{opt}}$, which minimize the loss function (5), with corresponding discretionary adjustments equal to zero. The set $\mathcal{G}^{\text{opt}}$ can contain multiple reaction functions and there can be combinations of g and ϵ_t that achieve the same loss as any $g \in \mathcal{G}^{\text{opt}}$. Importantly, since we do not restrict

[12]In this paper we generally do not exploit the structure of $\mathcal{G}$. As an example one could consider $\mathcal{G}$ to be defined as the class of linear functions of y_t and X_t, which would cover a large class of linear policy rules found in the literature.

the class $\mathcal{G}$ the optimal reaction functions can be arbitrarily complex, hence setting the corresponding discretionary adjustments to zero merely normalizes the optimal solution.

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

Definition 1. (Optimization failure)
An optimization failure is any pair of choices g and ϵ_t, such that either $g \notin \mathcal{G}^{\text{opt}}$ and/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. If there are known commitments that the policy maker has to fulfill the assumption can be adjusted to take such commitments into account.[13] We do not explore this possibility in the main text, but in Appendix B we show that the methodology of this paper continues to apply for constrained loss functions.

4 Detecting optimization failures

In this section we take the perspective of a researcher interested in assessing whether the policy maker's choices g^0 and ϵ_t^0 are optimal, i.e. in detecting a possible *optimization failure*, and we derive necessary conditions under which the policy maker's choices are (sub-)optimal. We assume that the researcher does not know the functions $f(\cdot,\cdot;g)$ and $g^0(\cdot,\cdot)$, nor the variables X_t or the discretionary adjustments ϵ_t^0.

Of course, if the function $f(\cdot,\cdot;g)$ in model (3) was 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, or an approximation thereof. 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, there is no broad consensus on the equations that govern the economy, and therefore the assumption that the researcher knows the explicit structure of the economy is often too strong. Assessing whether the policy choices are appropriate, or optimal, in this context requires a different approach.

[13] To give a concrete example, suppose that the policy maker previously committed to keep policy instrument j fixed at $p_{j,t+h} = c_h$, for $h = 1,\dots,H'$, then we would simply define the set $\mathcal{G}^{\text{opt}}$ by

$$\mathbb{E}_t\|\mathcal{W}^{1/2}Y_t(g,0)\|^2 \leq \mathbb{E}_t\|\mathcal{W}^{1/2}Y_t(\tilde{g},\tilde{\epsilon}_t)\|^2,\ \forall\, g \in \mathcal{G}^{\text{opt}},\ \tilde{g} \in \mathcal{G}\backslash\mathcal{G}^{\text{opt}}, \tilde{\epsilon}_t \neq 0\,, \quad s.t. \quad p_{j,t+h} = c_h\ \forall \quad h = 1,\dots,H'$$

where $Y_t(g,\epsilon_t) = \mathcal{R}(g)p_t + f(y_t, X_t; g) + \Xi_t$ and $p_t = g(y_t, X_t) + \epsilon_t$. This defines the set of reaction functions $\mathcal{G}^{\text{opt}}$ that minimize the loss function given the constraints. Appendix B provides a more general treatment.

We proceed in two steps. First, we derive the OPP statistic to detect an optimization failure, which could be due to a systematic mistake —a non-optimal reaction function—, a discretionary mistake —a one-time mistake—, or both. Second, we show that a sequence of OPP statistics can be used to detect systematic optimization failures.

4.1 The OPP statistic

To detect optimization failures we study how the loss function (6) changes when we modify the policy choice to $p_t^0 + \delta_t$, where $\delta_t = (\delta_{1,t}, \dots, \delta_{K,t})'$ is a vector of policy perturbations. If the loss function is lower for some $\delta_t \neq 0$ we may conclude that the choice p_t^0 was not optimal.

To derive a perturbation that (i) has attractive properties and (ii) is computable in practice, we mimic the policy maker's problem around the expected equilibrium allocation $(\mathbb{E}_t Y_t^0, p_t^0)$ and solve

$$\min_{\delta_t} \mathbb{E}_t \tilde{Y}_t' \mathcal{W} \tilde{Y}_t \qquad \text{where} \qquad \tilde{Y}_t = \mathcal{R}^0(p_t^0 + \delta_t) + f(y_t^0, X_t; g^0) + \Xi_t \ . \tag{8}$$

This problem is a simple weighted least squares problem for which the solution is given by

$$\delta_t^* = -(\mathcal{R}^{0\prime} \mathcal{W} \mathcal{R}^0)^{-1} \mathcal{R}^{0\prime} \mathcal{W} \mathbb{E}_t Y_t^0 \ , \tag{9}$$

where $\mathbb{E}_t Y_t^0 = \mathbb{E}_t(\mathcal{R}^0 p_t^0 + f(y_t^0, X_t; g^0) + \Xi_t)$. The statistic δ_t^* is the policy perturbation that forms the basis of our detection method. We refer to δ_t^* as the *optimal policy perturbation*, or OPP. It depends on three quantities: the dynamic causal effects $\mathcal{R}^0$ under g^0, the forecasts of the policy maker $\mathbb{E}_t Y_t^0$, which are conditional on the policy choice $p_t^0 = g^0(y_t^0, X_t) + \epsilon_t^0$ and the weighting matrix $\mathcal{W}$.

The OPP is proportional to the gradient of the loss function evaluated at p_t^0, i.e., $\delta_t^* \propto \mathcal{R}^{0\prime} \mathcal{W} \mathbb{E}_t Y_t^0$, and the intuition that was discussed for the gradient in Section 2 carries over to the OPP. The re-scaling gives the OPP a useful regression interpretation which we discuss below.

The properties of δ_t^* are summarized in the following proposition.

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

1. $\delta_t^* \neq 0$ *implies that* $g^0 \notin \mathcal{G}^{\text{opt}}$ *and/or* $\epsilon_t^0 \neq 0$*;*

2. $p_t^0 + \delta_t^* = p_t^*$*, where* $p_t^* = \arg\min_{p_t \in \mathbb{R}^K} \mathbb{E}_t \| \mathcal{W}^{1/2}(\mathcal{R}^0 p_t + f(y_t, X_t; g^0) + \Xi_t)\|^2$.

All proofs are provided in Appendix D.

The proposition shows that if δ_t^* is not equal to zero there must be a policy mistake as either the reaction function g^0 is non-optimal, a discretionary mistake was made $\epsilon_t^0 \neq 0$, or both. Further, δ_t^* has an economic meaning, and it can be interpreted as the magnitude of the deviation from optimality. Specifically, the OPP is the discretionary adjustment to the policy instruments that would correct the optimization failure, given the policy maker's reaction function g^0.[14]

The results for the OPP statistic in proposition 1 are theoretical as in practice δ_t^* is not observable and needs to be estimated. We discuss inference for δ_t^* in detail in Section 6, but two key points are worth clarifying now.

First, in this paper we work under the assumption that the forecasts $\mathbb{E}_t Y_t^0$ are available to the researcher. This is not unreasonable as many policy makers, such as central banks or fiscal agencies, routinely provide conditional forecasts for their target variables. For now we assume that the published forecasts correspond to $\mathbb{E}_t Y_t^0$, but in Section 6 we allow for the possibility that the forecasts are mis-specified.

Second, the coefficient matrix $\mathcal{R}^0$ captures the dynamic causal effects of p_t on the targets Y_t under g^0. The causal effects of policy are typically not published by the policy maker and thus need to be estimated by the researcher in order to compute δ_t^*. Estimating $\mathcal{R}^0$ will require a constant policy regime for some period prior to time t, the existence of some identification strategy to avoid endogeneity problems and some regularity conditions. These conditions are similar to those found in the treatment-effect or sufficient-statistic approaches, and we spell them out explicitly in Section 6.

Intuition: the OPP as a regression coefficient

Notice how the OPP statistic

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

looks like the formula of a Weighted Least-Square (WLS) regression. In fact, the OPP statistic has a simple regression interpretation, as δ_t^* is minus the coefficient estimate of a regression of $\mathbb{E}_t Y_t^0$ on $\mathcal{R}^0$, weighted by the preference matrix $\mathcal{W}$. To get the intuition behind this regression interpretation, assume for now that all targets are weighted equally so $\mathcal{W} = I$,

[14] Technically speaking, the OPP statistic is the first-step of a Gauss-Newton algorithm, an algorithm based on the Newton line search algorithm and designed to minimize loss functions. The Gauss-Newton approximation of the Hessian consists in approximating the Hessian with first-derivatives only (e.g., Nocedal and Wright, 2006). When the control variable (here p_t) has a linear effect on target variables (as we assumed in model (3)), the Gauss-Newton algorithm converges in one step, that is the OPP is the discretionary adjustment that would exactly corrects the optimization failure. In more general non-linear cases discussed in Appendix A, the Gauss-Newton algorithm is known to converge under some additional assumptions (loosely speaking, the non-linearities are mild), implying that the OPP improve the original policy choice p_t^0.

and recall that the set of dynamic causal effects $\mathcal{R}^0$ captures the effect of a change in the policy instrument on the policy maker's objectives. The goal of the optimal perturbation δ_t^* is then to use the causal effects ($\mathcal{R}$) in order to minimize the squared deviations of Y_t, weighted by the policy maker's preferences. This is nothing but a regression of Y_t^0 on $\mathcal{R}$, except one with a minus sign in front of the coefficient estimate since the goal is not to best fit the path for Y_t^0, but instead to best "undo" Y_t^0. Now since the future values of Y_t^0 are unknown, the policy maker aims to minimize expected squared deviations, and hence the OPP aims to undo $\mathbb{E}_t Y_t^0$, the expected path for Y_t^0 conditional on the policy choice p_t^0. This gives the WLS-type formula.

Intuition: the OPP as a score test statistic

The OPP can be interpreted as a population score test statistic, or Lagrange multiplier statistic, for testing the null hypothesis $H_0 : g^0 \in \mathcal{G}^{\text{opt}}, \epsilon_t = 0$. Under the null, the gradient (or score) should be zero and thus δ_t^* as well. The benefit of the score test is that it only requires the estimation of the parameters that are not fixed under H_0, in our case the dynamic causal effects and the forecasts. This feature is particularly useful when considering the reaction function g, as by fixing the reaction function at $g = g^0$ under the null, we avoid having to estimate g, which typically cannot be done without strong modeling assumptions.[15]

4.2 Testing the optimality of the policy rule

The OPP statistic δ_t^* is a practical tool to 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 states a sufficient condition that guarantees that a non-zero OPP comes from a discretionary mistake. If we can reject this sufficient condition, 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 that an economy described by equations* (3) *and* (4)*, Assumption 1 holds, and the discretionary adjustment satisfy* $\mathbb{E}(\epsilon_t^0|\mathcal{F}_t) = 0$. *We have that under* $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$*,*

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

[15]Interestingly, this approach has often been exploited in the econometrics literature in settings where the parameter of interest cannot be consistently estimated without strong assumptions (e.g. Stock and Wright, 2000; Kleibergen, 2002, 2005; Andrews and Mikusheva, 2015). A prominent example is the weak instruments problem in GMM, where the score test, or versions thereof are able to avoid the strong instrument assumption by fixing the structural parameters under the null and relying on the score of the objective function to conduct inference. Exactly the same intuition applies in our setting, where the reaction function is the parameter that cannot be estimated without strong assumptions.

The proposition formalizes 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. Loosely speaking, the OPP statistic should resemble a policy shock (e.g., Ramey, 2016) and be orthogonal to any variable in $\mathcal{F}_t$.

Formally, Proposition 2 provides a testable moment condition that enables the detection of systematic failures by studying the relationship between δ_t^* and the elements of the information set $\mathcal{F}_t$. The implementation of such test, i.e. the evaluation of the moment condition (10), requires a sample OPP statistics δ_t^* for different time periods as the population moment needs to be replaced by its sample equivalent. In Section 6 we provide the details for implementing this test. This involves replacing δ_t^* by its estimate and handling possible endogeneity problems to evaluate the moment condition.

4.3 Testing subsets of the policy plan

As we will see in the inference section, estimating the full matrix of causal effects $\mathcal{R}^0$ may be difficult in many practical situations. Instead, the 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 shows that the logic of the OPP test carries to that setting as well. Intuitively, since the OPP framework only aims at testing a necessary condition, 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 $S = (S_a', S_b')$ denote an orthogonal selection matrix and define

$$p_{a,t} = S_a p_t \ , \qquad \text{and} \qquad p_{b,t} = S_b p_t \ ,$$

where S_a is the $K_a \times K$ selection matrix that determines the linear combinations of the policy plan $p_{a,t} = S_a p_t$ that the researcher wants to evaluate. The other linear combinations $p_{b,t} = S_b p_t$ are not of interest and are only implicitly defined, with $(K - K_a) \times K$ selection matrix S_b such that S is an orthogonal matrix.

We give two examples in the context of monetary policy where the policy plan is the expected interest rate path $p_t = (i_{t|t}, \ldots, i_{t|t+H})'$. First, suppose that the researcher is only interested in testing the short rate, then we can take

$$S_a = \begin{bmatrix} 1 & 0 & \ldots & 0 \end{bmatrix} \qquad \text{such that} \qquad p_{a,t} = S_a p_t = i_{t|t} \ .$$

Second, Eberly, Stock and Wright (2019) summarize Fed policy in terms of the short rate $i_{t|t}$ and the slope of the yield curve $s_t = i_{t|t+H} - i_{t|t}$. We can accommodate this choice by considering

$$S_a = \begin{bmatrix} 1 & 0 & \dots & 0 \\ -1 & 0 & \dots & 1 \end{bmatrix} \qquad \text{such that} \qquad p_{a,t} = S_a p_t = \begin{bmatrix} i_{t|t} \\ s_t \end{bmatrix} .$$

To detect optimization failures for the subset of policies $p_{a,t}$ we proceed similarly as before by mimicking the policy maker's problem around the expected equilibrium allocation $(\mathbb{E}_t Y_t^0, p_t^0)$. The difference is now that we only perturb the policy instruments $p_a^0 = S_a p_t^0$. In particular, we consider

$$\min_{\delta_{a,t}} \mathbb{E}_t \tilde{Y}_t' \mathcal{W} \tilde{Y}_t \qquad \text{where} \qquad \tilde{Y}_t = \mathcal{R}_a^0 (p_{a,t}^0 + \delta_{a,t}) + \mathcal{R}_b^0 p_{b,t}^0 + f(y_t^0, X_t; g^0) + \Xi_t \ ,$$

where $\mathcal{R}_a^0 = \mathcal{R}^0 S_a'$ and $\mathcal{R}_b^0 = \mathcal{R}^0 S_b'$. The solution for $\delta_{a,t}$ defines the subset-OPP

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

The subset OPP $\delta_{a,t}^*$ has the same theoretical properties as the full vector OPP, which are summarized in the following proposition.

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

1. $\delta_{a,t}^* \neq 0$ *implies that* $g^0 \notin \mathcal{G}^{\text{opt}}$ *and/or* $\epsilon_{a,t}^0 \neq 0$*, where* $\epsilon_{a,t}^0 = S_a \epsilon_t^0$*;*
2. $p_{a,t}^0 + \delta_{a,t}^* = p_{a,t}^*$*, where* $p_{a,t}^* = \arg\min_{p_{a,t} \in \mathbb{R}} \mathbb{E}_t \| \mathcal{W}^{1/2} (\mathcal{R}_a^0 p_{a,t} + \mathcal{R}_b^0 p_{b,t}^0 + f(y_t, X_t; g^0) + \Xi_t) \|^2$

And, given that the discretionary adjustment satisfy $\mathbb{E}(\epsilon_{a,t}^0 | \mathcal{F}_t) = 0$*, we have that under* $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$*,*

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

The proposition shows that Propositions 1 and 2 continue to apply when we consider subset tests. Note that the implication in part 1 $g^0 \notin \mathcal{G}^{\text{opt}}$ is driven by the fact that $g_a^0 = S_a g^0$ is not optimally set, which implies $g^0 \notin \mathcal{G}^{\text{opt}}$.

The intuition underlying Proposition 3 is similar as before. The main benefit of this extension is that it accommodates a broad array of tests that a researcher might want to consider in practice, such as simply testing the short rate, or any other combination of the policy plan. The requirements are the same as before: the conditional forecasts $\mathbb{E}_t Y_t^0$ must be available, or an estimate thereof, and the causal effects $\mathcal{R}_a^0$ must be estimable.

In the next section, we work under the assumption that δ_t^* (or $\delta_{a,t}^*$) is known to the researcher. This will allow us to motivate and clarify the working of propositions 1 and 2 (or proposition 3) in a stylized example. In Section 6, we will extend our approach to "test" macroeconomic policies in practice, which requires estimating the causal effects and taking into account that the policy maker may not be able to provide the oracle forecasts $\mathbb{E}_t Y_t^0$.

5 Illustration: the OPP in a structural macro model

In this section, we illustrate Propositions 1 and 2 —the core properties of the OPP statistic— in the context of a structural macro model. We consider a stylized New Keynesian model where the policy maker is the central bank that sets the short term interest rate (e.g. Galí, 2015). 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

$$\begin{aligned} \pi_t &= \mathbb{E}_t \pi_{t+1} + \kappa x_t + e_t^s \ , \quad (12) \\ x_t &= \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}(i_t - \mathrm{E}_t \pi_{t+1} - r_t^e) \ , \quad (13) \end{aligned}$$

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^e the efficient real interest rate and e_t^s 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 operates under discretion and considers the loss function

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

with λ the weight on output gap fluctuations. Under discretion, the central bank's instrument vector p_t reduces to only one instrument, i_t the interest rate at time t. In Appendix C we show that a similar analysis can be conducted for a central bank who commits, starting at time t, to an entire interest rate plan $p_t = (i_t, i_{t+1}, \ldots)'$ in order to minimize the loss function $\mathbb{E}_t \sum_{h=0}^{\infty}(\pi_{t+h}^2 + \lambda x_{t+h}^2)$.

The standard approach

Denote by i_t^0 the policy implemented by the policy maker. The traditional approach to evaluate i_t^0 is to contrast this policy decision with that implied by that of a planner choosing 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 policy rule $g^{\text{opt}} \in \mathcal{G}^{\text{opt}}$ given by

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

with $\phi_\pi = \frac{\kappa\sigma}{\lambda}$ (e.g., Galí, 2015).[16] 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:

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

As highlighted in (17), the policy choice i_t^0 can be non-optimal for two reasons. First, with $\gamma^0 > 0$ the central bank is not following the optimal policy rule and is reacting too strongly to movements in the inflation gap. In our general notation this implies that g^0 is not optimal. Second, with $\epsilon_t^0 \neq 0$ the central bank is making a discretionary error. Note that ϵ_t^0 is like a monetary shock in common macro parlance (Ramey, 2016).

Calculating the gradient (and the OPP) requires the two statistics $\mathrm{E}_t Y_t^0$ and $\mathcal{R}^0$. Under

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

i_t^0 —the prescription of the non-optimal rule (17)— and assuming determinacy, the vector $\mathrm{E}_t Y_t^0$ is given by

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

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

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

so that we can compute the OPP

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

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

We can discard optimality when $\delta_t^* \neq 0$, which is the case when $\gamma^0 \neq 0$ or $\epsilon_t^0 \neq 0$, that is when the central bank does *not* follow the optimal Taylor rule (16).

At the optimum the impulse response $\mathcal{R}^0$ is orthogonal to the conditional expectation $\mathbb{E}_t Y_t^0$ and we have $\mathcal{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 = 0$. Intuitively, when $\gamma^0 = 0$ and $\epsilon_t^0 = 0$ there is no combination of the causal effects $\mathcal{R}^{0\pi}$ and $\mathcal{R}^{0x}$ that can further stabilize $\mathbb{E}_t Y_t^0$ (in an ℓ^2 norm sense), as any additional stabilization of inflation would be more than compensated by additional variability in the output gap. In contrast, when $\gamma^0 > 0$ or when $\epsilon_t^0 > 0$ the central bank is over-stabilizing inflation, and lowering the interest rate can lower the loss function, i.e., better stabilize inflation and unemployment as a whole.

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 policy rule is g^0 (and not necessarily an optimal rule g^{opt}).[17] To see that , we note that minimizing $\mathcal{L}_t$ 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 $i_t^0 + \delta_t^*$.

Importantly, this example also illustrates how the OPP statistic alone cannot be used to

[17]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}.

isolate the reason for the optimization failure. We could have $\delta_t^* \neq 0$, because $\epsilon_t^0 \neq 0$, $\gamma^0 \neq 0$ or both. In other words, the OPP is silent about the source of the optimization failure, i.e., whether it is due to a one-time mistake ϵ_t^0 or to a systematic mistake, i.e., a non-optimal reaction function $g^0 \neq g^{\text{opt}}$ as with $\gamma^0 \neq 0$. To separate these two sources of optimization failure, we can however rely on Proposition 2 and use a sequency of OPP statistics, as we illustrate next.

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 policy rule g^0. 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 the variables in the reaction function, in this case inflation, on the OPP statistic.

Since the reaction function is linear we can do so using the second stage model

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

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.[18] In the New Keynesian model the cost-push shock e_t^s provides a valid instrument as equation (18) implies that e_t^s only affects δ_t^* via π_t^0.

Using this instrument the coefficient b is identified by

$$b = \frac{\mathbb{E}(e_t^s \delta_t^*)}{\mathbb{E}(e_t^s \pi_t^0)} \ . \tag{19}$$

Under $H_0 : g^0 = g^{\text{opt}}$ we have that $\gamma^0 = 0$ and the expression for δ_t^* in equation (18), implies that $\mathbb{E}(e_t^s \delta_t^*) = 0$ and hence $b = 0$. In contrast, under the alternative $H_1 : g^0 \neq g^{\text{opt}}$ we have that $\gamma^0 \neq 0$ and we can derive $b = -\gamma^0 \phi_\pi \omega \neq 0$.

In practice, a sequence of δ_t^*'s is required to replace the expectations in (19) by their sample averages. The resulting estimate then allows to separate whether the optimization failures come from γ^0 —a non-optimal policy rule— or from $\{\epsilon_t^0\}$ —a sequence of discretionary policy failures—. 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. In contrast, when the reaction function is optimal, (18) implies $\delta_t^* \propto \epsilon_t^0$, and the OPP statistic resembles a monetary shock such that $\mathbb{E}(e_t^s \epsilon_t^0) = 0$ and thus $b = 0$.

[18] 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.

6 Testing macro-economic policy: OPP inference

In this section we discuss inference for the OPP, for generality we focus on the subset OPP statistic that was discussed in section 4.3. In this setting the researcher aims to test whether a $K_a \times 1$ subset of the policy plan, e.g. $p_{a,t} = S_a p_t$, was optimally set and the corresponding subset OPP statistic was given by

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

whose properties are summarized in Proposition 3. Note that this statistic includes as a special case the OPP statistic which is obtain by choosing $S_a = I_K$, such that $p_{a,t} = p_t$.

The computation of the subset OPP requires two 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 does not know the causal effects $\mathcal{R}^0_a$ and (ii) the optimal forecasts $\mathbb{E}_t Y^0_t$ cannot be computed by the policy maker.

Taking these constraints into consideration, we now develop an approach to "testing" the optimality of macro policies.[19] Specifically, 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 $\mathbb{E}_t Y^0_t$.

Specifically, the causal effects $\mathcal{R}^0_a$ need to be estimated by the researcher and thus face estimation uncertainty. Second, the policy makers' forecasts can only approximate the conditional expectation $\mathbb{E}_t Y^0_t$.[20] Because of these two sources of uncertainty — parameter and model uncertainty—, the researcher could wrongly conclude that there was an optimization failure.[21] To guard against such a risk, we derive confidence bands around the OPP estimate. These bands allow the researcher to state the level of confidence attached with a detection of an optimization failure.

Based on estimates for $\delta^*_{a,t}$, we discuss how Proposition 2 can be implemented in practice to detect systematic optimization failures. We discuss how endogeneity biases can be avoided, and provide a detailed implementation guide for the class of linear reaction functions.

[19] 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.

[20] 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 and/or (ii) the measure underlying $\mathbb{E}_t$ may be incorrectly specified, e.g. a case of distribution mis-specification.

[21] We take a conservative approach here in the sense that we aim to guard against incorrect rejections of optimality. Depending on the researcher's taste or objective one could argue that incorrectly not-rejecting optimality is also undesirable. However, similarly to hypothesis testing one cannot generally guard against both types of errors.

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

In this section we discuss the estimation of the dynamic causal effects. We emphasize that the material in this section is standard and we only include it to be explicit about the necessary assumptions for estimating $\mathcal{R}_a^0$ consistently and obtaining an asymptotically valid approximation to its limiting distribution. This approximating distribution is then used to construct confidence bounds around the OPP.

Overall the assumptions imposed are similar to those found in other treatment-effect or sufficient statistic type approaches. Invariably they include a constant policy regime, an identification assumption (e.g. existence of valid instruments in our case) and a set of regularity conditions.

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 $\mathcal{R}_a^0$.[22] 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

$$\begin{aligned} Y_s^0 &= \mathcal{R}_a^0 p_{a,s}^0 + \underbrace{\mathcal{R}_b^0 p_{b,s}^0 + f(y_s^0, X_s; g^0) + \Xi_s}_{\zeta_s} \\ \underbrace{\begin{bmatrix} p_{a,s}^0 \\ p_{b,s}^0 \end{bmatrix}}_{p_t^0} &= \underbrace{\begin{bmatrix} g_a^0(y_s^0, X_s) \\ g_b^0(y_s^0, X_s) \end{bmatrix}}_{g^0(y_s^0, X_s)} + \underbrace{\begin{bmatrix} \epsilon_{a,s}^0 \\ \epsilon_{b,t} \end{bmatrix}}_{\epsilon_s^0} \end{aligned} . \tag{20}$$

To estimate $\mathcal{R}_a^0$ based on model (20) 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 typically does not have access to $f(y_s^0, X_s; g^0)$ nor $p_{b,t}$. Both problems imply that the policy paths $p_{a,s}^0$ are correlated with the error term ζ_s. This implies that the researcher cannot use ordinary least squares to consistently estimate $\mathcal{R}_a^0$.

To solve the endogeneity problem we rely on instrumental variables. Similar to Barnichon and Mesters (2020) we note that 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 $p_{a,s}^0$. This makes $\epsilon_{a,s}^0$ an attractive instrumental variable. While the discretionary component 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. For instance, for monetary policy the high frequency series of Gürkaynak, Sack and Swanson (2005) can be exploited

[22] We slightly abuse notation as formally not all elements of Y_s are observable for $s > t - H$, these elements are in practice omitted from the analysis.

after adjusting them for a possible information effect, see Nakamura and Steinsson (2018) and Miranda-Agrippino and Ricco (2019). Also, Ramey and Zubairy (2018) and Alesina, Favero and Giavazzi (2019) provide potential instrument series for fiscal policy.

In general, we postulate that the researcher has access to a sequence of instruments $\{z_s\}$, where z_s has dimension $L \times 1$, with $L \geq K_a$, that correlate only with the policy choice $p^0_{a,s}$.

Assumption 3. (Exogenous variation)
The instrumental variables z_s satisfy

1. $\mathbb{E}(z_s \zeta_s) = 0$ *for all* s
2. $\frac{1}{n}\sum_{s=t_0}^{t} \mathbb{E}(z_s p^{0\prime}_{a,t})$ *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 $p^0_{a,s}$.[23]

We use the instruments to define the following moment estimator for the $K_a M(H+1) \times 1$ vector $r^0_a \equiv \text{vec}(\mathcal{R}^0_a)$.

$$\hat{r}_{a,n} = (Q'_{a,n} \hat{D}_n Q_{a,n})^{-1} Q'_{a,n} \hat{D}_n Z' Y^0 \qquad \text{and} \qquad \hat{r}_{a,n} = \text{vec}(\widehat{\mathcal{R}}_{a,n}) \ , \tag{21}$$

where $Q_{a,n} = Z' P^0_a$, $P^0_a = (P^{0\prime}_{a,t_0}, \ldots, P^{0\prime}_{a,t})'$, with $P^0_{a,s} = p^{0\prime}_{a,s} \otimes I_{M(H+1)}$, $Z = (Z'_{t_0}, \ldots, Z'_t)'$, with $Z_s = z'_s \otimes I_{M(H+1)}$ and $Y^0 = (Y^{0\prime}_{t_0}, \ldots, Y^{0\prime}_t)'$. For the weighting matrix $\hat{D}_n$ different choices can be considered including $\hat{D}_n = (n^{-1} Z' Z)^{-1}$ which leads to the two-stage least squares estimator.

To give an example, suppose that $K_a = 1$ and $M = 1$, then $r_{a,n}$ corresponds to the dynamic causal effect of the scalar policy $p^0_{a,s}$ on $y_{1,s}, \ldots, y_{1,s+H}$. The estimator is then equivalent to the commonly used LP-IV estimator discussed in Stock and Watson (2018) which is based on the local projection framework of Jordà (2005). The difference is that we simultaneously estimate all dynamic causal effects, because to conduct inference on the OPP the confidence region around the entire estimate $\hat{r}_{a,n}$ is required.[24]

With assumptions 2-3 in place and some standard regularity conditions listed in Appendix A under assumption 5, we obtain a limiting distribution for $\hat{r}_n$ as formalized in the following propositions

[23] We restrict our exposition to the case of strong instruments (e.g. Assumption 3 part 2.). Such assumption may be too strong for some applications. In such cases confidence regions for $\mathcal{R}^0$ should be constructed using weak instrument robust methods, see Andrews, Stock and Sun (2019) for a comprehensive review.

[24] The estimate $\hat{r}_{a,n}$ can also be referred to as the vector of impulse responses of a one-unit change in $p^0_{a,s}$ on $y_{1,s}, \ldots, y_{1,s+H}$. We use the terminology of dynamic causal effects as it generalizes more naturally to the non-linear setting, see Appendix A for details.

Proposition 4. *Given assumptions 2-3 and regularity conditions 5 stated in Appendix A we have that*

$$\Omega_{a,n}^{-1/2}\sqrt{n}(\hat{r}_{a,n} - r_a^0) \xrightarrow{d} N(0, I) \qquad \text{and} \qquad \widehat{\Omega}_{a,n} - \Omega_{a,n} \xrightarrow{p} 0 \ ,$$

where

$$\widehat{\Omega}_{a,n} \equiv (Q_{a,n}'\hat{D}_n Q_{a,n}/n)^{-1} Q_{a,n}'\hat{D}_n \hat{V}_n \hat{D}_n Q_{a,n} (Q_{a,n}'\hat{D}_n Q_{a,n}/n)^{-1}$$

with $\hat{V}_n$ *any consistent estimate for the asymptotic variance* $V_n = \mathrm{Var}(n^{-1/2}\sum_{s=t_0}^{t} Z_s\zeta_s)$*, e.g.* $\hat{V}_n - V_n \xrightarrow{p} 0$.

This asymptotic approximation $r^0 \sim N(\hat{r}_n, \widehat{\Omega}_{a,n})$ is used below to construct confidence bounds around the subset OPP statistic which avoids that we reject optimality because of estimation error in the causal effects.

Due to possible serial correlation in the error term we suggest to use a heteroskedasticity and serial correlation robust estimator for $\hat{V}_n$. Several choices exists and we refer to the recent review of Lazarus et al. (2018) for an overview.

In practice the following trade-off arises. To accommodate that the causal effects pertain to a stable regime it is attractive to rely on a short sampling period to estimate $\mathcal{R}_a^0$. The unfortunate consequence is that generally the variance of the estimator (e.g. $\widehat{\Omega}_{a,n}$) will increase, which as we show below, will imply that the researcher needs to be more conservative when rejecting optimality. Therefore a careful assessment of the stability of the policy regime is important.

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 $\widehat{Y}_{t|t}$ which can be regarded as an approximation to $\mathbb{E}_t Y_t^0$. To avoid that we reject optimality using the OPP because of such approximation error, we outline two methods for approximating the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}$.

First, some policy makers, in addition to providing their point forecasts, provide an associated distribution $F_{Y_t^0}$ that characterizes the uncertainty around the forecasts. A simple solution is then to assume that this distribution corresponds to the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}$. While this might seem like a strong assumption, recall that many policy makers spend a lot of effort at assessing the risk around their baseline conditional forecasts. They often experiment with different parameter settings and model specifications in order to asses the effects of model uncertainty.[25]

[25] To give an example, note that most policy documents typically include an "Assessment of Forecast

Alternatively, the researcher can construct its own approximation to the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}$. The difficulty in practice is that historical misspecification errors $\{\mathbb{E}_s Y_s^0 - \widehat{Y}_{s|s}\}_{s=t_0}^t$ are not observable, and we cannot exploit such sequence to predict the distribution of $\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}$. Instead, we take a conservative approach and rely on the historical forecast errors to construct upper-bounds for the confidence interval for $\mathbb{E}_t Y_t^0$.

Specifically, we have

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

First, we assume that the misspecification error $\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}$ follows a normal distribution.[26] Second, we upper bound the variance of $\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}$ with the variance of the forecast errors, which are observable.[27]

With these assumptions in place we approximate the distribution of $\mathbb{E}_t Y_t^0$ by

$$\mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t} \sim \underbrace{N\left(0, \widehat{\Sigma}_{t|t}\right)}_{F_{Y_t^0}} , \tag{22}$$

where $\widehat{\Sigma}_{t|t}$ is an estimate for the mean squared forecast error $\Sigma_{t|t} = \mathbb{E}_t(Y_t - \widehat{Y}_{t|t})(Y_t - \widehat{Y}_{t|t})'$ that we obtain from the historical forecast errors

$$\widehat{\Sigma}_{t|t} = \frac{1}{n}\sum_{s=t_0}^{t}(Y_s - \widehat{Y}_{s|s})(Y_s - \widehat{Y}_{s|s})' . \tag{23}$$

In specific applications there might be additional information available about historical misspecification errors. This would allow to further sharpen the approximation (22).

Regardless whether $F_{Y_t^0}$ is provided by the policy maker, or constructed by the researcher, we exploit such distribution to construct confidence bounds around the OPP.

Uncertainty" section, see for instance the Fed Tealbook or the Bank of England fan-charts, where such model uncertainty assessments are made.

[26]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.

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

6.3 Confidence interval for the OPP

Having characterized approximating distributions for $\mathcal{R}_a$ and $\mathbb{E}_t Y_t^0$ we are able to construct a confidence interval for the subset OPP, of which the full vector OPP is a special case. In particular, we use the distributions for $\hat{r}_{a,n}$ and $\mathbb{E}_t Y_t^0$ to approximate the distribution of $\delta_{a,t}^* = -(\mathcal{R}_a^{0\prime}\mathcal{W}\mathcal{R}_a^0)^{-1}\mathcal{R}_a^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0$. To do so we compute by simulation

$$\begin{gathered}\left\{\delta_{a,t}^{(j)}, j = 1, \ldots, B\right\}, \qquad \text{with} \qquad \delta_{a,t}^{(j)} = -\left(\mathcal{R}_a^{(j)\prime}\mathcal{W}\mathcal{R}_a^{(j)}\right)^{-1}\mathcal{R}_a^{(j)\prime}\mathcal{W}Y_{t|t}^{(j)}, \\ r_a^{(j)} \sim N\left(\hat{r}_{a,n}, \widehat{\Omega}_{a,n}\right), \qquad Y_{t|t}^{(j)} \sim \widehat{Y}_{t|t} + U^{(j)}, \qquad U^{(j)} \sim F_{Y_t^0}, \end{gathered} \tag{24}$$

where B is the number of independent draws from the approximating distributions. In our empirical work we report the median and the upper and lower bounds of the simulated distribution $\{\delta_{a,t}^{(j)}, j = 1, \ldots, B\}$. Whenever the confidence bounds exclude zero we conclude that either g_a^0, $\epsilon_{a,t}^0$, or both were not optimal.

6.4 A Brainard conservatism principle for the OPP

An interesting point is that, the mean of the distribution $\{\delta_{a,t}^{(j)}, j = 1, \ldots, B\}$, say $\widehat{\delta}_{a,t}$, does not correspond to $\hat{\delta}_{a,t}^* = (\widehat{\mathcal{R}}_{a,n}^{0\prime}\mathcal{W}\widehat{\mathcal{R}}_{a,n})^{-1}\widehat{\mathcal{R}}_{a,n}^{\prime}\mathcal{W}\widehat{Y}_{t|t}$. The latter would be an intuitive plug-in estimator, at least from a frequentist perspective. However, in finite sample the variance of the causal effect estimates shows up in the inverse of the OPP.

In particular, we have

$$\widehat{\delta}_{a,t} \quad = \quad -(\widehat{\mathcal{R}}_{a,n}^{\prime}\mathcal{W}\widehat{\mathcal{R}}_{a,n} + \widehat{\Gamma}_{a,n})^{-1}\widehat{\mathcal{R}}_{a,n}^{\prime}\mathcal{W}\widehat{Y}_{t|t} \tag{25}$$

where $\widehat{\Gamma}_{a,n} = \sum_{i=1}^{M(H+1)}\sum_{j=1}^{M(H+1)}\mathcal{W}_{ii}\mathcal{W}_{jj}\widehat{\Omega}_{a,n,(i,j)}$, with $\mathcal{W}_{ii}$ the ith diagonal element of $\mathcal{W}$, and $\widehat{\Omega}_{a,n,(i,j)}$ denotes the (i,j) block of $\widehat{\Omega}_{a,n}$ that is of dimension $K \times K$.

Since the matrix $\widehat{\Gamma}_{a,n}$ is positive definite by construction, we get $|\widehat{\delta}_{a,t}| < |\hat{\delta}_{a,t}^*|$ implying that with parameter uncertainty, the mean OPP estimate is smaller (in absolute value) than the optimal policy perturbation without parameter uncertainty. In other words, the adjustment $\widehat{\Gamma}_n$ can be thought of as capturing an attenuation bias coming from measurement error in the dynamic causal effect estimates.

This result is analog to the seminal Brainard (1967) conservatism principle. Brainard's principle states that in the face of parameter uncertainty, a policy maker should be more conservative in its use of the policy instruments and refrain from fulling minimizing the loss function. A similar logic is at work in our context: a researcher that faces uncertainty in its estimate of the effects of policy (uncertainty in the causal effects) needs to be more conservative – on average – when aiming to reject that the current policy choice is non-optimal.

6.5 Inference for systematic failures

Next, we discuss the details for determining whether a systematic optimization failure occurred over the sampling period $s = t_0, \ldots, t$. To do so, we rely on Proposition 3 which states that if $g^0 \in \mathcal{G}^{\text{opt}}$ we have that $\mathbb{E}(\delta^*_{a,s}|\mathcal{F}_s) = 0$.

To test $\mathbb{E}(\delta^*_{a,s}|\mathcal{F}_s) = 0$ we restrict the class of reaction functions $\mathcal{G}$ to be linear, which covers the majority of reaction functions considered in the literature. To implement the test we require estimates for a sequence of OPP statistics: δ^*_s for $s = t_0, \ldots, t$. These estimates can be computed from equation (25) for each period s. Based on these estimates we consider

$$\widehat{\delta}_{a,s} = Bw_s + \eta_s \ , \qquad s = t_0, \ldots, t \ , \tag{26}$$

where $\widehat{\delta}_{a,s}$ is the estimate for $\delta^*_{a,s}$ given in (25), the $d_w \times 1$ vector w_s includes a constant and any subset of the variables y^0_s and X_s that the researcher suspects to be overlooked, B is the $K \times d_w$ coefficient matrix and η_t is the error term. Note that in general the reaction function can be arbitrarily complex and the researcher may not know, or have access to, all relevant variables. Fortunately, this is not necessary as the optimality of the rule g^0 is rejected if $B \neq 0$, for any subset of variables w_s. In other words, finding $B \neq 0$ implies that $g^0 \notin \mathcal{G}^{\text{opt}}$ but the converse is not necessarily true. The error term η_s includes the estimation error of the OPP statistic $\widehat{\delta}_s - \delta^*_s$ and possibly other omitted variables.

The consistent estimation of B generally requires an identification strategy as even under $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$ and Assumption 1 we have that $\delta^*_{a,s} = \epsilon^0_{a,s}$ 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 can be used. We postulate that the researcher has access to a sequence z^δ_s with dimension $L_\delta \geq d_w$ and impose the following assumption.

Assumption 4. (Systematic inference)

1. *$\mathcal{G}$ is a linear function class,*[28] *with $\mathcal{G}^{\text{opt}} \subset \mathcal{G}$, and Assumption 1 holds for $s = t_0, \ldots, t$;*

[28] That is $\mathcal{G} = \{g : g(y, x) = a + By + Cx, a \in \mathbb{R}^K, B \in \mathbb{R}^{K\times M}, C \in \mathbb{R}^{K\times \dim(x)}\}$.

2. $\mathbb{E}(z_s^\delta \eta_s) = 0$ *for all* s *and* $\frac{1}{n}\sum_{s=t_0}^{t} \mathbb{E}(z_s^\delta w_s')$ *has uniformly full column rank;*

The first part of the assumption restricts the class of reaction functions to be linear, imposes that at least one optimal reaction function lies within this class and ensures that this reaction function is optimal over the sampling period. The second part imposes that the instrumental variables z_t^δ are exogenous and relevant.

We consider the following moment estimator for $b = \text{vec}(B)$

$$\hat{b}_{a,n} = (Q_n^{\delta\prime} \hat{D}_n^\delta Q_n^\delta)^{-1} Q_n^{\delta\prime} \hat{D}_n^\delta d_a \tag{27}$$

where $Q_n^\delta = Z^{\delta\prime} W$, $W = (W_{t_0}', \ldots, W_t')'$, with $W_s = w_s' \otimes I_K$, $Z^\delta = (Z_{t_0}^{\delta\prime}, \ldots, Z_t^{\delta\prime})'$, with $Z_s^\delta = z_s^{\delta\prime} \otimes I_K$ and $d_a = (\widehat{\delta}_{a,t_0}', \ldots, \widehat{\delta}_{a,t}')'$. For the weighting matrix $\hat{D}_n^\delta$ different choices can be considered including again $\hat{D}_n^\delta = (n^{-1} Z^{\delta\prime} Z^\delta)^{-1}$ which leads to the two-stage least squares estimator.

Given Assumption 4 and a set of regularity conditions stated in Assumption 6 in Appendix A the following proposition formalizes the detection of systematic optimization failures using the OPP.

Proposition 5. *Given model* (26)*, Assumption 4 and regularity conditions 6 stated in Appendix A, we have that*

$$\textit{if} \quad n\hat{b}_{a,n}' \widehat{\text{Var}}(\hat{b}_{a,n})^{-1} \hat{b}_{a,n} > \chi^2_{K_a d_w, 1-\alpha} \qquad \textit{we reject } H_0 : g^0 \in \mathcal{G}^{\text{opt}}$$

with confidence level α*. Here* $\hat{b}_{a,n}$ *is defined in* (27)*,* $\chi^2_{K_a d_w, 1-\alpha}$ *is the* $1-\alpha$ *critical value of the* χ^2*-distribution with* $K_a d_w$ *degrees of freedom and*

$$\widehat{\text{Var}}(\hat{b}_{a,n}) \equiv (Q_n^{\delta\prime} \hat{D}_n^\delta Q_n^\delta / n)^{-1} Q_n^{\delta\prime} \hat{D}_n \hat{V}_n^\delta \hat{D}_n^\delta Q_n^\delta (Q_n^{\delta\prime} \hat{D}_n^\delta Q_n^\delta / n)^{-1} \ .$$

with $\hat{V}_n^\delta$ *any consistent estimate for the asymptotic variance* $V_n^\delta = \text{Var}(n^{-1/2} \sum_{s=t_0}^{t} Z_s^\delta \eta_s)$*, e.g.* $\hat{V}_n^\delta - V_n^\delta \xrightarrow{p} 0$*.*

The proposition formalizes the use of the Wald statistic for detecting optimization failures. In particular, if the Wald statistic $n\hat{b}_n' \widehat{\text{Var}}(\hat{b}_n)^{-1} \hat{b}_n$ exceeds the critical value we can reject the null hypothesis that the systematic reaction function g^0 is optimal. In practice, a researcher may also want to compute the t-statistics corresponding to the variables w_s to get a sense of which variables are determining the violation of the moment condition $\mathbb{E}(\delta_{a,s}^* | \mathcal{F}_s) = 0$. To this extent, the Wald statistic of Proposition 5 merely formalizes such inference.

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. Could the policy maker have lowered the loss function with a different policy choice at time t, i.e., was there an optimization failure at time t?

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. Are the optimization failures over the sampling period due to discretionary mistakes or does it a point to a systematic failure, i.e., to a non-optimal reaction function? In other words, over the past n periods, could the policy maker have achieved a better outcome by using a different reaction function?

We aim to answer these questions for monetary policy decisions in the United States. For answering the first two questions we concentrate on three case studies: the Fed funds rate policies as 1990M6, as of 2008M4 on the eve of the Great Recession, and as of 2010M4 in the middle of the Great recession. Before describing these exercises, we first detail the data series that we used, the policy instruments that we considered, and the specifics of the causal effects estimation.

7.1 Implementation details

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

$$\mathcal{L}_t = \mathbb{E}_t \left\|\Pi_t\right\|^2 + \mathbb{E}_t \left\|U_t\right\|^2 \ , \tag{28}$$

where $\Pi_t = (\pi_t - \pi^*, \dots, \pi_{t+H} - \pi^*)'$ the vector of inflation gaps and $U_t = (u_t - u_t^*, \dots, u_{t+H} - u_{t+H}^*)'$ the vector of unemployment gaps. In line with the Fed's "balanced approach" and as assumed by the Board staff in the Tealbook, we put equal weight on stabilizing inflation and unemployment ($\lambda = 1$) as in the optimal policy simulations of the Fed's Tealbook. We also take an horizon of $H = 5$ years and consider discount rates $\beta_h = 1$ for all h.

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, at least since 1985, the US economy has evolved in a stable monetary regime (e.g., Clarida, Galí and Gertler, 2000).

The Fed has a wide range of tools at its disposal, the current fed funds rate, forward guidance about the path of the fed funds rate, asset purchases (QE) and possibly others. Following Eberly, Stock and Wright (2019) we consider two specific policy instruments: (i) the current level of the fed funds rate, or fed funds rate policy for short, and we take $p_{a,t} = i_{t|t}$ as the policy instrument that we aim to test, 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 as in Eberly, Stock and Wright (2019) we consider these instruments separately as the sampling periods over which they are used are different. 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.

Testing the optimality of $p_{a,t}$ requires estimates for $\mathcal{R}_a^0$, for which we rely on the instrumental variable methods outlined in Section 6.1. Following Kuttner (2001) and Eberly, Stock and Wright (2019) we use as instruments the monetary policy surprises measured around the FOMC announcements with 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 $\widehat{R}_{a,n}$ using equation (21).

For the conditional forecasts $\widehat{Y}_{t|t}$, we use the median FOMC forecast reported in the Survey of Economic Projections (SEP).[29] To capture the uncertainty around these point forecasts, we use the Board staff assessment of forecast uncertainty, as reported in the Tealbook.

Based on the dynamic causal effect estimates and the forecasts we compute the mean

[29] 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).

subset OPP statistics $\widehat{\delta}_{a,t}$ as in equation (25) and construct confidence bands using equation (24). We consider different choices for the subset OPP including (i) selecting only the short rate $\widehat{\delta}_{i,t}$, (ii) selecting only the slope policy $\widehat{\delta}_{\Delta,t}$ and (iii) the joint subset OPP that selects both the short rate and the slope policy $\widehat{\delta}_{(i,\Delta),t}$. To illustrate the effects of uncertainty we also report the corresponding naive OPP estimates $\hat{\delta}^*_{a,t} = (\widehat{\mathcal{R}}^{0\prime}_{a,n}\mathcal{W}\widehat{\mathcal{R}}_{a,n})^{-1}\widehat{\mathcal{R}}'_{a,n}\mathcal{W}\widehat{Y}_{t|t}$ which ignore the uncertainty in the causal effects estimates.

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 expected unemployment to run over its target, but it also expected inflation to run over its target. The Fed was confronted with a classic inflation-unemployment trade-off: while it would have liked to lower the fed funds rate to fight excess unemployment, it was prevented to do so by the high and on-going inflation (Bluebook, June 2006). The question for the OPP is thus whether the level of the fed funds rate optimally balanced this trade-off.

Computing the OPP

Figure 1 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 $\mathbb{E}_t\Pi^0_t$. 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 $\widehat{\mathcal{R}}^{0\pi}_i$. The right column reports the same information for unemployment: $\mathbb{E}_tU^0_t$ and $\widehat{\mathcal{R}}^{0u}_i$.

For illustration purposes, in this first case study we omit confidence bands and treat the causal effect estimates and forecasts as fixed. As shown in Table 1, absent uncertainty, the OPP for the Fed funds rate policy is slightly negative at $\hat{\delta}^*_{i,t} = -0.21$ppt.

A trade-off across mandates

To understand the reason for that small optimization failure, it is helpful to isolate the contribution of each mandate —inflation or unemployment— to the dual mandate OPP $\hat{\delta}^*_{i,t}$. Specifically, to highlight how the policy maker must balance a trade-off across mandates, we can re-write $\hat{\delta}^*_{i,t}$ as a weighted-average of the OPP for each mandate with

$$\hat{\delta}^*_{i,t} = (1-\omega)\hat{\delta}^{\pi*}_{i,t} + \omega\hat{\delta}^{u*}_{i,t} \tag{29}$$

with $\delta_{i,t}^{w*} = -(\widehat{\mathcal{R}}_i^{w\prime}\widehat{\mathcal{R}}_i^{w})^{-1}\widehat{\mathcal{R}}_i^{w\prime}\mathbb{E}_t W_t^0$ the OPP for a single mandate —inflation ($\lambda = 0$) or unemployment ($\lambda = \infty$)— with $W_t = (w_t - w^*, \dots, w_{t+H} - w^*)'$ for $w = \pi$ or u, and with

$$\omega = \frac{1}{1 + \hat{\kappa}^2/\lambda} \tag{30}$$

a scalar weight that depends on the ratio of society's preference between the two mandates (λ), and the central bank's instrument "average" ability to transform unemployment into inflation $\hat{\kappa} = \|\widehat{\mathcal{R}}_i^{\pi*}\|/\|\widehat{\mathcal{R}}_i^{u*}\|$.[30] The single mandate OPP $\hat{\delta}_{i,t}^{w*}$ captures the optimal perturbation needed to best stabilize only of the two mandates.

Expression (29) highlights how a non-zero OPP —an optimization failure— can come from a failure to appropriately balance the two mandates, and the weight ω captures how a policy maker should balance the two mandates. The weight depends on a policy instrument's ability to influence one mandate versus another (unemployment versus inflation) as well as the policy maker's preference for stabilizing one mandate versus another. In the case of the fed funds rate, we get $\hat{\kappa} \approx .07$, because the fed funds rate is much more effective at moving unemployment than moving inflation. As a result, with $\lambda = 1$, we get $\omega \approx .08$ and the dual-mandate OPP $\hat{\delta}_{i,t}^{*}$ is (ceteris paribus) tilted towards the stabilization of unemployment, i.e., towards the OPP for unemployment $\hat{\delta}_{i,t}^{u*}$.

Table 1 indicates that the central bank should run a more contractionary policy to lower inflation ($\hat{\delta}_{i,t}^{\pi*} \approx +1.2 > 0$), but a more expansionary policy ($\hat{\delta}_{i,t}^{u*} \approx -.3 < 0$) to lower unemployment. With $\omega = .08$, these conflicting goals roughly balance each other out, and the dual-mandate OPP is negative with $\delta_{i,t}^{*} \approx -0.2$. That being said, taking estimation and model uncertainty into account, Table 1 shows that the error band for $\hat{\delta}_{i,t}$ includes zero, and we cannot discard that the fed funds rate was set optimally. We will discuss more explicitly the effect of uncertainty in the next case study.

7.3 Fed funds rate policy as of April 2008

Narrative

In the second case study, we evaluate the fed fund 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.[31] At that meeting, the fed funds rate

[30] Note that the same relation holds in our example based on a New-Keynesian model. In that case κ reduces to the slope of the Phillips curve.

[31] 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)

was lowered by .25ppt to 2 percent, and it remained at that level until October 2008, i.e., until the collapse of Lehman brothers.

As is clear from the April Tealbook and forecast narratives reported by the FOMC, the central bank was facing two conflicting issues in April 2008: (i) a marked deterioration in the growth outlook due declining housing prices and tensions in the financial market, and (ii) upside risks to inflation coming from "persistent surprises to energy and commodity prices" (Kohn, SEP report, April 2008). While the deterioration in credit conditions pushed the central bank into easing monetary policy, the upside risk to inflation called for a more prudent approach with no additional easing.

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?[32]

Computing the OPP

Figure 2 has the same structure as Figure 1 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 2. 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. Table 1 shows that *both* OPPs are negative with $\hat{\delta}^{u*}_{i,t} \approx -.7$ *and* $\hat{\delta}^{\pi*}_{i,t} \approx -.2$, and absent uncertainty the dual-mandate OPP calls for lowering the Fed funds rate by an additional 60 basis points, e.g. $\hat{\delta}^{*}_{i,t} = -0.6$, so a 50 basis points cut with rounding at the nearest quarter percentage point.

A trade-off across horizons

Before discussing the effect of uncertainty on this conclusion,we briefly mention the reason for having a negative value for $\hat{\delta}^{\pi*}_{i,t}$. Indeed, $\hat{\delta}^{\pi*}_{i,t} < 0$ 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, which would seem to call for a higher interest rate, not a lower one.

The reason for this result is that the Fed must also balance a trade-off across horizons. In particular, Figure 2 shows that the fed funds rate policy has a very delayed effect of

following the failure of Lehman Brothers in September 2008.

[32] 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.

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.[33] The main difference compared to June 1990 is that in April 2008, the positive inflation gap was expected to be much more transitory and even to turn negative after three years. As a result, in contrast to June 1990, we get $\delta^{\pi *}_{i,t} \simeq -0.2 < 0$: the fed funds rate should be lowered to undo that 2011 negative inflation gap.

Inference about an optimization failure

The discussion so far has ignored estimation and model uncertainty. Taking estimation uncertainty into account, Table 1 shows that the OPP mean estimate drops from $\hat{\delta}^*_{i,t} = -0.6$ to about $\widehat{\delta}_{i,t} = -0.4$. This is due to the Brainard attenuation bias discussed in section 6.4, implying that with estimation uncertainty the researcher must be more conservative when aiming to reject optimality.

Table 1 also reports the 68% error bands as computed based on equation (24). We find that the error bands include zero. In other words, there is a more than 32 percent chance that the chosen policy $i^0_{t|t}$ actually balances the true conditional expectations for inflation and unemployment, *even though* the FOMC point forecasts suggest otherwise. In words, uncertainty is too large, and we cannot discard optimality despite the expected increase in unemployment: the expected increase in unemployment was not strong enough and too uncertain to justify 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 use its slope instrument $\Delta i_{t+10\text{yr}|t}$ to stabilize the economy.

Figure 3 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 $\widehat{\delta}_{\Delta,t} = -0.92$. This time, the deviations from targets are so large that we can clearly discard optimality at the 68 percent confidence level.[34]

Unlike in June 1990, there is no trade-off between inflation and unemployment: the expected path of the inflation gap is flat around zero, the OPP is driven almost exclusively

[33] This precisely captures a common wisdom of central banking: central banks should "look through" transitory inflationary episodes.

[34] A few caveats to this conclusion: (i) our approach does not highlight which specific policies would have been able to induce such shift in the slope, and (ii) the exercise for the slope instrument is done with the benefit of hindsight, since our evidence on the effect of the slope instrument comes precisely from that time period.

by the large expected gap in unemployment ($\hat{\delta}^{u*}_{\Delta,t} = -1.81$).

7.5 Testing the Fed reaction function

In this last section, we study the systematic component of the Fed's policy and explore whether the Fed's reaction function is optimal, under the assumption of a constant policy regime over 1990-2018.

Our test is based on Proposition 2, or more specifically its subset counterpart Proposition 3, that is we will look 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 rate is set according to 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-negative effect on the OPP statistic.[35]

Specifically, using either one of our two OPP sequences —the OPP for the fed funds rate and the OPP for slope policy—, we estimate the regression

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

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, and the bottom row presents the Wald test of joint significance, as described in Proposition 5. 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 Summary and future research

In this paper we focused on detecting optimization failures for given macroeconomic policy decisions. We now summarize the framework and provide several directions for future research.

[35] 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.

Summary

The starting point was given by a *known* loss function and a *generic* model:

$$\mathcal{L}_t = \mathbb{E}_t \|\mathcal{W}^{1/2} Y_t\|^2 \qquad \text{where} \qquad Y_t = \mathcal{R}(g) p_t + f(y_t, X_t; g) + \Xi_t \ .$$

The policy maker aimed to minimize $\mathcal{L}_t$ using the reaction function g^0 and possible discretionary adjustments ϵ_t^0 which implied the proposed policy plan $p_t^0 = g^0(y_t, X_t) + \epsilon_t^0$. To test whether these choices were optimal we relied on the OPP statistic

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

which depends on the dynamic causal effects $\mathcal{R}^0$, the forecasts $\mathbb{E}_t Y_t^0$ and the weighting matrix $\mathcal{W}$. We showed that if $\delta_t^* \neq 0$ there was an optimization failure and if $\mathbb{E}_t(\delta_t^* | \mathcal{F}_t) \neq 0$ there was a systematic optimization failure. To empirically evaluate these conditions we estimated the causal effects $\mathcal{R}^0$ using instrumental variables and approximated the distribution of the mis-specification in the forecasts $\widehat{Y}_{t|t} - \mathbb{E}_t Y_t^0$. This allowed to determine the distribution of δ_t^* and determine whether $\delta_t^* = 0$ and/or $\mathbb{E}_t(\delta_t^* | \mathcal{F}_t) = 0$. The latter moment condition was evaluated using a second stage IV regression.

The key benefit of the OPP approach is that it does not require knowledge of any specific underlying model as it only requires conditional forecasts, which policy makers routinely provide, as well as estimates of the causal effects of the policy instruments, i.e., consistent estimates of $\mathcal{R}$. This is attractive in practical situations, where there is no universally agreed upon underlying model and where policy makers routinely use instincts and judgment calls to make their policy decisions. By constructing confidence bounds around the OPP the detection approach is robust to parameter uncertainty and model mis-specification. This allows to assign a level of confidence attached to a rejection of optimality.

A key requirement for the OPP approach is that the causal effects $\mathcal{R}$ need to be estimable. This requires: (i) a stable policy regime (for some period), (ii) some identification strategy (e.g. instrumental variables) and (iii) some regularity conditions. In certain scenarios satisfying these assumptions may not be possible, for instance if the policy has never been implemented or if the policy regime has recently changed.

Future research

There are several ways in which the OPP approach can be improved. First, improvements in the quality of forecasts and causal effect estimates directly improve the ability to detect optimization failures by shrinking the confidence bands of the OPP. Policy makers should provide an explicit *and* correct quantitative assessment of the model uncertainty around the forecasts. Currently, not all policy makers provide a rigorous assessment of the model uncertainty attached to their forecasts. For instance, while the Bank of England which reports fan charts around its central forecasts, the Fed only provides a qualitative assessment of the risks around its forecasts. In the context of fiscal policy, governments often only publish their point 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. Similarly, more accurate estimates of the causal effects, possibly allowing for time-variation or state dependence could improve the detection of optimization failures.

Second, the assumptions that we imposed could be relaxed. For instance, in the Appendix we outline how the linearity assumption (with respect to p_t) could be relaxed. This requires more advanced estimation methods for the dynamic causal effects, but the general approach of the OPP is shown to hold in the presence of non-linearities. Further, the instrumental variable based identification approach could be replaced by any other identification strategy that the researcher deems appropriate.

Third, 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.[36] 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.[37] 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 with a double objective of high trend GDP growth and low income inequality, a government interested in exchange rate management, or a low/middle income country interested in foreign-exchange reserve management.[38]

[36] 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.

[37] To some extent, a "forecast deficit targeting" approach is already followed by the EU commission, as countries have to justify how they *expect* to bring the deficit back under the 3 percent ceiling. However, there is no objective criterion defining the "appropriate" pace of corrective measures, similarly to the imprecision of the "looking good" criterion in the context of forecast inflation targeting.

[38] The optimal level of foreign-exchange reserve is an important area of research for both developing and low-income economies. The 3-months of import rule advocated by the IMF has no strong theoretical justification and it does take into account time and country specificities (e.g., Jeanne and Ranciere, 2011; Barnichon, 2009).

References

Alesina, Alberto, Carlo Favero, and Francesco Giavazzi. 2019. *Austerity: When It Works and When It Doesn't.* Princeton University Press.

Andrews, Isaiah, and Anna Mikusheva. 2015. "Maximum likelihood inference in weakly identified dynamic stochastic general equilibrium models." *Quantitative Economics*, 6(1): 123–152.

Andrews, Isaiah, James Stock, and Liyang Sun. 2019. "Weak Instruments in IV Regression: Theory and Practice." *Annual Review of Economics.* forthcoming.

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

Barnichon, Regis. 2009. "The optimal level of reserves for low-income countries: Self-insurance against external shocks." *IMF staff Papers*, 56(4): 852–875.

Barnichon, Regis, and Geert Mesters. 2020. "Identifying modern macro equations with old shocks." *The Quarterly Journal of Economics.* forthcoming.

Barro, Robert J., and David B. Gordon. 1983. "Rules, discretion and reputation in a model of monetary policy." *Journal of Monetary Economics*, 12(1): 101 – 121.

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

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

Chari, Varadarajan V, Lawrence J Christiano, and Patrick J Kehoe. 1994. "Optimal fiscal policy in a business cycle model." *Journal of polı́tical 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.

Gavin, William T., and Rachel J. Mandal. 2003. "Evaluating FOMC forecasts." *International Journal of Forecasting*, 19(4): 655 – 667.

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, 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.

Jeanne, Olivier, and Romain Ranciere. 2011. "The optimal level of international reserves for emerging market countries: A new formula and some applications." *The Economic Journal*, 121(555): 905–930.

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

Kleibergen, Frank. 2002. "Pivotal Statistics for Testing Structural Parameters in Instrumental Variables Regression." *Econometrica*, 70(5): 1781–1803.

Kleibergen, Frank. 2005. "Testing Parameters in GMM Without Assuming that They Are Identified." *Econometrica*, 73(4): 1103–1123.

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

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

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.

Lazarus, Eben, Daniel J. Lewis, James H. Stock, and Mark W. Watson. 2018. "HAR Inference: Recommendations for Practice." *Journal of Business & Economic Statistics*, 36(4): 541–559.

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.

Miranda-Agrippino, Silvia, and Giovanni Ricco. 2019. "The Transmission of Monetary Policy Shocks." *AEJ: Macroeconomics.* Conditional Acceptance.

Nakamura, Emi, and Jón Steinsson. 2018. "High frequency identification of monetary non-neutrality: The information effect." *Quarterly Journal of Economics.*

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.

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

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

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. 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 Jonathan H. Wright. 2000. "GMM with Weak Identification." *Econometrica*, 68(5): 1055–1096.

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.

Figure 1: Fed funds rate policy in June 1990

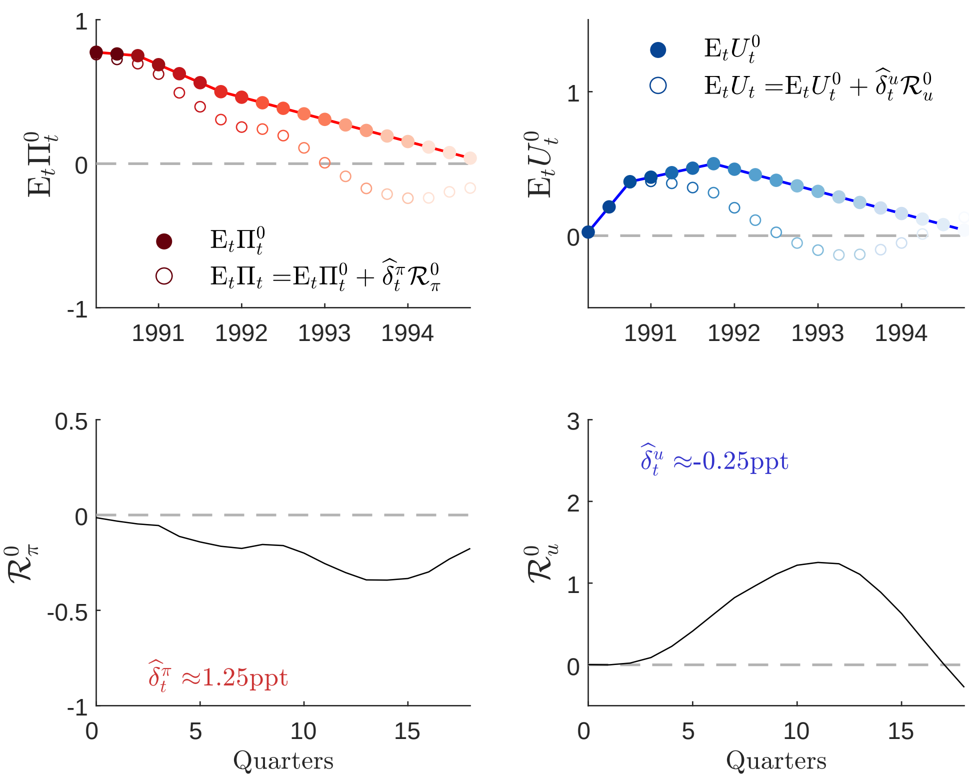

Notes: Top panel: median FOMC forecasts for the inflation and unemployment gaps as of 1990-M6. Bottom panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock. In red (blue) is the OPP $\delta^{*\pi}$ (δ^{*u}) for a strict inflation (unemployment) targeter.

Figure 2: Fed funds rate policy in April 2008

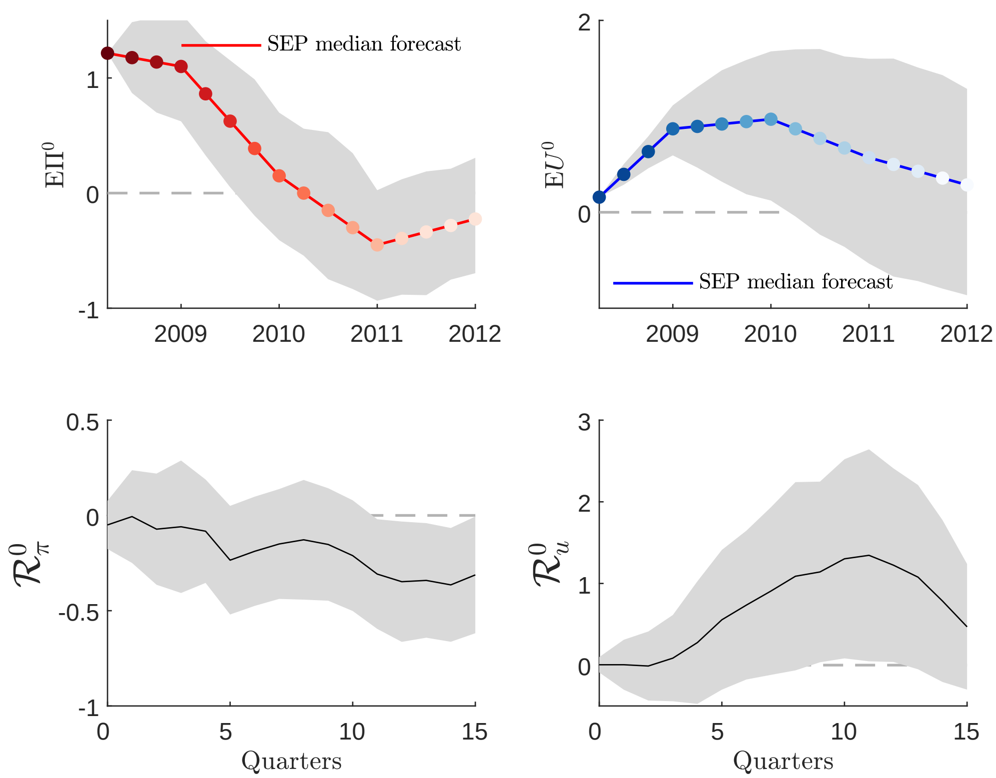

Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2008-M4 (in red and blue) along with the 68 percent confidence bands. Bottom panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock with the 95 percent confidence intervals.

Figure 3: Slope policy in April 2010

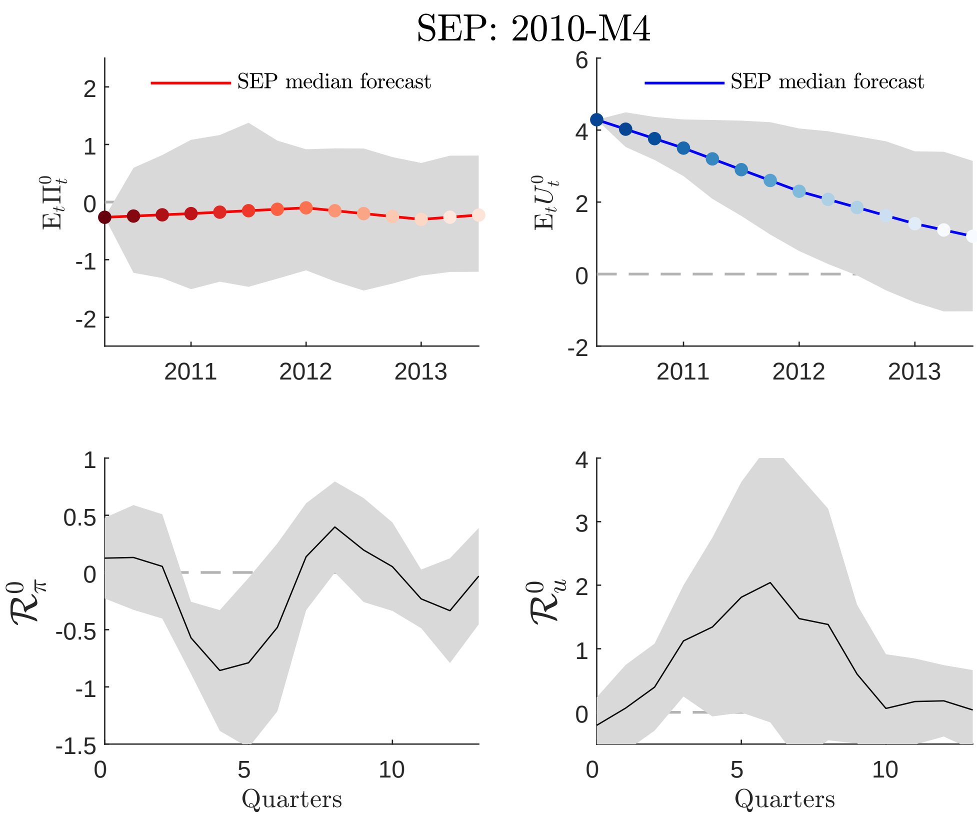

Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2010-M4 (in red and blue) along with the 68 percent confidence bands uncertainty. Bottom panel: impulse responses of the inflation and unemployment gaps to a slope policy shock with the 95 percent confidence intervals.

Figure 4: A sequence of OPP for Fed monetary policy (1990-2018)

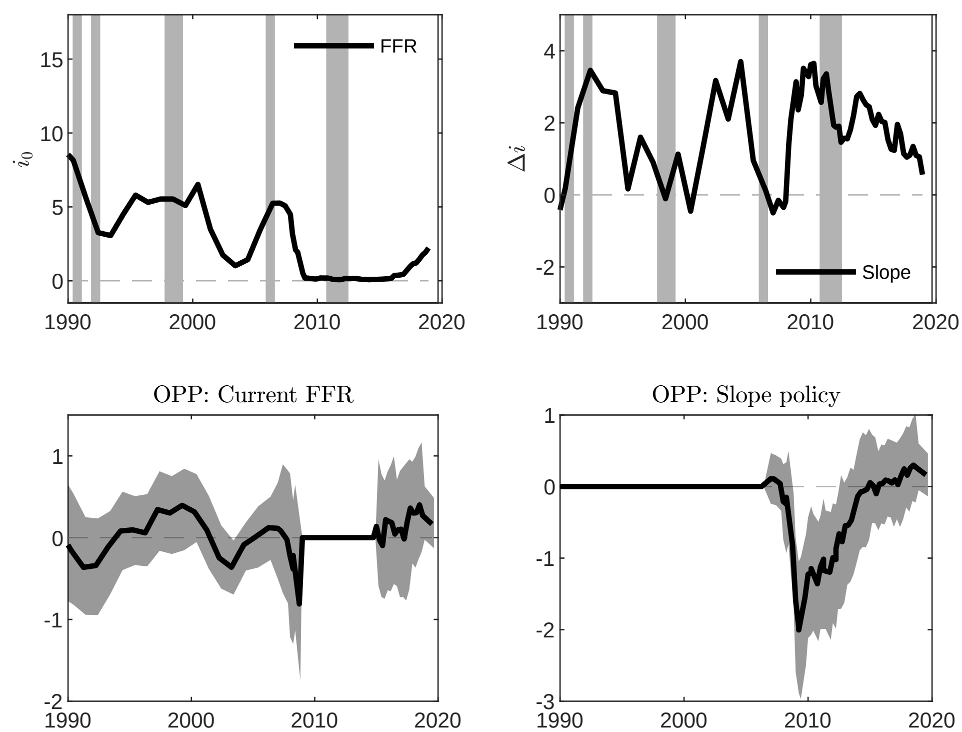

Notes: Top panels: the fed funds rate ("FFR", left-panel) and the difference between the 10-year bond yield and the fed funds rate ("slope", right panel). Grey bars denote NBER recessions. Bottom panels: OPP for the fed funds rate at time t (left-panel) and OPP for the slope instrument at time t (right-panel). 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
$\widehat{\delta}_{i,t}$	-0.13 [-.30,.54]	-0.39 [-.79,.13]	$\widehat{\delta}_{\Delta,t}$	-0.92 [-1.50,-.37]
$\hat{\delta}^{*}_{i,t}$	-0.2	-0.6	$\hat{\delta}^{*}_{\Delta,t}$	-1.6
$\hat{\delta}^{\pi*}_{i,t}$	1.2	-0.2	$\hat{\delta}^{\pi*}_{\Delta,t}$	-0.0
$\hat{\delta}^{u*}_{i,t}$	-0.3	-0.7	$\hat{\delta}^{u*}_{\Delta,t}$	-1.8

Notes: $\widehat{\delta}_{i,t}$ and $\widehat{\delta}_{\Delta,t}$ denote the mean estimates for the dual-mandate OPP ($\lambda = 1$) 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. Further, $\hat{\delta}^{*}_{i,t}$ and $\hat{\delta}^{*}_{\Delta,t}$ denotes the naive plug-in OPP estimates for the short rate and slope policies that ignore uncertainty. Similarly, $\hat{\delta}^{\pi*}_{s,t}$ and $\hat{\delta}^{u*}_{t}$ denote the naive OPPs for a strict inflation targeter ($\lambda = 0$) and a strict unemployment targeter ($\lambda = \infty$), for $s = i, \Delta$.

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	1990-2006	1990-2006	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.

Appendix A: Regularity conditions

To derive the limiting distribution of our estimate for $\mathcal{R}^0$, stacked in the vector $\hat{r}_n$, we require a set of regularity conditions – defined in terms of dependence and moment assumptions – that ensure the applicability of a law of large numbers and a central limit theorem.

In particular, for Propositions 4 and 5 we posit the following standard regularity conditions that allow for heterogeneity and weak dependence, see White (2000), Theorem 5.23, for more discussion.

Assumption 5. (Regularity conditions for Proposition 4)
The following assumptions hold

1. $\{(z_s', p_s^{0'}, \zeta_s')\}$ *is an* α*-mixing sequence with mixing coefficients of size* $-a/(a-2)$*, for* $a > 2$*;*
2. $\mathbb{E}|z_{i,s}\zeta_{j,s}|^a < \Delta < \infty$ *and* $\mathbb{E}|z_{i,s}p^0_{j,s}|^{(a/2)+\rho} < \Delta < \infty$ *for all* i, j, s *and some* $\rho > 0$*;*
3. $V_n = \mathrm{Var}(n^{-1/2}\sum_{s=t_0}^{t} Z_s\zeta_s)$ *is uniformly positive definite and there exists* $\hat{V}_n$*, symmetric and positive definite, such that* $\hat{V}_n - V_n \xrightarrow{p} 0$*;*
4. $\hat{D}_n - D_n \xrightarrow{p} 0$ *where* $D_n = O(1)$ *and is symmetric and uniformly positive definite.*

Assumption 6. (Regularity conditions for Proposition 5)
We assume that

1. $\{(z_s^{\delta'}, w_s', \eta_s')\}$ *is an* α*-mixing sequence with mixing coefficients of size* $-a_\delta/(a_\delta - 2)$*, for* $a_\delta > 2$*;*
2. $\mathbb{E}|z^\delta_{i,s}\eta_{j,s}|^{r_\delta} < \Delta_\delta < \infty$ *and* $\mathbb{E}|z^\delta_{i,s}w_{j,s}|^{r_\delta/2+\rho_\delta} < \Delta_\delta < \infty$ *for all* i, j, s *and some* $\rho_\delta > 0$*;*
3. $V_n^\delta = \mathrm{Var}(n^{-1/2}\sum_{s=t_0}^{t} Z_s^\delta\eta_s)$ *is uniformly positive definite and there exists* $\hat{V}_n^\delta$*, symmetric and positive definite, such that* $\hat{V}_n^\delta - V_n^\delta \xrightarrow{p} 0$*;*
4. $\hat{D}_n^\delta - D_n^\delta \xrightarrow{p} 0$ *where* $D_n^\delta = O(1)$ *and is symmetric and uniformly positive definite.*

Appendix B: 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:

$$Y_t = f(p_t, y_t, X_t; g) + \Xi_t \ , \qquad p_t = g(y_t, X_t) + \epsilon_t \ , \tag{31}$$

where $f(p_t, y_t, X_t; g)$ 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

$$\mathfrak{R}_t^0 \equiv \left. \frac{\partial f(p_t + \delta_t, y_t^0, X_t; g^0)}{\partial \delta_t'} \right|_{\delta_t = 0} \ . \tag{32}$$

It is easy to see that under model (3) we have that $\mathfrak{R}_t^0 = \mathcal{R}^0$. 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 $\mathbb{E}_t Y_t^0 = \mathbb{E}_t(f(p_t^0, y_t^0, X_t; g^0) + \Xi_t)$. The GOPP reduces to the OPP under the partially linear model (3). 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 7. *Let* $y_t \in \mathcal{Y}$ *and* $X_t \in \mathcal{X}$ *be random vectors and let* $\mathcal{D}$ *be an open convex subset of* $\mathbb{R}^K$. *We assume that*

1. *the function* $\mathfrak{R}_t(p_t, y_t, X_t; g) \equiv \partial f(p_t, y_t, X_t; g)/\partial p_t$ *exists for all* $y_t \in \mathcal{Y}$, $X_t \in \mathcal{X}$, $p_t \in \mathcal{D}$ *and* $g \in \mathcal{G}$, *and* $\text{rank}(\mathfrak{R}_t^0) = K$

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 $Y_t(g, \epsilon_t) = f(p_t, y_t, X_t; g) + \Xi_t$ *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. 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 (32) involves $\mathfrak{R}_t^0 = \mathfrak{R}_t(p_t^0, y_t^0, X_t; g^0)$, 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 6. *Given model* (31) *and Assumption 7, we have that* $\delta_t^{g*} \neq 0$ *implies* $g^0 \notin \mathcal{G}^{\text{opt}}$ *and/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 constrained optimal choice $p_t^* = \arg\min_{p_t \in \mathcal{D}} \|f(p_t, y_t, X_t; g^0) + \Xi_t\|^2$. 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 8. *We assume that*

1. $\mu_{\min} > 0$, *where* $\mu_{\min}$ *is the smallest eigenvalue of* $\mathfrak{R}_t(p_t^*, y_t, X_t; g^0)' \mathcal{W} \mathfrak{R}_t(p_t^*, y_t, X_t; g^0)$ *uniformly over* $y_t \in \mathcal{Y}$ *and* $X_t \in \mathcal{X}$

2. $\|(\mathfrak{R}_t(p_t, y_t X_t; g^0) - \mathfrak{R}_t(p_t^*, y_t, X_t; g^0))' \mathcal{W} \mathbb{E}_t[f(p_t^*, y_t, X_t; g^0) + \Xi_t]\| \leq c \|p_t - p_t^*\|$, *with constant* $c < \mu_{\min}$ *for all* $(p_t, y_t, X_t) \in \mathcal{D} \times \mathcal{Y} \times \mathcal{X}$.

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 7. *Given Assumptions 7 and 8 we have there exists* $e > 0$ *such that for all* $p_t^0 \in \mathcal{N}(p_t^*, e)$ *we have*[39]

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

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

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.[40] Proposition 7 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 7 show that for nonlinear models the GOPP can only bring the policy choice closer to the constrained minimum p_t^*.

While the Assumption 7 imposed virtually no restrictions on the functional form of $\mathfrak{R}_t(p_t, y_t, X_t; g^0)$, Assumption 8 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 3, 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 7 and 8 we impose

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

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

Assumption 9. $\mathfrak{R}_t$ *is independent of* p_t*.*

When compared to assumptions 7 and 8 assumption 9 rules out any dependence of the derivatives on the policy choice. Given this assumption we obtain the following result.

Proposition 8. *Given Assumptions 7 and 9, 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.

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 6 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.[41] 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 (similar to Assumption **??**). Importantly, the regularity conditions will include further smoothness assumptions on f, implying that in order to conduct inference using the GOPP statistic requires more assumptions on the model, e.g. besides Assumption 7. should we be more specific on the smoothness conditions, or just not say anything. Right now this last sentence is a bit cryptic.

Apart from the estimation of the derivative function $\mathfrak{R}_t$ inference for the GOPP proceeds exactly the same as discussed in Section 6.

Appendix C: OPP with given constraints

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).

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

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

$$Cp_t = c \;, \tag{33}$$

where C is a known $\#r \times K$ matrix with full row rank and c is a known $\#r \times 1$ vector. We assume that the policy choice of the policy maker $p_t^0 = g^0(y_t, X_t) + \epsilon_t^0$ satisfies the constraints. The constrained class of optimal policy choices is defined, similarly as in Assumption 1.

Assumption 10. Existence of constrained optimum
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 $Y_t(g, \epsilon_t) = \mathcal{R}(g)p_t + f(y_t, X_t; g) + \Xi_t$.

The assumption defines a constrained class of reaction functions $\mathcal{G}^{\text{copt}}$ that minimize the loss function. Again, this class can contain multiple rules.

In this setting, the standard OPP approach needs to account for the restrictions. In particular, we now perturb the policy maker's problem taking into account the constraints and consider

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

This problem can be easily solved using the Lagrange function to give the constrained OPP (COPP) δ_t^{c*} that is given by

$$\delta_t^{c*} = \delta_t^* - (\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}C'\left(C(\mathcal{R}^{0\prime}\mathcal{W}\mathcal{R}^0)^{-1}C'\right)^{-1}(C\delta_t^* - c) \;, \tag{34}$$

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. The constrained OPP has the same properties as the unconstrained OPP but now with respect to Assumption 10.

Proposition 9. *Given an economy defined by equations* (3) *and* (4)*, and policy constraints* (33)*, we have that under Assumption 10:*

1. $\delta_t^{c*} \neq 0$ *implies that* $g^0 \notin \mathcal{G}^{\text{copt}}$ *and/or* $\epsilon_t^0 \neq 0$ *;*
2. $p_t^0 + \delta_t^{c*} = p_t^{c*}$*, where* $p_t^{c*} = \arg\min_{p_t \in \mathbb{R}} \mathbb{E}_t\|\mathcal{W}^{1/2}(\mathcal{R}^0p_t + f(y_t, X_t; g^0) + \Xi_t)\|^2$ *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 (34) subjected to nonlinear or inequality constraints can still be solved numerically if the inputs $\mathcal{R}^0$ and $\mathbb{E}_t Y_t^0$ are known or estimable.

Appendix D: The OPP in the New Keynesian model

In this appendix we revisit the illustrative New Keynesian model of Section 5. We show that the OPP approach can be used in the same way to test a policy maker who considered the loss function

$$\mathcal{L}_0 = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda x_t^2)$$

and aims to minimize this loss function with respect to the entire path $p_0 = (i_0, i_1, \ldots)'$. The equations that govern the economy are similar as before and restated for convenience

$$\begin{aligned} \pi_t &= \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + e_t^s \ , \\ x_t &= \mathbb{E}_t x_{t+1} - \tfrac{1}{\sigma}(i_t - \mathrm{E}_t \pi_{t+1} - r_t^e) \ . \end{aligned} \tag{35}$$

This setting is often referred to as the optimal monetary policy problem under commitment (e.g. Galí, 2015). We stress that the OPP test is applicable in this scenario if there are no outstanding commitments from before time $t = 0$. If there are known commitments from before time $t = 0$ these can be incorporated using the extension outlined in Appendix B.

Galí (2015), page 106, shows that the loss minimizing solution is supported by a reaction function of the form

$$i_t = r_t^e - (1 - \delta)\left(1 - \frac{\sigma\kappa}{\lambda}\right)\hat{p}_t \ ,$$

where

$$\delta \equiv \frac{1 - \sqrt{1 - 4\beta a^2}}{2a\beta} \ , \qquad a \equiv \frac{\lambda}{\lambda(1+\beta) + \kappa^2} \ ,$$

and $\hat{p}_t \equiv p_t - p_{-1}$ is the (log) deviation between the price level and an implicit target given by the price level prevailing one period before the policy maker chooses its optimal plan.

Using the relationship $x_t = -\frac{\kappa}{\lambda}\hat{p}_t$ this reaction function can be conveniently restated in terms of the output gap as

$$i_t = g^{\text{opt}}(r_t^e, x_t) = r_t^e + \underbrace{(1 - \delta)\left(\frac{\lambda}{\kappa} - \sigma\right)}_{\gamma} x_t \ . \tag{36}$$

Based on equations (35) and (36), we start by showing that the elements of the OPP statistic $\delta^*_{i,t}$ for $i = 0, 1, \ldots$ are indeed zero under the optimal reaction function (36).

To define the (infinite) dimensional matrix of dynamic causal effects under g^{opt} let $Y_0 = (\pi_0, x_0, \pi_1, x_1, \ldots)'$, $p_0 = (i_0, i_1, \ldots)$,

$$A = \begin{bmatrix} \beta + \frac{\kappa\sigma^{-1}}{1+\gamma\sigma^{-1}} & \frac{\kappa}{1+\gamma\sigma^{-1}} \\ \frac{\sigma^{-1}}{1+\gamma\sigma^{-1}} & \frac{1}{1+\gamma\sigma^{-1}} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} -\kappa\sigma^{-1} \\ -\sigma^{-1} \end{bmatrix} .$$

Then we can denote the causal effects matrix by

$$\mathcal{R}(g^{\text{opt}}) \equiv \frac{\partial Y_0}{\partial p_0'} = \begin{bmatrix} B & AB & A^2B & A^3B & \ldots & \ldots \\ 0 & B & AB & A^2B & \ddots & \ddots \\ 0 & 0 & B & AB & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

The forecasts under g^{opt} are given by

$$\mathbb{E}_0 x_0 = -\frac{\kappa\delta}{\lambda} e_0^s \,, \qquad \mathbb{E}_0 \pi_0 = \delta e_0^s \,, \qquad \mathbb{E}_0 x_t = -\frac{\kappa\delta^{t+1}}{\lambda} e_0^s \,, \qquad \mathbb{E}_0 \pi_t = (\delta^{t+1} - \delta^t) e_0^s \,, \qquad t > 0$$

The weighting matrix is given by

$$\mathcal{W} = \text{diag}(\mathcal{W}_0, \mathcal{W}_1, \ldots) \qquad \mathcal{W}_t = \text{diag}(\beta^t, \lambda\beta^t)$$

Now that we have defined all elements, we can verify that the OPP statistic is indeed zero under the optimal reaction function. To do so, note that $\delta^*_t = (\mathcal{R}'(g^{\text{opt}})\mathcal{W}\mathcal{R}(g^{\text{opt}}))^{-1}\mathcal{R}'(g^{\text{opt}})\mathcal{W}\mathbb{E}_0 Y_0$ is an infinite dimensional vector, where $\mathcal{R}'(g^{\text{opt}})\mathcal{W}\mathcal{R}(g^{\text{opt}}) \succ 0$. Hence, we study the gradient term $\mathcal{R}'(g^{\text{opt}})\mathcal{W}\mathbb{E}_0 Y_0$ element wise to verify that the OPP is indeed zero. The first element is given by

$$\begin{aligned} B'\mathcal{W}_0\mathbb{E}_0 y_0 &= [-\kappa\sigma^{-1}, -\sigma^{-1}] \begin{bmatrix} 1 & 0 \\ 0 & \lambda \end{bmatrix} \begin{bmatrix} \delta \\ -\frac{\kappa\delta}{\lambda} \end{bmatrix} e_0^s \\ &= \left(-\kappa\delta\sigma^{-1} + \sigma^{-1}\lambda\frac{\kappa\delta}{\lambda} \right) e_0^s = 0 \end{aligned}$$

The second element is given by

$$B'A'\mathcal{W}_0\mathbb{E}_0 y_0 + B'\mathcal{W}_1\mathbb{E}_0 y_1 = [-\kappa\sigma^{-1}, -\sigma^{-1}] \begin{bmatrix} \beta + \frac{\kappa\sigma^{-1}}{1+\gamma\sigma^{-1}} & \frac{\sigma^{-1}}{1+\gamma\sigma^{-1}} \\ \frac{\kappa}{1+\gamma\sigma^{-1}} & \frac{1}{1+\gamma\sigma^{-1}} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \lambda \end{bmatrix} \begin{bmatrix} \delta \\ -\frac{\kappa\delta}{\lambda} \end{bmatrix} e_0^s$$

$$+ [-\kappa\sigma^{-1}, -\sigma^{-1}] \begin{bmatrix} \beta & 0 \\ 0 & \beta\lambda \end{bmatrix} \begin{bmatrix} \delta^2 - \delta \\ -\frac{\kappa\delta^2}{\lambda} \end{bmatrix} e_0^s$$

$$= \left(-\beta\delta\kappa\sigma^{-1} + \beta\delta\kappa\sigma^{-1}\right) e_0^s = 0$$

The third element is given by

$$B'A'A'\mathcal{W}_0\mathbb{E}_0 y_0 + B'A'\mathcal{W}_1\mathbb{E}_0 y_1 + B'\mathcal{W}_2\mathbb{E}_0 y_2 =$$

$$[-\kappa\sigma^{-1}, -\sigma^{-1}] \begin{bmatrix} \beta + \frac{\kappa\sigma^{-1}}{1+\gamma\sigma^{-1}} & \frac{\sigma^{-1}}{1+\gamma\sigma^{-1}} \\ \frac{\kappa}{1+\gamma\sigma^{-1}} & \frac{1}{1+\gamma\sigma^{-1}} \end{bmatrix} \begin{bmatrix} \beta + \frac{\kappa\sigma^{-1}}{1+\gamma\sigma^{-1}} & \frac{\sigma^{-1}}{1+\gamma\sigma^{-1}} \\ \frac{\kappa}{1+\gamma\sigma^{-1}} & \frac{1}{1+\gamma\sigma^{-1}} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \lambda \end{bmatrix} \begin{bmatrix} \delta \\ -\frac{\kappa\delta}{\lambda} \end{bmatrix} e_0^s +$$

$$[-\kappa\sigma^{-1}, -\sigma^{-1}] \begin{bmatrix} \beta + \frac{\kappa\sigma^{-1}}{1+\gamma\sigma^{-1}} & \frac{\sigma^{-1}}{1+\gamma\sigma^{-1}} \\ \frac{\kappa}{1+\gamma\sigma^{-1}} & \frac{1}{1+\gamma\sigma^{-1}} \end{bmatrix} \begin{bmatrix} \beta & 0 \\ 0 & \beta\lambda \end{bmatrix} \begin{bmatrix} \delta^2 - \delta \\ -\frac{\kappa\delta^2}{\lambda} \end{bmatrix} e_0^s +$$

$$[-\kappa\sigma^{-1}, -\sigma^{-1}] \begin{bmatrix} \beta^2 & 0 \\ 0 & \beta^2\lambda \end{bmatrix} \begin{bmatrix} \delta^3 - \delta^2 \\ -\frac{\kappa\delta^3}{\lambda} \end{bmatrix} e_0^s =$$

$$\left(-\kappa\sigma^{-1}\beta^2\delta - \frac{\kappa^2\sigma^{-2}\beta\delta}{1+\gamma\sigma^{-1}} - \frac{\beta\kappa\delta\sigma^{-1}}{1+\gamma\sigma^{-1}}\right) e_0^s$$

$$+ \left(-\kappa\sigma^{-1}\beta^2\delta^2 + \kappa\sigma^{-1}\beta^2\delta + \frac{\kappa^2\sigma^{-2}\beta\delta}{1+\gamma\sigma^{-1}} + \frac{\beta\kappa\delta\sigma^{-1}}{1+\gamma\sigma^{-1}}\right) e_0^s$$

$$+ \kappa\sigma^{-1}\beta^2\delta^2 e_0^s = 0$$

The other elements follow similarly. Hence, we can verify that under the optimal rule the elements of the OPP statistic δ_t^* are equal to zero.

Next, we show that when we deviate from the optimal rule the OPP will become non-zero. To keep the illustration analytically tractable we consider a policy maker who made a discretionary mistake in time period $t = 0$. More elaborate deviations from optimality can be more easily verified numerically using simulation methods. The proposed policy path $p_t^0 = (i_0^0, i_1^0, \ldots)$ is then given by

$$i_0^0 = r_0^e + (1-\delta)\left(\frac{\lambda}{\kappa} - \sigma\right) x_0 + \epsilon_0^0$$

with $\epsilon_t^0 \neq 0$ and for all $t > 0$ we have

$$i_t^0 = r_t^e + (1-\delta)\left(\frac{\lambda}{\kappa} - \sigma\right) x_t$$

In this scenario the dynamic causal effects remain the same, but the forecasts now become

$$\mathbb{E}_0 x_0 = -\frac{\kappa\delta}{\lambda} e_0^s - \frac{\kappa^2\delta\sigma^{-1}}{\lambda(\gamma\sigma^{-1}+1)}\epsilon_0^0 \qquad \mathbb{E}_0\pi_0 = \delta e_0^s - \frac{(\delta\kappa - 1)\sigma^{-1}}{\gamma\sigma^{-1}+1}\epsilon_0^0$$

and for $t > 0$ we have recursively defined

$$\mathbb{E}_0 x_t = \mathbb{E}_0 x_{t-1} \qquad \mathbb{E}_0\pi_t = -\frac{\lambda}{\kappa}\mathbb{E}_0 x_t + \frac{\lambda}{\kappa}\mathbb{E}_0 x_{t-1}$$

We again study the gradient term $\mathcal{R}'(g^{\text{opt}})\mathcal{W}\mathbb{E}_0 Y_0$ element wise. The first element becomes

$$\begin{aligned} B'\mathcal{W}_0\mathbb{E}_0 y_0 &= [-\kappa\sigma^{-1}, -\sigma^{-1}]\begin{bmatrix} 1 & 0 \\ 0 & \lambda \end{bmatrix}\begin{bmatrix} \delta \\ -\frac{\kappa\delta}{\lambda} \end{bmatrix} e_0^s + [-\kappa\sigma^{-1}, -\sigma^{-1}]\begin{bmatrix} 1 & 0 \\ 0 & \lambda \end{bmatrix}\begin{bmatrix} -\frac{(\delta\kappa-1)\sigma^{-1}}{\gamma\sigma^{-1}+1} \\ \frac{\kappa^2\delta\sigma^{-1}}{\lambda(\gamma\sigma^{-1}+1)} \end{bmatrix}\epsilon_0^0 \\ &= \left(-\kappa\delta\sigma^{-1} + \sigma^{-1}\lambda\frac{\kappa\delta}{\lambda}\right) e_0^s + \left(\kappa\frac{(\delta\kappa-1)\sigma^{-2}}{\gamma\sigma^{-1}+1} - \frac{\kappa^2\delta\sigma^{-2}}{\gamma\sigma^{-1}+1}\right)\epsilon_0^0 = \frac{-\kappa\sigma^{-2}}{\gamma\sigma^{-1}+1}\epsilon_0^0 \neq 0 \end{aligned}$$

And hence since $\mathcal{R}'(g^{\text{opt}})\mathcal{W}\mathcal{R}(g^{\text{opt}}) \succ 0$, we have that $\delta_{0,t}^* \neq 0$ and the OPP thus detects the optimization failure. The same can be verified for the other elements of δ_t^*.

Appendix E: 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

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

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

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

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

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

for $p_t^{\text{opt}} = g^{\text{opt}}(y_t^0, X_t)$,[42] we have that $\delta_t^* \neq 0$ implies that $p_t^0 \neq p_t^{\text{opt}}$ which can arise because $g^0 \notin \mathcal{G}^{\text{opt}}$ and/or $\epsilon_t^0 \neq 0$.

Part 2. Let $p_t = p_t^0 + \delta_t$, and plug this in the gradient to get

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

Setting the gradient to zero and solving for δ_t gives

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

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 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 4. The assumptions 2, 3 and **??** correspond exactly to assumptions (i)-(v) in Theorem 5.23 of White (2000). Hence, the proof of White (2000) Theorem 5.23 applies. □

Proof of Proposition 5. Model (26) and Assumption 4 parts 2-6 correspond to exactly to assumptions (i)-(v) in Theorem 5.23 of White (2000). The Theorem implies that

$$\widehat{\text{Var}}(\hat{b}_n)^{-1/2}\sqrt{n}(\hat{b}_n - b) \xrightarrow{d} N(0, I_{Kd_w})$$

Now Assumption 4 part 1 implies that $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$ we have that $b = 0$ and thus

$$nb_n\widehat{\text{Var}}(\hat{b}_n)^{-1}b_n \xrightarrow{d} \chi^2_{Kd_w}$$

under H_0. Hence, we reject $H_0 : g^0 \in \mathcal{G}^{\text{opt}}$ for any level of confidence α when $nb_n\widehat{\text{Var}}(\hat{b}_n)^{-1}b_n > \chi^2_{Kd_w,1-\alpha}$. □

Proof of Proposition 6. By Assumption 7 part 1, the loss function $\mathbb{E}_t\|f(p_t, y_t, X_t; g) + \Xi_t\|^2$ is continuously differentiable on $\mathcal{D}$, thus by Lemma 4.3.1 in Dennis and Schnabel (1996) and Assumption 7 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

[42] 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.

gradient condition $\frac{\partial}{\partial p_t}\mathbb{E}_t\|f_t(p_t, y_t, X_t; g^{\text{opt}}) + \Xi_t\|^2|_{p_t=p_t^{\text{opt}}} = 0$. 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 $Y_t^0 = f(p_t^0, y_t, X_t; g^0) + \Xi_t$, which since $\mathfrak{R}_t^0$ has full column rank (see Assumption 7) 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 7. For convenience let $\mathbb{E}_tY_t^* = \mathbb{E}_tf(p_t^*, y_t, X_t; g^0) + \mathbb{E}_t\Xi_t$. Let κ be a fixed constant in $(1, \mu_{\min}/c)$ and note that such constant exists as $c < \mu_{\min}$ by assumption 2.1. Note that $\mathfrak{R}_t(p_t^0, y_t, X_t; g^0)'\mathcal{W}\mathfrak{R}_t(p_t^0, y_t, X_t; g^0)$ is non-singular and thus there exists a constant $\epsilon_1 > 0$ such that

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

Let

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

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 8 - 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 8 - part 2 that

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

Combining the bounds gives

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

which can be simplified using the definition of ϵ to obtain

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

This completes the proof. □

Proof of Proposition 8. Note that since $\mathfrak{R}_t$ is independent of p_t under assumption 9 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_t around $\delta_t = 0$ we have

$$f_t(p_t^0 + \delta_t, X_t^y; \theta(\psi^0)) = f_t(p_t^0, X_t^y; \theta(\psi^0)) + \mathfrak{R}_t\delta_t$$

which implies that the policy problem is strictly convex in δ_t. Now recall from proposition 1 that the first order conditions at $p_t + \delta_t$ are equal to $\mathfrak{R}_t'\mathbb{E}_t f_t(p_t^0 + \delta_t, X_t^y; \theta(\psi^0)) + \mathbb{E}_t W_t = 0$, which using the expansion can be rewritten as

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

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