



Optimal Policy Perturbations

Régis Barnichon
Geert Mesters

April 2020

Barcelona GSE Working Paper Series

Working Paper n° 1171

Optimal Policy Perturbations*

Regis Barnichon^(a) and *Geert Mesters*^(b)

^(a) Federal Reserve Bank of San Francisco and CEPR

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

April 6, 2020

Abstract

Model mis-specification remains a major concern in macroeconomics, and policy makers must often resort to heuristics to decide on policy actions; combining insights from multiple models and relying on judgment calls. Identifying the most appropriate, or optimal, policy in this manner can be challenging however. In this work, we propose a statistic—the *Optimal Policy Perturbation* (OPP)—to detect “optimization failures” in the policy decision process. The OPP does not rely on any specific underlying economic model, and its computation only requires (i) forecasts for the policy objectives conditional on the policy choice, and (ii) the impulse responses of the policy objectives to shocks to the policy instruments. We illustrate the OPP in the context of US monetary policy decisions. In forty years, we only detect one period with major optimization failures; during 2010-2012 when unconventional policy tools should have been used more intensively.

JEL classification: C14, C32, E32, E52.

Keywords: Macroeconomic stabilization, Optimal policy, Impulse Responses, Sufficient statistics, Forecast targeting.

*We thank Martin Eichenbaum, Ethan Ilzetzki, Jinil Kim, Sylvain Leduc, Huiyu Li, Emi Nakamura, Nicolas Petrosky-Nadeau, Adam Shapiro, Jon 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.

Although contention about the appropriate model of the economy and its microeconomic foundations continues, macroeconomic policy decisions have to be made.

Blanchard and Fischer (1989)

1 Introduction

Despite impressive recent progress in structural macro modeling, model misspecification remains a major concern, and policy makers often rely on heuristics to determine how to optimally balance conflicting objectives; combining insights from different models and relying on judgement calls.¹ This practical approach has benefits in terms of robustness to model mis-specification, but a major downside is that it can be difficult to identify the most appropriate course of policy: without a comprehensive quantitative framework some opportunities for improvement could be left on the table.

In this work, we propose a statistic – the *Optimal Policy Perturbation*, OPP – to detect “optimization failures” in the policy decision process, failures that could arise whenever the underlying economic model is unknown or too complex to write down, as is often the case in practice. The OPP can be applied to a broad range of policy problems, notably macroeconomic stabilization objectives, such as a central bank with a dual inflation-unemployment mandate or a government interested in smoothing business cycle fluctuations but concerned about excessive deficits.

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 objective function can depend on multiple objectives over arbitrary horizons and can have multiple policy instruments as argument. 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 choice can lower the loss function.

The specific policy perturbation that we consider, the OPP, only depends on (i) the

¹See e.g., Svensson (2003), Blanchard (2016).

expectations for the policy objectives —the targets— conditional on the policy choice, and (ii) the impulse responses (IRs) of the targets to shocks to the policy instruments. In fact, calculating the OPP boils down to an OLS regression where we regress the conditional expectations for the targets on their impulse responses to policy shocks. Intuitively, recall that the IRs capture how perturbations to the policy instruments affect the targets. When the policy is optimal, the IRs should be orthogonal to the expected paths of the mandates: there is no combination of the IRs, i.e., no adjustment to the instruments, that can better stabilize the expected paths of the mandates. In contrast, when the policy is not optimal, a regression of the conditional expectations for the targets on the IRs can, under a linearity assumption, determine how to use the IRs to minimize the sum-of-squared deviations of the targets, i.e., how to optimally adjust the instruments.

Specifically, the OPP does not rely on any specific underlying economic model, but the properties of OPP depends on how policy affects the targets. When the effect is linear, the OPP is optimal in the sense that adjusting the current policy choice by the OPP leads to a policy choice that is guaranteed to minimize the loss function.² When the policy effects are mildly nonlinear, or when the current policy choice is in the neighborhood of the optimum, adjusting the policy by the OPP brings the policy closer to the optimum. Finally, in the most general setting, i.e., even with highly nonlinear effects of policy, a non-zero OPP implies that the current policy choice is non-optimal.³

To assess the optimality of a policy using the OPP, we must confront two practical issues: (i) the impulse responses need to be estimated and thus face estimation uncertainty, and (ii) the conditional expectations are oracle forecasts, i.e., optimal forecasts that can generally only be approximated by the policy makers. Because of these two sources of error —IR estimation error and model misspecification error—, one could wrongly conclude that a policy is not optimal. To guard against such risks, we derive confidence bands around the

²This case covers most macro models, such as vector autoregressive models, linearized DSGE models and linear state space models.

³In particular, this most general case is immune to the Lucas critique whereby policy perturbations could prompt regime changes.

OPP. These bands allow the researcher to state the level of confidence attached with any assessment of optimality.

To showcase the usefulness of the OPP and illustrate its practical implementation, we revisit past US monetary policy decisions. We discuss the results from two different exercises. First, we consider a number of past Fed decisions where policy was not always set optimally. These examples illustrate how our methodology can identify deviations from optimality, the reasons for these deviations and the effects of parameter uncertainty and model misspecification on these conclusions. Second, we use our framework to systematically re-assess past Fed policies, and we study the optimality of the Fed policy over 1980-2018 without imposing any modeling restriction on the Fed’s decision making process.

To implement these exercises we require forecasts for inflation and unemployment and estimates of the impulse responses to monetary policy shocks. We obtain conditional FOMC forecasts from historical records of monetary reports to Congress from 1980 until 2018. These projections are conditional on the Fed following an optimal policy, as judged by the FOMC members, which allows us to measure the distance to optimality of the Fed’s actions back to 1980. We then group the Fed’s monetary policy instruments into two groups: a first one captures conventional monetary policy and operates through the fed funds rate; and a second one, available since 2007, captures a broad class of unconventional monetary policies that operate through the slope of the yield curve, as in Eberly, Stock and Wright (2019). We estimate the impulse responses of interest with local projection instrumental variable methods (Jordà, 2005; Stock and Watson, 2018), using external instruments derived from changes in asset prices around FOMC announcements (Kuttner, 2001; Gürkaynak, Sack and Swanson, 2005).

We find that the Fed monetary policy has been remarkably close to optimal. Specifically, we cannot discard optimality of the fed funds rate in all but three periods, the early 1980s when policy was too tight, in 2003 when the fed funds rate could have been lowered further⁴,

⁴As we discuss in the main text, this conclusion holds for a central bank with a dual inflation-unemployment mandate. With additional objectives, e.g., limiting excessive growth in domestic debt, we cannot reject that the fed funds rate was set optimally.

and in 2008 when the Fed should have lowered the fed funds rate faster. At other times, the magnitude of the deviations from optimality are not only non-significant but also relatively minor, averaging only about a quarter percentage point. Since the Fed only moves the short term rate in quarter percentage points, this implies that, except for these three time periods, the fed funds rate OPPs typically round to zero over the past 40 years. In contrast, for the slope instrument, which only operates after 2007, 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 down in line with optimality.

While the detection of optimization failures has received limited attention in the literature, the policy perturbation approach has important links to several large literatures.

First, the OPP goes back to the work of Tinbergen (1952) in that we do not evaluate the full-fledged social welfare function, but instead focus on a simpler high-level macro welfare function, as specified by the policy maker. While the micro-foundation of the loss function can be crucial in certain policy contexts, representing the preferences of the policy makers with a quadratic loss function can yield great benefits when, as is often the case in macroeconomic applications, the underlying economic model is so complicated that policy makers must rely on heuristic methods to approximate the optimal policy. The OPP preserves the simplicity and transparency of the Tinbergen rule, but generalizes it to any dynamic context.⁵

In a public finance context, the OPP shares important similarities with the sufficient statistic approach (e.g. Chetty, 2009), in that both methods exploit the fact that the “welfare” consequences of a policy can be derived from (estimable) high-level elasticities, in our case the IRs to policy innovations. The two approaches differ in that the sufficient statistic approach builds on a micro-founded welfare function, whereas the OPP approach starts from

⁵The Tinbergen rule states that “to reach any number of independent policy objectives, the government needs at least an equal number of policy instruments”. The OPP shows that in a dynamic world where policy *transmits* to the economy, this is not as simple, because any instrument will move the targets over multiple horizons, leading to within-mandate trade-offs, i.e., trade-offs across horizons.

a high-level quadratic loss function, as specified by the policy maker. Since the OPP does not rely on any envelope condition (unlike the sufficient statistic approach, e.g., Chetty, 2009, , section 4.1), the OPP method does not require knowing the underlying model. As we illustrate in this paper, this is of particular interest in the context of macro-economic stabilization, which has been little studied by the sufficient-statistic literature.

Second, the OPP is complementary to the large literature on optimal policy making, see for many examples the textbooks of Ljungqvist and Sargent (2004), Woodford (2003) and Bénassy-Quéré et al. (2018). Unlike the optimal policy literature focused on deriving optimal policy rules —formulas for setting the policy instrument as a function of other variables—, the OPP is a statistic designed to detect optimization failures. Its strength is that it does not rely on any underlying model, but this is also its limitation as its lack of theoretical foundation limits its applicability as a general policy rule. However, the OPP can be seen as a key element of the forecast-targeting rules used in monetary policy making (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*b*, 2019).⁶ However, a limitation is that the “looking good” criterion is imprecise and leaves the policy maker uncertain about the optimality of the policy choice. The OPP provides a precise quantitative condition for optimality, and can thus be seen as providing a quantitative foundation for forecast-targeting rules.⁷

A large literature on optimal policy has aimed at deriving simple and robust optimal policy rules in the form of simple interest rate rules that deliver good performances across a

⁶As 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*b*).

⁷For instance, the OPP is immediately applicable to the Fed “policy rule” described by Bernanke (2015): “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 optimally balancing the expected paths of inflation and unemployment.

broad range of models.⁸ In practice policy makers do not rely on Taylor-type rules. Instead, a more accurate representation of the reaction function is in terms of a forecast-targeting rule. For instance, central bank provide forecasts that converge to the policy objective at some “appropriate” rate. EU countries who deviate from the 3 percent Maastricht deficit rule do not abide by that rule, but instead provide deficit forecasts that converge to (less than) 3 percent at some “appropriate” rate.⁹ Our paper fits into this practical representation of policy making, by providing a statistic that can help understand, quantify and communicate what constitutes an “appropriate” convergence rate.

From a general macro-econometrics perspective, while the structural impulse response literature (e.g. Ramey, 2016) has focused on estimating impulse responses to shocks in order to guide model building (e.g., Christiano, Eichenbaum and Evans, 2005), our paper provides a novel and important role for impulse response estimates: as a testbed for the optimality of policy. And while econometricians shied away from using optimal control theory to devise policy rules since the Lucas critique, the OPP shows that econometrics can play a major role in the optimal policy literature. Two core econometric tasks —economic forecasting and structural IR estimation— are central to the OPP, and progress in IR estimation precision and in forecasting performance will directly improve our ability to detect optimization failures.

The remainder of this paper is organized as follows. In the next section we introduce the environment in which the policy maker and the researcher operate. Section 3 presents the OPP statistic and discusses its properties and underlying intuition. The confidence bands for the OPP statistic are derived in Section 4. In Section 5 we apply our methodology to study monetary policy decisions from the US. Section 6 concludes.

⁸As Galí (2015) puts it “an interest rate rule is generally considered “simple” if it makes the policy instrument a function of observable variables only, and does not require any precise knowledge of the exact model or the values taken by its parameters. The desirability of any given simple rule is thus given to a large extent by its robustness, i.e., its ability to yield a good performance across different models and parameter configurations.”

⁹In the case of Portugal for instance, the Council of the European Union stated that “in order to bring the headline government deficit below the 3 percent-of-GDP reference value by 2015 in a credible and sustainable manner, Portugal was recommended to bring the headline deficit to 5,5 of GDP in 2013, 4,0% of GDP in 2014 and 2,5% of GDP in 2015.” (Economic and Financial Affairs Council, 12 July 2016)

2 Environment

In this section we describe the environment and the policy maker's problem. 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.

The policy maker acts under discretion and has M mandates that span H horizons. The horizon H is arbitrary and can be considered infinite. Let $y_{m,t+h}$ denote the value of the variable that corresponds to mandate m and horizon h . The current time period is indexed by t and the target value for $y_{m,t+h}$ is denoted by y_m^* . To incorporate preferences across mandates and horizons, let λ_m be the fixed preference parameter for mandate m and let β_h be the discount factor for horizon h . The policy maker considers a loss function of the form

$$\mathcal{L}_t = \mathbb{E}_t \sum_{h=0}^H \sum_{m=1}^M \lambda_m \beta_h (y_{m,t+h} - y_m^*)^2, \quad (1)$$

where the expectation is conditional on the time t information set \mathcal{F}_t , e.g. $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_t)$, and with respect to the true distribution of the target deviations $y_{m,t+h} - y_m^*$.¹⁰

The K policy instruments that are available to the policy maker are denoted by $p_t = (p_{1,t}, \dots, p_{K,t})'$. For instance, if the policy maker is a central bank $p_{k,t}$ may correspond to the short term nominal interest rate i_t , or alternatively $p_{k,t}$ may correspond to a forward guidance announcement. In general, the objective of the policy maker is to minimize the loss function (1) using its instruments p_t .

For convenience, we stack all re-weighted deviations from the targets in the $M(H+1) \times 1$ vector $Y_t = [\sqrt{\lambda_j \beta_h} (y_{m,t+h} - y_m^*)]_{m=1, \dots, M, h=0, \dots, H}$ and simply refer to this vector as the targets. We consider the following generic model for the targets

$$Y_t = f_t(p_t, X_t, \xi_t), \quad (2)$$

¹⁰The target value y_m^* can equally well be considered horizon dependent, e.g. $y_{m,t+h}^*$ instead of y_m^* .

where X_t is a vector of \mathcal{F}_t measurable random variables,¹¹ and ξ_t is an arbitrary vector of future shocks that are unknown at time t . The function f_t is assumed to be differentiable with respect to the policy choice p_t .

With these definitions we summarize the policy makers' problem as the following least squares problem¹²

$$\min_{p_t \in \mathcal{D}} \mathbb{E}_t \|Y_t\|^2, \quad \text{with} \quad Y_t = f_t(p_t, X_t, \xi_t), \quad (3)$$

where $\mathcal{D} \subset \mathbb{R}^K$. Problem (3) can be thought of as a static representation of a dynamic policy problem. Forecast horizons and mandates enter in a symmetric fashion (after reweighting by preferences and discount factors) and the policy maker has K instruments to hit the $(H + 1)M$ static targets collected in Y_t . Since K is generally smaller than $(H + 1)M$, the solution to the policy problem implies a set of trade-offs across horizons and mandates.

The possibility of optimization failures

The policy makers' *perceived* solution to problem (3) is denoted by p_t^0 and we do not make any assumptions on how the policy choice is determined.

In practice, the policy makers' choice p_t^0 may not be optimal—a case of optimization failure in the policy decision process—for a number of reasons.

The most immediate one is simply a policy mistake. This is possible, but this is not the main idea of this paper. Another reason is that the policy maker does not have the correct model of the economy, i.e., she uses a mis-specified representation of $f_t(\cdot)$. This is possible, and we will consider this possibility later in Section 4, but this is also not the main idea of this paper.

Instead, our idea is that the policy maker has access to the true underlying model f_t *but* at a very high cost. Specifically, the underlying model is so complex that it cannot be

¹¹For instance, X_t can include expected, current and lagged values of y_t or of any other set of relevant variables.

¹²We consider the usual vector norm where for any $a \in \mathbb{R}^n$ we have that $\|a\| = \sqrt{\sum_{t=1}^n a_t^2}$.

written down explicitly and computing $\mathbb{E}_t f_t(p_t, X_t, \xi_t)$ for a given p_t is a very costly process: it requires many iterations of smaller models, discussions among staff members and judgement calls.¹³

With a high cost of computing f_t , the policy maker is only able to compute $\mathbb{E}_t f_t(p_t, X_t, \xi_t)$ for a small set of values for p_t , i.e., for a small number of alternative policy strategies. This heuristic approach amounts to an incomplete grid search that can leave the policy maker unsure of having set the optimal policy, even if she has access to the true model.¹⁴ To help improve the policy decision process, we will propose a statistic that can help detect optimization failures at no additional computing cost.

3 Optimal Policy Perturbations

In this paper, we take the perspective of a researcher interested in assessing whether the policy maker's choice p_t^0 is optimal, i.e. in detecting a possible *optimization failure*.

To assess optimality, our simple starting idea is to study how the loss function (1) 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 δ_t we may conclude that the choice p_t^0 was not optimal.

While many perturbations are possible, in this work we propose the Optimal Policy Perturbation, or OPP for short, which is defined as

$$\delta_t^* = -(\mathcal{R}_t' \mathcal{R}_t)^{-1} \mathcal{R}_t' \mathbb{E}_t Y_t^0, \quad (4)$$

¹³A telling example are the massive resources devoted by the Board of Governors to (i) assess the current state of the economy (historically called the Greenbook, Part 2), and (ii) construct forecasts for key macroeconomic variables (historically called the Greenbook, Part 1). Every FOMC cycle (8 times a year), the writing of the Greenbook involves a substantial share of the Board staff.

¹⁴In the context of monetary policy for instance, the central bank staff can only consider a limited number of alternative interest paths; e.g., the Bluebook of the Fed lists only 3 or 4 alternative monetary strategies for consideration for the FOMC members.

where

$$\mathcal{R}_t = \left. \frac{\partial f_t(p_t^0 + \delta_t, X_t, \xi_t)}{\partial \delta_t'} \right|_{\delta_t=0} \quad \text{and} \quad \mathbb{E}_t Y_t^0 = \mathbb{E}_t f_t(p_t^0, X_t, \xi_t) .$$

The $M(H+1) \times K$ matrix \mathcal{R}_t denotes the impulse responses of the targets to the policy perturbation δ_t evaluated at $\delta_t = 0$, while $\mathbb{E}_t Y_t^0$ are the forecasts for the targets given the policy choice p_t^0 .

Note that computing the OPP does not require knowing $f_t(\cdot)$ and instead only requires two typically known or estimable quantities: (i) the forecasts $\mathbb{E}_t Y_t^0$ which are routinely published by policy makers, and (ii) the impulse response matrix \mathcal{R}_t , which can be estimated. Indeed, policy makers, such as central banks or fiscal authorities, routinely provide their forecasts to make their policy decisions transparent (e.g., the Survey of Economic Projections from the Fed’s FOMC members) and there is a large literature on the estimation of impulse response (e.g., Ramey, 2016). We discuss inference for the OPP in Section 4, but for now we assume that δ_t^* is known to the researcher.

3.1 Intuition

Before stating the formal properties of the OPP, we provide a heuristic discussion that clarifies why δ_t^* is an attractive perturbation to consider. We discuss the intuition behind the OPP from two perspectives; first from an optimization perspective and then from an econometrics perspective. Finally, we provide some illustrations to clarify the workings of the OPP.

An optimization perspective

The policy maker’s problem (3) can be seen as an optimization problem where the objective function is costly to compute, making a grid search prohibitively expensive. In this context, the idea of the OPP is to build on the optimization literature and exploit the attractive properties of the Gauss-Newton (GN) optimization algorithm, a line-search algorithm designed to minimize quadratic loss functions (e.g., Nocedal and Wright, 2006). In fact, the

OPP δ_t^* is precisely the first iteration of a GN algorithm $p_t = p_t^0 + \delta_t^*$ with p_t^0 the starting vector and p_t the updated vector.

Like many optimization algorithms, the GN algorithm is based on the property that at the optimum, the gradient of the loss function should be zero, or in our case

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

which implies that the OPP must satisfy $\delta_t^* = -(\mathcal{R}_t' \mathcal{R}_t)^{-1} \mathcal{R}_t' \mathbb{E}_t Y_t^0 = 0$ at the optimum. This orthogonality condition will form the basis of our approach to detect optimization failures under very mild regularity conditions.

With more assumptions, specifically when the effect of policy is close to linear or when the initial policy choice is close to being optimal, we can go further and use the OPP to improve the initial policy choice. Specifically, the Gauss-Newton algorithm is based on a standard Newton line-search algorithm,¹⁵ but with the specificity that the Hessian is approximated with first-derivatives only. When that Hessian approximation is good, which happens when non-linearities are mild or when we are already close to the optimum, the GN algorithm is known to converge to a minimum, in which case the OPP will improve the initial policy choice.

To illustrate this point, imagine that p_t^0 is close to the optimum or that non-linearities are mild so that we can approximate f_t with the first order Taylor expansion of f_t around $\delta_t = 0$:

$$Y_t \approx Y_t^0 + \mathcal{R}_t \delta_t \tag{5}$$

When that approximation is exact, policy has linear effects and the GN algorithm actually converges in one step to the minimum:¹⁶ the OPP will make policy optimal. For intermediary cases where the approximation (5) is good but not exact, the OPP will not make policy

¹⁵In a Newton line-search algorithm, the iteration is of the form $p_t = p_t^0 + \delta_t$ with $\delta_t = -(\nabla^2 \mathcal{L}_t)^{-1} \nabla \mathcal{L}_t$ with $\nabla \mathcal{L}_t$ the gradient of the loss function evaluated at p_t^0 and $\nabla^2 \mathcal{L}_t$ the corresponding Hessian.

¹⁶In the linear case, the second derivatives of the loss function with respect to δ_t are zero, and the GN Hessian approximation is exact.

optimal but it will bring it closer to optimality.

An econometrics perspective

From an econometrics perspective, note how the expression for the OPP looks like the formula of an OLS regression. In particular, δ_t^* is minus the coefficient estimate of a regression of $\mathbb{E}_t Y_t^0$ on \mathcal{R}_t . To get the intuition behind this OLS interpretation, go back to the case where policy has a linear effect on the targets so that

$$Y_t = Y_t^0 + \mathcal{R}_t \delta_t. \quad (6)$$

In that case, the set of impulse responses \mathcal{R}_t captures the effect of a change in the policy instrument on the policy maker's objectives.¹⁷ The goal of the optimal perturbation δ_t^* is then to use the impulse response (\mathcal{R}_t) in order to minimize the squared deviations of Y_t . This is nothing but an OLS regression of Y_t^0 on \mathcal{R}_t , 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 value of Y_t^0 is 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 OLS-type formula (4). A benefit of this regression interpretation is that it will help understand not only when but also why a policy choice is not optimal, i.e., which trade-offs may have been overlooked when setting policy.

Illustration

To make the discussion more concrete, we consider three simple scenarios that illustrate the workings of the OPP. We consider a policy maker with a single instrument.

For now, the policy maker has a single mandate: to close the gap y over the next H

¹⁷ In other words, this setting accommodates commonly used linear models such as vector autoregressive models, linearized DSGE models and linear state space models, as well as non-linear models for small perturbations δ_t around the policy choice p_t^0 .

periods. The loss function is thus $\mathcal{L}_t = \mathbb{E}_t \|Y_t\|^2$ with $Y_t = (y_t - y^*, \dots, y_{t+H} - y^*)'$ a vector of future gaps. For instance, the policy maker could be a central bank with an inflation target of 2 percent, in which case Y_t is a vector of inflation gaps. In that case, the policy maker's instrument could be the short term interest rate.

Figure 1 depicts graphically two hypothetical scenarios. The top row reports the policy makers' expected future gaps conditional on the policy choice p_t^0 , that is $\mathbb{E}_t Y_t^0$, and the bottom row reports the impulse response \mathcal{R}_t^y for each scenario. The expected path for the gaps is identical in the two scenarios—policy makers expect y to run over its target for some time—, but the impulse response \mathcal{R}_t^y differs across scenarios. In the first scenario (left-column) the instrument has a delayed effect on the gap variable— \mathcal{R}_t^y is relatively back-loaded—, but in the second scenario (right-column) the instrument has a relatively rapid effect on the gap variable— \mathcal{R}_t^y is relatively front-loaded—.

The goal of the optimal adjustment is to explore whether a small change to the policy rate could further stabilize the path of the target variable y . In other words, the goal of the optimal adjustment is to use the impulse response \mathcal{R}_t^y —the effect on y of an infinitesimal change in the instrument—in order to minimize the squared deviations of the future gaps.

In the first scenario, the instrument's effect is so delayed that there is little more the policy maker can do to stabilize the short-lived y . In other words, the impulse response is orthogonal to the expected path of the gap ($\mathcal{R}_t^{y'} \mathbb{E}_t Y_t^0 = 0$), and there is no adjustment to the policy rate that can better stabilize the target variable. In that case, $\delta_t^* = 0$, and we cannot discard optimality.

In the second scenario, the impulse response shows a faster reaction—the instrument has a faster effect on y —, and the impulse response actually overlaps with the expected path of the gap. In that case, $\delta_t^* \neq 0$, and the policy is not optimal: the policy maker can use the impulse response \mathcal{R}_t^y to bring y down faster by raising the level of its instrument. As illustrated in Figure 1, adjusting the policy choice by $\delta_t^* > 0$ will effectively trade off a smaller positive gap now for a slightly negative gap later—a trade-off across horizons—and in doing so will lower the sum-of-squared gaps. In the linear case, the adjustment needed

to make the policy optimal is given by the OLS regression of $\mathbb{E}_t Y_t^0$ on $-\mathcal{R}_t^y$. After this adjustment, the expected path for the gap becomes $\mathbb{E}_t Y_t = \mathbb{E}_t Y_t^0 + \delta_t^* \mathcal{R}_t^y$, and $\mathbb{E}_t Y_t$ will be (by construction) orthogonal to \mathcal{R}_t^y : the necessary condition for optimality ($\mathcal{R}_t^{y'} \mathbb{E}_t Y_t = 0$) will then be satisfied.

Figure 2 considers a third scenario where the policy maker has a dual mandate to stabilize two targets Y and Z , that is a loss function of the form $\mathcal{L}_t = \mathbb{E}_t \|Y_t\|^2 + \lambda \mathbb{E}_t \|Z_t\|^2$. While one can think of Y and Z as inflation and unemployment in the case of a central bank with a dual mandate, this example works generally for any policy maker with two mandates.

The first mandate y (left-column, in red) has the same expected path as in the two earlier scenarios, while the second mandate z (right-column, in blue) sees a similar transitory positive gap. As in the first scenario, there is little the policy maker can do about $\mathbb{E}_t Y_t^0$ with \mathcal{R}_t^y (the instrument takes too long to have any effect), but in this case the policy maker can use \mathcal{R}_t^z to reduce the positive z gap faster. In this case, the trade-off is not so much a trade-off across horizon but rather a trade-off between mandates: by lowering z faster now (top-right panel), the policy maker will incur a larger y gap later (top-left panel).¹⁸

3.2 Properties of OPP

In this section we formalize the properties of the OPP as defined in equation (4). In particular, we give the (increasingly stronger) conditions under which the OPP 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.

(i) Discarding optimality

In order to use the OPP 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.

¹⁸In other words, the policy is not optimal because the policy maker could lower its loss function by trading off one mandate versus another.

Assumption 1. Let $X_t \in \mathcal{X}$ and $\xi_t \in \Xi$ be random vectors and \mathcal{D} an open convex subset of \mathbb{R}^K . We assume that

1. $f_t : \mathcal{D} \times \mathcal{X} \times \Xi \rightarrow \mathbb{R}^{M(H+1)}$ is continuous and there exists a random variable Z_t such that $\|f_t\| \leq Z_t$ uniformly with $\mathbb{E}(Z_t) < \infty$.
2. there exists a function $R_t \equiv \partial f_t / \partial p_t$ such that uniformly we have (a) R_t is independent of ξ_t , (b) R_t has full column rank and (c) $\|R_t\| \leq \Delta$ for some finite constant Δ .

These assumptions are mild. In words they imply that the underlying true model is continuously differentiable with respect to the policy choice and that there exists a dominating measure which effectively ensures that we can interchange differentiation and integration. Further, the impulse responses R_t do not depend on future shocks, have full column rank and are bounded. These conditions are easily satisfied for commonly encountered impulse responses in macro economics. Note that the OPP formula involves $\mathcal{R}_t = R_t(p_t^0, X_t)$, as the OPP is computed using the impulse responses evaluated at p_t^0 .

The assumptions ensure that \mathcal{L}_t is continuously differentiable with respect to p_t and that the policy problem has a well defined solution on the set \mathcal{D} . Formally, we define the optimal policy solution, denoted by p_t^{opt} , as follows

$$\mathbb{E}_t \|f_t(p_t^{\text{opt}}, X_t, \xi_t)\|^2 \in \inf_{p_t \in \mathcal{D}} \mathbb{E}_t \|f_t(p_t, X_t, \xi_t)\|^2$$

which exists under Assumption 1. The following proposition formalizes the notion of discarding optimality using the OPP.

Proposition 1. *Given Assumption 1, we have that $\delta_t^* \neq 0$ implies $p_t^0 \neq p_t^{\text{opt}}$.*

All proofs are presented in the appendix.

The proposition implies that p_t^0 is not optimal if δ_t^* is not equal to zero. Note that the converse is not implied: it is not necessarily true that p_t^0 is optimal if $\delta_t^* = 0$. For this we need additional conditions as we show below.

Importantly, assumption 1 allows the impulse responses R_t to be an essentially arbitrary function of the policy choice p_t . This implies that Proposition 1 is immune to the Lucas (1976) critique, in the sense that it allows for feedback effects from the policy choice to the impulse responses, and hence *detecting optimization* failures can be carried out in a robust manner. Intuitively, to discard optimality we only have to consider the effect \mathcal{R}_t of an *infinitesimally* small perturbation for which regime change is not an issue.

(ii) Improving policy

Next, we formalize the conditions under which the perturbation δ_t^* can bring p_t^0 closer to the optimum p_t^{opt} . To do so let μ_{\min} be the smallest eigenvalue of $R_t(p_t^{\text{opt}}, X_t)' R_t(p_t^{\text{opt}}, X_t)$ uniformly over $X_t \in \mathcal{X}$, and note that $\mu_{\min} > 0$ by Assumption 1 part 2(b). With this notation we make the following additional assumption.

Assumption 2. *We assume that*

1. $\|(R_t(p_t, X_t) - R_t(p_t^{\text{opt}}, X_t))' \mathbb{E}_t f_t(p_t^{\text{opt}}, X_t, \xi_t)\| \leq c \|p_t - p_t^{\text{opt}}\|$, with constant $c < \mu_{\min}$ for all $(p_t, X_t) \in \mathcal{D} \times \mathcal{X}$
2. R_t is Lipschitz continuous with respect to p_t on \mathcal{D} with parameter γ .

The first part of the assumption ensures that the loss function is not too nonlinear in the neighborhood of p_t^{opt} . The second part imposes a smoothness condition on the impulse responses.

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

Proposition 2. *Given Assumptions 1 and 2 we have there exists $\epsilon > 0$ such that for all $p_t^0 \in \mathcal{N}(p_t^{\text{opt}}, \epsilon)$ we have¹⁹*

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

The proposition states that, if the policy choice of the policy maker is in the neighborhood of the optimal policy, the OPP will bring p_t^0 closer to the optimum. Importantly, as we show

¹⁹The neighborhood $\mathcal{N}(p_t^{\text{opt}}, \epsilon)$ is defined in the usual way: $\mathcal{N}(p_t^{\text{opt}}, \epsilon) = \{p_t \in \mathcal{D} : \|p_t - p_t^{\text{opt}}\| < \epsilon\}$.

in the proof of the proposition, the “size” ϵ of the neighborhood $\mathcal{N}(p_t^{\text{opt}}, \epsilon)$ 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 R_t —, the smaller the neighborhood has to be.²⁰

While the Assumption 1 imposed virtually no restrictions on the functional form of $R_t(p_t, X_t)$, Assumption 2 restricts the effect of p_t on R_t . Therefore Proposition 2 is not entirely immune to the Lucas (1976) critique. Instead it formalizes a notion of modest policy interventions where small perturbations, i.e. those close to the optimum, can improve the policy choice. Such modest policy interventions have been previously characterized, within the context of a structural model by Leeper and Zha (2003). Assumption 2 quantifies the idea of modest policy interventions and generalizes it to a large class of models.

(iii) Getting to the optimal policy

Finally, we describe the conditions under which the OPP brings us directly to the optimal policy choice. This happens when policy has a linear effect on the targets. Note that commonly used models in empirical macroeconomics, such as vector autoregressive models and linear Gaussian state space models, fall into this category (see Ljungqvist and Sargent, 2004, for many examples).

To facilitate a comparison with the previous assumptions, we formalize this last requirement as follows:

Assumption 3. *R_t is independent of p_t .*

When compared to assumptions 1 and 2 assumption 3 rules out any dependence of the impulse responses on the policy choice. Given this assumption we obtain the following result.

Proposition 3. *Given Assumptions 1 and 3, we have $p_t^0 + \delta_t^* = p_t^{\text{opt}}$*

In words, if there are no regime changes following a policy change, adjusting the policy choice by the optimal policy perturbation minimizes the loss function. This follows from

²⁰Note that for any specific model f_t the neighborhood can be determined exactly.

Assumption 3, because in that case the policy problem (3) is strictly convex and there exists a unique minimizer p_t^{opt} which can be reached in one-step regardless of the starting point p_t^0 .

4 Inference for OPP

The computation of the OPP requires two statistics: (i) the impulse responses \mathcal{R}_t , and (ii) the conditional expectations $\mathbb{E}_t Y_t^0$. While the previous section treated these statistics as given, in practice (i) the researcher does not know the true impulse responses \mathcal{R}_t and (ii) the optimal forecasts $\mathbb{E}_t Y_t^0$ cannot be computed by the policy maker.

In this section we extend our framework to capture these practical constraints. First, we take into account that the impulse responses need to be estimated by the researcher and thus face estimation uncertainty. Second, we take into account that the policy makers' forecasts can only approximate the conditional expectation $\mathbb{E}_t Y_t^0$ (e.g. Hansen and Sargent, 2008).²¹ Because of these two sources of error —IR estimation error and conditional expectation error—, the researcher could wrongly conclude that there was an optimization failure.²² To guard against such a risk, we derive confidence bands around the OPP. These bands allow the researcher to state the level of confidence attached with detection of an optimization failure.

Importantly, we do not take a stance on how impulse responses should be calculated or how forecast errors should be assessed. Different policy problems are likely to require different methods that to some extent will depend on data availability. Moreover, there already exists a large econometric literature that discuss impulse response estimation (e.g. Stock and Watson, 2016; Ramey, 2016) and forecast evaluation (e.g. Elliot and Timmermann, 2016). Instead, we merely provide a general methodology that illustrates how to combine

²¹The conditional expectations $\mathbb{E}_t Y_t^0$ can generally only be approximated for two reasons: (i) the model f_t 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.

²²We 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. 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.

impulse response and model mis-specification uncertainty to conduct inference for the OPP.

Impulse response uncertainty

To capture impulse response uncertainty, let $r_t = \text{vec}(\mathcal{R}_t)$ and let \hat{r}_t denote the vector of impulse response estimates of the researcher. We assume that we can approximate the distribution of r_t by

$$\hat{r}_t \sim N(r_t, \Omega_t) , \quad (7)$$

where Ω_t is the variance matrix of all impulse responses: across all horizons and mandates. The normality assumption can be justified from a frequentist perspective by asymptotic arguments for different estimators when taking the sample size to infinity.²³ Further, we posit that there exists a consistent estimate for Ω_t that we denote by $\hat{\Omega}_t$.²⁴

Alternatively, when considering a Bayesian approach to estimation, as is often done in macroeconomics, we may postulate that the researcher is able to obtain the approximation

$$r_t | \mathcal{Y}_t \sim N(\tilde{r}_t, \tilde{\Omega}_t) ,$$

where \mathcal{Y}_t denotes the data used and \tilde{r}_t and $\tilde{\Omega}_t$ are the posterior mean and variance, respectively. We continue our discussion using the frequentist approximation (7), but emphasize that the Bayesian perspective follows similarly.

Model misspecification uncertainty

To capture model misspecification uncertainty, let $\hat{Y}_{t|t}$ denote the forecast of the policy maker that we regard as an approximation to the conditional expectation $\mathbb{E}_t Y_t^0$. In practice, we

²³Suppose that the researcher uses $\{y_{1,s}, \dots, y_{M,s}, p_s, w_s\}_{s=t-n}^t$, with w_s control variables or instruments, to estimate the impulse responses. Then several estimators, such as those based on local projections and structural vector autoregressive models, satisfy $\sqrt{n}(\hat{r}_t - r_t) \xrightarrow{d} N(0, \text{Avar}(\hat{r}_t))$. This implies that if we let $\Omega_t = \frac{1}{\sqrt{n}} \text{Avar}(\hat{r}_t)$ we obtain (7).

²⁴These assumptions are mild. For instance, in our empirical work below we rely on local projections with instrumental variables to estimate the impulse responses (e.g. Jordà, 2005; Stock and Watson, 2018). The corresponding estimator \hat{r}_t is approximately normal under standard textbook assumptions (White, 2000) and Newey and West (1994) methods deliver a consistent estimator $\hat{\Omega}_t$ for Ω_t .

do not observe the historical misspecification errors $\{\mathbb{E}_s Y_s^0 - \widehat{Y}_{s|s}\}_{s=t-n}^{t-H}$ as the true model is unknown, 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 $E_{t|t} = \mathbb{E}_t Y_t^0 - \widehat{Y}_{t|t}$ follows a normal distribution.²⁵ Second, we upper bound the variance of $E_{t|t}$ with the variance of the forecast errors, which are observable.²⁶

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

$$\mathbb{E}_t Y_t^0 \sim N\left(\widehat{Y}_{t|t}, \widehat{\Sigma}_{t|t}\right) , \quad (8)$$

where $\widehat{\Sigma}_{t|t}$ is an (upper-bound) 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

$$\widehat{\Sigma}_{t|t} = \frac{1}{n-H} \sum_{s=1}^{n-H} (Y_s - \widehat{Y}_{s|s})(Y_s - \widehat{Y}_{s|s})' . \quad (9)$$

In specific applications there might be additional information available about historical misspecification errors. This would allow to further sharpen the approximation (8).

²⁵Since 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.

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

Confidence interval for the OPP

Having characterized approximating distributions for \mathcal{R}_t and $\mathbb{E}_t Y_t^0$ we are able to construct a confidence interval for the OPP. In particular, we use the distributions (7) and (8) to approximate the distribution of $\delta_t^* = -(\mathcal{R}_t' \mathcal{R}_t)^{-1} \mathcal{R}_t' \mathbb{E}_t Y_t^0$. To do so we compute by simulation

$$\left\{ \delta_t^{(j)}, j = 1, \dots, B \right\}, \quad \text{with} \quad (10)$$

$$\delta_t^{(j)} = - \left(\mathcal{R}_t^{(j)'} \mathcal{R}_t^{(j)} \right)^{-1} \mathcal{R}_t^{(j)'} Y_{t|t}^{(j)}, \quad r_t^{(j)} \sim N \left(\hat{r}_t, \hat{\Omega}_t \right), \quad Y_{t|t}^{(j)} \sim N \left(\hat{Y}_{t|t}, \hat{\Sigma}_{t|t} \right),$$

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 $\{\delta_t^{(j)}, j = 1, \dots, B\}$.

A Brainard conservatism principle for the OPP

An interesting point is that, the mean of the distribution $\{\delta_t^{(j)}, j = 1, \dots, B\}$, say $\hat{\delta}_t$, does not correspond to $(\hat{\mathcal{R}}_t' \hat{\mathcal{R}}_t)^{-1} \hat{\mathcal{R}}_t' \hat{Y}_{t|t}$. The latter would be an intuitive plug-in estimator, at least from a frequentist perspective, as asymptotically the impulse response estimates are consistent. However, in finite sample the variance of the impulse response estimates shows up in the inverse of the OPP.

In particular, with approximations (7) and (8) we have

$$\hat{\delta}_t = (\hat{\mathcal{R}}_t' \hat{\mathcal{R}}_t + \hat{\Omega}_t)^{-1} \hat{\mathcal{R}}_t' \hat{Y}_{t|t}, \quad (11)$$

where $\hat{\Omega}_t$ can be thought of as capturing an attenuation bias coming from measurement error in the impulse response 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 fully 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 IRs) needs to be more conservative – on average – when aiming to reject that the current policy choice is non-optimal.

5 Optimal perturbations for monetary policy

In this section we apply our optimal perturbation framework to study monetary policy decisions in the United States. Here the policy maker is the Fed which aims to fulfill a number of mandates like stable inflation and stable unemployment and possibly others such as financial stability. The central bank has a number of instruments: the current short term interest rate, the expected path of the short-term rate —forward-guidance— as well as more unconventional policies such as the size of the balance sheet, Quantitative Easing, LSAP or maturity management (see e.g., Eberly, Stock and Wright, 2019). Inspired by Eberly, Stock and Wright (2019), we group the Fed’s instruments into two broad categories, defined from their intended effect on the yield curve: (i) the current fed funds rate —conventional monetary policy—, which affects the short-end of the yield curve and (ii) policies that affect the slope of the yield curve —slope policy—.

We conduct two different exercises. First, to illustrate the workings of our method, we consider a number of past Fed decisions where policy was not always set optimally.²⁷ These examples illustrate how our methodology can clarify the monetary trade-offs at play and help identify opportunities for policy improvements.

Second, we use our framework to systematically re-assess past Fed policies, and we study the optimality of the Fed policy over 1980-2018 without imposing any modeling restriction on the Fed’s decision making process.²⁸

Before describing these exercises, we first detail the datasets and methods that we use to calculate the optimal perturbations. The OPP requires two pieces of information to assess

²⁷In our analysis of past Fed decisions, it is important to keep in mind that we have the benefit of hindsight in that our estimates of the effects of policy changes (the impulse responses) are drawn from evidence that was not necessarily available at the time.

²⁸This is in contrast to earlier studies that rely on a model to construct forecasts for inflation and unemployment and thus implicitly assume that their model was the forecasting model that the Fed adopted in the past (e.g. Rudebusch and Svensson, 1999; Sack, 2000; Rudebusch, 2002).

the optimality of the Fed policy over time: (i) the Fed’s projections for the target variables (e.g., inflation and unemployment) conditional on policy makers’ desired policy path and (ii) the impulse responses of the target variables to the monetary perturbations.

5.1 Data and impulse response estimation

FOMC projections over 1980-2018

Since the passage of the Full Employment and Balanced Growth Act of 1978 (also known as Humphrey-Hawkins), federal law requires the Federal Reserve Board to submit written reports to Congress containing discussions of “the conduct of monetary policy and economic developments and prospects for the future”. As part of this report, the Fed provides a summary of Federal Open Market Committee (FOMC) participants’ projections for the unemployment rate and inflation (among other variables): the Survey of Economic Projections (SEP).

The data were manually extracted from digital records from the archives of the House Financial Services Committee. We obtained bi-annual SEP data for the median forecasts of FOMC members for inflation and unemployment at one- and two-year ahead horizons over the period 1980-2006.²⁹ After 2006, SEP data are published four times a year and additionally include the median forecasts at a three-year ahead horizon. In addition, 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.³⁰ The SEP reports the long-run forecasts of inflation and unemployment only after 2007, so for the pre-2007 period we use real-time estimates of the natural rate of unemployment constructed by (Orphanides and Williams, 2012) and long-run inflation expectations from the Federal Reserve Board “PTR” variable, which is a measure a long-run

²⁹The price index underlying the inflation measure has changed over time, ranging from the GNP deflation, CPI to PCE in the more recent period. Using a linear model of the form $\pi_t^{pce} = \alpha + \beta\pi_t^x + \varepsilon_t$ with x denoting the underlying price index, we adjusted the different inflation measures to make them consistent with a PCE-based measure.

³⁰We obtain very similar results using instead a convergence time of 10 years.

inflation expectations derived from the Survey of Professional Forecasters (SPF). Since the SEP projections are annual, we linearly interpolate them in order to project them on the quarterly impulse responses to monetary shock.

Estimation of impulse responses

Following the works of Gürkaynak, Sack and Swanson (2005), Lakdawala (2019) and more closely Eberly, Stock and Wright (2019), we estimate the impulse responses to (i) innovations to the current Fed funds rate, (ii) innovations to the slope of the yield curve, where the slope is the spread between the yield on 10-year Treasuries and the Fed funds rate.

In our analysis we assume that the impulse responses are constant over the 1980-2018 period, i.e. we have that $\mathcal{R}_t = \mathcal{R}$ for all t . While this assumption is made for simplicity here,³¹ it is actually consistent with the widely held belief that, at least since 1985, the US economy has evolved in a stable monetary regime with anchored inflation expectations.

The impulse responses for the different targets $y = \pi, u$ horizons h and perturbation types $k = i_0, \Delta i$, with i_0 indicating the fed funds rate and Δi the slope of the yield curve, are defined as follows.

$$\mathcal{R}_{k,h}^y = \mathbb{E}(y_{t+h} | \varepsilon_t^k = 1, c_t) - \mathbb{E}(y_{t+h} | \varepsilon_t^k = 0, c_t) ,$$

where ε_t^k corresponds to a structural monetary policy shock of type $k = i_0, \Delta i$ and c_t is a vector of control variables. These impulses are the elements of the matrix \mathcal{R} defined above.

To estimate the impulse responses we use local projections with external instruments (Jordà, 2005; Stock and Watson, 2018). For the Fed funds shock, the instrument is the difference between the target decision and the expectation implied by current-month Federal funds futures contracts, constructed as described by Kuttner (2001). For the slope shock, the instrument identifies policy induced changes in the slope of the Interbank/Treasury term structure, holding constant changes in the Fed funds rate. To this end, the slope instrument

³¹While outside the scope of this application section, one could envision estimating time-varying (Primiceri, 2005) or state dependent impulse responses (e.g., Tenreyro and Thwaites, 2016).

is the residual from a regression of announcement-window changes in the ten-year on-the-run Treasury yield onto the fed funds rate shock.

The local projections, for $y = \pi, u$, $h = 0, \dots, H$ and $k = i_0, \Delta i$, are given by

$$y_{t+h} = x_t^k \mathcal{R}_{k,h}^y + c_t' \gamma_{k,h}^y + \nu_{t+h}^{k,y} , \quad (12)$$

where $x_t^k = i_{0,t}, \Delta i_t$ the fed funds rate or the slope of the yield curve, and c_t a set of control variables consistent of lags of y and x . We estimate the impulse responses using observations from 1990 until 2007 for the fed funds rate instrument and observations from 2007 until 2018 for the slope instrument, consistent the time frame over which the Fed used forward-guidance (e.g. Eberly, Stock and Wright, 2019). We compute the variance matrix of the impulse response estimates using Newey and West (1994).

5.2 Illustrative case studies

We first consider some concrete examples where we assess the optimality of past Fed decisions. For clarity purposes, in the first examples we ignore the effects of uncertainty and treat the impulse responses as given and the FOMC forecasts as capturing the conditional expectations. That is we act as if δ_t^* is known with certainty. Later examples discuss the role of uncertainty. For our analysis, we take an horizon of $H = 5$ years and consider discount rates $\beta_h = 1$ for all h .

5.2.1 On the optimality of the level of the fed funds rate

One mandate: We start with a central bank with a unique inflation mandate and a unique instrument; the fed funds rate. In this scenario the policy makers' problem is given by

$$\min_{i_{0,t} \in \mathbb{R}} \mathcal{L}_t , \quad \mathcal{L}_t = \mathbb{E}_t \|\Pi_t\|^2 , \quad (13)$$

where $\Pi_t = (\pi_t - \pi^*, \dots, \pi_{t+H} - \pi^*)'$ is the vector of inflation gaps with target $\pi^* = 2$.

The optimal perturbation is given by the scalar $\delta_t^{*\pi} = -(\mathcal{R}^{\pi'}\mathcal{R}^\pi)^{-1}\mathcal{R}^{\pi'}\mathbb{E}_t\Pi_t^0$, which measures the adjustment to the fed funds rate that minimizes the sum-of-squares of the expected inflation gaps.

Figure 3 depicts graphically all the information needed to assess the optimality of the fed funds rate in 1990-M6. The top-left panel reports the expected paths for inflation conditional on the current policy choice, including the current level of the fed funds rate. Note how lighter colors denote forecasts at more distant horizons. The bottom-left panel reports the impulse responses of inflation to a 1ppt innovation to the current fed funds rate, that is it reports \mathcal{R}^π .

To illustrate the regression interpretation, Figure 4 combines all the previous information in one scatter plot of the inflation forecast ($\mathbb{E}_t\Pi_t^0$) against the negative of the impulse responses of inflation ($-\mathcal{R}^\pi$). To denote the time dimension, recall that we use lighter colors to indicate horizons farther away in the future: that way, we can visualize how the forecast and the impulse response co-move over the forecast horizon h . There is no new information in Figure 4, but the scatter plot captures all the information needed to compute the OPP and thereby discard (or not) optimality. Indeed, recall that the OPP for a strict inflation targeter is given by a regression of the conditional forecast for inflation ($\mathbb{E}_t\Pi_t^0$) on the (negative) IR of inflation to a policy shock ($-\mathcal{R}^\pi$). Thus, the slope of the best linear fit for that scatter plot is $\delta_t^{*\pi}$, the OPP for a strict inflation targeter (dashed-red line).³²

The scatter plot helps understand the determinants of the regression line, and thus which trade-offs may have been overlooked when setting policy. Informally speaking, dots close to the x-axis mean that the target is already close to zero and thus with little room for improvement, while dots close to the y-axis cannot be influenced by the central bank's instrument. Thus, the room for improvements comes from the dots around the 45 degree line where the deviations are not only substantial but also “influentiable” with the instrument.

³²The OPP is given by the regression of $\mathbb{E}_t\Pi_t^0$ on the *negative* of \mathcal{R}^π , because the goal of the OPP is not to best fit the expected path for inflation with the impulse response, but instead to best “undo” it. By having $-\mathcal{R}^\pi$ on the x-axis, the slope of the best fit line is directly the OPP, which helps with the interpretability of the figure.

Since the Fed traditionally moves in steps of quarter percentage points, the shaded grey area depicts the region where we cannot discard optimality, because the optimal adjustment is below 12.5 basis points (and thus rounds to zero). In other words, whenever the regression line, overlaps with the grey region, we cannot discard that the current fed funds rate is set optimally.

In 1990-M6, the FOMC expected inflation to run over its target for some time, as depicted in the top-left panel of Figure 3. For a strict inflation targeter, the 1990-M6 fed funds rate was not set optimally, because a higher fed funds rate—a positive perturbation to the policy rate ($\delta_t^{\pi} > 0$)—can lower the expected loss function by better stabilizing inflation.³³ In this case, the overlooked trade-off is a trade-off across horizons, i.e., a “dynamic” trade-off: By raising interest rates, the Fed could have traded lower inflation at shorter horizons with too low inflation at longer horizons.³⁴

Two mandates: We now consider a central bank with a dual inflation-unemployment mandate, and we consider the policy problem

$$\min_{i_{0t} \in \mathbb{R}} \mathcal{L}_t, \quad \mathcal{L}_t = \mathbb{E}_t \|\Pi_t\|^2 + \lambda \mathbb{E}_t \|U_t\|^2, \quad (14)$$

where $U_t = (u_t - u_t^*, \dots, u_{t+H} - u_{t+H}^*)'$ is the vector of unemployment gaps.

In line with the Fed’s “balanced approach”, we put equal weight on stabilizing inflation and unemployment ($\lambda = 1$) as in the optimal policy simulations of the Tealbook.

With a dual mandate, the situation is more complicated because in addition to the dynamic trade-off highlighted above, policy makers must also take into account possible conflicts between the different mandates; in this case a “static” trade-off between inflation and unemployment. In fact, we can re-write δ_t^* , the dual-mandate OPP, as a weighted-

³³The dashed red line in Figure 4 lies outside the grey area, so that the magnitude of the OPP (the slope of the red line) is large enough to discard optimality of the fed funds rate (for a hypothetical strict inflation targeter).

³⁴The dynamic trade-off comes from the blunt nature of the monetary instrument: when the Fed moves the fed funds rate, it affects the whole path of macro variables, not only one horizon.

average of the OPP for each mandate, i.e.,

$$\delta_t^* = \omega \delta_t^{*\pi} + (1 - \omega) \delta_t^{*u} ,$$

where $\delta_t^{*\pi} = -(\mathcal{R}^{\pi'} \mathcal{R}^{\pi})^{-1} \mathcal{R}^{\pi'} \mathbb{E}_t \Pi_t^0$ and $\delta_t^{*u} = -(\mathcal{R}^{u'} \mathcal{R}^u)^{-1} \mathcal{R}^{u'} \mathbb{E}_t U_t^0$ are the optimal perturbations for a single mandate (inflation or unemployment) and $\omega = \frac{1}{1+\lambda/\kappa}$ is a scalar weight that depends on the ratio of society's preference between the two mandates (λ),³⁵ and the central bank's instrument "average" ability to transform unemployment into inflation $\kappa = \|\mathcal{R}^{\pi}\|/\|\mathcal{R}^u\|$.³⁶ Thus, the dual-mandate OPP δ_t^* captures two types of trade-offs: (i) a dynamic trade-off depending on the relative ability of the central bank to influence each mandate at different horizons, e.g., longer vs. shorter horizons, and (ii) a static trade-off depending on the relative ability of the central bank to influence one mandate versus another.³⁷

Going back to our 1990-M6 example, taking the expected path of unemployment (top-right panel of Figure 3, blue colors) into consideration completely changes the picture. In fact, in this case, we can no longer discard that the fed funds rate was actually set optimally despite substantial deviations of inflation and unemployment from target. The reason is the standard static trade-off between inflation and unemployment, and this can be seen directly from the regression lines in Figure 4: it is desirable to run a more contractionary policy ($\delta_t^{*\pi} \approx +1.25 > 0$, the red line slopes upwards) to lower inflation *and* at the same time run a more expansionary policy ($\delta_t^{*u} \approx -.25 < 0$, the blue line slopes downwards) to lower unemployment, which is not possible. Graphically, this can be seen by the fact that the blue line and the red line display opposite slopes. Since the black line (with slope equal to the dual-mandate OPP) is an average between these red and blue lines, it ends up with a flat

³⁵We have $\omega \xrightarrow{\lambda \rightarrow 0} 1$ (a single inflation mandate), and similarly $\omega \xrightarrow{\lambda \rightarrow \infty} 0$ (a single unemployment mandate).

³⁶Note that κ reduces to the slope of the Phillips curve in the baseline New-Keynesian model (Galí, 2015).

³⁷The estimated impulse responses to a fed funds rate shock imply $\kappa \approx .07$ and $\omega \approx .08$, because the fed funds rate is much more effective at moving unemployment than at moving inflation. As a result, the overall OPP δ_t will be heavily tilted towards the stabilization of unemployment, and unless the deviations of inflation are large and correlated with \mathcal{R}^{π} (such that $\delta_t^{*\pi}$ is large), the dual-mandate OPP will be driven to a large extent by the OPP for unemployment.

slope that rounds to zero, i.e., inside the grey cone of optimality.

In words, while the Fed 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 excess inflation. This was explicitly acknowledged in the 1990-M6 Bluebook.

Three mandates: Our framework can accommodate any number of mandates, and it is instructive to consider adding a financial stability mandate, as has been the subject of numerous debates since the great recession (e.g., Svensson, 2017a).

To capture that third mandate, imagine that the Fed wants to stabilize credit growth around some target, guided for instance by numerous works that point to the link between fast credit growth and the severity of subsequent recessions (e.g., Jordà, Schularick and Taylor, 2013). Specifically, we consider the policy problem

$$\min_{i_{0t} \in \mathbb{R}} \mathcal{L}_t, \quad \mathcal{L}_t = \mathbb{E}_t \|\Pi_t\|^2 + \mathbb{E}_t \|U_t\|^2 + \lambda_B \mathbb{E}_t \|B_t\|^2, \quad (15)$$

where λ_B is the preference parameter for the financial stability target and B_t denotes the path of the growth rate of non-financial domestic debt, in deviation from its target. For illustration purposes, we fix that target at 5.6 percent, the average debt growth rate over the 1990-2000 period, and we set $\lambda_B = .25$, the ratio of the variance of inflation to the variance of debt growth of 1990-2000.

Figure 5 plots the same set of plots as Figure 3 but with the additional third mandate, and we consider the Fed in 2003-M6, a time when credit expansion was particularly strong, and the Fed might have wanted to preemptively lean against the wind. Note that the 2003-M6 meeting marked the end of an easing cycle for the Fed.

Looking at the individual optimal adjustment for each mandates (depicted in the middle panels), we can see that the OPP for the “traditional” mandates ($\delta_t^{*\pi}$ and δ_t^{*u}) are aligned: Inflation was too low and unemployment too high, so that both mandates call for lower rates (this is in contrast to the previous 1990-M6 case study). As a result, $\delta_t^{*(2)}$, the OPP for a dual inflation-unemployment targeter, rounds to $-.75 < 0$ percentage points. At the same time

however, credit growth was very strong and called for a higher interest rate ($\delta_t^b \approx .25 > 0$) in order to more quickly bring back credit growth in line with target.³⁸ A policy maker with three mandates is thus confronted with another static trade-off in that the OPP for credit growth is positive but the OPPs for inflation and unemployment are negative.

Because of this static trade-off the three-mandate OPP ($\delta_t^{*(3)} \approx -.25$) is lower (in absolute value) than the dual-mandate OPP ($\delta_t^{*(2)}$): it still calls for lowering the fed funds rate to raise inflation and lower unemployment, but at the same time it tries to limit the side-effects on credit growth (as shown by the counter-factual forecasts displayed in the top panels).

5.2.2 On the effect of uncertainty on the OPP

To illustrate the effects of IR and model uncertainty on the OPP and on a researcher's ability to discard optimality, we consider the Fed as of 2008-M4. This particular example is interesting, because it is 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.³⁹ An interesting question in hindsight is thus whether the 2008-M4 fed funds rate was optimal. In other words, can we conclude that the Fed should have lowered its fed funds rate earlier *given* the FOMC forecasts of the time and *given* the uncertainty attached to its forecasts.

We consider a central bank with a traditional inflation-unemployment dual mandate as in policy problem (14), and Figure 6 has the same structure as our previous plots except that the top panels now also report the 68 percent confidence intervals calculated from past forecast errors over the previous 10 years. In the scatter plot, we also added the confidence intervals

³⁸Echoing an argument made by Svensson (2017a), the fed funds rate is a blunt instrument to control credit growth. The peak effect of the policy change only happens after 3 years (middle-right panel of Figure 5), by which time credit growth was already expected to be back to target. As a result, the overall reduction in the loss function (15) is quantitatively small. Stated differently, while a higher fed funds rate does reduce credit growth, it does so with considerable delay, and thus offers little benefits. In contrast, the cost in terms of higher unemployment is substantial. Thus, the cost-benefit ratio of leaning against the wind is relatively high.

³⁹By the end of 2008 however, unemployment had reached 7.3 percent, and the Fed had dropped the fed funds rate by almost 2ppt (to the zero lower bound) in the span of only three months (September-December) following the failure of Lehman Brothers in September 2008.

for the OPP; the yellow and grey cones. The yellow cones show the 68 percent confidence interval for the OPP given IR estimation uncertainty only. Then, the grey cones show the 68 percent confidence interval for the OPP given IR estimation uncertainty *and* model specification uncertainty. If that cone does not overlap with the x-axis, we can conclude that there is a failure to optimize with 68 percent probability.

Ignoring model specification uncertainty for now, a quick inspection of the median SEP forecast and the poor unemployment outlook suggests that policy in 2008-M4 was not optimal: the Fed should have cut rates more aggressively. The OPP confirms this impression and the median value for the OPP rounds to -0.5ppt (dash-black line), driven precisely by this expected increase in unemployment.⁴⁰ In other words, the OPP signals that the Fed should have lowered the fed funds rate faster than it did, given the median SEP forecast for unemployment.

However, note the large uncertainty around the unemployment forecast. In fact, in this case forecast uncertainty is so large that we cannot discard optimality despite the substantial expected increase in unemployment. The 68% grey cones overlap with the x-axis, meaning that there is a more than 32 percent chance that the chosen policy p_t^0 actually balances the (unobserved) conditional expectations for inflation and unemployment, *even though* the FOMC point forecasts suggest otherwise.

To summarize, as of 2008-M4 the expected increase in unemployment was not strong enough and too uncertain to justify a more aggressive monetary stimulus and a faster drop in the policy rate.

In this context, 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 to stabilize the economy. Thus, Figure 7 displays the usual set of plots but for 2010-M4 and for the slope instrument, so that the middle panels show the impulse

⁴⁰The positive inflation gap is seen as transitory and as a result displays little correlation with the delayed effect of the fed funds rate on inflation. As a result, there is little the Fed can do about that positive inflation gap, and the OPP for inflation rounds to zero. Note how this case captures a common wisdom of central banking: central banks should “look through” transitory inflationary episodes.

responses of inflation and unemployment to a 1ppt increase in the slope of the yield curve. This time, the deviations from targets are so large that we can clearly discard optimality at the 68 percent confidence level. In fact, the OPP calls for lowering the slope instrument by an additional 150 basis.⁴¹

5.3 A retrospective analysis of US monetary policy

In this section, we systematically study the optimality of the Fed policy over 1980-2018, in effect repeating our case studies of section 5.2 for every single FOMC meeting since 1980 for which we have the conditional forecasts.

Figure 8 displays in the top row the time series of the two instruments that we considered: the fed funds rate and the slope of the yield curve (the difference between the 10-year bond yield and the fed funds rate). The bottom panels report the corresponding OPPs along with their confidence intervals, as implied by both impulse response and model uncertainty. In addition, to summarize the distance to optimality in terms of “welfare loss”, the bottom panels of Figure 8 also show in shades of red the additional loss $\Delta\mathcal{L}_t$ incurred by the observed deviation of optimality, expressed in ppt of average extra inflation and unemployment gaps over the next 5 years. This welfare loss metric is useful to put into perspective the magnitude of the welfare cost of an optimality deviation.⁴²

While the fed funds rate has not been set exactly at its optimal level since 1980, the optimal adjustment (in absolute value) is overall relatively small averaging only 25 basis points over the full sample. There are only three larger misses; large misses in the early 1980s, a relatively small but significant miss in 2003 (a case we already discussed in the

⁴¹A 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.

⁴²We compute

$$\Delta\mathcal{L}_t(\hat{\delta}_t) = \mathcal{L}_t(\hat{\delta}_t) - \mathcal{L}_t(0) , \quad \text{with} \quad \mathcal{L}_t(\delta_t) = \mathbb{E}_t(Y_t^0 + \mathcal{R}_t\delta_t)'(Y_t^0 + \mathcal{R}_t\delta_t) \quad (16)$$

where $\mathcal{L}_t(\hat{\delta}_t)$ is the value of the loss function under the optimal policy (i.e., after the actual policy has been corrected with the optimal discretionary adjustment $\hat{\delta}_t$ and where $\mathcal{L}_t(0) = \mathbb{E}_t(Y_t^{0'}Y_t^0)$ is the value of the loss function under the policy actually implemented at time t .

previous section) and a miss at the onset of the great recession when the fed funds rate could have been brought down more rapidly to the zero lower-bound (although that OPP is not significant, as we discussed above).⁴³

The largest miss is in terms of the slope instrument, which could have been used more aggressively during the financial crises and its aftermath, a conclusion echoing that of Eberly, Stock and Wright (2019). The OPP drops rapidly to -2 percentage points and only slowly revert back to zero. In fact, the OPP for the slope instrument remains significantly different from zero over the whole 2009-2013 period, indicating that for the period 2008-2012 a more extensive use of slope policies would have brought the slope of the term structure closer to optimal.⁴⁴

6 Conclusion

In this paper we propose a methodology for detecting optimization failures in the policy decision process. The proposed OPP statistic determines if the proposed policy choice is non-optimal, and it can highlight which trade-offs have been overlooked.

Although little touched on in this paper, the OPP can also be used by the policy makers themselves in order to articulate their views and communicate their policy decisions, either to their peers (in the context of deliberations among FOMC members for instance) or to private actors (market participants for instance). In particular, the transparent regression interpretation can help policy makers articulate and quantify their policy prescriptions around three central concepts (i) their preference between different objectives, (ii) their assessment of the economic outlook, and (iii) their views on the effects of policy.⁴⁵

⁴³The large miss in the early 80s should be taken with caution. Our approach is valid as long as there are no regime changes and our IR estimates capture the transmission of monetary policy in the current regime. This approach is reasonable in the post-1990 era of anchored inflation expectations, but it is more contentious in the early 1980s period, where the Fed was actively trying to anchor inflation expectations at a lower level. In that context, our finding of an “overly tight” fed funds rate in the early 80s could just be capturing the Fed’s efforts to engineer a regime switch.

⁴⁴We re-iterate that our approach does not highlight which specific policies would have been able to induce such shift in the slope.

⁴⁵These three concepts are the three central sources of disagreement among policy makers (See e.g., Orphanides, 2019, for a recent example).

While we focused on monetary policy making to illustrate the OPP,⁴⁶ our framework can be used to inform the policy responses in many other contexts where optimality is hard to assess. 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. Similar to Taylor rules in the context of monetary policy, 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.⁴⁷ 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.⁴⁸ The OPP would provide a quantitative criterion to formalize how a policy maker should balance (in expectation) fiscal discipline with (say) 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.⁴⁹

⁴⁶In the context of monetary policy, the OPP statistic could also be trivially adapted to alternatives to inflation targeting, such as average inflation targeting or nominal GDP targeting, which are actively debated as the Fed is reconsidering its policy framework.

⁴⁷In 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.

⁴⁸To 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.

⁴⁹The 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).

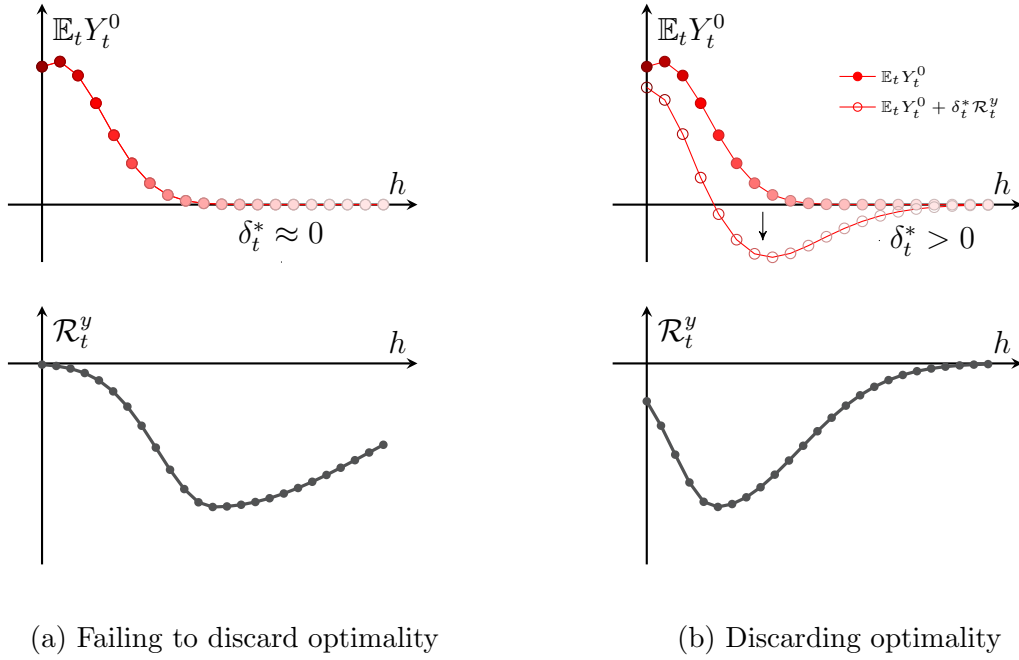
References

- Barnichon, Regis.** 2009. “The optimal level of reserves for low-income countries: Self-insurance against external shocks.” *IMF staff Papers*, 56(4): 852–875.
- Bénassy-Quéré, Agnès, Benot Coeuré, Pierre Jacquet, and Jean Pisani-Ferry.** 2018. *Economic Policy: Theory and Practice. Second edition.* Oxford University Press.
- Blanchard, Olivier J.** 2016. “Do DSGE Models Have a Future?” *Policy Brief PB*, 16–11.
- Blanchard, Olivier J., and Stanley Fischer.** 1989. *Lectures on macroeconomics.* MIT press.
- Brainard, William C.** 1967. “Uncertainty and the Effectiveness of Policy.” *The American Economic Review*, 57(2): 411–425.
- Chetty, Raj.** 2009. “Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods.” *Annu. Rev. Econ.*, 1(1): 451–488.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 2005. “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy.” *Journal of Political Economy*, 113(1): 1–45.
- 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.
- Elliot, Graham, and Allan Timmermann.** 2016. *Economic Forecasting.* Princeton University Press.
- Galí, Jordi.** 2015. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications.* Princeton University Press.
- 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.** 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.

- Jordà, Òscar, Moritz Schularick, and Alan M Taylor.** 2013. “When credit bites back.” *Journal of Money, Credit and Banking*, 45(s2): 3–28.
- 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.
- Lakdawala, Aeimit.** 2019. “Decomposing the Effects of Monetary Policy Using an External Instruments SVAR.” *Journal of Applied Econometrics*. forthcoming.
- Leeper, Eric M., and Tao Zha.** 2003. “Modest policy interventions.” *Journal of Monetary Economics*, 50(8): 1673–1700.
- Ljungqvist, Lars, and Thomas J. Sargent.** 2004. *Recursive Macroeconomic Theory*. Cambridge, Massachusetts: MIT Press.
- Lucas, Robert Jr.** 1976. “Econometric policy evaluation: A critique.” *Carnegie-Rochester Conference Series on Public Policy*, 1(1): 19–46.
- Newey, Whitney K., and Kenneth D. West.** 1994. “Automatic Lag Selection in Covariance Matrix Estimation.” *The Review of Economic Studies*, 61(4): 631–653.
- Nocedal, Jorge, and Stephen Wright.** 2006. *Numerical optimization*. Springer Science & Business Media.
- Orphanides, Athanasios.** 2019. “Monetary Policy Strategy and its Communication.” Working paper, presented at Jackson Hole Conference August 22, 2019.
- Orphanides, Athanasios, and John C Williams.** 2012. “Monetary Policy Mistakes and the Evolution of Inflation Expectations.” In *The Great Inflation: The Rebirth of Modern Central Banking*. 255–288. University of Chicago Press.
- Primiceri, Giorgio. E.** 2005. “Time Varying Structural Vector Autoregressions and Monetary Policy.” *The Review of Economic Studies*, 72: 821–852.
- Ramey, Valerie.** 2016. “Macroeconomic Shocks and Their Propagation.” In *Handbook of Macroeconomics*, ed. J. B. Taylor and H. Uhlig. Amsterdam, North Holland: Elsevier.
- Rudebusch, Glenn, and Lars E.O. Svensson.** 1999. “Policy rules for inflation targeting.” In *Monetary policy rules*. 203–262. University of Chicago Press.
- Rudebusch, Glenn D.** 2002. “Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty.” *The Economic Journal*, 112(479): 402–432.
- Sack, Brian.** 2000. “Does the fed act gradually? A VAR analysis.” *Journal of Monetary Economics*, 46(1): 229 – 256.
- Scheffe, Henry.** 1953. “A Method for Judging all Contrasts in the Analysis of Variance.” *Biometrika*, 40(1/2): 87–104.

- Stock, James H., and Mark W. Watson.** 2016. "Chapter 8 - Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics." In . Vol. 2 of *Handbook of Macroeconomics*, , ed. John B. Taylor and Harald Uhlig, 415 – 525. Elsevier.
- Stock, James H., and Mark W. Watson.** 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal*, 128(610): 917–948.
- Svensson, Lars E.O.** 1999. "Inflation targeting as a monetary policy rule." *Journal of Monetary Economics*, 43(3): 607 – 654.
- Svensson, Lars E.O.** 2003. "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules." *Journal of Economic Literature*, 41(2): 426–477.
- Svensson, Lars E.O.** 2017a. "Cost-benefit analysis of leaning against the wind." *Journal of Monetary Economics*, 90: 193–213.
- Svensson, Lars E.O.** 2017b. "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.
- Tenreyro, Silvana, and Gregory Thwaites.** 2016. "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions." *American Economic Journal: Macroeconomics*, 8(4): 43–74.
- Tinbergen, Jan.** 1952. *On the Theory of Economic Policy*. North-Holland Publishing Company, Amsterdam.
- White, H.** 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.

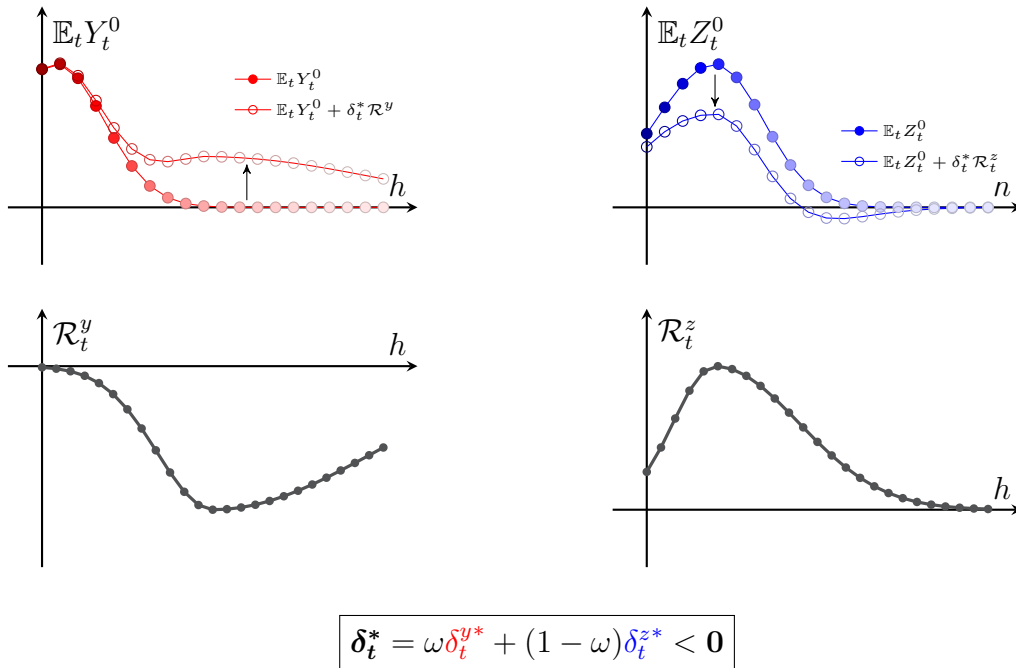
Figure 1: Illustration of the OPP: a single mandate



$$\hat{\delta}_t^* = - \left(\mathcal{R}_t^{y'} \mathcal{R}_t^y \right)^{-1} \mathcal{R}_t^{y'} \mathbb{E}_t Y_t^0$$

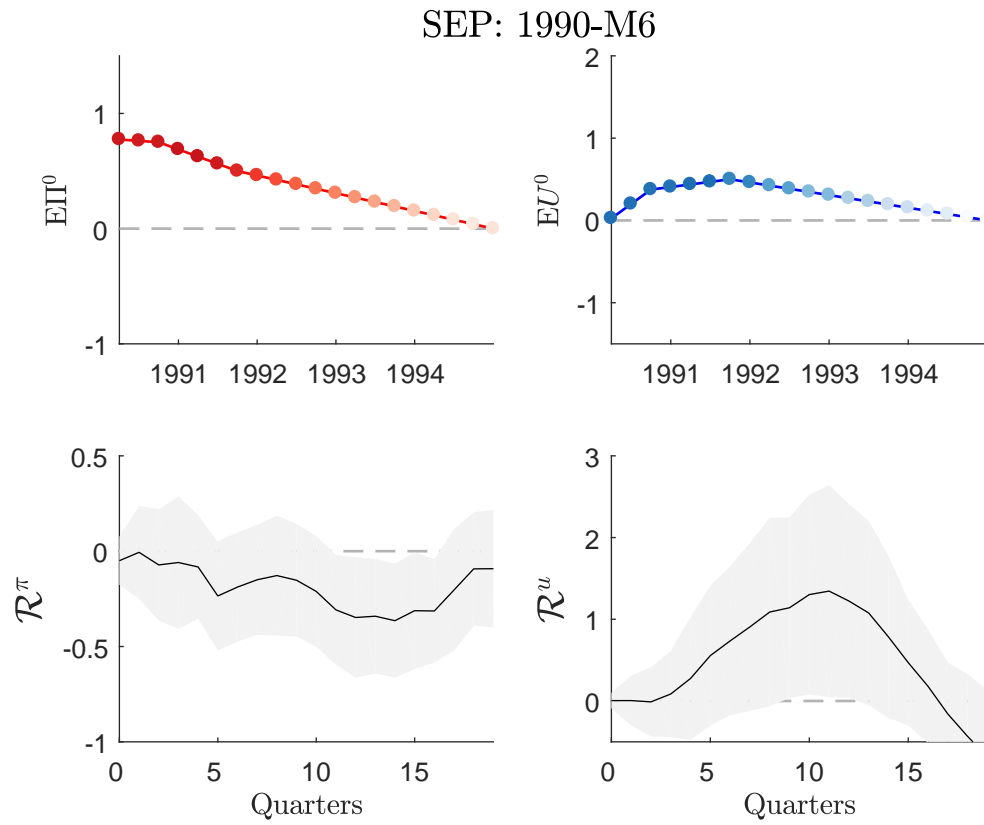
Notes: Top panel: Expected paths for the policy objective Y_t under the initial policy choice (filled circles) and under the policy modified with the OPP δ_t^* (empty circles). Bottom row: impulse responses of the policy objective y_t to a 1ppt innovation to the policy instrument.

Figure 2: Illustration of the OPP: a dual mandate




Notes: Top panel: Expected paths for the two policy objectives Y_t (in red) and U_t (in blue) under the initial policy choice (filled circles) and under the policy modified with the OPP δ_t^* (empty circles). Bottom row: impulse responses of the policy objective y_t and z_t to a 1ppt innovation to the policy instrument. The OPP δ_t^* is a weighted average of δ_t^{y*} and δ_t^{z*} , the OPPs corresponding to the single mandate y or z , with the weight $\omega = \frac{1}{1+\lambda/\kappa}$ where $\kappa = \|\mathcal{R}_t^y\|/\|\mathcal{R}_t^z\|$.

Figure 3: A static inflation-unemployment trade-off




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

Figure 4: A regression interpretation of the OPP




Notes: In red, a scatter plot of the median FOMC inflation gap forecast ($E\Pi^0$) as of 1990-M6 against the impulse response of the inflation gap (\mathcal{R}^π) with the best linear fit (red dashed line). The slope of that regression line is $\delta^{*\pi}$, the OPP for a strict inflation targeter. Same information in blue for the unemployment gap. The dashed black line depicts the best-linear fit after including all points, and the slope of that line is δ^* , the OPP for a dual targeter.

Figure 5: Dynamic trade-offs and leaning against the wind



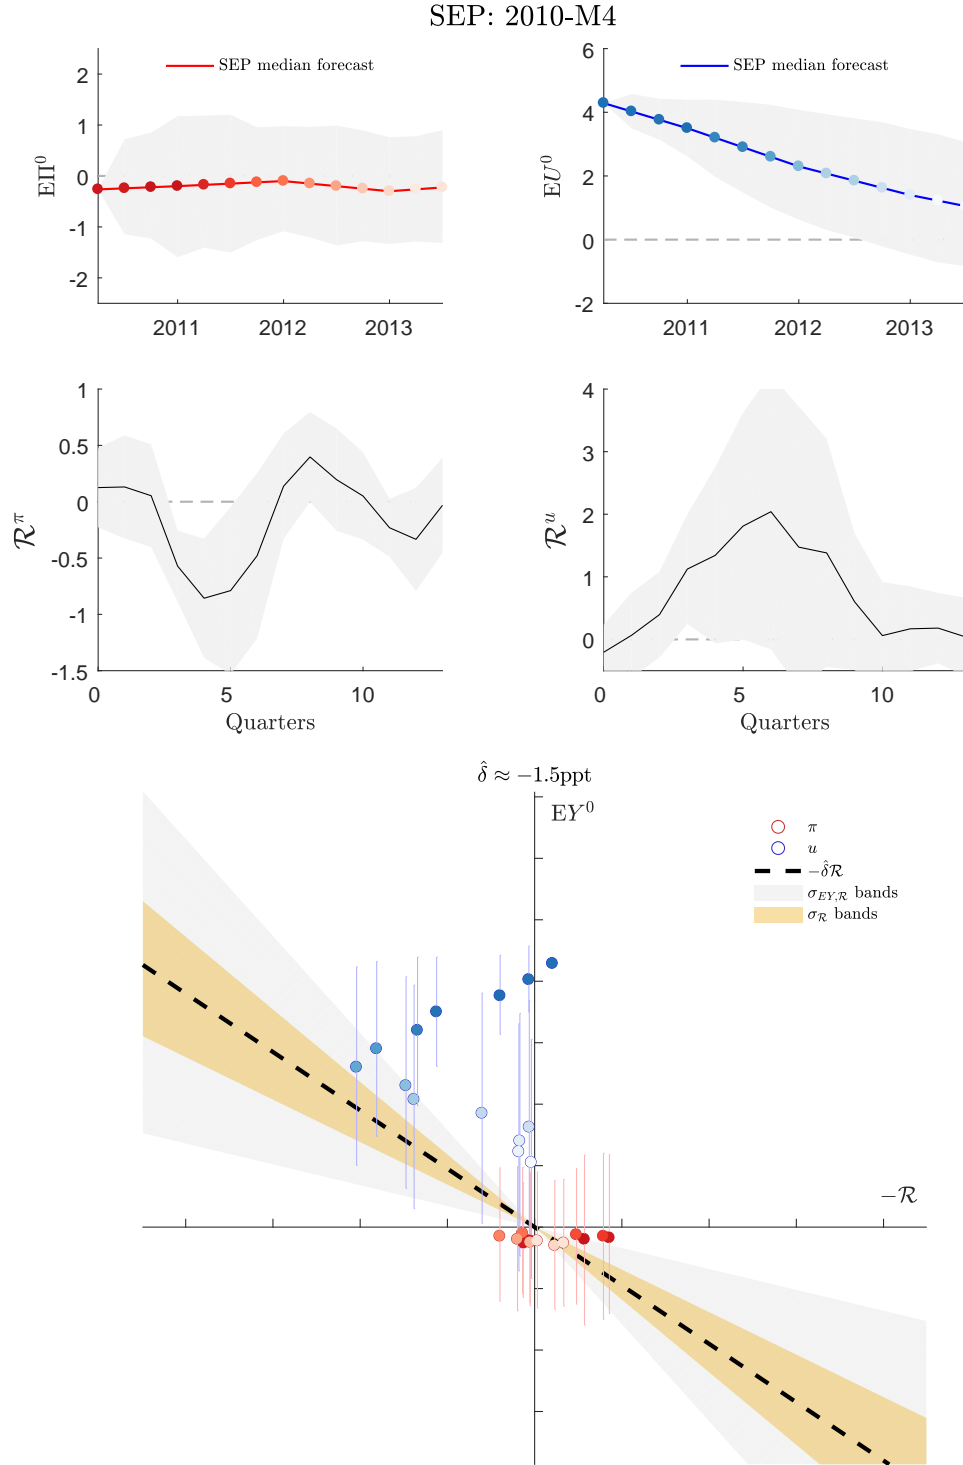
Notes: Top panel: median SEP forecasts for the inflation, unemployment and debt growth gaps as of 2003-M6. Middle panel: impulse responses of the inflation, unemployment, and debt growth gaps to a fed funds rate shock. In red (resp. blue, green) is the OPP $\delta^{*\pi}$ (resp. δ^{*u} , δ^{*b}) for a strict inflation (resp. unemployment, debt growth) targeter. Bottom panel: scatter plot of the median SEP forecasts for inflation (red), unemployment (blue) and debt growth (green) against the corresponding impulse responses. The dot-dashed black line depicts the best-linear fit after including inflation and unemployment points, and the slope of that line is $\delta^{*(2)}$, the OPP for a dual inflation-unemployment targeter. The thick-dashed line depicts the best-linear fit after including all points, and the slope of that line is $\delta^{*(3)}$, the OPP for a central bank with three mandates.

Figure 6: Taking uncertainty into account: fed funds rate policy



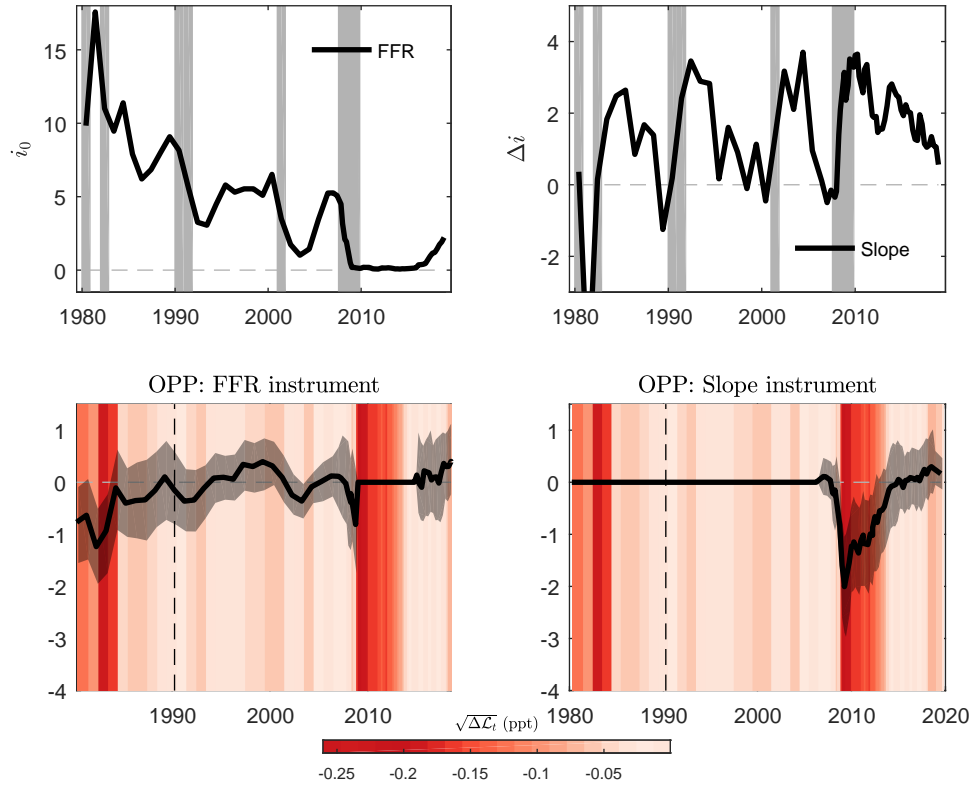
Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2008-M4 (in red and blue) along with the 67 percent confidence bands. Middle panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock with the 95 percent confidence intervals. Bottom panel: scatter plot of the median SEP inflation and unemployment gap forecasts (EY^0) against (minus) the impulse responses (IR) of the inflation and unemployment gaps ($-\mathcal{R}$) with the best linear fit (dashed-black line) with average slope $\hat{\delta}$. The beige area areas depicts the 95 confidence intervals from IR estimation uncertainty, and the light grey shaded areas depicts the 67 confidence intervals from IR estimation and model mis-specification uncertainty.

Figure 7: Taking uncertainty into account: slope policy



Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps as of 2010-M4 (in red and blue) along with the 67 percent confidence bands uncertainty. Middle panel: impulse responses of the inflation and unemployment gaps to a slope policy shock with the 95 percent confidence intervals. Bottom panel: scatter plot of the median SEP inflation and unemployment gap forecasts (EY^0) against (minus) the impulse responses (IR) of the inflation and unemployment gaps ($-\mathcal{R}$) with the best linear fit (dashed-black line) with average slope $\hat{\delta}$. The beige area areas depicts the 95 confidence intervals from IR estimation uncertainty, and the light grey shaded areas depicts the 67 confidence intervals from IR estimation and model mis-specification uncertainty.

Figure 8: OPP for the Fed instruments (1980-2018)



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 67 percent probability. The shades of red represent the magnitude of the additional loss incurred by deviating from optimality, with darker colors indicating larger losses. The additional loss $\sqrt{\Delta\mathcal{L}_t}$ is expressed in ppt of average extra inflation and unemployment gaps over the next 5 years.

Online Appendix (not for publication)

Proof of Proposition 1. Note that

$$\begin{aligned}
\frac{\partial}{\partial p_t} \mathcal{L}_t &= \frac{\partial}{\partial p_t} \mathbb{E}_t \|f_t(p_t, X_t, \xi_t)\|^2 \\
&= \lim_{h \rightarrow 0} \frac{1}{h} (\mathbb{E}_t \|f_t(p_t + h, X_t, \xi_t)\|^2 - \mathbb{E}_t \|f_t(p_t, X_t, \xi_t)\|^2) \\
&= \lim_{h \rightarrow 0} \mathbb{E}_t \left(\frac{\|f_t(p_t + h, X_t, \xi_t)\|^2 - \|f_t(p_t, X_t, \xi_t)\|^2}{h} \right) \\
&= \lim_{h \rightarrow 0} \mathbb{E}_t \left(\frac{\partial}{\partial p_t} \|f_t(p_t(h), X_t, \xi_t)\|^2 \right),
\end{aligned}$$

where $p_t(h) \in (p_t, p_t + h)$ exists by the continuity of f_t on \mathcal{D} by the mean value theorem. We have for all $p_t \in \mathcal{D}$ that

$$\begin{aligned}
\frac{\partial}{\partial p_t} \|f_t(p_t, X_t, \xi_t)\|^2 &= 2 \left(\frac{\partial f(p_t, X_t, \xi_t)}{\partial p'_t} \right)' f_t(p_t, X_t, \xi_t) \\
&= 2R_t(p_t, X_t)' f_t(p_t, X_t, \xi_t) \\
&\leq 2\|R_t(p_t, X_t)\| \|f_t(p_t, X_t, \xi_t)\| \\
&\leq 2\Delta Z_t,
\end{aligned}$$

where the second equality follows as R_t is independent of ξ_t , the first inequality from Cauchy-Schwartz and the second inequality from Assumption 1 and 2.(c). Note that Δ is a finite constant and $\mathbb{E}_t(Z_t) < \infty$, thus we can use the dominated convergence theorem to conclude that almost surely

$$\begin{aligned}
\lim_{h \rightarrow 0} \mathbb{E}_t \left(\frac{\partial}{\partial p_t} \|f_t(p_t(h), X_t, \xi_t)\|^2 \right) &= \mathbb{E}_t \left(\lim_{h \rightarrow 0} \frac{\partial}{\partial \delta_t} \|f_t(p_t(h), X_t, \xi_t)\|^2 \right) \\
&= \mathbb{E}_t \left(\frac{\partial}{\partial \delta_t} \|f_t(p_t, X_t, \xi_t)\|^2 \right) \\
&= 2R'_t \mathbb{E}_t f(p_t, X_t, \xi_t).
\end{aligned}$$

Hence, we have that the loss function $\mathbb{E}_t \|f_t(p_t, X_t, \xi_t)\|^2$ is continuously differentiable on \mathcal{D} , thus by Lemma 4.3.1 in Dennis and Schnabel (1996) the optimal policy p_t^{opt} satisfies the gradient condition $\frac{\partial}{\partial p_t} \mathbb{E}_t \|f_t(p_t, X_t, \xi_t)\|^2|_{p_t=p_t^{\text{opt}}} = 0$. Hence, if p_t^0 is optimal we must have that $\mathcal{R}'_t \mathbb{E}_t Y_t^0 = 0$, which since $\mathcal{R}_t = R_t(p_t^0, X_t)$ has full column rank implies that $\delta_t^* = -(\mathcal{R}'_t \mathcal{R}_t)^{-1} \mathcal{R}'_t \mathbb{E}_t Y_t^0$ must satisfy $\delta_t^* = 0$ if p_t^0 is optimal. \square

Proof of Proposition 2. For convenience let $\mathbb{E}_t Y_t^{\text{opt}} = \mathbb{E}_t f_t(p_t^{\text{opt}}, X_t, \xi_t)$. 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 $R_t(p_t^0, X_t)'R_t(p_t^0, X_t)$ is non-singular and thus there exists a constant $\epsilon_1 > 0$ such that

$$\|(R_t(p_t^0, X_t)'R_t(p_t^0, X_t))^{-1}\| \leq \frac{\kappa}{\mu_{\min}} \quad \forall p_t^0 \in \mathcal{N}(p_t^{\text{opt}}, \epsilon_1) .$$

Let

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

Now consider

$$\begin{aligned} p_t^0 + \delta_t^* - p_t^{\text{opt}} &= p_t^0 - p_t^{\text{opt}} - (\mathcal{R}_t' \mathcal{R}_t)^{-1} \mathcal{R}_t' \mathbb{E}_t Y_t^0 \\ &= -(\mathcal{R}_t' \mathcal{R}_t)^{-1} [\mathcal{R}_t' \mathbb{E}_t Y_t^0 - (\mathcal{R}_t' \mathcal{R}_t)(p_t^{\text{opt}} - p_t^0)] \\ &= -(\mathcal{R}_t' \mathcal{R}_t)^{-1} [\mathcal{R}_t' \mathbb{E}_t Y_t^{\text{opt}} - \mathcal{R}_t' (\mathbb{E}_t Y_t^{\text{opt}} - \mathbb{E}_t Y_t^0 - \mathcal{R}_t(p_t^{\text{opt}} - p_t^0))] . \end{aligned}$$

Now the Lipschitz assumption (e.g. Assumption 2.2) implies that

$$\|\mathbb{E}_t Y_t^{\text{opt}} - \mathbb{E}_t Y_t^0 - \mathcal{R}_t(p_t^{\text{opt}} - p_t^0)\| \leq \frac{\gamma}{2} \|p_t^{\text{opt}} - p_t^0\|^2 ,$$

see Lemma 4.1.12 in Dennis and Schnabel (1996). Recall from proposition 1 that $R_t(p_t^{\text{opt}}, X_t)' \mathbb{E}_t Y_t^{\text{opt}} = 0$, and thus we have by assumption 2.1 that

$$\|\mathcal{R}_t' \mathbb{E}_t Y_t^{\text{opt}}\| \leq c \|p_t^0 - p_t^{\text{opt}}\| .$$

Combining the bounds gives

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

which can be simplified using the definition of ϵ to obtain

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

This completes the proof. □

Proof of Proposition 3. Note that since \mathcal{R}_t is independent of p_t the derivative of every el-

ement of \mathcal{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

$$f_t(p_t^0 + \delta_t, X_t, \xi_t) = f_t(p_t^0, X_t, \xi_t) + \mathcal{R}_t \delta_t$$

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 $\mathcal{R}'_t \mathbb{E}_t f_t(p_t^0 + \delta_t, X_t, \xi_t) = 0$, which using the expansion can be rewritten as

$$\mathcal{R}'_t (\mathbb{E}_t f_t(p_t^0, X_t, \xi_t) + \mathcal{R}_t \delta_t) = 0$$

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