

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

February 14, 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 by studying the optimality of past and present US monetary policy decisions.

*We thank Martin Eichenbaum, Ethan Ilzetzki, Sylvain Leduc, Huiyu Li, Nicolas Petrosky-Nadeau, Adam Shapiro and seminar participants at the Bank of England, McGill, Ohio State, the University of Ottawa, the Cleveland Fed, the St Louis Fed and the 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. Mesters acknowledge support from the Spanish Ministry of Science and Technology: Grant ECO2015-68136-P, Marie Skłodowska-Curie Actions: FP7-PEOPLE-2012-COFUND Action grant agreement no: 600387, the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075) and the Analysis of Big Data in Economics and Empirical Applications - 2016 BBVA Foundation Grants for Scientific Research.

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 the 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 the OPP 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. If the policy instruments are set optimally, perturbations to the instruments should have no first-

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

order effect on the policy maker’s loss function. If this is not the case, the instruments can be adjusted until that first-order condition is satisfied, and this “optimal policy perturbation” is a measure of the distance between the current and the optimal policy.

The OPP only depends on (i) the policy maker’s forecasts for the policy objectives — the targets— conditional on the policy choice, and (ii) the impulse responses of the targets to shocks to the policy instruments. In fact, calculating the OPP boils down to an OLS regression where we regress the forecasts for the targets on their impulse responses to policy shocks. Intuitively, recall that 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 lower the loss function. In contrast, when the policy is not optimal, a regression of the forecasts on the IRs will determine how to use the IRs to minimize the sum-of-squared deviations of the targets, i.e., how to optimally adjust the instruments. This optimal adjustment is the OPP.

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—, the researcher could wrongly conclude that a policy is not optimal. To guard against such risks, we derive confidence bands around the 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 further², 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

²As 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.

used more aggressively to bring the slope down in line with optimality.

This paper relates to a number of influential 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.³

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 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 (2011) 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

³The 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.

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).⁴ 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.⁵

A large literature on optimal monetary 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 broad range of models.⁶ 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.⁷

⁴As 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 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*).

⁵For 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.

⁶As 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.”

⁷In 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)

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.

Finally, from a general macro perspective, this paper highlights an important link between the structural impulse response literature (e.g. Ramey, 2016) and the optimal policy literature (e.g., Svensson, 2010). While the impulse response literature has traditionally 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.

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 baseline OPP statistic and discusses the 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.

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 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 Ω_t , e.g. $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot|\Omega_t)$, and with respect to the true distribution of the target deviations $y_{m,t+h} - y_m^*$.⁸

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. A generic model for the targets is given by $Y_t = f(p_t, X_t)$, where $f(\cdot, \cdot)$ is assumed to be differentiable with respect to its first argument. The function $f(p_t, X_t)$ maps the policy choice p_t and an arbitrary set of relevant variables X_t to the targets Y_t . For instance, X_t may include lagged values of the targets and any other variables that determine Y_t including expectations for future variables.

With these definitions we rewrite the policy makers' problem as the following least squares problem⁹

$$\min_{p_t \in \mathbb{R}^K} \mathbb{E}_t \|Y_t\|^2 \quad \text{with} \quad Y_t = f(p_t, X_t). \quad (2)$$

Problem (2) 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, see Tinbergen (1952).

⁸The target value y_m^* can equally well be considered horizon dependent, e.g. $y_{m,t+h}^*$ instead of y_m^* .

⁹We consider the usual vector norm where for any $a \in \mathbb{R}^n$ we have that $\|a\| = \sqrt{\sum_{t=1}^n a_t^2}$.

The policy makers’ *perceived* solution to problem (2) is denoted by p_t^0 . In practice, the policy makers’ choice p_t^0 may not be optimal—a case of optimization failure in the policy decision process—, because the underlying model $f(\cdot, \cdot)$ is unknown or too “complicated” to allow the policy maker to determine whether p_t^0 is optimal. Specifically, even if the policy maker has access to the true underlying model, that model is often too complex to write down and $f(p_t, X_t)$ is very costly to compute, being the result of many model iterations and judgement calls (Manganelli, 2009).¹⁰ With a high cost of computing $f(p_t, X_t)$, the policy maker is only able to assess the value of $f(p_t, X_t)$ for a small set of values for p_t , i.e., for a small number of alternative policy strategies. This approach amounts to an incomplete grid search that can leave the policy maker unsure of having set the optimal policy.¹¹

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. Importantly, we do not assume that the researcher has access to the underlying model $f(p_t, X_t)$.

To assess optimality, we study the implications for the loss function (1) of modifying 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 one could consider a number arbitrary deviations from the policy p_t^0 and evaluate whether any of them lower the loss function, a more efficient approach consists in finding a well-chosen perturbation that exploits the first-order condition for optimality.

Specifically, we will use as perturbation the first step of a Gauss-Newton optimization algorithm, an algorithm that only relies on the first and second derivatives of the objective

¹⁰A 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.

¹¹In 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 to discuss.

function.¹² These methods are attractive in our setting because the first and second derivatives of the loss function only involve two typically known or estimable quantities: (i) the forecast of the policy objectives conditional on the policy choice, and (ii) the IR of the policy objectives to shocks to the policy instruments.¹³

Under a “no regime-switching” assumption described below, we can derive an optimal adjustment that makes policy optimal in one iteration, which means that the adjustment is the “Optimal Policy Perturbation”, i.e., the distance between the current policy choice and the optimal policy. The OPP thus allows us to not only detect optimization failures but also to identify the reason for these failures and the correction needed to go back to optimality.

Even without the “no regime-switching” assumption, our method can be used to detect optimization failures in the policy decision process. In that case however, we cannot quantify the magnitude of the optimization failure.

3.1 Measuring distance to optimality

To measure the distance to optimality of the current policy choice p_t^0 , we make the following assumption.

Assumption 1 (No Regime-Switching). *The $M(H + 1) \times K$ matrix of impulse responses*

$$\mathcal{R} = \frac{\partial Y_t}{\partial \delta_t'} \quad (3)$$

is fixed and does not depend on δ_t .

Assumption 1 imposes that the impulse responses are constant over time and do not depend on δ_t (or p_t). As we detail in the appendix, the time-invariance restriction is only imposed for convenience and the main ideas of this paper continue to apply for time-varying impulse responses \mathcal{R}_t , except that the data requirements for estimating \mathcal{R}_t are stronger.

¹²Specifically, the Gauss-Newton algorithm consists in approximating the objective function with its second-order Taylor expansion, and finding the minimum of that approximate function.

¹³For (i), policy makers, such as central banks, routinely provide their forecasts to make their policy decisions transparent, e.g., the Survey of Economic Projections from FOMC members. For (ii), we can draw on a large literature to estimate impulse responses to policy shocks.

The “independence from δ_t ” restriction is important however. It imposes that the policy innovations δ_t do not lead agents to revise their expectation about the policy maker’s reaction function (which affects their decision making process).

Stated differently, Assumption 1 states that the OPP is not affected by the Lucas critique, a reasonable assumption for small deviations from optimality. Indeed, if the policy innovations are sufficiently small —modest interventions in the sense of Leeper and Zha (2003)—, we can reasonably consider that the economy remains in the same regime before and after a policy innovation, and assumption 1 is satisfied.¹⁴

With Assumption 1 in place we mimic the policy makers’ problem and determine whether the current policy choice is optimal. The formal statement and solution are summarized in the following theorem.

Theorem 1. *Given Assumption 1 and under the assumption that \mathcal{R} and $\mathbb{E}_t Y_t^0$, with $Y_t^0 = f(p_t^0, X_t)$, are given with $\mathcal{R}'\mathcal{R} \succ 0$, the Optimal Policy Perturbation (OPP) is given by*

$$\delta_t^* = \arg \min_{\delta_t \in \mathbb{R}^K} \mathbb{E}_t \|Y_t\|^2 = -(\mathcal{R}'\mathcal{R})^{-1} \mathcal{R}' \mathbb{E}_t Y_t^0. \quad (4)$$

The proof is deferred to the appendix.

Under assumption 1, if δ_t^* is not equal to zero the policy maker’s policy choice p_t^0 is not optimal as the policy maker could adjust to $p_t^0 + \delta_t^*$ in order to achieve a lower expected loss $\mathbb{E}_t \|Y_t\|^2$. Therefore δ_t^* is a statistic that measures the “distance to optimality” of the policy choice p_t^0 making δ_t^* the *Optimal Policy Perturbation* (OPP).

A few comments are in order.

First, the OPP only requires the forecasts of the policy maker $\mathbb{E}_t Y_t^0$ (conditional on his optimal policy choice p_t^0) and the impulse responses \mathcal{R} .

Second, 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} . To get the intuition behind this OLS interpretation, first notice that under Assumption 1, the effect of

¹⁴According to Leeper and Zha (2003), a *modest* policy intervention does not significantly shift agents’ beliefs about policy regime and does not induce the changes in behavior that Lucas (1976) emphasizes.

a small policy perturbation on Y_t is given by¹⁵

$$Y_t = \mathcal{R}\delta_t + Y_t^0,$$

which means that the set of impulse responses \mathcal{R} captures the effect of a change in the policy instrument on the policy maker’s objectives.¹⁶ The goal of the optimal perturbation δ_t^* is then to use the impulse response (\mathcal{R}) in order to minimize the squared deviations of Y_t . This is nothing but an OLS regression of Y_t^0 on \mathcal{R} , except one with a minus sign in front of the coefficient estimate since the goal is not to best fit the path for Y_t^0 , but instead to best “undo” it. Since the future value of the objectives is unknown, the policy maker instead tries to undo $\mathbb{E}_t Y_t^0$, the expected path for Y_t^0 conditional on the policy choice p_t^0 , and we obtain the OLS-type formula (4).

Third, expression (4) makes clear that our approach relies on a *direct* mapping between the policy perturbations and the target variables, as opposed to specifying a recursive model for the targets.¹⁷ This is attractive in our context, because it allows us to represent the policy maker’s problem as a static problem, whereby policy makers are not just facing M dynamic targets but instead $M(H + 1)$ static targets. Under that light, we can view the impulse responses as capturing the joint effects of the instruments on all these static targets, and optimal policy involves cutting trade-offs across mandates *and* horizons. As we will see in the application section, this static representation will prove useful to understand and visualize the trade-offs at play at any point in time.

¹⁵This can be seen by taking a first-order expansion of $f(p_t^0 + \delta_t, X_t)$ around $\delta_t = 0$, which gives $f(p_t^0 + \delta_t, X_t) \approx f(p_t^0, X_t) + \mathcal{R}\delta_t$. Under Assumption 1, this approximation is exact since the higher-order terms are zero.

¹⁶In other words, Assumption 1 accomodates commonly used linear models suchs as vector autoregressive models and linear state space models, as well as non-linear models for small perturbations δ_t around the policy choice p_t^0 .

¹⁷Recursive models such as VAR models or linear state space models are more common in the optimal policy literature see for examples Chow (1972, 1973); Sack (2000); Rudebusch (2002); Orphanides (2003); Swanson (2004) or the textbook treatment of Ljungqvist and Sargent (2004). Note that any linear recursive model can be re-written to satisfy (5) simply by adjusting the definition for Π_t^0 and U_t^0 .

3.2 Rejecting optimality

Assumption 1 imposes that \mathcal{R} , the responses to policy changes, are invariant to policy perturbations δ and thus not subject to the Lucas critique. For large deviations from optimality, the OPP can be large and thereby lead to a change in regime. While this prevents the use of the OPP as a distance-to-optimality measure, we now show it is still possible to use our approach to assess optimality.

To allow for regime changes following discretionary monetary interventions, we relax assumption 1 and allow the impulse responses to depend explicitly on the policy perturbation.

Assumption 2. *The $M(H+1) \times K$ matrix of impulse responses*

$$\mathcal{R}(\delta_t) \equiv \mathcal{R}(p_t^0 + \delta_t, X_t) = \frac{\partial Y_t}{\partial \delta_t'} \quad (5)$$

where $\mathcal{R}(\delta_t)$ is assumed differentiable with respect to δ_t .

The notation in assumption 2 leaves the dependence on p_t^0 and X_t implicit. The dependence of the impulse responses on δ_t is crucial as it allows policy responses to change with the policy perturbation, our object of interest.

With this relaxation we modify Theorem 1 as follows.

Theorem 2. *Given Assumption 2 and under the assumption that $\mathcal{R}(\delta_t)$ and $\mathbb{E}_t Y_t^0$ are given with $Y_t^0 = f(p_t^0, X_t)$ and with $\mathcal{R}(0)' \mathcal{R}(0) \succ 0$ and $\frac{\partial \text{vec}(\mathcal{R}(\delta_t)')}{\partial \delta_t'} \Big|_{\delta_t=0} \succ 0$, the Improving Policy Perturbation (IPP) is given by*

$$\delta_t^{**} = \arg \min_{\delta_t \in \mathbb{R}^K} \mathbb{E}_t \|Y_t\|^2 = - \left(\mathcal{R}(0)' \mathcal{R}(0) + (\mathbb{E}_t Y_t^{0'} \otimes I_{M(H+1)}) \frac{\partial \text{vec}(\mathcal{R}(\delta_t)')}{\partial \delta_t'} \Big|_{\delta_t=0} \right)^{-1} \mathcal{R}(0)' \mathbb{E}_t Y_t^0 \quad (6)$$

The consequence of relaxing assumption 1 is twofold.

First, the policy perturbation δ_t^{**} does no longer measure the distance to optimality. However, it does indicate whether the current policy choice is optimal: if $\delta_t^{**} \neq 0$ it is still

possible to reduce the loss by choosing $p_t^1 = p_t^0 + \delta_t^{**}$. The update δ_t^{**} is merely the first step in the Gauss-Newton algorithm.¹⁸

Second, in general, computing the policy perturbation δ_t^{**} would require large data samples. The extra difficulty comes from the fact that one needs to estimate both the unknown function $\mathcal{R}(\cdot)$ as well as its derivative $\mathcal{R}'(\cdot)$. Given the availability of instrumental variables, one could estimate $\mathcal{R}(\cdot)$ using non-parametric IV methods to estimate (see Newey, 2013), but these methods require large sample sizes and are therefore not often adopted for studying causal effects in aggregate macro time series.

However, we can considerably simplify the problem by only trying to assess whether δ^{**} is zero or not, i.e., whether the current policy choice satisfies the first-order necessary condition for optimality. This can be done by simply considering the right-most component of equation (6): $\mathcal{R}(0)' \mathbb{E}_t Y_t^0$.¹⁹ The corollary below formalizes this approach.

Corollary 1. *Given Assumption 2 and under the assumption that $\mathbb{E}_t Y_t^0$ with $Y_t^0 = f(p_t^0, X_t)$ is given, that $\mathcal{R}(\delta_t)$ is given with $\mathcal{R}(0)' \mathcal{R}(0) \succ 0$ and $\left. \frac{\partial \text{vec}(\mathcal{R}(\delta_t)')}{\partial \delta_t'} \right|_{\delta_t=0} \succ 0$, the policy choice p_t^0 is optimal if*

$$\zeta_t^* = 0, \quad \text{where} \quad \zeta_t^* = \mathcal{R}(0)' \mathbb{E}_t Y_t^0. \quad (7)$$

A limitation of ζ_t^* is that it cannot inform about the distance to the optimal policy. It can only be used to detect optimization failures, but it cannot be used to assess the magnitude of the failure. However, since the OPP will only likely fail to satisfy Assumption 1 for large values, we can use the more robust method to discard optimality.

¹⁸If it were possible to evaluate $\mathbb{E}_t Y_t^1 = \mathbb{E}_t f(p_t^1, X_t)$ then the optimal policy might be found by iterating the Gauss-Newton algorithm (provided the local quadratic approximation of the loss function is reasonable). While this is not possible for a researcher outside the policy institute, this could be possible for the policy maker's staff, since they construct the forecasts $\mathbb{E}_t Y_t$ and they can iterate the updating procedure until convergence.

¹⁹This amounts to switching from a Gauss-Newton optimization iteration (which requires the second derivative and adds extra difficulties in computing δ_t^{**}), to a gradient-descent iteration that only requires the first derivative estimated at the current policy choice.

4 Inference for OPP

The computation of the OPP requires two statistics: (i) the impulse responses \mathcal{R} , 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} 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 $\mathbb{E}_t Y_t^0$. 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.²⁰ 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.

Impulse response uncertainty

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

$$\hat{r} \sim N(r, \Omega) , \quad (8)$$

where Ω is the variance matrix of all impulse responses: across all horizons and mandates.

The normality assumption can be justified by asymptotic arguments for different estimators when taking the sample size to infinity.²¹ Further, we posit that there exists a consistent

²⁰We 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.

²¹Suppose 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 $\sqrt{n}(\hat{r} - r) \xrightarrow{d} N(0, \text{Avar}(\hat{r}))$. This implies that if we let

estimate for Ω that we denote by $\hat{\Omega}$.²²

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 do not observe the historical misspecification errors $\{\mathbb{E}_s Y_s^0 - \hat{Y}_{s|s}\}_{s=t-H}^{t-1}$ as the true model is unknown, and we cannot exploit such sequence to predict the distribution of $\mathbb{E}_t Y_t^0 - \hat{Y}_{t|t}$. 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

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

First, we assume that the misspecification error $E_{t|t} = \mathbb{E}_t Y_t^0 - \hat{Y}_{t|t}$ follows a normal distribution.²³ Second, we upper bound the variance of $E_{t|t}$ with the variance of the forecast errors, which are observable.²⁴

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

$$\mathbb{E}_t Y_t^0 \approx N \left(\hat{Y}_{t|t}, \hat{\Sigma}_{t|t} \right) , \quad (9)$$

where $\hat{\Sigma}_{t|t}$ is an (upper-bound) estimate for the mean squared forecast error $\Sigma_{t|t} = \mathbb{E}_t(Y_t - \mathbb{E}_t Y_t^0)^2$.²⁵ we obtain (8).

²²These 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} is approximately normal under standard textbook assumptions (White, 2000) and Newey and West (1994) methods deliver a consistent estimator $\hat{\Omega}$ for Ω .

²³Since 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.

²⁴This requires the assumption that the covariance between the future error and the misspecification error is zero.

$\widehat{Y}_{t|t})(Y_t - \widehat{Y}_{t|t})'$ that we estimate using the historical forecast errors

$$\widehat{\Sigma}_{t|t} = \frac{1}{n-H} \sum_{s=1}^{n-H} (Y_s - \widehat{Y}_{s|s})(Y_s - \widehat{Y}_{s|s})' . \quad (10)$$

Confidence bands for the OPP

We can then use the distributions (8) and (9) to approximate the distribution of $\delta_t^* = -(\mathcal{R}'\mathcal{R})^{-1}\mathcal{R}'\mathbb{E}_t Y_t^0$. We compute

$$\left\{ \delta_t^{(j)}, j = 1, \dots, B \right\} , \text{ where} \quad (11)$$

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

where B is the number of independent draws that we take equal to 10,000 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 the fact that $\widehat{\delta}_t$, the mean of the distribution $\{\delta_t^{(j)}, j = 1, \dots, B\}$, does not correspond to $(\widehat{\mathcal{R}}'\widehat{\mathcal{R}})^{-1}\widehat{\mathcal{R}}'\widehat{Y}_{t|t}$, a naive estimator of the OPP based on an OLS regression of the point forecast $\widehat{Y}_{t|t}$ on the IR point estimate $\widehat{\mathcal{R}}$.

Instead, we have

$$\widehat{\delta}_t = (\widehat{\mathcal{R}}'\widehat{\mathcal{R}} + \widehat{\Omega})^{-1}\widehat{\mathcal{R}}'\widehat{Y}_{t|t} , \quad (12)$$

where $\widehat{\Omega}$ 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 fulling minimizing the loss

function. A similar logic is at work in our context: a researcher that faces uncertainty in its estimate of the effects of policy (uncertainty in the 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.²⁵ 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.²⁶

Before describing these exercises, we first detail the datasets and methods that we use to

²⁵In 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.

²⁶This 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).

calculate the optimal perturbations. The OPP requires two pieces of information to assess 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 1980-2006.²⁷ 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.²⁸ 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

²⁷The 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.

²⁸We obtain very similar results using instead a convergence time of 10 years.

expectations from the Federal Reserve Board “PTR” variable, which is a measure a long-run 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.

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 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

$$y_{t+h} = x_t^k \mathcal{R}_{k,h}^y + c_t' \gamma_{k,h}^y + \nu_{t+h}^{k,y}, \quad (13)$$

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 (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

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

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 1 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 middle-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 .

Finally, the bottom panel 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 the bottom panel, 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.²⁹

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.

Since the Fed only moves in quarter percentage points, the shaded grey area depicts the

²⁹The 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.

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 1. 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.³⁰ 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.³¹

Note the magnitude of the optimization failure for an hypothetical inflation targeter. As shown in the middle-left panel, the OPP amounts to 1.25ppt a large economic deviation for the fed funds rate.

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

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

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-

³⁰The dashed red line 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).

³¹The 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) and the central bank's instrument "average" ability to transform unemployment into inflation $\kappa = \|\mathcal{R}^\pi\|^2 / \|\mathcal{R}^u\|$.³²

Thus, the dual-mandate OPP δ_t^* depends on two 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 (as captured by $\delta_t^{*\pi}$ or δ_t^{*u}), and (ii) a static trade-off depending on the relative ability of the central bank to influence one mandate versus another (as captured by ω).³³

Going back to our 1990-M6 example, taking the expected path of unemployment (top-right panel of Figure 1, blue colors) into consideration completely changes the picture. In fact, in this case, we can no longer discard optimality 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 coefficients $\delta_t^{*\pi}$ and δ_t^{*u} : it is desirable to run a more contractionary policy ($\delta_t^{*\pi} \approx +.25 > 0$, 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. And since the black line

³²Note that $\omega \xrightarrow{\lambda \rightarrow 0} 1$ (a single inflation mandate), and similarly $\omega \xrightarrow{\lambda \rightarrow \infty} 0$ (a single unemployment mandate). In the appendix, we show that κ reduces to the slope of the Phillips curve in the baseline New-Keynesian model (Galí, 2015)

³³The 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 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.

(with slope equal to the dual-mandate OPP) is an average between these red and blue lines, it ends up with a flat 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

$$\min_{i_{0t} \in \mathbb{R}} \mathcal{L}_t, \quad \mathcal{L}_t = \mathbb{E}_t \|\Pi_t\|^2 + \mathbb{E}_t \|U_t\|^2 + \lambda_B \mathbb{E}_t \|B_t\|^2, \quad (16)$$

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 2 plots the same set of plots as Figure 1 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 $-.75$ 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.³⁴ 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 ($\delta_t^{*(3)}$) is slightly lower than the dual-mandate OPP ($\delta_t^{*(2)}$). Lowering the fed funds rate by $\delta_t^{*(2)} = -.75$ ppt to stabilize inflation and unemployment would raise credit growth by too much in 2006 and overall increase the expected loss. The three-mandate OPP ($\delta_t^{*(3)} \approx -.25$ ppt) strikes a balance between these conflicting goals: it calls for lowering the fed funds rate to raise inflation and lower unemployment, while limiting 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.³⁵ An interesting question in hindsight is thus whether the 2008-M4 fed funds rate was optimal.

³⁴Echoing 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), by which time credit growth was already expected to be back to target. As a result, the overall reduction in the loss function (16) is relatively small, only 5 percent (middle-right panel). This point is an important determinant of the overall OPP with three mandates ($\delta_t^{*(3)}$). 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.

³⁵By 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.

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 (15), and Figure 3 has the same structure as our previous plots except that the top panels now also report the 67 percent confidence intervals calculated from past forecast errors over the previous 10 years.

While the median SEP forecast indicate a worsening outlook for the unemployment gap, note the large uncertainty surrounding that forecast. Thus, while casual inspections of the unemployment point forecast and the impulse response suggests that the policy was not optimal –the Fed should have lowered the Fed funds rate—,³⁶ the large uncertainty around that forecast could substantially weaken that conclusion.

The confidence bands around the OPP can precisely (and quantitatively) capture this tension. The yellow cone shows the 68 percent confidence interval for the OPP: if that cone does not overlap with the x-axis, we can conclude that there is a failure to optimize with 68 percent probability (i.e., given sampling uncertainty). In terms of magnitude, the median value for the OPP rounds to -0.5ppt (dash-black line), driven almost exclusively by the expected increase in unemployment.³⁷ 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, taking into account model mis-specification and the possibility that the median SEP forecast is not the conditional expectation can radically change this conclusion. Perhaps, the Fed model was so misspecified and thus so far off the conditional expectation $\mathbb{E}_t Y_t^0$, that we cannot reject optimality given the level of model uncertainty. The grey cone report the 68 percent central interval around the median OPP, taking into account both model and IR

³⁶In other words, the Fed should have acted earlier by cutting rates more aggressively to counter the coming crisis.

³⁷The 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. Interestingly, this case captures a common wisdom of central banking: central banks should “look through” transitory inflationary episodes.

uncertainty. In this case, model uncertainty is so large that we cannot discard optimality despite the expected increase in unemployment. In other words, we cannot discard optimality because there is a more than 33 percent chance that the conditional expectations $\mathbb{E}_t \Pi_t^0$ and $\mathbb{E}_t U_t^0$ are such that the policy optimally balances the expected paths of inflation and unemployment. 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 4 displays the usual set of plots but for 2010-M4 and for the slope instrument, so that the middle panels show the impulse 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.³⁸

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 5 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

³⁸A 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.

panels of Figure 5 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.³⁹

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 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).⁴⁰

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.⁴¹

³⁹We compute

$$\Delta\mathcal{L}_t(\hat{\delta}_t) = \mathcal{L}_t(\hat{\delta}_t) - \mathcal{L}_t(0) , \quad \text{with} \quad \mathcal{L}_t(\delta_t) = \mathbb{E}_t(Y_t^0 + \mathcal{R}_t\delta_t)'(Y_t^0 + \mathcal{R}_t\delta_t) \quad (17)$$

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 .

⁴⁰The 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.

⁴¹We re-iterate that our approach does not highlight which specific policies would have been able to induce such shift in the slope.

6 Conclusion

In this paper we propose a methodology for determining the distance to optimality of a given policy choice. The proposed optimal policy perturbation statistic determines whether the proposed policy choice was optimal and it highlights which trade-offs have been overlooked in settings where the underlying economic model is unknown or too complex to write down, as is often the case in practice.

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 static 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.⁴²

While we focused on monetary policy making to illustrate the OPP,⁴³ 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 deficit 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.⁴⁴ 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.⁴⁵ The OPP would provide a

⁴²These three concepts are the three central sources of disagreement among policy makers (See e.g., Orphanides, 2019, for a recent example).

⁴³In 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.

⁴⁴In 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. And in response to this trade-off, some commentators proposed a *cyclically-adjusted* 3 percent rule.

⁴⁵To some extent, a “forecast deficit targeting” approach is already followed by the EU commission, as

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.⁴⁶

Going beyond the OPP, we note that the OPP is simply the first iteration of a Newton-Raphson algorithm. A natural and important extension of the present paper is thus to explore the conditions under which a Newton algorithm could be used to find the optimal policy. Indeed, finding the optimal policy amounts to finding the zero of the derivative of the loss function. When the underlying model is unknown or very complex, the derivative of the loss function will typically (i) only be calculated approximately, and (ii) be very costly to compute, an exhaustive grid search strategy is difficult to implement.⁴⁷ In this context, a Newton algorithm could help find the optimal policy. While a researcher outside the central bank cannot iterate the algorithm beyond one step, because it cannot compute $f(p_t, X_t)$ for different policy choices, a researcher inside the central bank could iterate further and get substantially closer to the optimality policy with a few additional iterations.⁴⁸ This idea—replacing the current “incomplete grid search” approach to policy making with a much more efficient Newton-type algorithm is an important avenue for future research.

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.

⁴⁶The 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).

⁴⁷Recall that the construction of the forecast at central banks or finance ministries involves repeated iterations between models and judgement and lengthy discussions between/among policy makers and their supporting staff.

⁴⁸Under some regularity conditions on the underlying true model, the rate of convergence of the Newton algorithm is quadratic.

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, Benoœuré, Pierre Jacquet, and Jean Pisani-Ferry.** 2018. *Economic Policy: Theory and Practice. Second edition.* Oxford University Press.
- Blanchard, Olivier.** 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.
- Chow, Gregory C.** 1972. “Optimal Control of Linear Econometric Systems with Finite Time Horizon.” *International Economic Review*, 13(1): 16–25.
- Chow, Gregory C.** 1973. “Effect of Uncertainty on Optimal Control Policies.” *International Economic Review*, 14(3): 632–645.
- 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.
- Eberly, Janice C., James H. Stock, and Jonathan H. Wright.** 2019. “The Federal Reserve’s Current Framework for Monetary Policy: A Review and Assessment.” Working paper.
- Galí, Jordi.** 2015. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications.* Princeton University Press.
- 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.
- 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, L., and T. 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.
- Manganelli, Simone.** 2009. "Forecasting With Judgment." *Journal of Business & Economic Statistics*, 27(4): 553–563.
- Newey, Whitney K.** 2013. "Nonparametric Instrumental Variables Estimation." *The American Economic Review*, 103(3): 550–556.
- Newey, Whitney K., and Kenneth D. West.** 1994. "Automatic Lag Selection in Covariance Matrix Estimation." *The Review of Economic Studies*, 61(4): 631–653.
- Orphanides, A.** 2019. "Monetary Policy Strategy and its Communication." Working paper, presented at Jackson Hole Conference August 22, 2019.
- Orphanides, Athanasios.** 2003. "Monetary policy evaluation with noisy information." *Journal of Monetary Economics*, 50(3): 605 – 631.
- 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.
- Paul, Pascal.** 2019. "The Time-Varying Effect of Monetary Policy on Asset Prices." *The Review of Economics and Statistics*. Forthcoming.
- Primiceri, G. 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.
- Ramey, Valerie A., and Sarah Zubairy.** 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data." *Journal of Political Economy*, 126.
- Rudebusch, Glenn, and Lars EO 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.** 2018. “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments.” *The Economic Journal*, 128(610): 917–948.
- Svensson, Lars.** 2010. “Inflation Targeting, Handbook of Monetary Economics.” *Handbook of Monetary Economics*, edition, 1: 1237–1302.
- 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 EO.** 2017*a*. “Cost-benefit analysis of leaning against the wind.” *Journal of Monetary Economics*, 90: 193–213.
- Svensson, Lars EO.** 2017*b*. “What Rule for the Federal Reserve? Forecast Targeting.” National Bureau of Economic Research.
- Svensson, Lars EO.** 2019. “Monetary policy strategies for the Federal Reserve.” 4–5.
- Swanson, Eric T.** 2004. “Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rules.” *Macroeconomic Dynamics*, 8: 27–50.
- 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.** 2011. *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.

Appendix

The OPP at work in a simple macro model

In this section, we illustrate the working of the OPP in a simple but general economic model. This allows to provide the intuition behind the construction of the OPP and also to explicitly discuss the conditions underlying Assumption 1.

For ease of exposition, it is useful to consider a static model involving one macro variable y , one policy variable p and one exogenous variable x . All the important results are captured by this simple model. It is straightforward to generalize what follows to a multi-variate dynamic setting.

A general model can be written as

$$\begin{cases} y = f(p, x; \hat{g}(\cdot)) & \text{(M) curve} \\ p = g(y, x) + \varepsilon & \text{(P) curve} \end{cases} \quad (18)$$

The first equation is the “macro equation” —the (M) curve— that describes how the macro variable depends through the function $f(\cdot)$. For instance, when the policy instrument is the interest rate, this equation is the (IS) curve. The second equation is the policy equation —the (P) curve—. The policy choice p is determined by (i) $g(y, x)$, the reaction function of the policy maker, which captures the systematic response of the policy maker to economic developments, and by (ii) a policy shock ε that captures random fluctuations in policy due to taste shocks, mistakes in the policy process, or rounding (quarter ppt rounding for the fed funds rate for instance), among other possibilities. For instance, in the context of monetary policy the reaction function is often represented as a simple interest rate Taylor rule. In practice, the reaction function $g(\cdot)$ can be arbitrarily complicated and is unknown to the public,⁴⁹ and agents approximate that reaction function with some function $\hat{g}(\cdot)$. The function $f(\cdot)$ can depend on $\hat{g}(\cdot)$ —the approximate reaction function can influence the macro equation—, because agents may base their decisions on their expectation of the policy

⁴⁹As Svensson (2019) puts it in the context of monetary policy, “the reaction function, meaning how the policy choice responds to information available to the policy maker is far too complex to write down. The policy choice will respond to all relevant information that shifts the forecast of inflation and unemployment. This is a long and changing list.” Note that the reaction function is often too complex to write down, even for the policy maker. For instance, the reaction function of the Fed is the outcome of series of steps — model iteration/combination, lengthy discussions among staff members and judgement calls—, that cannot be easily written down as a Taylor-rule type of function. Instead, the Fed reaction function is better thought of as akin to a forecast targeting rule in the sense of Svensson (2019). As 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”.

maker's reaction function.

We denote by (y^0, p^0) the equilibrium allocation in this model. Our goal in this paper is to assess whether the policy choice

$$p^0 = g(y^0, x) + \varepsilon$$

is the optimal choice, defined (in this simpler case) as the policy setting p that minimizes $(y - y^*)^2$ where y^* is the policy maker's objective allocation.

To assess the optimality of p^0 , we consider a small perturbation δ to the policy choice p^0 , so that the optimization problem is

$$\min_{\delta} (f(g(y, x) + \varepsilon + \delta))^2.$$

Since f is unknown, we take a second-order approximation of the loss function for δ around 0 and solve instead the optimization problem

$$\min_{\delta} \left[(y^0)^2 + 2y^0 \mathcal{R}(0) + 2y^0 \left. \frac{d\mathcal{R}}{d\delta} \right|_{\delta=0} \delta^2 \right].$$

where $\mathcal{R} = \frac{dy}{d\delta}$. The first-order condition implies

$$\delta^* = - \left(\mathcal{R}(0)^2 + y^0 \left. \frac{d\mathcal{R}}{d\delta} \right|_{\delta=0} \right)^{-1} \mathcal{R}(0)y^0. \quad (19)$$

Expression (19) is the first-iteration of a Gauss-Newton algorithm, and the two requirements to compute δ^* are (i) y^0 , and (ii) the impulse response at $\delta = 0$; $\mathcal{R}(0)$. To calculate \mathcal{R} , we differentiate

$$y = f(g(y) + \varepsilon + \delta; \hat{g}(.))$$

with respect to δ and get

$$dy = f'_p(p^0) [g'_y(y^0)dy + d\hat{g}(.)]$$

which gives

$$\mathcal{R} = \frac{dy}{d\delta} = \frac{f'_p(p^0)}{1 - g'_y(y^0)f'_p(p^0)} + \frac{d\hat{g}(.)}{d\delta} \frac{1}{1 - g'_y(y^0)f'_p(p^0)}$$

or in more compact form

$$\mathcal{R} = \mathcal{R}(p^0, x, \delta)$$

so that the impulse response function (i) is state-dependent, depending on both the value of the exogenous variable x as well as value of the policy choice p^0 , and (ii) is not necessarily invariant to the discretionary policy adjustment δ .

Assumption 1 in the main text imposes that \mathcal{R} is fixed, which allows us first to estimate \mathcal{R} from policy shocks and simple VAR-IV or LP-IV estimation methods,⁵⁰ and second to interpret the OPP $\delta^* = -(\mathcal{R}^2)^{-1} \mathcal{R}y^0$ as a distance-to-optimality measure.

Assumption 1 excludes the possibilities of (i) and (ii).

Excluding (i) requires that non-linearities in the underlying model are small. This is not a conceptual limitation for the OPP and our approach in general, because the interpretation of the OPP as a distance to optimality measure remains correct even when $\mathcal{R} = \mathcal{R}(p^0, x)$.⁵¹ Allowing for (i) imposes however stronger requirements in terms of the data as more variation is needed to estimate $\mathcal{R} = \mathcal{R}(p^0, x)$. Provided there is enough variation however, we can draw on recent advances in macro-econometrics to estimate such state-dependent impulse responses (e.g. Tenreyro and Thwaites, 2016; Ramey and Zubairy, 2018).⁵²

Excluding (ii), invariance to discretionary policy perturbations, is an important restriction, because $\frac{d\mathcal{R}}{d\delta}|_{\delta=0}$ enters the formula for δ^* . As a result, when $\frac{d\mathcal{R}}{d\delta}|_{\delta=0} \neq 0$, δ^* no longer captures the magnitude of the deviation from optimality.⁵³ Intuitively, (ii) corresponds to the Lucas critique: a discretionary change in policy through δ can lead agents to revise their representation of the reaction function and induce a regime shift — a change in the macro equation $f(\cdot)$. In our context, excluding (ii) is reasonable however, as long as we are considering a small perturbation to p^0 , i.e., as long as the deviation from optimality is not too large. Intuitively, if the deviation from optimality is not substantially larger than a typical policy shock, a perturbation δ_t to p_t will not lead agents to revise $\hat{g}(\cdot)$. Thus, as long as δ^* remains modest in the sense of Leeper and Zha, the optimal policy perturbation will not lead to a regime switch and assuming \mathcal{R} independent of δ is reasonable.⁵⁴

⁵⁰While \mathcal{R} captures the response to a policy perturbation δ instead of a policy shock ε , the corresponding impulse response is the same since ε and δ enter $f(\cdot)$ linearly (such that $\frac{df}{d\delta} = \frac{df}{d\varepsilon}$).

⁵¹Indeed, as long as \mathcal{R} does not depend on δ , the second-order approximation of the loss function is exact and the Gauss-Newton algorithm converges in one step, i.e., δ^* is the distance between p^0 and the optimal policy choice p^* .

⁵²An additional source of variation for \mathcal{R} is time-variation in the underlying structure of the economy (the function f is time-dependent and becomes f_t) or in the policy maker's reaction function (the function g is time-dependent and becomes g_t), due to changes in the policy maker's preference, view about the underlying economic model or decision making process. The impulse response becomes \mathcal{R}_t . Similarly to state-dependence in the impulse response, this case can be addressed with the appropriate methodology, such as time-varying parameter VAR or time-varying LP (e.g., Primiceri, 2005; Paul, 2019).

⁵³With $\mathcal{R}(\delta)$, the second-order expansion of the optimization problem is only an approximation, and the Gauss-Newton algorithm takes more than one iteration to converge to the optimal policy.

⁵⁴Provided that the different regimes have been visited in the past, it is possible in theory to estimate how the impulse response function would change with the policy perturbation. To give one example in the context of monetary policy, if there exists two regimes —unanchored vs. anchored inflation expectations— and if the economy visited these two regimes in the past, one could model and estimate $\mathcal{R}(\delta)$. This requires large data sample however.

Proofs

Proof of Theorem 1. Consider the loss function $\mathbb{E}_t\|Y_t\|^2 = \mathbb{E}_t\|f(p_t^0 + \delta_t, X_t)\|^2$. Under assumption 1 the Taylor series around $\delta_t = 0$ is given by

$$\mathbb{E}_t\|Y_t\|^2 = \mathbb{E}_t\|Y_t^0\|^2 + 2\mathbb{E}_t Y_t^{0'} \mathcal{R} \delta_t + 2\delta_t' \mathcal{R}' \mathcal{R} \delta_t$$

as $\partial \mathbb{E}_t\|Y_t\|^2 / \partial \delta_t|_{\delta_t=0} = 2\mathcal{R}' \mathbb{E}_t Y_t^0$ and $\partial^2 \mathbb{E}_t\|Y_t\|^2 / \partial \delta_t \partial \delta_t'|_{\delta_t=0} = 2\mathcal{R}' \mathcal{R}$. The first derivative with respect to δ_t is given by

$$\frac{\partial \mathbb{E}_t\|Y_t\|^2}{\partial \delta_t} = 2\mathcal{R}' \mathbb{E}_t Y_t^0 + 2\mathcal{R}' \mathcal{R} \delta_t = 0$$

Solving using $\mathcal{R}' \mathcal{R} \succ 0$ implies the result. \square

Proof of Theorem 2. Consider the loss function $\mathbb{E}_t\|Y_t\|^2 = \mathbb{E}_t\|f(p_t^0 + \delta_t, X_t)\|^2$. Under assumption 2 the second order Taylor approximation around $\delta_t = 0$ is given by

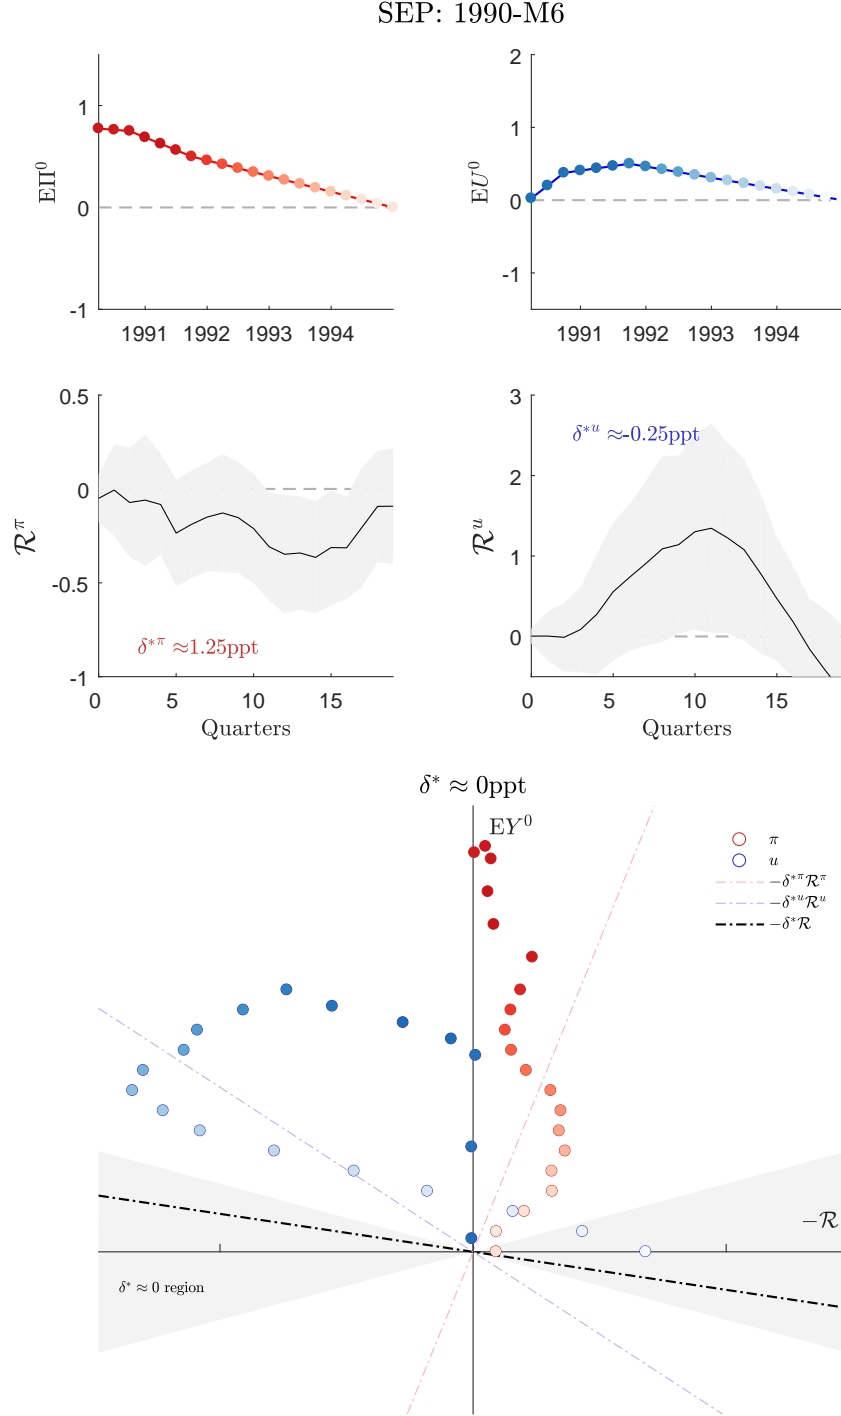
$$\mathbb{E}_t\|Y_t\|^2 \approx \mathbb{E}_t\|Y_t^0\|^2 + 2\mathbb{E}_t Y_t^{0'} \mathcal{R}(0) \delta_t + 2\delta_t' \left(\mathcal{R}(0)' \mathcal{R}(0) + (\mathbb{E}_t Y_t^{0'} \otimes I_{M(H+1)}) \frac{\partial \text{vec}(\mathcal{R}(\delta_t'))}{\partial \delta_t'} \bigg|_{\delta_t=0} \right) \delta_t$$

as $\partial \mathbb{E}_t\|Y_t\|^2 / \partial \delta_t = 2\mathcal{R}(\delta_t)' \mathbb{E}_t Y_t = 2(\mathbb{E}_t Y_t' \otimes I_{M(H+1)}) \text{vec}(\mathcal{R}(\delta_t'))$ and $\partial^2 \mathbb{E}_t\|Y_t\|^2 / \partial \delta_t \partial \delta_t' = 2(\mathcal{R}(\delta_t)' \otimes I_{M(H+1)}) \text{vec}(\mathcal{R}(\delta_t')) + 2(\mathbb{E}_t Y_t' \otimes I_{M(H+1)}) \frac{\partial \text{vec}(\mathcal{R}(\delta_t'))}{\partial \delta_t'}$. The first derivative with respect to δ_t is given by

$$2\mathcal{R}(0)' \mathbb{E}_t Y_t^0 + 2 \left(\mathcal{R}(0)' \mathcal{R}(0) + (\mathbb{E}_t Y_t^{0'} \otimes I_{M(H+1)}) \frac{\partial \text{vec}(\mathcal{R}(\delta_t'))}{\partial \delta_t'} \bigg|_{\delta_t=0} \right) \delta_t = 0$$

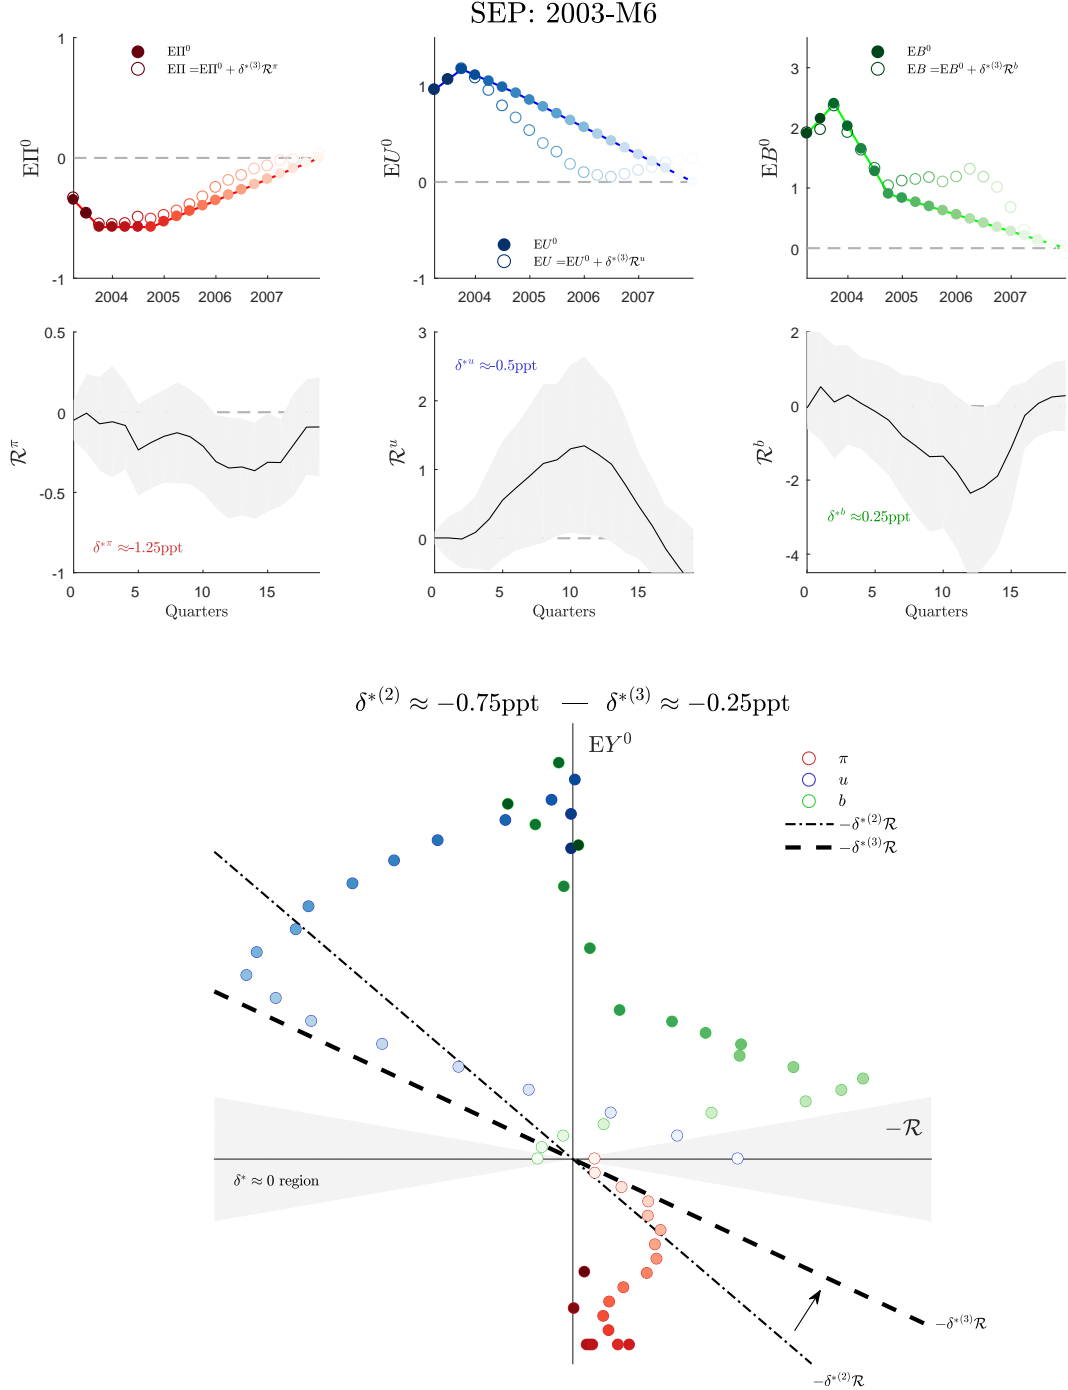
Solving using $\mathcal{R}' \mathcal{R} \succ 0$ implies the result. \square

Figure 1: A static inflation-unemployment trade-off



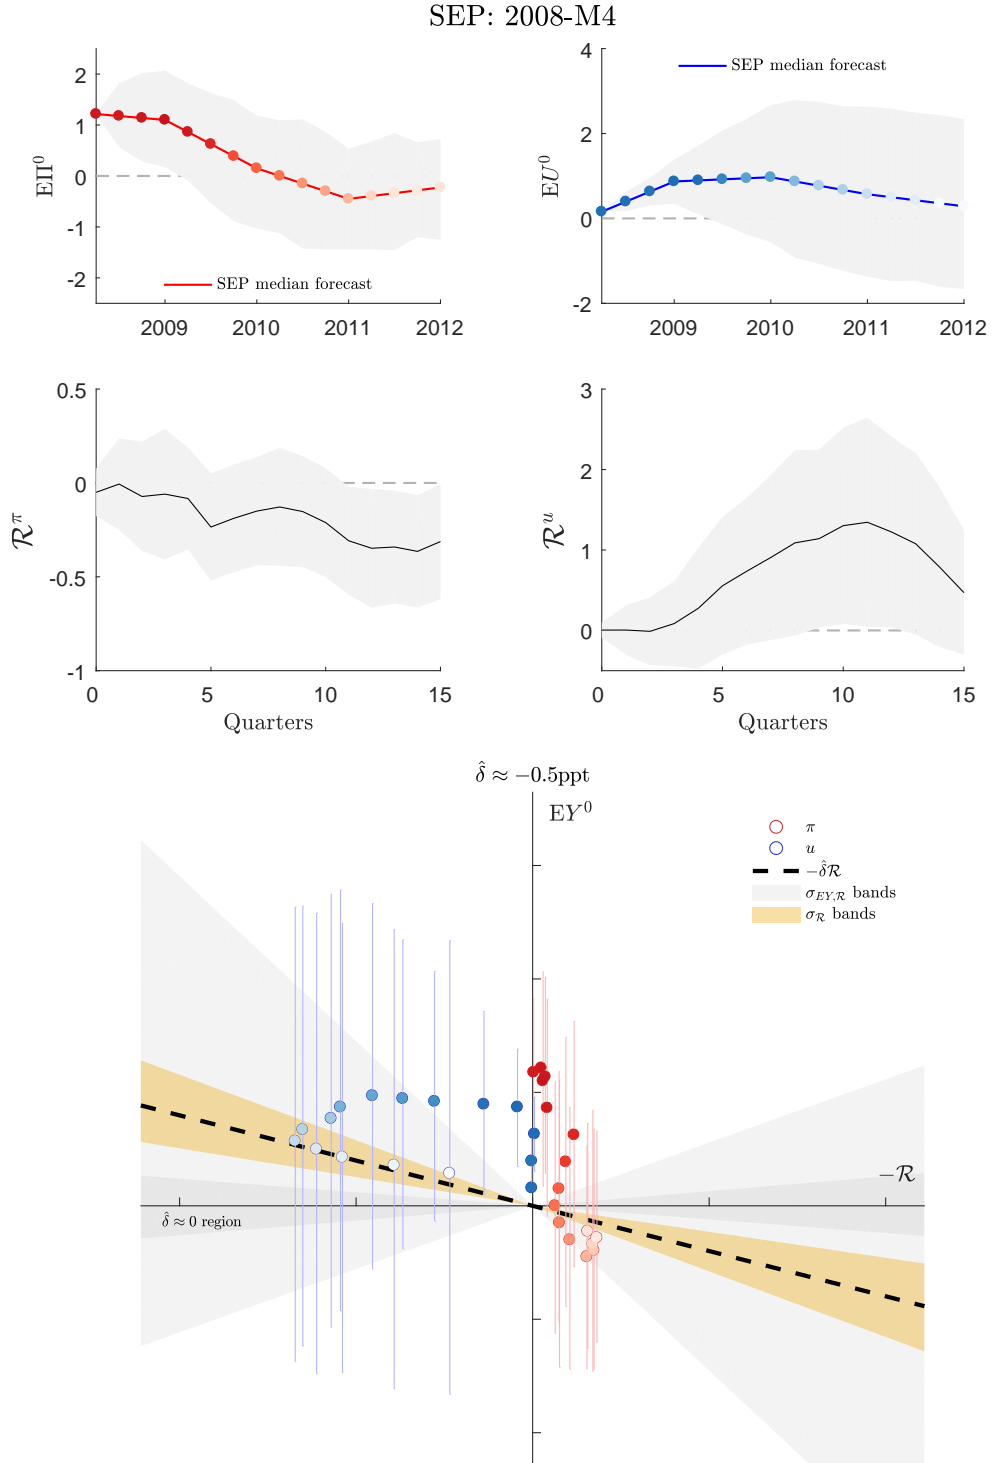
Notes: Top panel: median FOMC forecasts for the inflation and unemployment gaps. Middle panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock. In red (blue) is the OPP $\delta^{*\pi}$ (δ^{*u}) for a strict inflation (unemployment) targeter. Bottom panel: in red, a scatter plot of the median FOMC inflation gap forecast ($E\Pi^0$) 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 2: Dynamic trade-offs and leaning against the wind




Notes: Top panel: median SEP forecasts for the inflation, unemployment and debt growth gaps. 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 3: Taking uncertainty into account: fed funds rate policy



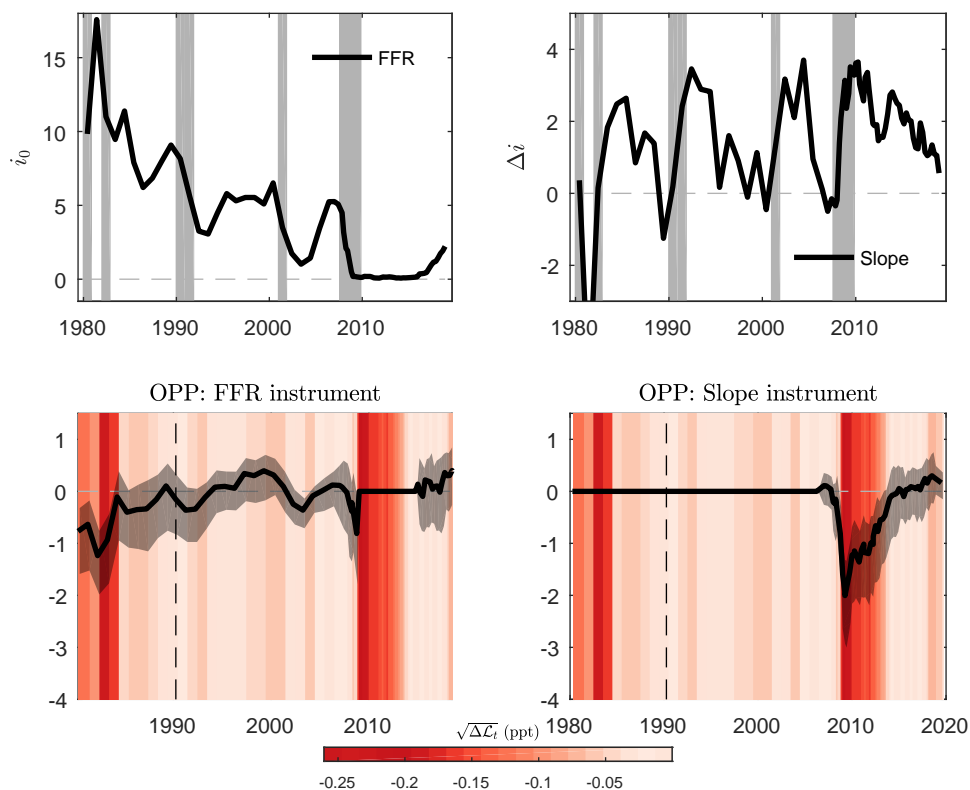
Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps (in red and blue) along with the 67 percent confidence bands. Middle panel: impulse responses of the inflation and unemployment gaps to a fed funds rate shock with the 95 percent confidence intervals. Bottom panel: scatter 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 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 4: Taking uncertainty into account: slope policy



Notes: Top panel: Median SEP forecasts for the inflation and unemployment gaps (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 5: 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.