

Web-appendix for:

A SUFFICIENT STATISTICS APPROACH FOR
MACRO POLICY EVALUATION

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

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

^(b) Universitat Pompeu Fabra and BSE

March 28, 2023

Abstract

In this web-appendix, we provide the following additional results:

- S0: Details: OPP implementation
- S1: Example: NK model with commitment
- S2: General nonlinear OPP framework
- S3: General convex loss functions
- S4: Details: Constrained OPP
- S5: Estimating robust preference parameters
- S6: Additional results for the empirical study

References to lemmas, equations, etc..., which start with a “S” are references to this document.

References, which consist of only a number refer to the main text.

S0 OPP implementation

Section 5 in the main text sketches how to estimate the sufficient statistics in practice, and then how to implement the OPP policy evaluation framework. This section provides more details on the construction of the two sufficient statistics and the OPP.

We will pay particular attention to the forecast construction step, which depends on the specific perspective and the question at hand. In the main text, we took the perspective that (ϕ^0, ϵ_t^0) was the policy maker's choice and $\mathbb{E}_t \mathbf{P}_t^0$ was known. Here we discuss other settings that can be of interest. The organization is the same as in the main text; we discuss forecasting, impulse response estimation, uncertainty assessment and OPP construction. In addition, we will explicitly discuss how to apply the OPP for policy recommendation (or improvement).

S0.1 Constructing the baseline forecasts

We start with providing more details on how to construct approximation of $\mathbb{E}_t \mathbf{Y}_t^0$ and $\mathbb{E}_t \mathbf{P}_t^0$: the expected paths implied by the baseline policy (ϕ^0, ϵ_t^0) . These oracle forecasts are defined in Lemma 1 in terms of $\mathbf{S}_t = (\mathbf{X}'_{-t}, \boldsymbol{\Xi}'_t)'$, which includes the initial conditions \mathbf{X}_{-t} and the non-policy news shocks $\boldsymbol{\Xi}_t$, and the policy news shocks ϵ_t^0 . Restated for convenience:

$$\begin{aligned}\mathbb{E}_t \mathbf{Y}_t^0 &= \Gamma_y^0 \mathbf{S}_t + \mathcal{R}_y^0 \epsilon_t^0 \\ \mathbb{E}_t \mathbf{P}_t^0 &= \Gamma_p^0 \mathbf{S}_t + \mathcal{R}_p^0 \epsilon_t^0\end{aligned}\tag{S1}$$

We will distinguish between three cases: (i) the researcher can directly download both $\mathbb{E}_t \mathbf{P}_t^0$ and $\mathbb{E}_t \mathbf{Y}_t^0$ (ii) $\mathbb{E}_t \mathbf{P}_t^0$ is known but $\mathbb{E}_t \mathbf{Y}_t^0$ must be approximated or (iii) both $\mathbb{E}_t \mathbf{P}_t^0$ and $\mathbb{E}_t \mathbf{Y}_t^0$ need to be approximated.

S0.1.1 Downloading forecasts

The simplest yet not always feasible way in which a researcher can obtain a baseline forecast is to use the forecasts that are provided by the policy maker. Indeed, several macro policy makers make their forecasts for the policy objectives publicly available and these can then be directly used to compute the OPP and evaluate/improve the expected policy path $\mathbb{E}_t \mathbf{P}_t^0$ of the policy maker. For instance, in our empirical work we use the inflation and unemployment forecasts of the Fed to evaluate the Fed's monetary policy decisions. Alternatively, the researcher could use professional forecasts, such as those from the Survey of Professional Forecasters (SPF) or the Blue Chip forecasts.

S0.1.2 Approximating $\mathbb{E}_t \mathbf{Y}_t^0$ with given $\mathbb{E}_t \mathbf{P}_t^0$

In the case where the evaluation of policy is done by the policy maker, or the policy maker's staff, we can consider that the expected path $\mathbb{E}_t \mathbf{P}_t^0$ is known. This is the case that we discussed in the main text.

Since the state \mathbf{S}_t is generally not observable, we denote by \mathbf{Z}_t a (possibly large) vector of observable variables that can be used to approximate \mathbf{S}_t . The best linear prediction for \mathbf{Y}_t^0 in terms of $\mathbb{E}_t \mathbf{P}_t^0$ and \mathbf{Z}_t can be recovered from

$$\mathbf{Y}_t^0 = \mathbf{B}_{yz}^0 \mathbf{Z}_t + \mathbf{B}_{yp}^0 \mathbb{E}_t \mathbf{P}_t^0 + \mathbf{U}_t^y, \quad (\text{S2})$$

where the maps \mathbf{B}_{yz}^0 and \mathbf{B}_{yp}^0 include the best linear prediction coefficients and the error \mathbf{U}_t^y is orthogonal to $(\mathbf{Z}_t, \mathbb{E}_t \mathbf{P}_t^0)$ by construction.

The model contains an equation for each vector of objectives y_{t+h} , for $h = 0, 1, \dots$. In practice, the forecast will be truncated at some maximum horizon H and some in sample period will be used to estimate the coefficients. This implies that the forecasting model that we consider in practice is given by

$$\mathbf{Y}_{s:s+H}^0 = \mathbf{B}_{yz}^{0,t} \mathbf{Z}_s + \mathbf{B}_{yp}^{0,t} \mathbb{E}_s \mathbf{P}_{s:s+H}^0 + \mathbf{U}_{s:s+H}^y \quad \text{for } s = t_0, \dots, t, \quad (\text{S3})$$

where we have truncated the paths at horizon H , so that $\mathbf{Y}_{s:s+H}^0 \equiv (y'_t, y'_{t+1}, \dots, y'_{t+H})'$.

Based on some prior sampling period —over which the policy rule ϕ^0 was implemented— we can estimate the model parameters $\mathbf{B}_{yz}^{0,t}$ and $\mathbf{B}_{yp}^{0,t}$. These parameters can be structured (e.g. VAR, DFM, DSGE, etc) and they can be estimated using standard regression methods, possibly allowing for shrinkage and penalization to improve the model fit. The estimated model parameters are denoted by $\widehat{\mathbf{B}}_{yz}^{0,t}$ and $\widehat{\mathbf{B}}_{yp}^{0,t}$.

The resulting forecasts for time period t are given by

$$\widehat{\mathbf{Y}}_{t:t+H}^0 = \widehat{\mathbf{B}}_{yz}^{0,t} \mathbf{Z}_t + \widehat{\mathbf{B}}_{yp}^{0,t} \mathbb{E}_t \mathbf{P}_{t:t+H}^0, \quad (\text{S4})$$

and the model uncertainty error is given by

$$\mathcal{U}_{t:t+H}^y = \widehat{\mathbf{Y}}_{t:t+H}^0 - \mathbb{E}_t \mathbf{Y}_{t:t+H}^0.$$

Combining (S1) and (S4) we see that the difference between the oracle forecast and the forecast of the policy maker comes from (i) imperfectly measuring \mathbf{S}_t and (ii) estimating the parameters.

Example S1 (Dynamic factor model). *Let \mathbf{X}_s denote a large panel of disaggregated macroeconomic and financial time series for time period s , with $s < t$. Following Stock and Watson*

(2002b), and many others, we can approximate the information in the high dimensional vector \mathbf{X}_s by a small number of common factors.

$$\mathbf{X}_s = \mathbf{\Lambda}f_s + \zeta_s ,$$

where the dimension of the common factors f_s is typically small. The factors can be recovered by principal components (e.g. Bai and Ng, 2002; Stock and Watson, 2002a) or by state space methods (e.g. Durbin and Koopman, 2012). With the estimated common factors in hand we can set $\mathbf{Z}_s = \hat{f}_s$ and estimate the regression coefficients in (S3) and construct the forecast (S4).

Example S2 (Vector Autoregressive model). Another popular way of constructing forecasts is by using vector autoregressive models, most notably large scale Bayesian VARs have shown good forecasting performance in macro settings (e.g. Banbura, Giannone and Reichlin, 2010). In our context, we can consider such models as well by setting $\mathbf{X}_s = (y_s^0, \mathbb{E}_s \mathbf{P}_{s:s+H}^0, w_s^0)'$, where y_s^0 are the policy objectives at time s , $\mathbb{E}_s \mathbf{P}_{s:s+H}^0$ is the expected policy path at time s and w_s^0 include other possible predictors. The VAR is then given by

$$\mathbf{X}_s = \mathbf{\Phi}_1 \mathbf{X}_{s-1} + \dots + \mathbf{\Phi}_p \mathbf{X}_{s-p} + \eta_s . \quad (\text{S5})$$

An important difference with a standard VAR is that we jointly model the entire path $\mathbb{E}_s \mathbf{P}_{s:s+H}^0$ together with the other observable variables. This is important as we want to allow the expected policy path, say $\mathbb{E}_s p_{s+h}^0$ to influence observations prior to $s+h$ directly.

To make this clear, consider a standard VAR, where the VAR consists of $\mathbf{X}_s = (y_s', p_s', w_s')'$ and to construct a forecast at time t we feed in the expected path $\mathbb{E}_t \mathbf{P}_t^0 = \mathbb{E}_t (p_t^0, p_{t+1}^0 \dots)'$.¹ Since the standard VAR does not allow the expected future policy path $\mathbb{E}_t p_{t+h}^0$ to affect contemporaneous values, it can only fit the expected path $\mathbb{E}_t \mathbf{P}_t^0$ by introducing future shocks in the forecast, and thus cannot provide a valid approximation for $\mathbb{E}_t \mathbf{Y}_t^0$.

Based on (S5) we can iterate the VAR forward to obtain expressions for y_s^0, y_{s+1}^0, \dots . These can be stacked to obtain a representation like (S3). Moreover, based on (S5) we can estimate the model parameters using OLS or more advanced shrinkage methods to ultimately compute the forecast (S4).

S0.1.3 Approximating $\mathbb{E}_t \mathbf{Y}_t^0$ and $\mathbb{E}_t \mathbf{P}_t^0$

When neither $\mathbb{E}_t \mathbf{Y}_t^0$ nor $\mathbb{E}_t \mathbf{P}_t^0$ is available, we can follow the same recipe as in the main text, with the distinction that we need to use the observable variables \mathbf{Z}_t to also infer the expected

¹In the VAR literature this route is often referred to as conditional forecasting, see Waggoner and Zha (1999) and Banbura, Giannone and Lenza (2015) for details.

policy path $\mathbb{E}_t \mathbf{P}_t^0$. The consequences of not knowing the expected policy path need to be taken into account. Specifically, policy evaluation here implies evaluating the policy that is currently implemented, as inferred by the researcher.

Starting from (S1) the researcher seeks variables \mathbf{Z}_t that best span \mathbf{S}_t —the state of the economy—and $\boldsymbol{\epsilon}_t^0$ —the policy news shocks. For any given approximating vector \mathbf{Z}_t we can write the forecasting models for approximating $\mathbb{E}_t \mathbf{P}_t^0$ and $\mathbb{E}_t \mathbf{Y}_t^0$ as

$$\mathbf{P}_t^0 = \mathbf{B}_{pz}^0 \mathbf{Z}_t + \mathbf{U}_t^p \quad \text{and} \quad \mathbf{Y}_t^0 = \mathbf{B}_{yz}^0 \mathbf{Z}_t + \mathbf{U}_t^y, \quad (\text{S6})$$

where the coefficients \mathbf{B}_{pz}^0 and \mathbf{B}_{yz}^0 are the best linear prediction coefficients that explain \mathbf{P}_t^0 and \mathbf{Y}_t^0 in terms of \mathbf{Z}_t . The error terms include the future errors as well as the approximation errors from not knowing the exact state of the economy and the policy news shocks.

It is instructive to compare (S6) to the forecasting model used when $\mathbb{E}_t \mathbf{P}_t^0$ known, i.e. equation (S2). The only difference is that $\mathbb{E}_t \mathbf{P}_t^0$ is now omitted and the researcher will end up evaluating the inferred $\mathbb{E}_t \mathbf{P}_t^0$ based on the model for \mathbf{P}_t^0 in (S6).

In practice, the forecast will be truncated and some in sample period will be used to estimate the coefficients. This implies that the forecasting models that we consider in practice are given by

$$\mathbf{P}_{s:s+H}^0 = \mathbf{B}_{pz}^{0,t} \mathbf{Z}_s + \mathbf{U}_{s:s+H}^p \quad \text{and} \quad \mathbf{Y}_{s:s+H}^0 = \mathbf{B}_{yz}^{0,t} \mathbf{Z}_s + \mathbf{U}_{s:s+H}^y. \quad (\text{S7})$$

Based on some prior sampling period, least squares methods can be used to estimate the model parameters $\mathbf{B}_{pz}^{0,t}$ and $\mathbf{B}_{yz}^{0,t}$. The estimated model parameters are denoted by $\widehat{\mathbf{B}}_{yz}^{0,t}$ and $\widehat{\mathbf{B}}_{yp}^{0,t}$.

The resulting forecasts for time period t are given by

$$\mathbf{P}_{t:t+H}^0 = \widehat{\mathbf{B}}_{pz}^{0,t} \mathbf{Z}_t \quad \text{and} \quad \mathbf{Y}_{t:t+H}^0 = \widehat{\mathbf{B}}_{yp}^{0,t} \mathbf{Z}_t, \quad (\text{S8})$$

and the model uncertainty errors are given by

$$\mathcal{U}_{t:t+H}^p = \widehat{\mathbf{P}}_{t:t+H}^0 - \mathbb{E}_t \mathbf{P}_{t:t+H}^0 \quad \text{and} \quad \mathcal{U}_{t:t+H}^y = \widehat{\mathbf{Y}}_{t:t+H}^0 - \mathbb{E}_t \mathbf{Y}_{t:t+H}^0.$$

The dynamic factor and VAR examples above apply easily to this setting.

S0.2 Impulse response estimation

To estimate \mathcal{R}_a^0 we can rely on the large macro-econometric literature that discusses the estimation of impulse responses to policy shocks. Starting from Lemma 1 we add and

subtract the realized series from $\mathbb{E}_t \mathbf{Y}_t^0$ and $\mathbb{E}_t \mathbf{P}_t^0$, to get

$$\mathbf{Y}_t^0 = \mathcal{R}_{a,y}^0 \boldsymbol{\epsilon}_{a,t} + \mathbf{V}_{a,t}^y \quad \text{and} \quad \mathbf{P}_t^0 = \mathcal{R}_{a,p}^0 \boldsymbol{\epsilon}_{a,t} + \mathbf{V}_{a,t}^p, \quad (\text{S9})$$

where $\mathbf{Y}_t^0 = (y_t^0, y_{t+1}^0, \dots)'$ and $\mathbf{P}_t^0 = (p_t^0, p_{t+1}^0, \dots)'$ are the observed policy objectives and policy instruments, which pertain to the policy regime ϕ^0 . The error terms $\mathbf{V}_{a,t}^y$ and $\mathbf{V}_{a,t}^p$ include all other time t structural shocks $\boldsymbol{\epsilon}_{a^\perp,t}$ and $\boldsymbol{\Xi}_t$, the initial conditions \mathbf{X}_{-t} and the future shocks $\mathbf{Y}_t^0 - \mathbb{E}_t \mathbf{Y}_t^0$ or $\mathbf{P}_t^0 - \mathbb{E}_t \mathbf{P}_t^0$, respectively. By construction these errors are orthogonal to the identifiable policy news shocks $\boldsymbol{\epsilon}_{a,t}$.

Equations (S9) can be viewed as local projections (Jordà, 2005) after realizing that $\boldsymbol{\epsilon}_{a,t}$ are time t measurable news shocks. In practice, we only estimate local projections up to some finite horizon H . The implicit assumption being that the impulse responses are indistinguishable from zero after this horizon.² Specifically, let $\mathbf{Y}_{t:t+H}^0 = (y_t^0, y_{t+1}^0, \dots, y_{t+H}^0)'$ and $\mathbf{P}_{t:t+H}^0 = (p_t^0, p_{t+1}^0, \dots, p_{t+H}^0)'$. Since, the error terms of the local projections in (S9) depend on the same underlying shocks they are correlated and we will estimate them jointly. Let $\mathcal{Y}_{t:t+H}^0 = (\mathbf{Y}_{t:t+H}^0, \mathbf{P}_{t:t+H}^0)'$ and consider

$$\mathcal{Y}_{t:t+H}^0 = R_{a,H}^0 \boldsymbol{\epsilon}_{a,t} + \mathbf{V}_{a,t:t+H},$$

where $R_{a,H}^0$ stacks the corresponding blocks of $\mathcal{R}_{a,y}^0$ and $\mathcal{R}_{a,p}^0$. Note that $\boldsymbol{\epsilon}_{a,t}$ is the subset of policy news shocks that can be identified and we assume that its finite dimensional with dimension K_a .

In general, to estimate \mathcal{R}_a^0 , we will assume that we have available the some prior sampling period over which the rule ϕ^0 was implemented. Let n denote the number of time periods. Ideally, we should not use observations past time period t to accurately reflect the information available at time t . That said, in practice it may be necessary to let go of this requirement as otherwise new policies cannot be evaluated in hindsight. Yet we stress that in such cases one must acknowledge that the policy maker at time t could not have computed the causal effect of the possible policy action.

Finally, in practice the policy news shocks are not observed directly, but some identifying strategy is required. Prominent examples include using zero-, long-run, or inequality restrictions (e.g. Sims, 1980; Blanchard and Quah, 1989; Faust, 1998; Uhlig, 2005), or by using past exogenous variations as instrumental variables (e.g. Mertens and Ravn, 2013; Stock and Watson, 2018), see Ramey (2016) for a detailed review. It is important to stress that our sufficient statistics approach can be implemented using any of these methods. In addition, it can be attractive to impose some smoothness restrictions on the local projections, either

²Alternatively, it could be that the loss function only places weight on the objectives corresponding to the first H periods.

directly as in Barnichon and Brownlees (2018), or by rewriting the LP in its vector autoregressive form (e.g. Plagborg-Møller and Wolf, 2021). See Li, Plagborg-Møller and Wolf (2022) for a recent comparison of impulse response estimation methods.

Regardless of the specific approach chosen it is important to stress that each approach will require specifying a specific reduced form model, i.e. control variables, lag lengths, and so on need to be specified. In addition, regularity conditions in the form of specific moment and dependence conditions are needed to ensure that the distribution of the estimators can be approximated reliably. Finally, the identification assumptions need to hold.

S0.3 Uncertainty assessment

To approximate the distribution of the OPP statistic and the adjustment we need to approximate the distribution of the impulse response estimates $\mathcal{U}_t^{\mathcal{R}^a} = \mathcal{R}_a^0 - \widehat{\mathcal{R}}_a^0$ and the distribution of model uncertainty: $\mathcal{U}_t^y = \mathbb{E}_t \mathbf{Y}_t^0 - \widehat{\mathbf{Y}}_t^0$ and possibly $\mathcal{U}_t^p = \mathbb{E}_t \mathbf{P}_t^0 - \widehat{\mathbf{P}}_t^0$. Recall that model misspecification arises from not being able to perfectly approximate the state of the economy and parameter estimation uncertainty.

For convenience we distinguish between two scenarios. First, consider the case where we use a single reduced form model to construct both the impulse responses and the forecasts. In this case conventional asymptotic theory or Bayesian methods can be used to obtain an estimate of the joint distribution of $(\mathcal{U}_t^{\mathcal{R}^a}, \mathcal{U}_t^y, \mathcal{U}_t^p)$. We denote the approximated distribution by \widehat{F} .

Second, consider the case where we compute subset OPPs using external forecasts, such as those obtained from the policy maker, or from other external sources. In such cases we will typically have to make the additional assumption that $\mathcal{U}_t^{\mathcal{R}^a}$ is independent of $(\mathcal{U}_t^y, \mathcal{U}_t^p)$ as we will not have any method for recovering the joint distribution. The distribution of the impulse response estimates can be obtained using standard methods. For the distribution of model uncertainty, it is sometimes also provided along with the published point forecasts. Alternatively, we can approximate it with the distribution of the historical forecast errors $\{\mathbf{Y}_s^0 - \widehat{\mathbf{Y}}_s^0, \mathbf{P}_s^0 - \widehat{\mathbf{P}}_s^0\}_{s=t^0}^t$ over some prior sampling period, along with a normality assumption. That said, this approach will give an *upper-bound* of the variance of model uncertainty.³

S0.4 Constructing OPPs

After obtaining the approximating distribution of the dynamic causal effects and the oracle forecasts, we can compute the distribution of any of the OPP statistics using simulation

³The variance of the forecast errors will upper-bound the variance of model uncertainty, because forecast errors mix two sources of uncertainty: (i) model uncertainty *and* (ii) future uncertainty.

methods for a given preference matrix \mathcal{W} . Specifically, we simulate dynamic causal effects and forecast misspecification errors from \widehat{F} and compute

$$\delta_{a,t}^j = -(\mathcal{R}_{a,y}^{j'} \mathcal{W} \mathcal{R}_{a,y}^j)^{-1} \mathcal{R}_{a,y}^{j'} \mathcal{W} \widehat{\mathbf{Y}}_t^j$$

where $\mathcal{R}_{a,y}^j = \widehat{\mathcal{R}}_{a,y}^0 + \mathcal{U}_t^{\mathcal{R}_{a,y},j}$ and $\widehat{\mathbf{Y}}_t^j = \widehat{\mathbf{Y}}_t^0 + \mathcal{U}_t^{y,j}$ with $\mathcal{U}_t^j \sim \widehat{F}$. Similarly for the constrained subset OPP statistic we compute for each draw

$$\begin{aligned} \delta_{a,t}^{c,j} &= \underset{\delta_{a,t}}{\operatorname{argmin}} (\widehat{\mathbf{Y}}_t^j + \widehat{\mathcal{R}}_{a,y}^j \delta_{a,t})' \mathcal{W} (\widehat{\mathbf{Y}}_t^j + \widehat{\mathcal{R}}_{a,y}^j \delta_{a,t}) \\ \text{s.t. } & C(\widehat{\mathbf{Y}}_t^j + \mathcal{R}_{y,a}^j \delta_{a,t}, \widehat{\mathbf{P}}_t^j + \mathcal{R}_{p,a}^j \delta_{a,t}) \geq \mathbf{c}, \end{aligned} \quad (\text{S10})$$

where $\widehat{\mathbf{P}}_t^j = \widehat{\mathbf{P}}_t^0 + \mathcal{U}_t^{p,j}$. Note that in the case where $\mathbb{E}_t \mathbf{P}_t^0$ is known we replace $\widehat{\mathbf{P}}_t^j$ by the known path $\mathbb{E}_t \mathbf{P}_t^0$.

For each statistic we sample for a large number of draws $j = 1, \dots, S_d$, and report the average statistics and the confidence interval for some $\alpha \in (0, 1)$. For instance for the subset OPP we have

$$\widehat{\delta}_{a,t} = \frac{1}{S_d} \sum_{j=1}^{S_d} \delta_{a,t}^j \quad \text{and} \quad \left[\delta_{a,t}^{(\alpha S_d)}, \delta_{a,t}^{((1-\alpha) S_d)} \right], \quad (\text{S11})$$

where $\delta_{a,t}^{(k)}$ denotes the (element wise) k th largest draw of $\{\delta_{a,t}^j, j = 1, \dots, S_d\}$. The same can be done for the constrained subset statistic. We will conclude that a policy $\mathbb{E}_t \mathbf{P}_t^0$ is not optimal whenever the confidence bands of $\delta_{a,t}^*$ or $\delta_{a,t}^{c*}$ exclude zero at any desired level of confidence.

With the distribution of the subset OPP in hand, we can use simulation to compute the distribution of the best adjustment to policy path, and compute the distribution of $\mathbb{E}_t \mathbf{P}_t^0 + \mathcal{R}_{a,p}^0 \delta_{a,t}$ and the associated paths of policy objectives $(\mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}_{a,y}^0 \delta_{a,t})$.

S0.5 OPP implementation for policy recommendation

An important property of the OPP and its variants is the ability to improve a non-optimal policy decision. Going forward, a policy maker could be interested in using this property to set policy with the OPP approach. Specifically, a policy maker could be interested in repeatedly (i.e., at each decision point) implementing OPP policy recommendations; i.e., in systematically using the OPP approach to set policy with minimal structural assumptions. Given the ubiquity of heuristics in macro decision making, this approach could be appealing to policy makers.

In this section, we discuss how one could operationalize such OPP-based decision making. Consider a policy maker who initially followed the rule ϕ_0 over some sample period and then

decide to add the OPP evaluation/improvement step to her decision making process. Since the OPP adjustment changes the policy rule (eq. (26) in the main text), the policy maker will see a change a change in her rule, to say ϕ_1 . To use the OPP repeatedly over time, it is important to keep the two sufficient statistics (the forecast and the impulse response) under the same rule. With impulse responses estimated under ϕ_0 , the baseline scenario must remain under the old rule ϕ_0 . Thus, the OPP can be used systematically, provided that one continues to construct forecasts based on the “old” rule ϕ_0 , and that the OPP recommendation are based on these forecasts.⁴ The forecasts under ϕ^0 can be constructed as discussed in section S0.1.

While it may seem surprising to construct forecasts based on an “outdated” rule, recall that the baseline policy need not be a policy to be evaluated or even implemented. The baseline policy can also be used as a tool to compute the optimal/improved policy path, where the baseline forecast serves to capture the characteristics of the time t problem. See our discussion on econometric-based optimal policy in the main text.

S1 Example: NK model with commitment

In this section we present the details for our sufficient statistics approach when applied to the optimal policy problem under commitment in the baseline New Keynesian model (e.g. Galí, 2015, Section 5.1.2). This extends the simple example of Section 2 for the case where the policy maker cares about future inflation and output deviations. The set-up can be regarded as a special case of our general framework.

The optimal policy problem under commitment is characterized by a policy maker that aims to minimize

$$\mathcal{L}_0 = \frac{1}{2} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + x_t^2) ,$$

with respect to π_t, π_{t+1}, \dots and x_t, x_{t+1}, \dots and subject to the constraints

$$\begin{aligned} \pi_t &= \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + \xi_t , \\ x_t &= \mathbb{E}_t x_{t+1} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \pi_{t+1}) , \end{aligned}$$

where β is the discount factor. Note that we can view this example as case a special case of our general policy problem (15) when taking \mathbb{E}_t as \mathbb{E}_0 , i.e. starting at $t = 0$. Clearly this is just an arbitrary normalization, i.e. the same would be achieved when starting from $\mathbb{E}_t(\cdot)$

⁴A policy maker operationalizing OPP recommendations will switch from following rule ϕ_0 to rule ϕ_1 . In the class of linear models considered, this information is common knowledge. Studying the problem of learning (e.g., Bullard and Mitra, 2002) the new policy rule, and more generally extending the sufficient statistics approach to incorporate learning, is an interesting avenue for future research.

and indexing the objectives by s for $s = t, t + 1, \dots$

The optimality conditions for this problem are given by

$$x_0 = -\kappa\pi_0 \quad \text{and} \quad x_t = x_{t-1} - \kappa\pi_t, \quad \forall t = 1, 2, \dots, \quad (\text{S12})$$

or

$$x_t = -\kappa\hat{p}_t \quad \forall t = 0, 1, 2, \dots,$$

where $\hat{p}_t = p_t - p_{-1}$ denotes the (log) deviation between the price level and an implicit target given by the price level prevailing one period before the central bank chooses its optimal plan (Galí, 2015, page 135).

A possible interest rate rule that (a) implements this optimal allocation and (b) leads to a unique equilibrium is given by

$$i_t = -[\phi_p + (1 - \delta)(1 - \kappa\sigma)] \sum_{k=0}^t \delta^k \xi_{t-k} - (\phi_p/\kappa)x_t$$

for any $\phi_p > 0$ (Galí, 2015, page 138). Note that this instrument rule is a special case of the generic policy rule (18). The coefficients in the rule are given by

$$\delta \equiv \frac{1 - \sqrt{1 - 4\beta a^2}}{2a\beta}, \quad \text{with} \quad a \equiv \frac{1}{1 + \beta + \kappa^2}.$$

The forecasts under the optimal allocation can be written as

$$\mathbb{E}_0\pi_0 = \delta\xi_0 \quad \mathbb{E}_0x_0 = -\kappa\delta\xi_0 \quad \mathbb{E}_0\pi_t = (\delta^{t+1} - \delta^t)\xi_0 \quad \mathbb{E}_0x_t = -\kappa\delta^{t+1}\xi_0 \quad (\text{S13})$$

for $t \geq 1$ (Galí, 2015, page 136).

Next, we rewrite this example in our general notation. Let $\mathbf{Y}_0 = (\pi_0, x_0, \pi_1, x_1, \dots)'$, $\mathbf{P}_0 = (i_0, i_1, \dots)'$, $\mathbf{\Xi}_0 = (\xi_0, \xi_1, \dots)'$ and denote by $\boldsymbol{\epsilon}_0 = (\epsilon_0, \epsilon_1, \dots)'$ the sequence of policy news shocks, which are equal to zero under the optimal rule (note that \mathbf{W}_t does not exist in this application). The loss function can be written as

$$\mathcal{L}_t = \frac{1}{2} \mathbb{E}_0 \mathbf{Y}'_0 \mathcal{W} \mathbf{Y}_0$$

where $\mathcal{W} = \text{diag}((\beta^0, \beta_1, \dots)' \otimes (1, 1)')$. The general model (16)-(18) becomes

$$\begin{aligned} \mathcal{A}_{yy} \mathbb{E}_0 \mathbf{Y}_0 - \mathcal{A}_{yp} \mathbb{E}_0 \mathbf{P}_0 &= \mathcal{B}_{y\xi} \mathbf{\Xi}_0 \\ \mathbb{E}_0 \mathbf{P}_0 - \mathcal{A}_{py} \mathbb{E}_0 \mathbf{Y}_0 &= \mathcal{B}_{p\xi} \mathbf{\Xi}_0 + \boldsymbol{\epsilon}_0 \end{aligned}$$

where the coefficient maps are given by

$$\mathcal{A}_{yy} = \begin{bmatrix} 1 & -\kappa & -\beta & 0 & \dots & \dots \\ 0 & 1 & -1/\sigma & -1 & 0 & \dots \\ 0 & 0 & 1 & -\kappa & -\beta & \ddots \\ 0 & 0 & 0 & 1 & -1/\sigma & \ddots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix} \quad \mathcal{A}_{yp} = \begin{bmatrix} 0 & 0 & 0 & \dots \\ 1/\sigma & 0 & 0 & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 1/\sigma & 0 & \dots \\ 0 & 0 & 0 & \ddots \\ 0 & 0 & 1/\sigma & \ddots \\ \vdots & \vdots & \ddots & \ddots \end{bmatrix}$$

$$\mathcal{A}_{py} = \begin{bmatrix} 0 & \phi_p/\kappa & 0 & 0 & \dots & \dots \\ 0 & 0 & 0 & \phi_p/\kappa & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & \phi_p/\kappa \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

and

$$\mathcal{B}_{y\xi} = \begin{bmatrix} 1 & 0 & 0 & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 0 & 1 & \ddots \\ 0 & 0 & 0 & \ddots \\ \vdots & \vdots & \ddots & \ddots \end{bmatrix} \quad \mathcal{B}_{p\xi} = \begin{bmatrix} \gamma_0 & 0 & 0 & \dots \\ \gamma_1 & \gamma_0 & 0 & \dots \\ \gamma_2 & \gamma_1 & \gamma_0 & \ddots \\ \vdots & \ddots & \ddots & \ddots \end{bmatrix}$$

where $\gamma_j = -[\phi_p + (1 - \delta)(1 - \kappa\sigma)]\delta^j$. It follows that \mathcal{R}_y , after some tedious manipulations, can be written —under the optimal policy rule— as

$$\mathcal{R}_y = \begin{bmatrix} \kappa/(\sigma v) & \kappa^2/(\sigma^2 v^2) + \kappa/\sigma v^2 + \kappa/(\sigma v) & \dots \\ 1/(\sigma v) & \kappa/(\sigma^2 v^2) + 1/(\sigma v^2) & \dots \\ 0 & \kappa/(\sigma v) & \dots \\ 0 & 1/(\sigma v) & \dots \\ 0 & 0 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

where $v = 1 - \phi_p/(\kappa/\sigma)$. Note that we only show the first two columns for ease of exposition. Given \mathcal{R}_y and the forecasts (S13) we can verify the equivalence condition, similar as shown

in equation (10) for the problem under discretion we have

$$\begin{aligned} \frac{\partial \mathcal{L}_0(\boldsymbol{\delta}_0)}{\partial \boldsymbol{\delta}_0} \Big|_{\boldsymbol{\delta}_t=0} &= \mathcal{R}'_y \mathcal{W} \mathbb{E}_0 \mathbf{Y}_0 \\ &= \begin{bmatrix} \kappa/(\sigma v) & \kappa^2/(\sigma^2 v^2) + \kappa/\sigma v^2 + \kappa/(\sigma v) & \dots \\ 1/(\sigma v) & \kappa/(\sigma^2 v^2) + 1/(\sigma v^2) & \dots \\ 0 & \kappa/(\sigma v) & \dots \\ 0 & 1/(\sigma v) & \dots \\ 0 & 0 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}' \begin{bmatrix} \delta \xi_0 \\ -\kappa \delta \xi_0 \\ \beta(\delta^2 - \delta) \xi_0 \\ -\beta \kappa \delta^2 \xi_0 \\ \vdots \\ \vdots \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ \vdots \end{bmatrix}. \end{aligned}$$

This mechanically shows that $\mathcal{R}'_y \mathcal{W} \mathbb{E}_0 \mathbf{Y}_0 = 0$ must also hold under optimality when considering the NK optimal policy problem under commitment.

In addition, note that as in the case under discretion, the impulse response matrix \mathcal{R}_y is sufficient to characterize the optimal targeting rule. Working under perfect foresight, the condition $\mathcal{R}'_y \mathcal{W} \mathbf{Y}_0 = 0$ corresponds exactly to the optimal targeting rule (S12).

S2 General nonlinear OPP framework

So far we have developed our sufficient statistics approach for policy evaluation in the context of linear models that can be written as in (16). In this section we explore for which other classes of models the statistics \mathcal{R}^0 and $\mathbb{E}_t \mathbf{Y}_t^0$ are sufficient to evaluate, and possibly improve, policy decisions. Key examples that we consider are models with state dependence (e.g. Auerbach and Gorodnichenko, 2013) and models with multiple policy regimes (e.g. Sims and Zha, 2006).

Recall that the properties of the OPP derive from the equivalence (i.e. Proposition 1)

$$\mathbb{E}_t \mathbf{P}_t^{\text{opt}} = \mathbb{E}_t \mathbf{P}_t^0 \iff \nabla_{\boldsymbol{\delta}_t} \mathcal{L}_t(\boldsymbol{\delta}_t) \Big|_{\boldsymbol{\delta}_t=0} = \mathcal{R}_y^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0 = \mathbf{0} \text{ , ,} \quad (\text{S14})$$

which we have shown to hold for linear models of the generic form (16).

Here we first provide a high level framework that exactly spells out the necessary conditions on the underlying economy that ensure that our sufficient statistics approach for policy evaluation applies.

After this we will make two specific points which are useful to place the OPP approach in the literature. First, as long as the model is linear conditional on time- t predetermined variables, the equivalence continues to hold and all our previous results hold. This case notably includes models of state dependent policy effects that are often considered in the

empirical literature (e.g. Auerbach and Gorodnichenko, 2013). Second, in a model with multiple regimes conditioned by the policy rule (e.g. Sims and Zha, 2006) as long as the economy can only be in a finite number of regimes, the equivalence no longer holds, but a non-zero gradient still implies that the proposed policy choice is non-optimal. In other words, the two statistics \mathcal{R}^0 and $\mathbb{E}_t \mathbf{Y}_t^0$ are still sufficient to evaluate a policy decision, but adjusting the policy choice with the OPP is no longer guaranteed to yield a superior policy decision.

B1: Generic nonlinear OPP

Consider an economy that can be represented at time t by the moment equations

$$\begin{cases} \mathbf{0} &= \mathbb{E}_t f(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t; \phi) \\ \mathbf{0} &= \mathbb{E}_t \phi(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t) \end{cases}, \quad (\text{S15})$$

where $f()$ and $\phi()$ are possibly nonlinear functions. The function $f()$ describes the general economy and takes as inputs the policy variables \mathbf{Y}_t , the policy instruments \mathbf{P}_t , the initial conditions \mathbf{X}_{-t} and the structural shocks $\boldsymbol{\Xi}_t$. The policy equation is characterized by ϕ , which takes similar inputs. A key difference with the set up in the main text is that we allow f to depend directly on the policy rule ϕ . This generalization directly corresponds to the generic set-up in Lucas (1976), see equations 16 and 17 in his paper, which allows the policy rule to alter the function $f()$ in an arbitrary way.

The optimal allocation, for the loss function $\mathcal{L}_t = \mathbb{E}_t \mathbf{Y}_t' \mathcal{W} \mathbf{Y}_t$, in this economy is characterized by

$$\min_{\mathbf{Y}_t, \mathbf{P}_t, \phi} \mathcal{L}_t \quad s.t. \quad (\text{S15})$$

where the difference is that the optimal allocation now also depends on ϕ , the choice for the policy rule that may affect the structure in the economy. The optimal expected policy path is again denoted by $\mathbb{E}_t \mathbf{P}_t^{\text{opt}}$.

To build up to our sufficient statistics approach consider a policy maker who can introduce exogenous surprises in the policy rule ϕ . As before we denote such policy shocks by $\boldsymbol{\epsilon}_t$ and we postulate that each element of \mathbf{P}_t corresponds to a specific element in $\boldsymbol{\epsilon}_t$. The economy becomes

$$\begin{cases} \mathbf{0} &= \mathbb{E}_t f(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t; \phi) \\ \mathbf{0} &= \mathbb{E}_t \phi(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t, \boldsymbol{\epsilon}_t) \end{cases}. \quad (\text{S16})$$

In this setting a policy choice is determined by the function $\phi \in \Phi$, where Φ denotes an arbitrary function class and a sequence of policy shocks $\boldsymbol{\epsilon}_t$. Let the proposed policy choice be denoted by $(\phi^0, \boldsymbol{\epsilon}_t^0)$, which implicitly characterizes the expected policy path $\mathbf{P}_t^{\epsilon_0}$ and allocation $\mathbb{E}_t \mathbf{Y}_t^0$.

The following high-level assumption exactly determines the class of models for which our sufficient statistics approach continues to work.

Assumption S1. *There exists a non-empty subset $\Phi^{\text{opt}} \subset \Phi$ such that*

1. *for all $\phi \in \Phi^{\text{opt}}$ and $\phi = \phi^0$ we have a unique and determinate equilibrium given by*

$$\begin{cases} \mathbb{E}_t \mathbf{Y}_t &= h_y(\mathbf{X}_{-t}, \boldsymbol{\Xi}_t, \boldsymbol{\epsilon}_t; \phi) \\ \mathbb{E}_t \mathbf{P}_t &= h_p(\mathbf{X}_{-t}, \boldsymbol{\Xi}_t, \boldsymbol{\epsilon}_t; \phi) \end{cases}.$$

where $h_y(\cdot)$ is continuously differentiable with respect to all $\boldsymbol{\epsilon}_t \in \mathcal{E}$, where \mathcal{E} is an open convex subset of \mathbb{R}^∞ .

2. $\mathcal{L}_t(\phi, \mathbf{0}) \leq \mathcal{L}_t(\tilde{\phi}, \tilde{\boldsymbol{\epsilon}}_t)$ for all $\phi \in \Phi^{\text{opt}}$, $\tilde{\phi} \in \Phi \setminus \Phi^{\text{opt}}$ and $\tilde{\boldsymbol{\epsilon}}_t \in \mathcal{E}$.

The first part of the assumption imposes the existence of a unique equilibrium under the optimal policy rules ϕ^{opt} and the proposed policy rule ϕ^0 . The second part defines the optimal rules as those that minimize the loss function. The key part in the assumption is that the loss function is minimized by ϕ^{opt} with $\boldsymbol{\epsilon}_t = 0$. That is, under the optimal rule it is not possible to further lower the loss function by introducing exogenous policy news shocks.

For nonlinear models like (S15) this assumption would need to be verified on a case by case basis. To better understand Assumption S1, it is helpful to consider our baseline (linear) framework and discuss how Assumption S1 is satisfied in this simpler case.

First, for the first part of Assumption S1, note that Assumption 2 in the main text ensures that a unique equilibrium exists under ϕ^0 . In addition, Assumption S1 also require that a unique equilibrium exists under ϕ^{opt} . The reason is that we need to ensure that the gradient of \mathbf{Y}_t with respect to $\boldsymbol{\epsilon}_t$ exists under the optimal policy. In the linear model this is not necessary as the optimal policy allocation does not depend on ϕ , see equation (17). With $f(\cdot)$ depending on ϕ however, this is no longer the case and the optimal policy does depend on ϕ . In that case, we also need a “well-behaved” optimal rule, i.e., that ϕ^{opt} implies a unique equilibrium. The second part of Assumption S1 holds in our linear framework as the optimal allocation $\mathbb{E}_t \mathbf{P}_t^{\text{opt}}$ does not depend on $\boldsymbol{\epsilon}_t$ and can be attained by a rule of form (18). For instance, we may take $\mathcal{A}_{pp} = \mathcal{A}_{pw} = \mathcal{B}_{px} = \mathcal{B}_{p\xi} = 0$ and $\mathcal{A}_{py} = \mathcal{R}^0 \mathcal{W}$, to obtain the optimal targeting rule $\mathcal{R}^0 \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0 = 0$.

With assumption S1 in hand we can proceed with our policy evaluation, similar as in the main text we are interested in evaluating whether $\mathbb{E}_t \mathbf{P}_t^0 = \mathbb{E}_t \mathbf{P}_t^{\text{opt}}$. As a test statistic we rely on the generalized OPP statistic given by

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

where

$$\mathfrak{R}_y^0 = \left. \frac{\partial h_y(\mathbf{X}_{-t}, \mathbf{\Xi}_t, \boldsymbol{\epsilon}_t; \phi)}{\partial \boldsymbol{\epsilon}_t} \right|_{\phi=\phi^0, \boldsymbol{\epsilon}_t=\boldsymbol{\epsilon}_t^0}.$$

Two comments are in order. First, note that for the linear model considered in the main text we have $h_y(\mathbf{X}_{-t}, \mathbf{\Xi}_t, \boldsymbol{\epsilon}_t; \phi^0) = \mathcal{R}_y^0 \boldsymbol{\epsilon}_t^e + \Gamma_y^0 \mathbf{S}_t$, see Lemma 1, such that we immediately have $\mathfrak{R}_y^0 = \mathcal{R}_y^0$. Second, we do not claim that \mathfrak{R}_y^0 can generally be estimated using conventional econometric methods. Clearly, when the function $h_y(\cdot)$ is unknown this is complicated as non-parametric methods will need to be used, which given the typically limited time series observations available may yield uninformative causal effects. The point here is simply to illustrate the theoretical limit of our approach.

We have the following key result.

Theorem S1. *Given model (S16) under assumption S1 if $\boldsymbol{\delta}_t^{g*}$ exists we have that*

$$\boldsymbol{\delta}_t^{g*} \neq \mathbf{0} \quad \Rightarrow \quad \mathbb{E}_t \mathbf{P}_t^0 \neq \mathbb{E}_t \mathbf{P}_t^{\text{opt}}.$$

Proof. By Assumption S1 part 1, the loss function \mathcal{L}_t is continuously differentiable on \mathcal{E} , thus by Lemma 4.3.1 in Dennis and Schnabel (1996) and Assumption S1 the optimal policy $\mathbb{E}_t \mathbf{P}_t^{\text{opt}} = h_p(\mathbf{X}_{-t}, \mathbf{\Xi}_t, \mathbf{0}; \phi^{\text{opt}})$ for any $\phi^{\text{opt}} \in \Phi^{\text{opt}}$ satisfies the gradient condition $\nabla_{\boldsymbol{\epsilon}_t} \mathcal{L}_t|_{\phi=\phi^{\text{opt}}, \boldsymbol{\epsilon}_t=\mathbf{0}} = \mathbf{0}$. Hence, if $\mathbb{E}_t \mathbf{P}_t^0$ is optimal we must have that $\nabla_{\boldsymbol{\epsilon}_t} \mathcal{L}_t|_{\phi=\phi^0, \boldsymbol{\epsilon}_t=\boldsymbol{\epsilon}_t^0} = \mathfrak{R}_y^0 \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0 = \mathbf{0}$, with $\mathbb{E}_t \mathbf{Y}_t^0 = h_y(\mathbf{X}_{-t}, \mathbf{\Xi}_t, \boldsymbol{\epsilon}_t^0; \phi^0)$, which if $\boldsymbol{\delta}_t^{g*}$ exists implies that $\boldsymbol{\delta}_t^{g*}$ must satisfy $\boldsymbol{\delta}_t^{g*} \neq \mathbf{0}$ if $\mathbb{E}_t \mathbf{P}_t^0 \neq \mathbb{E}_t \mathbf{P}_t^{\text{opt}}$. The existence of $\boldsymbol{\delta}_t^{g*}$ depends on whether the inverse of $\mathfrak{R}^0 \mathcal{W} \mathfrak{R}^0$ exists, which can be verified for any given model. \square

The key purpose for stating Theorem S1 is twofold. First, the key underlying Assumption (S1) characterizes the class of models for which our sufficient statistics approach can be used for policy evaluation. Second, Theorem S1 is useful for characterizing optimal policy in complicated possibly nonlinear models. Indeed while the traditional approach to optimal policy — as documented in Section 2 for the baseline NK model — may become complex for nonlinear models, the gradient based approach underlying Theorem S1 may provide a workable necessary condition for optimal policy. That is determining $\boldsymbol{\delta}_t^{g*} \neq \mathbf{0}$ may be easier than solving the model and computing the optimal policy.

Next, we discuss the details for the specific examples — state dependence and multiple regimes — as mentioned in the main text.

B1: State dependence

Numerous works have documented evidence for various forms of state dependence in the effects of fiscal and monetary policy, where the state dependence is governed by some time-

t pre-determined variable that is independent of the policy decision.⁵ Our main results continue to hold in this setting, and the two statistics \mathcal{R}^0 and $\mathbb{E}_t \mathbf{Y}_t^0$ are sufficient to detect and correct non-optimal policy decisions. The only difference is that the statistic \mathcal{R}^0 needs to be conditioned on the state of the economy.

As an illustration, consider the specification of Auerbach and Gorodnichenko (2013) where the economy can be in two states, depending on the value of some state variable z_t . Under such state dependence the generic model for the economy remains conditionally linear, i.e.

$$\begin{cases} \mathcal{A}_{yy}(z_t)\mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{yw}(z_t)\mathbb{E}_t \mathbf{W}_t - \mathcal{A}_{yp}(z_t)\mathbb{E}_t \mathbf{P}_t &= \mathcal{B}_{yx}(z_t)\mathbf{X}_{-t} + \mathcal{B}_{y\xi}(z_t)\boldsymbol{\Xi}_t \\ \mathcal{A}_{ww}(z_t)\mathbb{E}_t \mathbf{W}_t - \mathcal{A}_{wy}(z_t)\mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{wp}(z_t)\mathbb{E}_t \mathbf{P}_t &= \mathcal{B}_{wx}(z_t)\mathbf{X}_{-t} + \mathcal{B}_{w\xi}(z_t)\boldsymbol{\Xi}_t \end{cases}, \quad (\text{S17})$$

where z_t is some predetermined time- t measurable variable and for $\mathcal{D} = \mathcal{A}, \mathcal{B}$

$$\mathcal{D}_{..}(z_t) = F(z_t)\mathcal{D}_{..(1)} + (1 - F(z_t))\mathcal{D}_{..(2)}$$

where $F(z_t)$ can be interpreted as a measure of probability of being in state 1 at time t based on some time t predetermined variable z_t .⁶

The optimal state dependent policy $\mathbb{E}_t \mathbf{P}_t^{\text{opt}}(z_t)$ can be defined as the solution for \mathbf{P}_t to

$$\min_{\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t} \mathcal{L}_t \quad \text{s.t.} \quad (\text{S17}). \quad (\text{S18})$$

The generic policy rule is given by

$$\mathcal{A}_{pp}(z_t)\mathbb{E}_t \mathbf{P}_t - \mathcal{A}_{py}(z_t)\mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{pw}(z_t)\mathbb{E}_t \mathbf{W}_t = \mathcal{B}_{px}(z_t)\mathbf{X}_{-t} + \mathcal{B}_{p\xi}(z_t)\boldsymbol{\Xi}_t,$$

where the definition for the maps $\mathcal{A}_{..}(z_t)$ and $\mathcal{B}_{..}(z_t)$ is the same as above. We collect all the coefficients of the rule in $\phi(z_t)$. The state dependent OPP under a given rule $\phi^0(z_t)$ is given by

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

where $\mathcal{R}_y^0(z_t) = F(z_t)\mathcal{R}_{y,(1)}^0 + (1 - F(z_t))\mathcal{R}_{y,(2)}^0$ captures the causal state dependent effect of $\boldsymbol{\epsilon}_t$ on \mathbf{Y}_t . The following corollary summarizes the properties of the state dependent OPP.

Corollary S1. *Given model (S17). Under the assumptions that (1) the optimal policy $\mathbb{E}_t \mathbf{P}_t^{\text{opt}}(z_t)$ is unique, and (2) the rule $\phi^0(z_t)$ underlying the proposed policy path $\mathbb{E}_t \mathbf{P}_t^0(z_t)$*

⁵See e.g., Auerbach and Gorodnichenko (2012, 2013); Ramey and Zubairy (2018); Barnichon, Debortoli and Matthes (2021) for studies on whether fiscal policy is more or less effective when the economy is in a high unemployment state, and Tenreyro and Thwaites (2016); Ascari and Haber (2021); Eichenbaum, Rebelo and Wong (2022) for studies on whether monetary policy is more or less effective when unemployment is high.

⁶A popular functional form for $F(\cdot)$ is $F(z_t) = \exp(-\gamma z_t) / [1 + \exp(-\gamma z_t)]$ with γ a tuning parameter.

leads to a unique and determinate equilibrium, we have that

1. $\delta_t^*(z_t) = 0 \iff \mathbb{E}_t \mathbf{P}_t^0(z_t) = \mathbb{E}_t \mathbf{P}_t^{\text{opt}}(z_t)$.
2. $\mathbb{E}_t \mathbf{P}_t^0 + \mathcal{R}_p^0(z_t) \delta_t^*(z_t) = \mathbb{E}_t \mathbf{P}_t^{\text{opt}}(z_t)$.

Proof of Corollary S1. The proof is identical to the proof of Proposition 1 after changing the maps \mathcal{A}_\cdot and \mathcal{B}_\cdot to $\mathcal{A}_\cdot(z_t)$ and $\mathcal{B}_\cdot(z_t)$, respectively. \square

B2: Multiple policy regimes

Next, we consider an economy with a finite number of policy regimes, where the model coefficients (equations (16)) can depend on the policy regime. To give a concrete and relevant example in monetary policy, inflation expectations can be “anchored” —fixed at some value— or “unanchored”, in that inflation expectations depend the state of the economy (expectation formation could be e.g., rational or adaptive), and the anchoring of inflation expectations likely depends on the central bank’s policy rule or objective function (e.g., Bernanke, 2007).

In this type of model, the policy rule can affect the coefficients of the non-policy block (16). In that case, we can still use the OPP to detect optimization failures, but there is no longer any guarantee (without additional structural assumption) that an OPP adjustment would improve policy. In other words, the two statistics \mathcal{R}^0 and $\mathbb{E}_t \mathbf{Y}_t^0$ are still sufficient to evaluate a policy decision, but they may not be sufficient to correct a non-optimal policy decision.

To formalize this, consider a generalization of (16) with

$$\begin{cases} \mathcal{A}_{yy}(\vartheta) \mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{yw}(\vartheta) \mathbb{E}_t \mathbf{W}_t - \mathcal{A}_{yp}(\vartheta) \mathbb{E}_t \mathbf{P}_t &= \mathcal{B}_{yx}(\vartheta) \mathbf{X}_{-t} + \mathcal{B}_{y\xi}(\vartheta) \boldsymbol{\Xi}_t \\ \mathcal{A}_{ww}(\vartheta) \mathbb{E}_t \mathbf{W}_t - \mathcal{A}_{wy}(\vartheta) \mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{wp}(\vartheta) \mathbb{E}_t \mathbf{P}_t &= \mathcal{B}_{wx}(\vartheta) \mathbf{X}_{-t} + \mathcal{B}_{w\xi}(\vartheta) \boldsymbol{\Xi}_t \end{cases}, \quad (\text{S19})$$

where $\mathcal{A}_\cdot(\vartheta)$ and $\mathcal{B}_\cdot(\vartheta)$ capture describe the economy under the policy regime ϑ .

As in the baseline model, the generic policy rule is

$$\mathcal{A}_{pp} \mathbf{P}_t^e - \mathcal{A}_{py} \mathbb{E}_t \mathbf{Y}_t - \mathcal{A}_{pw} \mathbb{E}_t \mathbf{W}_t = \mathcal{B}_{px} \mathbf{X}_{-t} + \mathcal{B}_{p\xi} \boldsymbol{\Xi}_t, \quad (\text{S20})$$

and we collect all the rule parameters in $\phi = \{\mathcal{A}_{pp}, \mathcal{A}_{py}, \mathcal{A}_{pw}, \mathcal{B}_{px}, \mathcal{B}_{p\xi}\}$.

The policy regime ϑ can depend on the policy rule ϕ , creating a feedback from the policy rule to the coefficients of the non-policy block. Assume that there exists only a finite number of regimes. Here we consider the case with two regimes for clarity of exposition and following our example on anchored/unanchored inflation expectations. The maps $\mathcal{D}_\cdot(\vartheta)$, for $\mathcal{D} = \mathcal{A}, \mathcal{B}$, are of the form

$$\mathcal{D}_\cdot(\vartheta) = \begin{cases} \mathcal{D}_\cdot(\vartheta_1) & \text{if } \phi \in \Phi_1 \\ \mathcal{D}_\cdot(\vartheta_2) & \text{if } \phi \in \Phi_2 \end{cases}, \quad (\text{S21})$$

where $\Phi_1 \cup \Phi_2 = \Phi$ and $\Phi_i \cap \Phi_j = \emptyset$.

The characterization of the optimal policy choice requires more care. Here we define the optimal policy as the policy rule (and associated policy regime) that ensures the lowest loss. We consider

$$\min_{\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t, \phi} \mathcal{L}_t \quad \text{s.t.} \quad (\text{S19})\text{-(S20)} , \quad (\text{S22})$$

where we note that since the coefficients of the policy equation directly affect the coefficients that describe the economy, we need to take into account the policy equation when defining the optimal policy. We denote the solution for \mathbf{P}_t to this minimization problem by $\mathbb{E}_t \mathbf{P}_t^{\text{opt}}(\vartheta^{\text{opt}})$ where ϑ^{opt} corresponds to ϕ^{opt} which is the minimizing ϕ . We will assume that ϕ^{opt} lies in the interior of some Φ_i , which rules out boundary solutions.

Given a policy proposal $\mathbb{E}_t \mathbf{P}_t^0$, implied by choices ϕ^0 and ϵ^0 , where ϕ^0 implies the regime $\vartheta^0 \in \{\vartheta_1, \vartheta_2\}$, the regime specific OPP statistic is

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

where $\mathbb{E}_t \mathbf{Y}_t^0$ is the expected allocation under $\mathbb{E}_t \mathbf{P}_t^0$. Note that computing the regime specific OPP only requires $\mathcal{R}^0(\vartheta^0)$ the causal effects under the proposed policy and does not require knowledge of the causal effects in any of the other regimes.

The following corollary establishes the main property of the regime specific

Corollary S2. *Given model (S19)-(S21), under the assumptions that (1) the optimal policy $\mathbb{E}_t \mathbf{P}_t^{\text{opt}}(\vartheta^{\text{opt}})$ is unique and the underlying rule ϕ^{opt} leads to a unique and determinate equilibrium, and (2) the rule ϕ^0 underlying the proposed policy path $\mathbb{E}_t \mathbf{P}_t^0(z_t)$ leads to a unique and determinate equilibrium, we have that*

$$\delta_t^*(\vartheta^0) \neq 0 \quad \Rightarrow \quad \mathbb{E}_t \mathbf{P}_t^0(\vartheta^0) \neq \mathbb{E}_t \mathbf{P}_t^{\text{opt}}(\vartheta^{\text{opt}}) .$$

The corollary implies that if the regime specific OPP is non-zero we have that the policy choice is non optimal.

Proof of Corollary S2. The Lagrangian for the optimal policy problem (S22) is given by

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

where $\mathcal{A}_\cdot(\vartheta)$ and $\mathcal{B}_\cdot(\vartheta)$ capture describe the economy under the policy regime ϑ determined by the policy rule $\phi = \{\mathcal{A}_{pp}, \mathcal{A}_{py}, \mathcal{A}_{pw}, \mathcal{B}_{px}, \mathcal{B}_{p\xi}\}$

The first order conditions for $\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t$ are given by

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

Importantly, with a finite number of regimes if ϕ^{opt} lies in the interior of some Φ_i , all the derivatives of the maps $\mathcal{A}_\cdot(\vartheta)$ and $\mathcal{B}_\cdot(\vartheta)$ with respect to the elements of ϕ are zero. Intuitively, an infinitely small change in a rule coefficient does not trigger a regime change. In that case, unless the economy is already perfectly stabilized,⁷ the first order conditions with respect to the elements of ϕ imply that $\boldsymbol{\mu}_3 = 0$. To see that, note that optimization with respect to, for instance, the j th element of the first row of $\mathcal{B}_{p\xi}$ gives

$$\mu_{3,j}\mathbb{E}\xi_t = 0$$

where $\mu_{3,j}$ is the corresponding element of the $\boldsymbol{\mu}_3$ vector. Unless $\mathbb{E}_t\xi_t = 0$,⁸ this implies $\mu_{3,j} = 0$. We can proceed similarly with the other coefficients of ϕ to show $\boldsymbol{\mu}_3 = 0$.

With $\boldsymbol{\mu}_3 = 0$, note that the optimization problem (S22) has the same first order conditions as the following fictitious problem: given some optimal policy rule ϕ^{opt} the policy maker can choose $\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t$ and a sequence of policy shocks $\boldsymbol{\epsilon}_t$ to minimize \mathcal{L}_t . Indeed, for that problem the Lagrangian writes

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

The first-order conditions with respect to $\boldsymbol{\epsilon}_t$ yield $\boldsymbol{\mu}_3 = 0$, which immediately establishes that the first-order conditions of the fictitious problem are the same as those of problem (S22). As in the proof of Proposition 1, under the stated assumption that ϕ^{opt} implies a unique equilibrium, the fictitious problem can also be stated as

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

⁷Technically, we exclude that all shocks until time t and all initial conditions are zero, a trivial case we can discard.

⁸Technically, the condition is much weaker: we only have to exclude that the economy is perfectly stabilized when $\mathbb{E}\mathbf{P}_t = 0$, i.e., that all shocks and all initial conditions are zero. A trivial case we can discard.

which leads to the first order condition

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

This establishes that $\mathcal{R}_y^{\text{opt}'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t = \mathbf{0}$ is a necessary condition for the original optimization problem (S22).

Next, consider the second fictitious problem: given the proposed policy rule ϕ^0 the policy maker can choose \mathbf{Y}_t , \mathbf{W}_t , \mathbf{P}_t and a sequence of policy shocks $\boldsymbol{\epsilon}_t$ to minimize \mathcal{L}_t . Using the exactly the same steps as above, and noting that ϕ^0 leads to a unique equilibrium, it follows that $\mathcal{R}_y^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t = \mathbf{0}$ is a necessary condition for optimality under this rule. This implies that if we find that $\mathcal{R}_y^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t \neq \mathbf{0}$, then $\mathbb{E}_t \mathbf{P}_t^0 \neq \mathbb{E}_t \mathbf{P}_t^{\text{opt}}$ as there can only be two cases. First, if ϕ^0 implies the same regime as ϕ^{opt} it follows immediately as $\mathcal{R}_y^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t = \mathcal{R}_y^{\text{opt}'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t = \mathbf{0}$ is a necessary condition. Second, if we find that ϕ^0 does not imply the same regime as ϕ^{opt} the result also follows immediately as ϕ^0 then does not minimize the loss function.

Finally, note that since $\boldsymbol{\delta}_t^*(\vartheta^0)$ is just a rescaling of $\mathcal{R}_y^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t$ the same conclusion carries over, and $\boldsymbol{\delta}_t^*(\vartheta^0) = \mathbf{0}$ is a necessary condition for optimality. Importantly however, unlike in the linear case, that condition is not sufficient to characterize the optimal policy. With of a feedback from ϕ to the maps $\mathcal{A}_\cdot(\vartheta)$ and $\mathcal{B}_\cdot(\vartheta)$, a policy satisfying $\mathcal{R}_y^{0'} \mathcal{W} \mathbb{E}_t \mathbf{Y}_t = \mathbf{0}$ could be a *local* minimum, when the regime ϑ^0 associated with rule ϕ^0 is not the regime ϑ^{opt} associated with the optimal rule ϕ^{opt} . \square

S3 General convex loss functions

In the main text we restricted ourselves to quadratic loss functions when testing the optimality of a given policy choice. In this section we show that the main idea – exploiting the gradient of the loss function to evaluate optimality– continues to apply for essentially any convex loss function. The only difference is that the evaluation of the gradient will require the full forecast densities instead of only the mean oracle forecasts.

To show this, let $\mathcal{L}_t(\mathbf{Y}_t; \theta)$ denote a loss function which is convex and differentiable with respect to \mathbf{Y}_t and may depend on preference parameters denoted by θ . The quadratic loss function (15) in the main text is a special case where θ includes preference parameters λ and discount factors β . Using the same generic model (16), we can summarize the policy maker’s problem by

$$\min_{\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t; \theta) \quad \text{s.t} \quad (16)$$

To evaluate whether a proposed policy choice $\mathbb{E}_t \mathbf{P}_t^0$ solves this problem we can follow the same steps as in the main text. Specifically, using the proof of Proposition 1 it follows immediately that the equivalence $\mathbb{E}_t \mathbf{P}_t^0 = \mathbb{E}_t \mathbf{P}_t^{\text{opt}} \iff \nabla_{\boldsymbol{\delta}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t(\boldsymbol{\delta}_t); \theta)|_{\boldsymbol{\delta}_t=0}$ continues to

hold. However, now the gradient with respect to $\boldsymbol{\delta}_t$ evaluated at $\boldsymbol{\delta}_t = \mathbf{0}$ (i.e. $\mathbb{E}_t \mathbf{P}_t^0$) is given by

$$\nabla_{\boldsymbol{\delta}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t(\boldsymbol{\delta}_t); \theta) |_{\boldsymbol{\delta}_t = \mathbf{0}} = \mathcal{R}_y^{0'} \times \nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t(\boldsymbol{\delta}_t); \theta) |_{\boldsymbol{\delta}_t = \mathbf{0}} .$$

Given that $\mathcal{L}_t(\mathbf{Y}_t; \theta)$ is convex with respect to \mathbf{Y}_t and \mathbf{Y}_t is an affine function of $\boldsymbol{\epsilon}_t$ in equilibrium we have that if $\nabla_{\boldsymbol{\delta}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t(\boldsymbol{\delta}_t); \theta) |_{\boldsymbol{\delta}_t = \mathbf{0}} \neq 0$ the policy choice $\mathbb{E}_t \mathbf{P}_t^0$ is not optimal.

To evaluate the gradient we need to compute the derivative $\nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t(\boldsymbol{\delta}_t); \theta) |_{\boldsymbol{\delta}_t = \mathbf{0}}$. Under a quadratic loss $\frac{1}{2} \mathbf{Y}_t' \mathcal{W} \mathbf{Y}_t$, this expression simplifies to $\nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t(\boldsymbol{\delta}_t); \theta) |_{\boldsymbol{\delta}_t = \mathbf{0}} = \mathcal{W} \mathbb{E}_t \mathbf{Y}_t^0$ as in the main text, but for a general convex loss we have

$$\nabla_{\mathbf{Y}_t} \mathbb{E}_t \mathcal{L}_t(\mathbf{Y}_t(\boldsymbol{\delta}_t); \theta) |_{\boldsymbol{\delta}_t = \mathbf{0}} = \int_{\mathbf{Y}_t^0} \nabla_{\mathbf{Y}_t} \mathcal{L}_t(\mathbf{Y}_t^0; \theta) p(\mathbf{Y}_t^0 | \mathcal{F}_t) d\mathbf{Y}_t^0 , \quad (\text{S23})$$

where $p(\mathbf{Y}_t^0 | \mathcal{F}_t)$ is the forecast density under the proposed policy choice $\mathbb{E}_t \mathbf{P}_t^0$. Thus, provided the forecast density is available, we can construct the OPP statistic and OPP-based tests as in the main text. The only difference is that there is no closed form expression for the gradient, and numerical or Monte Carlo integration methods will be necessary.

S4 Details: constrained OPP

We provide further details for the constrained OPP statistic of Section 4.3. Specifically, we generalize the set-up in the paper to allow for constraints on any endogenous variable and provide formal statements for the properties of the constrained OPP and the subset constrained OPP.

S4.1 General constrained OPP

We consider the general case where the constraints are of the form

$$C(\mathbb{E}_t \mathbf{Y}_t, \mathbb{E}_t \mathbf{W}_t, \mathbb{E}_t \mathbf{P}_t) \geq \mathbf{c} , \quad (\text{S24})$$

where $C(\cdot, \cdot, \cdot)$ is a continuous function that allows to capture a variety of constraints that the policy maker may face. We give some explicit examples below. The policy maker's problem becomes

$$\min_{\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t} \mathcal{L}_t \quad \text{s.t.} \quad (16) \quad \text{and} \quad C(\mathbb{E}_t \mathbf{Y}_t, \mathbb{E}_t \mathbf{W}_t, \mathbb{E}_t \mathbf{P}_t) \geq \mathbf{c} . \quad (\text{S25})$$

Clearly, with nonlinear constraints there is no closed form expression for the constrained OPP. Nevertheless, the optimality of a given policy path can be assessed by numerically

solving the program

$$\begin{aligned} \boldsymbol{\delta}_t^{c*} &= \underset{\boldsymbol{\delta}_t}{\operatorname{argmin}} (\mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}_y^0 \boldsymbol{\delta}_t)' \mathcal{W} (\mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}_y^0 \boldsymbol{\delta}_t) \\ \text{s.t. } & C(\mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}_y^0 \boldsymbol{\delta}_t, \mathbb{E}_t \mathbf{W}_t^0 + \mathcal{R}_w^0 \boldsymbol{\delta}_t, \mathbb{E}_t \mathbf{P}_t^0 + \mathcal{R}_p^0 \boldsymbol{\delta}_t) \geq \mathbf{c} \end{aligned} \quad (\text{S26})$$

where \mathcal{R}_y^0 , \mathcal{R}_w^0 and \mathcal{R}_p^0 are the dynamic causal effects of $\boldsymbol{\epsilon}_t$ on \mathbf{Y}_t , \mathbf{W}_t and \mathbf{P}_t under ϕ^0 . The problem (S26) is equivalent to a constrained least squares problem and can be easily solved numerically.

The constrained OPP retains the same properties as its linear counterpart above.

Corollary S3. *Given the generic model (16)-(18) and under assumptions that: (i) $\mathbf{P}_t^{\text{opt},c}$ is the unique solution to (S25) and (ii) the rule ϕ^0 underlying the proposed policy path $\mathbb{E}_t \mathbf{P}_t^0$ leads to a unique and determinate equilibrium, we have that*

1. $\boldsymbol{\delta}_t^{c*} = 0 \iff \mathbb{E}_t \mathbf{P}_t^0 = \mathbb{E}_t \mathbf{P}_t^{\text{opt},c}$
2. $\mathbb{E}_t \mathbf{P}_t^0 + \boldsymbol{\delta}_t^{c*} = \mathbb{E}_t \mathbf{P}_t^{\text{opt},c}$,

where $\boldsymbol{\delta}_t^{c*}$ is defined in (S26).

S4.2 Subset OPP with nonlinear inequality constraints

Next, we extend our approach to derive a constrained subset-OPP.

Let $\mathbf{P}_{a,t}^{e0}$ denote the subset or linear combination of the proposed policy path for which the corresponding policy shocks $\boldsymbol{\epsilon}_{a,t}^{e0}$ can be identified. The subsets of the causal effects \mathcal{R}_a^0 , $\mathcal{C}_{a,w\epsilon}^0$ and $\boldsymbol{\Theta}_{a,\epsilon}^0$ measures the effects of $\boldsymbol{\epsilon}_{a,t}^e$ on \mathbf{Y}_t , \mathbf{W}_t and \mathbf{P}_t , respectively. The constraints that the policy maker faces take the same general form as in (S24).

The subset constrained OPP can be obtained from the program

$$\begin{aligned} \boldsymbol{\delta}_{a,t}^{c*} &= \underset{\boldsymbol{\delta}_{a,t}}{\operatorname{argmin}} (\mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}_{a,y}^0 \boldsymbol{\delta}_{a,t})' \mathcal{W} (\mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}_{a,y}^0 \boldsymbol{\delta}_{a,t}) \\ \text{s.t. } & C(\mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{R}_{a,y}^0 \boldsymbol{\delta}_{a,t}, \mathbb{E}_t \mathbf{W}_t^0 + \mathcal{R}_{a,w}^0 \boldsymbol{\delta}_{a,t}, \mathbb{E}_t \mathbf{P}_t^0 + \mathcal{R}_{a,p}^0 \boldsymbol{\delta}_{a,t}) \geq \mathbf{c} \end{aligned} \quad (\text{S27})$$

The constrained subset OPP can also be easily computed using numerical optimization routines. The only inputs are again the forecasts and the subset dynamic causal effects.

The inequality constrained subset OPP has the following properties.

Corollary S4. *Given the generic model (16)-(18) and under assumptions that: (i) $\mathbb{E}_t \mathbf{P}_t^{\text{opt},c}$ is the unique solution to (S25) and (ii) the rule ϕ^0 underlying the proposed policy path $\mathbb{E}_t \mathbf{P}_t^0$ leads to a unique and determinate equilibrium, we have that*

1. $\delta_{a,t}^{c*} \neq 0 \implies \mathbb{E}_t \mathbf{P}_t^0 \neq \mathbb{E}_t \mathbf{P}_t^{\text{opt},c}$
2. $\mathcal{L}_t(\delta_{a,t}^*, \mathbf{0}) \leq \mathcal{L}_t(\mathbf{0}, \mathbf{0})$, i.e. the adjusted path $\mathbb{E}_t \mathbf{P}_t^0 + \mathcal{R}_{a,p}^0 \delta_{a,t}^{c*}$ lowers the loss.

The properties are similar to those that were obtained for the unconstrained subset OPP in the main text.

We now give a few examples for the types of constraints that can be accommodated in this framework.

Example S3 (Zero lower bound). *A well known constraint on the interest rate path of a central bank is the zero lower bound. We can implement this for $\mathbf{P}_t = (i_t, i_{t+1}, \dots, i_{t+H})'$ by considering the constraint*

$$\mathbf{C} \mathbb{E}_t \mathbf{P}_t \geq \mathbf{c}, \quad \text{with} \quad \mathbf{C} = \mathbf{I}_H, \quad \mathbf{c} = \mathbf{0}.$$

This is the specification that we adopt in our empirical work in section 6 for the zero-lower bound period.

Example S4 (Pre-commitments). *A richer example of a policy constraint is the following announcement by the FOMC after its 09/16/2020 meeting:*

‘... The Committee decided to keep the target range for the federal funds rate at 0 to 1/4 percent and expects it will be appropriate to maintain this target range until labor market conditions have reached levels consistent with the Committee’s assessments of maximum employment and inflation has risen to 2 percent and is on track to moderately exceed 2 percent for some time. ...’

To illustrate how to incorporate such pre-commitment we will compute a constrained OPP, where we impose the constraint that the policy rate remains at the zero lower bound until inflation is expected to lie at least 0.5 ppt above target for 1 year and unemployment is less than 0.5 ppt above its long-run level. These thresholds are only chosen as means of illustration.

Letting $\mathbf{P}_t = (i_t, i_{t+1}, \dots)'$ be the interest rate path, we can translate this commitment into a conditional constraint of the form

$$C(\mathbb{E}_t \mathbf{Y}_t, \mathbb{E}_t \mathbf{W}_t, \mathbb{E}_t \mathbf{P}_t) = \begin{cases} \mathbb{E}_t \mathbf{P}_t = \mathbf{0} & \text{if } \mathbb{E}_t u_{t+h} > 4.5 \forall h, \mathbb{E}_t \pi_{t+j} < 2.5 \ j = 1, \dots, 4 \\ \mathbb{E}_t \mathbf{P}_t & \text{else} \end{cases}.$$

Note that we could also write this expression using an indicator function for the conditioning event.

S5 Estimating robust preference parameters

Researchers outside of the policy maker’s research staff may not have access to the policy maker’s preferences \mathcal{W} . For this setting, we outline an approach for conducting *preference robust* OPP evaluation. The idea is to exploit a sequence of past policy decisions to find the preferences that gives the smallest deviations from optimality on average. This approach can thus be seen as considering a worse-case scenario for “rejecting” optimality.

Specifically, we write $\omega = \beta \otimes \lambda$, the elements of the preference matrix \mathcal{W} , as a function of the $d_\theta \times 1$, parameter vector θ , i.e., $\omega = \omega(\theta)$, with $d_\theta \leq K$, and we estimate θ by numerically solving ⁹

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \widehat{S}(\theta) , \quad \widehat{S}(\theta) = \left\| \frac{1}{n} \sum_{s=t_0}^t \widehat{\delta}_{a,s}(\theta) \right\|^2 , \quad (\text{S28})$$

where $\widehat{\delta}_{a,s}(\theta)$ corresponds to the mean OPP estimate as a function of θ .

With the estimated $\hat{\theta}$ in hand a researcher can base the optimality assessment on the simulated distribution of the OPP given $\hat{\theta}$, which ensures that deviations from optimality are not due to potentially arbitrary choices for the preference parameters.

S6 Additional results for the empirical study

In this section we discuss additional results for our empirical study on testing US monetary policy decisions. These results are complementary to those presented in Section 6 of the main text. In particular, the different subsections discuss the following aspects.

1. Impulse responses to policy shocks
2. Sensitivity to the preference parameter λ
3. Alternative dynamic causal effect estimates: SVAR inference
4. Testing the stability of the macro environment

S6.1 Impulse responses to policy shocks

In our empirical study we use shocks to the policy rate and to the slope of the yield curve.

To better understand the policy experiments that we are considering, Figure S1 plots the full set of impulse responses to (i) innovations to the fed funds rate, and (ii) innovations to the slope of yield curve. We can see that both policy experiments correspond to somewhat

⁹To give a specific example, suppose that $M_y = 2$ and $M_p = 2$, then we can take $\theta = (\theta_1, \theta_2)'$ and set $\omega(\theta) = (\theta_1^0, \theta_1^1, \dots)' \otimes (1, \theta_2)'$, which implies that $\beta = (\theta_1^0, \theta_1^1, \dots)'$ and $\lambda = (1, \theta_2)'$.

persistent changes to the policy instrument, similar to earlier estimates of impulse responses to monetary shocks (e.g., Eberly, Stock and Wright, 2020). In other words, these experiments will allow to evaluate the optimality of the short- to medium-end of the policy paths.

S6.2 The preference parameter λ

In the main text we used $\lambda = 1$ to compute the OPP-based tests. In this section we compute the OPPs for different choices of λ between $[0.2, 2]$. The results for different choices of λ are shown in Figure S2. We find that the short-rate OPP is not sensitive to the choice for λ . In fact, all of our main findings hold for all choices of λ and the differences are often small. For the slope OPP the findings are a bit different. Here low values of λ , say $\lambda = 0.2$, move the slope OPP towards zero. The reason is that the causal effects of inflation are estimated with greater uncertainty, such that putting more weight on the inflation mandate (a lower λ) lessens the “power” of our evaluation framework. That said, we stress that the data indicates that such low values for λ are unlikely, as the worst case λ that we computed using (S28) was determined at $\lambda = 0.6$.

S6.3 Shadow rate results

In this section, we provide an alternative approach to evaluating policy decisions. Instead of using the fed funds rate as the policy instrument, we use the shadow fed funds rate constructed by Wu and Xia (2016). The benefit of this approach is that it removes the non-negativity constraint on the fed funds rate. The impulse responses are estimated as in the main text, only substituting the fed funds rate with the shadow fed funds rate (Figure S3)

The policy evaluation results (Figure S4) are similar to the baseline results reported in the main text (Figure 1), with the (unsurprising) exceptions of the two zero lower bound periods: 2009-2015 and 2020-2021. Unlike the constrained short-rate OPP reported in the main text, the shadow-rate OPP turns very negative during the Great Recession: the shadow fed funds rate was about 0.75 percentage point too high over 2009-2013. This finding confirms our conclusion that unconventional monetary policy should have been used more aggressively during the great recession. Interestingly, during the COVID crisis, the shadow short rate OPP also turn substantially negative (-0.5 percentage points), but model uncertainty was so high at that time that the OPP evaluation is inconclusive. This is to be expected: if little is known in real time about the transmission of an unusual shock (like COVID) to the economy, one would expect policy evaluation to be particularly difficult. So while the shadow short-rate OPP turns negative in 2020, the results are not significantly different from zero at any reasonable confidence level.

S6.4 Sequential optimization and time inconsistency

As is well-known, sequential decision making creates the possibility of dynamic inconsistency in the expected policy paths: An expected policy path that is optimal as of time t may not be optimal viewed from a time decision problem as of time $t + 1$ (Kydlan and Prescott, 1977).

The OPP statistic —being precisely about the policy maker’s time t decision problem—naturally inherits the time inconsistency associated with sequential decision making.

However, we saw in the main text that we can isolate the time inconsistent component and measure its quantitative relevance. Specifically, we can decompose the time t OPP as

$$\delta_t^* = \delta_{t-1}^* + \underbrace{[\mathbf{0}, \mathcal{D}^0] \Delta \mathbb{E}_t \mathbf{Y}_{t-1}^0}_{\text{Information update}} + \underbrace{\Delta \mathcal{D}^0 \mathbb{E}_{t-1} \mathbf{Y}_{t-1}^0}_{\text{Preference shift}} \quad (\text{S29})$$

or after simplifying as

$$\delta_t^* = \delta_{t-1}^* + \Delta \mathbb{E}_t \mathbf{Y}_t^0 + \mathcal{D}^0 \mathbb{E}_{t-1} \Delta \mathbf{Y}_t^0$$

where $\Delta \mathbb{E}_t(\cdot) \equiv \mathbb{E}_t(\cdot) - \mathbb{E}_{t-1}(\cdot)$ is the information update operator, $\Delta \mathbf{Y}_t^0 \equiv \mathbf{Y}_t^0 - \mathbf{Y}_{t-1}^0$ is the path in first differences and $\mathcal{D} \equiv -(\mathcal{R}_y^{0'} \mathcal{W} \mathcal{R}_y^0)^{-1} \mathcal{R}_y^{0'} \mathcal{W}$.

Equation (S29) decomposes the time t OPP into three components: (i) the “lagged, i.e., time $t - 1$, OPP —past deviations from optimality that have not been corrected—, (ii) an information update term —new information revealed at t —, and (iii) a preference shift term. The preference shift terms captures the dynamic time inconsistency associated with sequential decision making.

Whether dynamic inconsistency is quantitatively important depends on the policy problem at hand. In the case of US monetary policy decisions, it appears to be a minor problem. To see that, Figure S5 reports the result of decomposition (S29) for the (unconstrained) short-rate OPP over 1990-2022. Notice that preference shift (right-bottom panel) contributes very little to OPP adjustment between periods. In other words, time inconsistency is quantitatively irrelevant for our OPP evaluation: the conclusions are identical whether we assume that the policy maker follows time consistent policies or not. Instead, we can see that most of the adjustments to the OPP come from two forces: past deviations from optimality that have not been corrected —the “lagged” OPP—, and new information. For instance, the onset of the great recession (in early 2008) led to a large information update that implied a large negative adjustment to the OPP (red line), followed three years later by a positive adjustment as the recovery was stronger than expected. Not surprisingly, the COVID crisis was also marked by large information updates.


References

- Ascari, Guido, and Timo Haber.** 2021. “Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy.” *The Economic Journal*, 132(641): 37–57.
- Auerbach, Alan J., and Yuriy Gorodnichenko.** 2012. “Measuring the Output Responses to Fiscal Policy.” *American Economic Journal: Economic Policy*, 4: 1–27.
- Auerbach, Alan J., and Yuriy Gorodnichenko.** 2013. “Output Spillovers from Fiscal Policy.” *American Economic Review*, 103(3): 141–46.
- Bai, Jushan, and Serena Ng.** 2002. “Determining the Number of Factors in Approximate Factor Models.” *Econometrica*, 70(1): 191–221.
- Banbura, Marta, Domenico Giannone, and Lucrezia Reichlin.** 2010. “Large Bayesian Vector Auto Regressions.” *Journal of Applied Econometrics*, 25(1): 71–92.
- Banbura, Marta, Domenico Giannone, and Michele Lenza.** 2015. “Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections.” *International Journal of Forecasting*, 31(3): 739–756.
- Barnichon, Regis, and Christian Brownlees.** 2018. “Impulse Response Estimation By Smooth Local Projections.” *Review of Economics and Statistics*.
- Barnichon, Regis, Davide Debortoli, and Christian Matthes.** 2021. “Understanding the Size of the Government Spending Multiplier: Its in the Sign.” *The Review of Economic Studies*, 89(1): 87–117.
- Bernanke, Ben.** 2007. “Inflation expectations and inflation forecasting.” Board of Governors of the Federal Reserve System (US).
- Blanchard, Olivier Jean, and Danny Quah.** 1989. “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review*, 79(4): 655–673.
- Bullard, James, and Kaushik Mitra.** 2002. “Learning about monetary policy rules.” *Journal of monetary economics*, 49(6): 1105–1129.
- Dennis, John E., and Robert B. Schnabel.** 1996. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Siam.
- Durbin, J., and S. J. Koopman.** 2012. *Time Series Analysis by State Space Methods; 2nd Edition*. Oxford:Oxford University Press.

- Eberly, Janice C., James H. Stock, and Jonathan H. Wright.** 2020. “The Federal Reserve’s Current Framework for Monetary Policy: A Review and Assessment.” *International Journal of Central Banking*, 16(1): 5–71.
- Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong.** 2022. “State Dependent Effects of Monetary Policy: The Refinancing Channel.” *American Economic Review*. forthcoming.
- Faust, Jon.** 1998. “The Robustness of Identified VAR Conclusions About Money.” *Carnegie-Rochester Conference Series on Public Policy*, 49: 207 – 244.
- Galí, Jordi.** 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*. Princeton University Press.
- Jordà, Oscar.** 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *The American Economic Review*, 95: 161–182.
- Kydland, Finn E., and Edward C. Prescott.** 1977. “Rules Rather than Discretion: The Inconsistency of Optimal Plans.” *Journal of Political Economy*, 85(3): 473–491.
- Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf.** 2022. “Local Projections vs. VARs: Lessons From Thousands of DGPs.” Working paper.
- Lucas, Robert Jr.** 1976. “Econometric Policy Evaluation: A critique.” *Carnegie-Rochester Conference Series on Public Policy*, 1(1): 19–46.
- Mertens, Karel, and Morten O. Ravn.** 2013. “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States.” *American Economic Review*, 103(4): 1212–1247.
- Plagborg-Møller, Mikkel, and Christian K. Wolf.** 2021. “Local Projections and VARs Estimate the Same Impulse Responses.” *Econometrica*, 89(2): 955–980.
- Ramey, Valerie.** 2016. “Macroeconomic Shocks and Their Propagation.” In *Handbook of Macroeconomics*, ed. J. B. Taylor and H. Uhlig. Amsterdam, North Holland:Elsevier.
- Ramey, Valerie A., and Sarah Zubairy.** 2018. “Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data.” *Journal of Political Economy*, 126.
- Sims, Christopher A.** 1980. “Macroeconomics and reality.” *Econometrica*, 1–48.
- Sims, Christopher A., and Tao Zha.** 2006. “Were There Regime Switches in U.S. Monetary Policy?” *American Economic Review*, 96(1): 54–81.

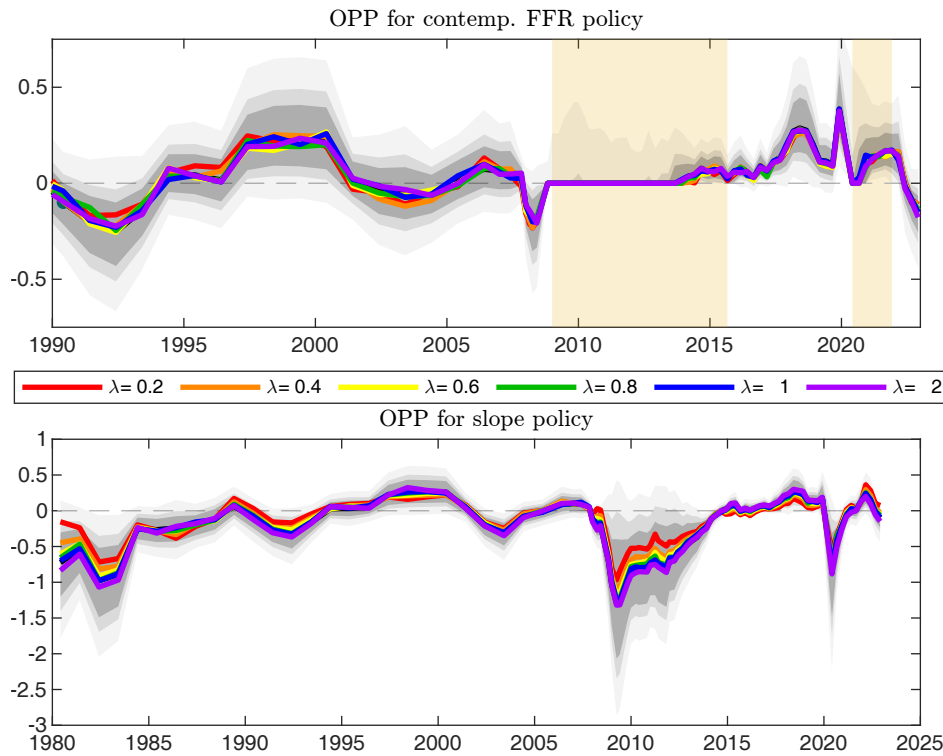
- Stock, James H., and Mark W. Watson.** 2002*a*. “Forecasting Using Principal Components from a Large Number of Predictors.” *Journal of the American Statistical Association*, 97(460): 1167–1179.
- Stock, James H., and Mark W. Watson.** 2002*b*. “Macroeconomic Forecasting Using Diffusion Indexes.” *Journal of Business & Economic Statistics*, 20(2): 147–162.
- Stock, James H., and Mark W. Watson.** 2018. “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments.” *The Economic Journal*, 128(610): 917–948.
- Tenreyro, Silvana, and Gregory Thwaites.** 2016. “Pushing on a String: US Monetary Policy Is Less Powerful in Recessions.” *American Economic Journal: Macroeconomics*, 8(4): 43–74.
- Uhlig, Harald.** 2005. “What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure.” *Journal of Monetary Economics*, 52(2): 381–419.
- Waggoner, Daniel F., and Tao Zha.** 1999. “Conditional Forecasts in Dynamic Multivariate Models.” *The Review of Economics and Statistics*, 81(4): 639–651.
- Wu, Jing Cynthia, and Fan Dora Xia.** 2016. “Measuring the macroeconomic impact of monetary policy at the zero lower bound.” *Journal of Money, Credit and Banking*, 48(2-3): 253–291.

Figure S1: IMPULSE RESPONSES TO POLICY INNOVATIONS




Notes: Impulse responses (IRs) of inflation, unemployment and the fed funds rate or the slope of the yield curve (10year-FFR spread) to a contemporaneous fed funds rate shock (left panel) or to a yield curve slope shock (right panel). Shaded bands denote the 68 and 90 percent confidence intervals.

Figure S2: OPP FOR DIFFERENT λ , 1990-2022



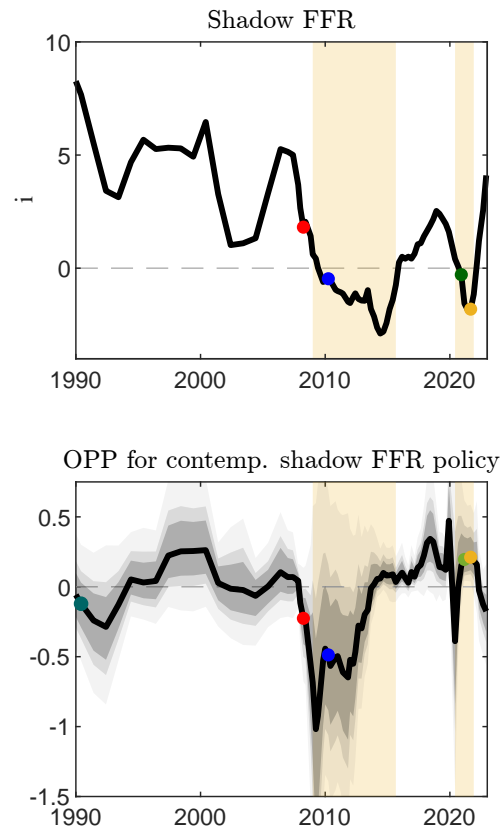
Notes: Top panel: short-rate OPP. The colored lines correspond to the OPP with λ 's between 0.2 and 2. Bottom panel: slope policy OPP. The colored lines correspond to the OPP with λ 's between 0.2 and 2. The grey areas capture impulse response and model uncertainty at 60%, 75% and 90% confidence (from darker to lighter shades) when using $\lambda = 1$ as in the main text.

Figure S3: IMPULSE RESPONSES TO POLICY INNOVATIONS



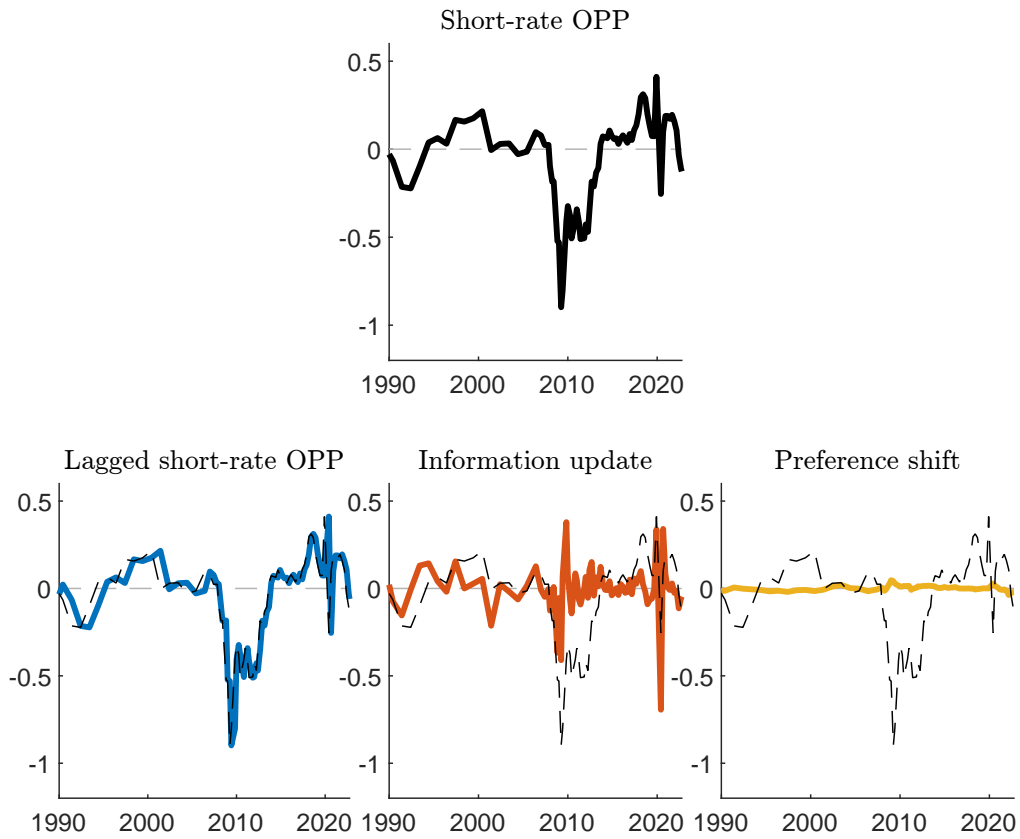
Notes: Impulse responses (IRs) of inflation, unemployment and the shadow fed funds rate (Wu and Xia, 2016) or the slope of the yield curve (10year-FFR spread) to a contemporaneous shadow fed funds rate shock (left panel) or to a yield curve slope shock (right panel). Shaded bands denote the 68 and 90 percent confidence intervals.

Figure S4: A SEQUENCE OF OPP FOR SHADOW FED FUNDS RATE POLICY (1990-2022)



Notes: Top panel: the shadow fed funds rate (Wu and Xia, 2016). The yellow shaded areas denote the zero-lower bound (ZLB) periods. Bottom panel: time series for the unconstrained shadow short-rate OPP (labeled “OPP for contemp. shadow FFR”, left panel) over 1990-2022 for a policy maker with a dual inflation–unemployment mandate ($\lambda = 1$). The grey areas capture impulse response and model uncertainty at 60%, 75% and 90% confidence (from darker to lighter shades). The case studies are marked as points: April 2008 (red), April 2010 (blue), March 2021 (green) and November 2021 (yellow).

Figure S5: A DECOMPOSITION OF THE SHORT-RATE OPP, 1990-2022



Notes: Top panel: the median short-rate OPP (black line). Bottom panels: decomposition of the short-rate OPP (dashed black line) into its three components: (i) the lagged OPP (left panel, blue line), (ii) the information update term (middle panel, red line), and (iii) the preference shift term (right panel, yellow line). See equation (S29). The preference shift terms captures the strength of dynamic time inconsistency coming from a sequential use of the OPP.