UNIVERSITY OF LIVERPOOL Management School

Working Paper in Economics

 # 202112

A Toolkit for Computing Constrained Optimal Policy Projections (COPPs)

Oliver de Groot
Falk Mazelis
Roberto Motto
Annukka Ristiniemi

A Toolkit for Computing Constrained Optimal Policy Projections (COPPs)*

Oliver de Groot† Falk Mazelis‡ Roberto Motto§
Annukka Ristiniemi¶

May 6, 2021

Abstract

This paper presents a toolkit for generating optimal policy projections. It makes five contributions. First, the toolkit requires a minimal set of inputs: only a baseline projection for target and instrument variables and impulse responses of those variables to policy shocks. Second, it solves optimal policy projections under commitment, limited-time commitment, and discretion. Third, it handles multiple policy instruments. Fourth, it handles multiple constraints on policy instruments such as a lower bound on the policy rate and an upper bound on asset purchases. Fifth, it allows alternative approaches to address the forward guidance puzzle. The toolkit that accompanies this paper is `Dynare` compatible, which facilitates its use. Examples replicate existing results in the optimal monetary policy literature and illustrate the usefulness of the toolkit for highlighting policy trade-offs. We use the toolkit to analyse US monetary policy at the height of the Great Financial Crisis. Given the Fed's early-2009 baseline macroeconomic projections, we find the Fed's planned use of the policy rate was close to optimal whereas a more aggressive QE program would have been beneficial.

Keywords: Optimal monetary policy, Commitment vs. discretion, Lower bound, Asset purchases, Forward guidance puzzle
JEL Classification: C61, C63, E52, E58

*COPPs toolkit repository: https://github.com/oliverdegroot/COPPs-toolkit. We wish to thank Camilo Marchesini for excellent research assistance and John Barrdear and Anton Nakov for invaluable contributions to an earlier version of the toolkit. We also wish to thank Anthony Diercks, Junior Maih, Geert Mesters, Ricardo Nunes, and participants at several seminars and workshops for valuable comments. The views expressed herein do not necessarily reflect those of the European Central Bank or the Eurosystem.

†University of Liverpool Management School & CEPR, Liverpool, UK (oliverdegroot@gmail.com)
‡European Central Bank, Frankfurt, Germany (falk.mazelis@ecb.europa.eu)
§European Central Bank, Frankfurt, Germany (roberto.motto@ecb.europa.eu)
¶European Central Bank, Frankfurt, Germany (annukka.ristiniemi@ecb.europa.eu)

1 Introduction

The setting of monetary policy has become decidedly more complicated in the last decade. The usual challenge of identifying adequate models of the economy and combining models with expert judgement has been accompanied by frequent encounters with the effective lower bound (ELB) on nominal interest rates. Central banks have responded to the ELB by adopting a range of non-standard measures such as forward guidance and large-scale asset purchases (or quantitative easing, QE).[1] But this has created the need to find an appropriate combination of these measures and to account for the fact that non-standard measures also face constraints, such as limits to asset purchases arising from legal or market-functioning considerations.

In this paper, we present a toolkit to compute Constrained Optimal Policy Projections (COPPs).[2] It offers a platform on which to elucidate policy trade-offs across multiple policy instruments facing real-world policymakers. Optimal policy projections under commitment were first introduced by Svensson (2005) and applied to a policy setting by Svensson and Tetlow (2005) using the Fed's FRB/US model. Optimal policy projections are related to the targeting rule approach to optimal policy (e.g., Svensson and Woodford 2004). Svensson (2010) describes optimal policy projections as the selection of projections for the target variables and policy instrument that "look best relative to the central bank's objectives". Our toolkit extends the existing approach in five directions.

First, the toolkit requires as inputs only the baseline projections for the target and instruments variables, and impulse responses for those variables to policy shocks. The toolkit does not require the specification of a model's structural equations to derive optimal policy projections. Impulse responses can come from a structural model such as a DSGE model or an identified SVAR but also from any other informed view about the transmission of policy.[3] As it requires only impulse responses to the policy instruments rather than a fully-fledged structural model, our toolkit makes it possible to use the large literature on monetary policy transmission. Hence, our toolkit differs from existing

[1] The list of policy instruments is rather larger depending on whether one wishes to distinguish, for example, between credit easing and QE, or to include liquidity operations into the mix. In addition, several central banks also have a macro-prudential mandate and macro-prudential instruments to deploy.

[2] We use the words *projection* and *path* interchangeably. One can also think of these projections as forecasts, although we avoid this terminology in the paper.

[3] While we work with impulse responses from structural models, in independent work Barnichon and Mesters (2020) take a complementary approach and calculate optimization failures of US monetary policy by instead estimating impulse responses from local projections to a federal funds "target" and "slope" shock. While their focus is on the effects of policy uncertainty, we focus on the consequences of time-inconsistency and expectations formation and we allow for multiple policy instruments and constraints.

algorithms to compute optimal policy (e.g., Dennis 2007) that require model-specific structural equations as input. The toolkit takes, as a starting point, a baseline projection of target and instrument variables. Among alternative projections of the instruments (computed using the impulse responses) that shift the target variables away from their baseline paths, the COPPs toolkit selects the alternative projection that minimizes the policymaker's objective function. One advantage of this approach is that it is agnostic about the structural shocks that drive the baseline.

Second, the toolkit can solve optimal policy projections under commitment, limited-time commitment, and discretion. The optimal policy projections literature cited above focused only on commitment. We extend optimal policy projections approach to analyse the discretion case. We also also intermediate situations between the case in which the policymaker commits indefinitely into the future and the case in which the policymaker re-optimizes every period. In particular, the toolkit allows for limited-time commitment in which a sequence of policymakers each serve a fixed term of a certain number of periods.[4] Each policymaker re-optimizes at the start of their term but can then commit for the remainder of their term.

Third, the toolkit can handle multiple policy instruments. In this paper, we study the interplay between standard interest rate policy (which incorporates forward guidance) and QE. The toolkit can be extended to more instruments and to study optimality problems outside the monetary policy realm. The optimal policy projections reviewed in Svensson (2010) focus on one instrument, the standard policy rate. Our toolkit is well suited to study current monetary policy frameworks that make use of several instruments.

Fourth, the toolkit can handle multiple constraints on policy instruments. We show applications in which we impose the ELB on the policy rate, an upper bound to asset purchases, and a maximum deviation of the policy instruments from their baseline path. We enforce these constraints using anticipated policy shocks. This follows the approach of Laséen and Svensson (2011) and Holden and Paetz (2012). Holden (2016) provides an efficient algorithm for implementing the procedure and Holden (2019) provides results on existence and uniqueness of equilibria. Other toolkits for solving models with occasionally binding constraints include the following: Guerrieri and Iacoviello (2015) use a pairwise-linear approach under perfect-foresight; Adjemian and Juillard (2011) use an extended path approach; whereas Eggertsson et al. (2020) solve models in a stochastic setting under

[4]Our limited-time commitment approach is closely related with Clymo and Lanteri (2020). Similar in spirit is the *loose commitment* approach of Debortoli and Nunes (2010) that assumes a policymaker can commit but with some exogenous probability succumbs to the temptation to re-optimize.

the assumption of a specific shock structure. However, none of these are designed to study optimal policy under both commitment and discretion, which is one of our contributions.[5]

Fifth, the COPPs toolkit allows for alternative approaches to address the forward guidance puzzle. Del Negro et al. (2012) documented that forward-looking DSGE models can generate puzzlingly strong effects of forward guidance. Optimal monetary policy prescribes policymakers to make promises about their future behaviour. Our approach to compute optimal policy relies on anticipated policy shocks, making explicit the dependence of outcomes on the forward-looking nature of economic agents.[6] There is a growing literature on resolutions of the forward guidance puzzle (e.g., McKay et al. 2016). However, most of these resolutions do not easily lend themselves to be employed into large-scale DSGE models. In fact, large-scale models used by policy institutions have typically not incorporated these approaches. We overcome this problem by introducing a consistent methodology that can be applied to mitigate the forward guidance puzzle in existing models. In particular, we modify expectations about policy announcements following de Groot and Mazelis (2020), capturing the notions of i) private-sector inattention (Gabaix 2020), ii) lack of credibility (Haberis et al. 2019), iii) finite planning horizons (Woodford 2019), and iv) learning (Cole 2020). Nakata et al. (2019) show that dampening the forward guidance puzzle has implications for optimal monetary policy. Our toolkit allows an assessment of the sensitivity of optimal policy projections to alternative approaches that modify expectations formation.

The toolkit that accompanies this paper is compatible with `Dynare` and provides an efficient method for conducting optimal policy in a suite of large-scale DSGE models. We provide several examples to show that the toolkit replicates existing results in the optimal monetary policy literature and to illustrate the practical usefulness of the toolkit and its innovative features.

Importantly, we use the toolkit to study the optimal mix of monetary policy instruments (forward guidance and QE) for the US at the height of the Global Financial Crisis. We take the perspective of early 2009 and use as baseline projections the macroeconomic variables documented in the March 2009 Greenbook. In that period, the Federal Reserve's staff were expecting the federal funds rate to remain at the ELB for about four years. We

[5] A contemporary paper to ours, Harrison and Waldron (2021), use a piecewise linear solution method to calculate optimal policy with occasionally binding constraints. While similar in some dimensions, their methods uses a different definition of time-consistent policy. They neither study limited commitment nor resolutions of the forward guidance puzzle as we do.

[6] Laséen and Svensson (2011) study the effect of anticipated policy rate paths in the context of optimal policy projections and find similarly puzzling effects resulting from the forward-looking nature of private sector agents.

employ two prominent models to compute impulse responses to policy shocks. The first model is Smets and Wouters (2007), which has the advantage of having been used in many applications in the literature but has the disadvantage of only one policy instrument—the policy rate. The second model is Sims and Wu (2020), which features QE as an additional instrument and is thus well suited to study this historical episode. Under the assumption of commitment, and correcting for the forward guidance puzzle, optimal policy projections suggest the Fed's use of the policy rate was close to optimal whereas a more aggressive QE program would have been beneficial.

The rest of the paper proceeds as follows. Section 2 presents the methodology. Section 3 highlights several features of the toolkit. Section 4 applies the toolkit to study the historical episode following the Great Financial Crisis in the US. Section 5 concludes.

2 Methodology

For clarity of exposition, we confine our attention to a model in which the policymaker has just two policy instruments. In the policy exercise in Section 4, these will be the short-term interest rate and QE. Extensions to more instruments are eminently possible.

This Section proceeds as follows. Subsection 2.1 defines the policymakers loss function and policy instruments. Subsection 2.2 defines the baseline projection. Subsection 2.3 demonstrates how to construct *alternative* policy projections. Subsection 2.4 demonstrates how to solve for *optimal* policy projections.

2.1 Policy preferences and policy instruments

The policymaker sets a $(n_x \times 1)$ vector of policy instruments, X_t, to minimize the quadratic loss function

$$\frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t L_t' Q L_t, \tag{1}$$

where $\mathbb{E}_0$ represents the mathematical expectations operator conditional on period 0 information, $\beta \in (0,1)$ is the discount factor, L_t is a $(n_l \times 1)$ vector of endogenous variables (henceforth, policy target variables), and Q is a symmetric positive semidefinite matrix comformable with L_t containing the policymaker's preference parameters. For convenience, all variables represent deviations from nonstochastic steady state values.

2.2 Baseline projection

We begin with a finite T-period [B]aseline projection, $\{L_t^B\}_{t=0}^T$ and $\{X_t^B\}_{t=0}^T$, of policy target and policy instrument variables, with dimensions $n_l \times (T+1)$ and $n_x \times (T+1)$, respectively.[7]

Remark 1. *The length of the baseline projection, T, must be sufficiently long such that all the policy preference and policy instrument variables have converged back to zero (the steady state) with a given tolerance.*

Our method is agnostic with regards to the origin of the baseline projections. They do not need to come from a model at the disposal of the researcher tasked with constructing the optimal policy projection. They need not even come from a formal structural model and can be purely judgemental in nature. Implicitly, we assume that there exists a model that could rationalize the baseline projections of the target and instrument variables in a consistent manner, but this assumption is very weak.

2.3 Alternative policy projections

To construct policy projections, we require unanticipated and anticipated impulse responses to policy instrument shocks. We illustrate this by considering a structural model described by the following system of equations[8]

$$\tilde{A}Y_t = B\mathbb{E}_t Y_{t+1} + \tilde{C}Y_{t-1} + \tilde{D}X_t, \tag{2}$$

$$X_t = FY_t + GM_t, \tag{3}$$

$$M_t = \mathcal{M}M_{t-1} + \mathcal{N}\tilde{V}_t, \tag{4}$$

where Y_t is an $n_y \times 1$ vector of endogenous variables, of which L_t is a subset.[9] The matrices $\tilde{A}, B, \tilde{C}, \tilde{D}, F$, and G contain the model's structural parameters, and are conformable with Y_t and X_t, as necessary. Equation (2) contains private sector equilibrium conditions and (3) contains policy rules.[10] Moreover, (3) is augmented with announced policy innovations.

[7]Later, it will be useful to work with vectorized versions of these objects. In general, we use the following notation: $Z \equiv \text{vec}\left(\{Z_t\}_{t=0}^T\right)$.

[8]The impulse responses to policy instrument shocks could, in principle, come from a time-series model such as an SVAR. However, the number of anticipated shocks that can be reliably estimated using an SVAR is typically small, which places an additional constraint on the policymaker. See D'Amico and King (2015) and Christoffel et al. (2020).

[9]It is without loss of generality that only X_t (and not X_{t-1} nor $\mathbb{E}_t X_{t+1}$) appears in (2) or that only Y_t appears in (3). These dependencies can be added with the use of auxiliary variables.

[10]Combining Equations (2)-(3) gives $AY_t = B\mathbb{E}_t Y_{t+1} + \tilde{C}Y_{t-1} + DM_t$, where $A \equiv \tilde{A} - \tilde{D}F$ and $D \equiv \tilde{D}G$.

Let V_t be a $(H+1) \times n_x$ innovation matrix, where the row denotes the horizon at which an innovation is realized and the column denotes the policy instrument to which the innovation applies. Thus, H is the maximum horizon for announced policy innovations. The model therefore contains a $n_x (H+1) \times 1$ vector of anticipated policy innovations, $\tilde{V}_t \equiv \text{vec}(V_t)$, which are $iid\,(0, \Omega_v)$, where $H \leq T$ is the horizon over which the policymaker can make policy announcements. One can think of these anticipated innovations as "forward guidance" about any of the policy instruments. The dimensions of G, $\mathcal{M}$, and $\mathcal{N}$ are $n_x \times n_x H(H+1)$, $n_x H(H+1) \times n_x H(H+1)$, and $n_x H(H+1) \times n_x (H+1)$, respectively.[11] Several remarks on the model are in order:

Remark 2. *By assumption, the conditions for existence of a rational expectations solution hold in this model.*

The optimal policy projections will be independent of the policy rule(s) described by (3). In that sense, these policy rules can be arbitrarily chosen so long as they obey a Taylor-type principle, ensuring determinacy.

Remark 3. *The model can provide a structural interpretation of the variables in the baseline projections.*

Thus far, we have presented the baseline projections and the model separately. This need not be the case. Model (2)-(3) could be used to construct the baseline projection with the addition of appropriate non-policy structural shocks. If the baseline projections did not come from the model, as would often be the case in a policy institution where the baseline forecast is a non-model based judgmental baseline, then the model can be used to find a sequence of structural innovations that recover the baseline path. Standard properties are required for identification in this case. However, the structural shocks that deliver the baseline path will not be used in constructing the optimal policy projection. It is sufficient to know the baseline paths of the target variables, not their drivers.

Remark 4. *The linear model remains valid when the baseline projections are subject to binding constraints.*

Consider, for example, the case of the effective lower bound (ELB) on the policy rate. Deviations from the policy rule (as a result of the ELB binding) can simply be interpreted as anticipated contractionary policy shocks in V_t. These anticipated shocks

[11] Non-policy structural shocks, such as technology shocks, have been removed for expositional clarity since they play no role in constructing alternative policy paths.

ensure that the actual and expected path of the policy rate does not fall below the ELB in a perfect-foresight simulation.[12] Thus, by adding anticipated shocks we transform a non-linear problem into a linear one.

To illustrate how anticipated shocks enter the model, we provide an example.

Example 5. *Suppose the policymaker i) has two policy instruments given by $X_t = [i_t, q_t]$; and ii) makes policy announcements up to horizon $H = 2$. Then*

$$
\begin{aligned}
M_t &= [m_{i00t}, m_{i01t}, m_{i11t}, m_{i02t}, m_{i12t}, m_{i22t}, m_{q00t}, m_{q01t}, m_{q11t}, m_{q02t}, m_{q12t}, m_{q22t}]', \\
\tilde{V}_t &= [v_{i0t}, v_{i1t}, v_{i2t}, v_{q0t}, v_{q1t}, v_{q2t}]', \quad G = [1,1,0,1,0,0] \otimes I_2, \\
\mathcal{M} &= \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \otimes I_2, \quad \mathcal{N} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \otimes I_2,
\end{aligned}
$$

where I_2 is the 2×2 identity matrix.

The [I]mpulse [R]esponse projections of this model can be written as

$$
\begin{aligned}
Y_0^{IR} &= \Phi \boldsymbol{v_{00}} + \mathcal{F}\Phi \boldsymbol{v_{10}} + \cdots + \mathcal{F}^{H-1}\Phi \boldsymbol{v_{H0}} + \mathcal{F}^{H}\Phi \boldsymbol{v_{H0}}, \\
Y_1^{IR} &= \mathcal{P}Y_0 + \Phi \boldsymbol{v_{10}} + \mathcal{F}\Phi \boldsymbol{v_{20}} + \cdots + \mathcal{F}^{H-1}\Phi \boldsymbol{v_{H0}} \\
&\vdots \\
Y_{H-1}^{IR} &= \mathcal{P}Y_{H-2} + \Phi \boldsymbol{v_{(H-1)0}} + \mathcal{F}\Phi \boldsymbol{v_{H0}}, \\
Y_H^{IR} &= \mathcal{P}Y_{H-1} + \Phi \boldsymbol{v_{H0}}, \\
Y_{H+\tau}^{IR} &= \mathcal{P}Y_{H-1+\tau} \quad \text{for} \quad 0 < \tau \leq T - H,
\end{aligned} \tag{5}
$$

where $Y_{-1}^{IR} = 0$ by design. Using Example 5, the notation $\boldsymbol{v_{ht}}$ denotes $[v_{iht}, v_{qht}]'$, i.e., these are all policy innovations announced in period t but implemented in period $h \geq t$. The

[12]In the context of perfect-foresight simulations, shock uncertainty does not play a role.

matrix $\mathcal{P}$ can be found by standard first-order solution techniques.[13] The other matrices are given as follows

$$\Phi = (A - B\mathcal{P})^{-1}\tilde{D}, \tag{6}$$

$$\mathcal{F} = (A - B\mathcal{P})^{-1}B. \tag{7}$$

Appendix A.1 provides the derivation of Equations (5)-(7). In matrix form, we can rewrite (5) as follows

$$Y^{IR} = \mathcal{A}\hat{V}_0, \quad \mathcal{A} = \mathbf{A}^{-1}\mathbf{C}, \tag{8}$$

where

$$\mathbf{A} = \begin{bmatrix} I & 0 & \cdots & 0 \\ -\mathcal{P} & I & \ddots & \vdots \\ & \ddots & \ddots & 0 \\ 0 & & -\mathcal{P} & I \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} \Phi & \mathcal{F}\Phi & \cdots & \mathcal{F}^H\Phi \\ 0 & \Phi & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathcal{F}\Phi \\ 0 & \cdots & 0 & \Phi \\ \hdashline 0 & \cdots & \cdots & 0 \end{bmatrix}, \tag{9}$$

$\hat{V}_t \equiv \text{vec}(V_t')$, and $Y^{IR} = [Y_0^{IR\prime} \cdots Y_T^{IR\prime}]'$ is a vector of length $n_y(T+1)$. When considering only the subvector of loss-function variables, L^{IR}, or policy instrument variables, X^{IR}, the appropriate submatrix of $\mathcal{A}$ is denoted $\mathcal{A}_L$ and $\mathcal{A}_X$, respectively.

Given a vector of time-0 anticipated policy shocks, $\hat{V}_0$, then Y^{IR} can be appropriately rearranged to give impulse response (IR) projections $\left\{Y_t^{IR}\right\}_{t=0}^T$. Finally, the [A]lternative policy projection is simply given by a linear sum of the baseline projection and the impulse response projections

$$\left\{Y_t^A\right\}_{t=0}^T = \left\{Y_t^B\right\}_{t=0}^T + \left\{Y_t^{IR}\right\}_{t=0}^T. \tag{10}$$

2.4 Optimal policy projections

Now that we have demonstrated how to construct an alternative policy projection, constructing an optimal policy projection involves choosing vector, V_0, that solves a given (constrained) optimization problem.

[13]The toolkit that accompanies this paper uses a `Dynare` compatible model file format. Thus, the model is automatically solved. Only matrix $\mathcal{F}$ needs to be constructed outside of `Dynare`.

2.4.1 Unconstrained commitment

This section considers the problem facing a policymaker that can fully commit to future actions. We consider a modified version of the loss function in (1), given by $\frac{1}{2}\sum_{t=0}^{T}\beta^t L_t' Q L_t$, where the infinite horizon problem has been replaced by a finite horizon problem. So long as T is sufficiently large, such that all relevant variables have converged back to zero, then this will serve as a good approximation. Since we are constructing perfect-foresight projections, we also drop the expectations operator.

The policy problem can be defined precisely as follows

$$\min_{\hat{V}_0} \frac{1}{2}\hat{V}_0' \mathcal{A}_L' \Omega \mathcal{A}_L \hat{V}_0 + L^{B\prime} \Omega \mathcal{A}_L \hat{V}_0, \tag{11}$$

where $\Omega \equiv \text{diag}\left(1, \beta, \ldots, \beta^T\right) \otimes Q$. Appendix A.2 derives (11) from the original problem in (1). The optimal set of policy announcements, $\hat{V}_0^*$, are given by

$$\hat{V}_0^* = -\left(\mathcal{A}_L' \Omega \mathcal{A}_L\right)^{-1} \mathcal{A}_L' \Omega L^B. \tag{12}$$

Remark 6. *Optimal policy projections under commitment coincide with the solution to optimal policy under commitment using Lagrangian methods.*

Remark 6 is important. It is common to solve the policymaker's problem by setting up a Lagrangian with the objective given by (1) and constraints given by (2) and deriving first-order conditions. The latter are often used to construct an instrument rule or a targeting rule. The toolkit that accompanies this paper includes several examples that show that optimal policy projection methods and Lagrangian methods are equivalent.[14] This is also true for constrained commitment (Section 2.4.2). The equivalence exists between optimal policy projections under discretion (Section 2.4.3) and Lagrangian methods for discretionary policy under some conditions. This topic is explored in detail below.

[14] However, optimal policy projections do not deliver a policy rule as a function of shocks or variables. The toolkit includes the following examples which are solved for both commitment and discretion using both methods: i) Three-equation New-Keynesian model (NK3) with a technology shock, ii) NK3 with an iid cost-push shock, iii) NK3 with a persistent cost-push shock, iv) NK3 with a discount factor shocks and the ELB, v) Sticky-wage New-Keynesian model, vi) New-Keynesian model with a hybrid Phillips Curve.

2.4.2 Constrained commitment

If there is an occasionally binding constraint on the policy instrument(s), then problem (11) is extended as follows

$$\min_{\hat{V}_0} \frac{1}{2}\hat{V}_0' \mathcal{A}_L' \Omega \mathcal{A}_L \hat{V}_0 + \left(L^B\right)' \Omega \mathcal{A}_L \hat{V}_0 \quad \text{s.t.} \quad X^{min} \leq X^B + \mathcal{A}_X \hat{V}_0 \leq X^{max}. \tag{13}$$

The constraints can take a number of different forms and can capture a number of different constraints facing a policymaker.[15]

Example 7. *Consider the setup from Example 5.*

1. *Effective lower bound, $\bar{i}$, on the policy rate or upper bound, $\bar{q}$, on asset purchases:*

$$X^{min} = vec\left(\begin{bmatrix} \bar{i} \cdot \mathbf{1}_{T+1}' \\ 0 \cdot \mathbf{1}_{T+1}' \end{bmatrix}\right), \quad X^{max} = vec\left(\begin{bmatrix} +\infty \cdot \mathbf{1}_{T+1}' \\ \bar{q} \cdot \mathbf{1}_{T+1}' \end{bmatrix}\right), \tag{14}$$

where $\mathbf{1}_n$ is a vector of ones of length n.

2. *Date-based forward guidance: Policy rate is constrained to remain at the ELB for κ periods, and to follow optimal policy thereafter:*

$$X^{min} = vec\left(\begin{bmatrix} \bar{i} \cdot \mathbf{1}_{T+1}' \\ -\infty \cdot \mathbf{1}_{T+1}' \end{bmatrix}\right), \quad X^{max} = vec\left(\begin{bmatrix} \bar{i} \cdot \mathbf{1}_{\kappa}', +\infty \cdot \mathbf{1}_{T+1-\kappa}' \\ +\infty \cdot \mathbf{1}_{T+1}' \end{bmatrix}\right) \tag{15}$$

3. *Small-deviations optimal policy:*

$$X^{min} = vec\left(\begin{bmatrix} i^B - c_i \\ q^B - c_q \end{bmatrix}\right), \quad X^{max} = vec\left(\begin{bmatrix} i^B + c_i \\ q^B + c_q \end{bmatrix}\right) \tag{16}$$

where $c_i, c_q \geq 0$.

In Example 7, *1.* incorporates an effective lower bound (ELB) on the policy rate, a non-negativity constraint on asset purchases, and an upper bound on asset purchases that may capture either quantity constraints or legal constraints. *2.* enforces the optimal path

[15] If $\bar{\Omega} \equiv \mathcal{A}_L' \Omega \mathcal{A}_L$ is positive definite then the quadratic function $\hat{V}_0' \bar{\Omega} \hat{V}_0$ is convex. In such a case, the quadratic programming problem in (13) is well defined and, if there exists a feasible solution, then there exists an optimal solution. In most cases, one would set up the problem such that the baseline (i.e., with $\hat{V}_0 = 0$) is feasible. These types of problems have been extensively studied in operations research. An interior-point method (see Gondzio 2012) is extremely efficient for solving this type of a problem and is available using MATLAB's `quadprog`. As a result, a large H does not pose a computational burden. This, however, does not hold for the case of constrained discretion, below.

to track a specific path (in this case the ELB for the policy rate) for a finite number of periods, capturing the notion of date-based forward guidance. *3.* allows the possibility that the policymaker is reluctant to make large policy changes from the baseline path. Thus, policy variables are constrained to lie within a corridor of the baseline. We will make use of these constraints in Section 4 when studying optimal policy in practice.

Finally, it is important to note that the set of periods for which the unconstrained commitment projection violates (a constraint like) the ELB is not necessarily the same as the set of periods for which the ELB strictly binds under constrained commitment. This is because the whole policy rate path is affected when the ELB is imposed.

2.4.3 Unconstrained discretion

In the absence of a commitment mechanism, the optimal policy projection described above may be time-inconsistent. To address this issue, in this section, we model the strategic interaction between policymakers at different points in time along the projection horizon. In particular, the policymaker in period-0 chooses $\hat{V}_0$ which applies from period 0 to H, the policymaker in period-1 chooses $\hat{V}_1$ which applies from 1 to H and so on. The period-H policymaker thus chooses a scalar $\hat{V}_H$.[16] Optimal time-consistent policy enforces the following set of conditions

$$\begin{aligned} \hat{V}^*_{0,(2:H-1)} &= \hat{V}^*_{1,(1:H-2)}, \\ \hat{V}^*_{1,(2:H-1)} &= \hat{V}^*_{2,(1:H-3)}, \\ &\vdots \\ \hat{V}^*_{H-1,(2)} &= \hat{V}^*_{H}, \end{aligned} \tag{17}$$

where the additional subscripts denote the elements of the vector. This says that the set of shocks chosen by the policymaker in period 0 are the same shocks as will be chosen by the policymaker that reoptimizes in period 1, and so on.[17]

[16]This timing is written from the perspective of the projection, or calendar time. From the perspective of the policymaker in period-1, for example, they are choosing a vector from 0 to $H-1$, while the policymaker in period-H is choosing a single period-0-type unanticipated surprise.

[17]This solution concept delivers a time-consistent equilibrium but is not a subgame perfect equilibrium. See Fershtman (1989) for a discussion. Subgame perfection is a stronger requirement than time consistency. Time consistency implies that the equilibrium strategies constitute an equilibrium only for subgames along the equilibrium path and not for all possible paths as the definition of subgame perfection requires. Dennis (2007) in solving for optimal discretion employs the stronger notion of time consistency, namely subgame perfection.

We use a recursive algorithm to solve the optimal policy projection under discretion. First, it is necessary to introduce additional notation. In particular, let us denote $Y^{IR}_{0:T-j} = \mathcal{A}^j_X \boldsymbol{v_{00}}$ which contains only the unanticipated period-0 policy innovations and impulse responses which run until period $T-j$, for $j = 0, \ldots, H$.

1. Set the initial guess $\hat{V}_0^{(0)}$.
2. For iteration k, calculate $L^{(k)} = L^B + \mathcal{A}_L \hat{V}_0^{(k)}$ and denote $L^{(k)}_{j:T}$ the projection from period j to T.
3. For each $j \in \{0, H\}$
   (a) Solve
   $$\min_{\boldsymbol{v_{00}}} \frac{1}{2} \boldsymbol{v_{00}'} \mathcal{A}^{j\prime}_L \Omega^j \mathcal{A}^j_L \boldsymbol{v_{00}} + L^{(k)\prime}_{j:T} \Omega^j \mathcal{A}^j_L \boldsymbol{v_{00}}, \tag{18}$$
   where $\Omega \equiv \text{diag}\left(1, \beta, \ldots, \beta^{T-j}\right) \otimes Q$, which gives
   $$\boldsymbol{v^*_{00}} = -\left(\mathcal{A}^{j\prime}_L \Omega^j \mathcal{A}^j_L\right)^{-1} \mathcal{A}^{j\prime}_L \Omega^j L^{(k)}_{j:T}. \tag{19}$$
   (b) Collect $\hat{V}^{\dagger}_{n_x j+1:n_x(j+1)} = \boldsymbol{v^*_{00}}'$.
4. If $\max\left(\left|\hat{V}_0^{\dagger}\right|\right) < \epsilon$, end. Else set $\hat{V}_0^{(k+1)} = \theta \hat{V}_0^{(k)} + (1-\theta)\hat{V}_0^{\dagger}$, where $\theta \in (0,1)$ is the updating parameter, set $k = k+1$ and return to step 2.

The intuition for the algorithm is as follows. We begin with a guess for the vector of policy announcements. Then, for each period in the projection horizon we solve for the optimal unanticipated shock. If this optimal unanticipated shock is zero, the policymaker has no incentive to deviate from the vector of policy announcements and the projection is a time-consistent equilibrium. If the optimal unanticipated shock at any period in the projection is non-zero, we update the initial guess for the vector of policy announcements.[18]

2.4.4 Constrained discretion

The algorithm for constrained discretion is the same as for unconstrained discretion, but for two minor changes. First, in Step 3(a), the additional constraint is as follows

$$X^{min}_{j:T} \leq X^{(k)}_{j:T} + \mathcal{A}^j_X \boldsymbol{v_{00}} \leq X^{max}_{j:T}. \tag{20}$$

[18] Unlike the commitment case, the discretionary problem is not necessarily convex and multiple equilibria may exist as pointed out by Blake and Kirsanova (2012). However, experimentation with alternative initial conditions for the algorithm in the examples below did not reveal multiple equilibria.

These can be the same types of constraints described in Example 7. Second, the solution in equation (19) needs to be replaced with a numerical solution.

2.4.5 Limited-time commitment

Under commitment, the policymaker commits for $D = H + 1$ periods (i.e., for the full projection horizon), whereas under discretion, policy is re-optimized every $D = 1$ periods. Under limited-time commitment, we assume that the projection horizon is made up of a sequence of non-overlapping policymakers that each commit for $1 < D < H + 1$ periods.[19] In this case, the time-consistency requirement given in (17), is generalized as follows

$$\begin{aligned} \hat{V}^*_{0,(D+1:H+1)} &= \hat{V}^*_{D,(1:H+1-D)}, \\ \hat{V}^*_{D,(D+1:H+1-D)} &= \hat{V}^*_{2D,(1:H+1-2D)}, \\ &\vdots \\ \hat{V}^*_{H+1-2D,(D+1:2D)} &= \hat{V}^*_{H+1-D}. \end{aligned} \tag{21}$$

The remainder of the algorithm for limited-time commitment follows closely the case of (un)constrained discretion.

2.5 Dampening the forward guidance puzzle

Here we sketch the methodology we use to modify expectations and attenuate the effects of future policy announcements. We closely follow de Groot and Mazelis (2020) in defining four approaches called I: Inattention, II: Credibility, III: Finite-planning horizon, and IV: Learning. The first three approaches are parameterized by a single parameter whereas the Learning approach requires two parameters. Each approach results in the impulse response loading matrix, $\mathbf{C}$, from Equation (9) being augmented as follows

$$\mathbf{C}^* = \mathbf{C} \odot Z^{(j)}, \tag{22}$$

for $j \in \{\text{I,II,III,IV}\}$, where $\odot$ denotes element-by-element multiplication and $Z^{(j)}$ is conformable with $\mathbf{C}$. Each element, z, of $Z^{(j)}$ exist in $[0, 1]$. Each z measures how much the response of an endogenous variable y at time $0 \leq t \leq h$ is attenuated by a shock announced in period-1 and to be realized in period-h. When $z = 0$, the shock has no effect on the endogenous variable in that period, and when $z = 1$, there is no attenuation. The

[19] For computational simplicity, we check whether $H+1$ is a multiple of D, and if it is not, we increase H.

foundations of $Z^{(j)}$ are found in de Groot and Mazelis (2020). Here, we simply provide a depiction of $Z^{(j)}$ in Figure 1 via a surface plot.

Figure 1: Attenuating forward guidance

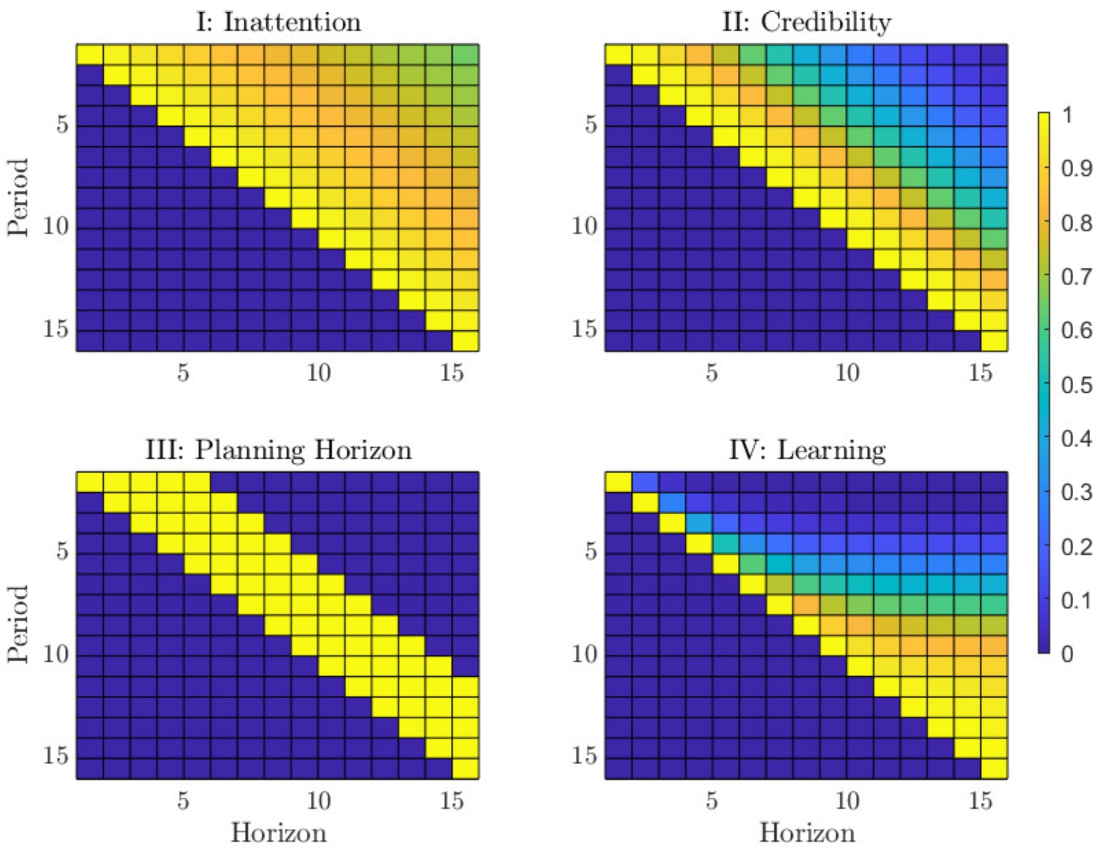

Note: Illustrative example of a Z matrix with $\alpha = 0.97$ (panel I and II), $N = 4$ (panel III), $\beta_1 = 0.5$, $\beta_2 = 5$ (panel IV).

The x-axis gives the horizon, the time at which a shock that is announced in period 1 hits, h, and the y-axis gives the current period, t. For example, the x-y quadrant [15,7] shows the attenuation in period 7 to a shock that is announced in period 1 and will be realized in period 15. The lower triangle is irrelevant given $\mathbf{C}$ is upper triangular. The colour scale depicts the values of z and ranges from 0 (dark blue) to 1 (yellow).[20] The diagonal contains the effect of the realization (when $t = h$) of each shock, hence the diagonal is always 1 (yellow).

Panel I depicts Inattention. Under Inattention, we assume a constant faction, α, of agents are fully attentive to announcements about future policy whereas $1 - \alpha$ are completely inattentive. Despite the constant fraction of inattentive agents, the effect of an announcement about a policy change farther in the future is more attenuated than

[20]The parameter values in this examples are illustrative.

the effect of an announcement about a near-term policy change.[21] This is depicted by the transition from green to yellow when moving leftwards in the panel. Moreover, as the realization of an announced policy change comes nearer, the effect of that announced policy changes becomes less attenuated. This is depicted by the transition from green to yellow when moving downwards in the panel.

Panel II depicts Credibility. Under Credibility, we assume that the fraction of agents that incorporate a policy announcement into their expectations is decaying in the horizon of the policy change. In particular, we assume α^h of agents incorporate a policy change h-periods in the future into their expectations. The attenuation pattern is qualitatively similar to Panel I except that the decay when moving towards the top-right corner of the panel is more rapid. For example, with $\alpha = 0.97$, the effect in period-1 to a policy change h-periods ahead (top-right corner) is attenuated by a factor of around 0.6 (green) under Inattention and by a factor of around 0.1 (blue) under Credibility.

Panel III depicts the results for the approach where private-sector agents have a fixed planning horizon. In this example, the planning horizon is set to 4. Thus, agents ignore any announcements that concern events more than 4 periods in the future.

Finally, Panel IV depicts Learning by private-sector agents. In period 1, forward guidance announcements are largely ignored, whether they are about policy in period 2 or 15 (the row is largely blue). However, by period 10, agents learn to understand or (pay attention) to the policymakers announcements and thus the row has turned largely yellow.

3 Examples

This section provides several demonstrations of how to use the proposed toolkit, using the canonical 3-equation New-Keynesian model of Galí (2008); Galí (2015) and single shocks to the natural rate and markup as scenarios. In our methodology the shocks are not relevant in themselves, but we use them here to create baseline projections to study the trade-offs faced by the central bank. The beauty of this set up is that it is easy to verify that the toolkit *exactly* replicates the optimal policy paths that would result from solving the first-order conditions of the policymaker's constrained optimization problem after setting up the Lagrangian.[22]

[21]This attenuation pattern can be seen with a simple example. If a fraction α of agents are attentive and form expectations as follows $m_t = \mathbb{E}_t m_{t+1}$, whereas a fraction $1 - \alpha$ of agents are inattentive and set $m_t = 0$, then by iterating forward and aggregating, we get $m_t = \alpha^h \mathbb{E}_t m_{t+h}$.

[22]The first-order conditions are derived in Appendix B. Replication code deriving optimal policy using our projection method and the Lagrangian method is available within the toolkit.

The model's social welfare function is given by

$$\sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \vartheta x_t^2 \right), \tag{23}$$

and the private sector conditions are given by

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + \epsilon_{p,t}, \tag{24}$$

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma} \left(i_t - \mathbb{E}_t \pi_{t+1} - r_{n,t} \right), \tag{25}$$

where both the cost-push shock, $\epsilon_{p,t}$, and the natural rate of interest, $r_{n,t}$, follow first-order autoregressive processes. To generate the baseline scenarios, the model is closed with a policy rule given by

$$i_t = r_n + \phi_\pi \pi_t + \phi_x x_t + \epsilon_{r,t}. \tag{26}$$

The assigned parameter values are as follows: $\sigma = 1$, $\varphi = 5$, $\theta = 3/4$, $\beta = 0.99$, $\bar{\alpha} = 1/4$, $\epsilon = 9$, $r_n = 1$, $\phi_\pi = 1.5$, and $\phi_x = 0.5/4$, where $\omega = (1 - \bar{\alpha}) / (1 - \bar{\alpha} + \bar{\alpha}\epsilon)$, and $\lambda = (1 - \theta)(1 - \beta\theta)\,\omega/\theta$, $\kappa = \lambda \left(\sigma + (\varphi + \bar{\alpha}) / (1 - \bar{\alpha}) \right)$, $\vartheta = \kappa/\epsilon$.

3.1 Classic results

Natural rate shocks. The baseline scenario in Figure 2 (solid black line) is the response to a 2% annualized natural rate shock with $\rho_{rn} = 0.9$ when policy is governed by (26). In this scenario the "divine coincidence" holds and optimal policy (dashed blue line) can close both the inflation and output gaps by allowing the policy rate to track the natural rate. Since optimal policy is time-consistent in this setting, the optimal policy paths for commitment and discretion are equivalent.

Cost-push shocks. To demonstrate that the toolkit can replicate optimal policy in a scenario in which the divine coincidence does not hold, Figure 3 exactly replicates the persistent cost-push shock from Galí (2008) Figure 5.1. Under discretion, the policymaker attempts to stabilize the medium-term output gap more than under commitment because the policymaker does not internalize the benefits in terms of near-term stability that results from allowing larger deviations of the output gap in the medium-term—the so-called "stabilization bias" associated with discretionary policy.[23]

[23] Appendix B, Figure C.2 extends the canonical model to include sticky wages and reproduces the results from Galí (2008) Figure 6.4 in response to technology shocks.

Figure 2: Natural rate shocks and the "divine coincidence"

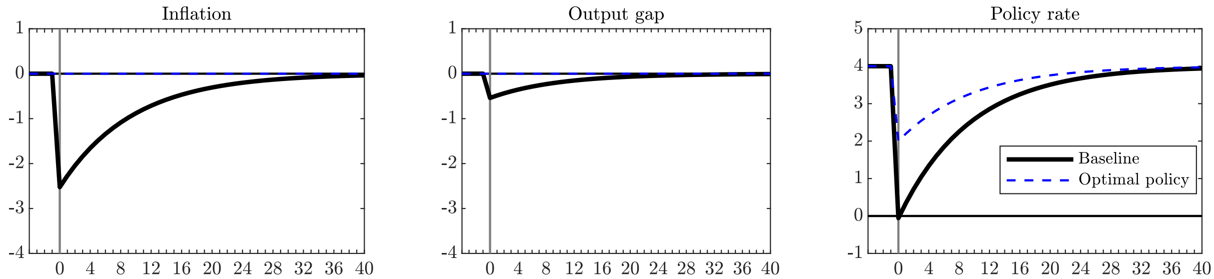

Note: All variables are given in percent. Inflation and the policy rate are annualized. The scenario is a 2% annualized natural rate shock with $\rho_{rn} = 0.9$. The "Optimal policy" line refers to both commitment and discretion since the two are equivalent in this scenario.

Figure 3: Cost-push shocks and "stabilization bias"

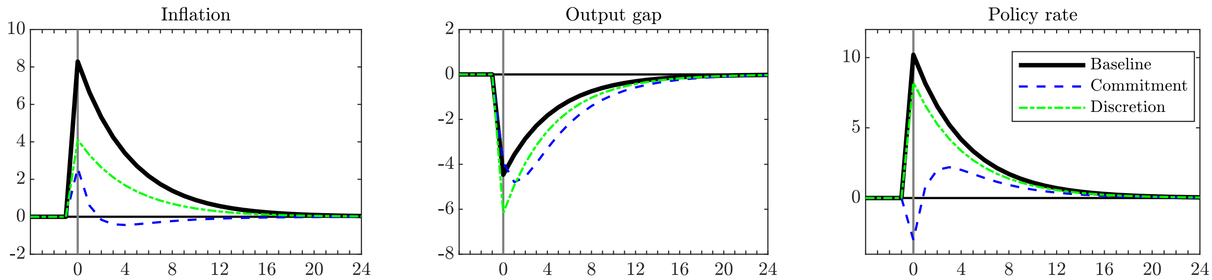

Note: All variables are given in percent. Inflation and the policy rate are annualized. The scenario is a 1% cost-push shock with $\rho_{\epsilon p} = 0.8$.

The zero lower bound. The previous two examples are linear-quadratic. Figure 4 introduces an occasionally binding constraint by revisiting the effect of a (very large) natural rate shock in the presence of the zero lower bound (ZLB).[24] Under the baseline, the ZLB binds for 5 quarters. Discretion improves upon the baseline by reducing the fall in inflation and output. A Taylor rule could replicate this discretionary outcome with a larger coefficient on inflation.[25] However, the striking improvement results from commitment, which is characterized by a promise to hold the policy rate at the ZLB for longer (increasing from 5 to 8 quarters) and a promise of an overshoot of inflation in the medium-term which lowers the ex-ante real rate, boosting aggregate demand.

[24]Computational times from the toolkit are reported Table D.1. This figure (with $H = 60$) took 2 seconds for commitment and 27 seconds for discretion. In general, the code is very efficient for calculating both unconstrained and constrained commitment. The computational demands are noticeably larger for discretion and is increasing in H. Lowering H to 40 and 20 reduces the computational time for discretion to 20 seconds and 9 seconds, respectively. The most computationally demanding optimal policy projection in this paper is Figure 13, which uses the Sims and Wu (2020) model, features two policy instruments, policy constraints, discretion, and $H = 40$ and takes 201 seconds.

[25]Setting $\phi_\pi = 150$ generates a Taylor-rule outcome similar to the discretionary outcome.

Figure 4: Natural rate shocks and the zero lower bound

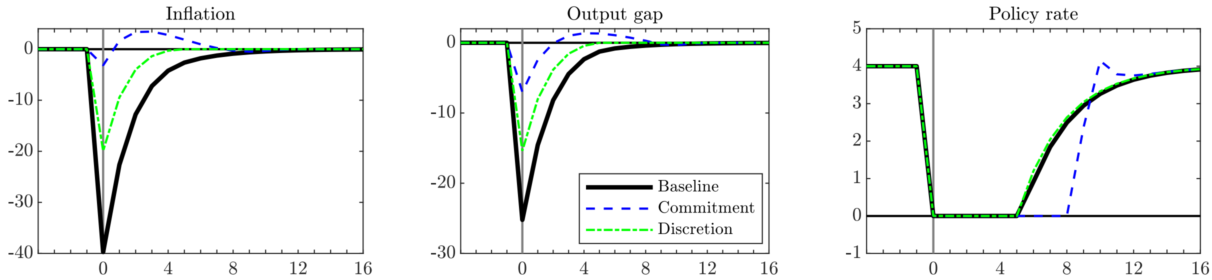

Note: All variables are given in percent. Inflation and the policy rate are annualized. The scenario is a 24% natural rate shock shock with $\rho_{rn} = 0.7$. With the baseline loss normalized to 1, the loss under commitment and discretion is 0.0253 and 0.2296, respectively.

3.2 Extensions

The toolkit is thus able to replicate classic results from the monetary policy literature. However, this toolkit allows the researcher to explore several further types of "constrained" optimal policy projections that provide valuable insight into the trade-offs faced by a policymaker. Here we highlight four.

Short H. As discussed in the previous section, the toolkit approximates optimal policy by replacing an infinite sequence of policy announcements with a finite horizon, H. The rule-of-thumb is to find the minimum horizon $\tilde{H}$ at which $\max \left| \left\{ L^B_{\tilde{H}}, X^B_{\tilde{H}} \right\} \right| <$ tol and ensure that $H >> \tilde{H}$. In the preceding examples, we set $H = 60$. Figure 5 demonstrates the consequence of H set too low in the natural rate shock scenario of Figure 2 under commitment. In particular, one observes that the policymaker sharply adjusts policy in period $H - 1$.[26]

However, despite being "incorrect", these simulations provide valuable insight into the trade-off faced by the policymaker. First, consider $H = 1$ (dashed blue line in Figure 5), in which the policymaker optimizes given just one deviation of policy in period 0. In this scenario, it is never optimal for both inflation and the output gap to be below target. The policymaker cuts the policy rate sharply, mitigating the fall in inflation and raising the output gap above target in period 0. Second, consider $H = 2$ (dot-dashed aqua line), in which the policymaker can both deviate from the policy rule in period 0 and 1. Again, the policymaker has the same incentive to cut the policy rate sharply in period 1, mitigating the period 1 fall in inflation and raising the output gap above target. However, this worsens the trade-off in period 0, raising the output gap too far above target. Thus, the optimal solution is to loosen policy aggressively in period 1 and fine tune policy in period 0. This process continues as we increase H.

[26]Since the canonical New-Keynesian model has no endogenous state variables, the economy returns to the baseline immediately in period H.

Figure 5: Short H

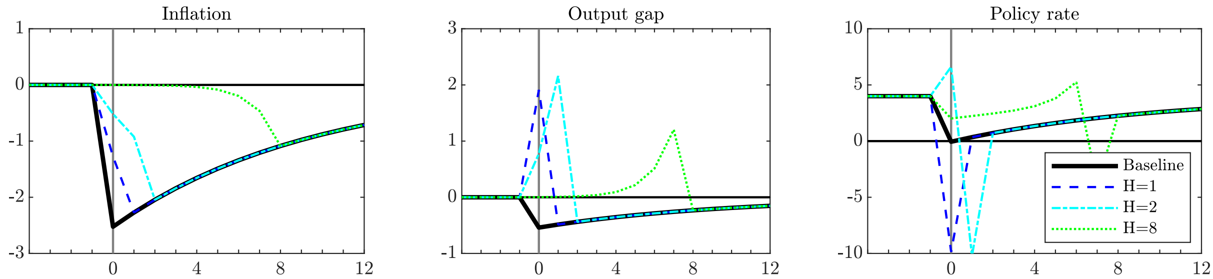

Note: All variables are given in percent. Inflation and the policy rate are annualized. The scenario is a 24% natural rate shock shock with $\rho_{rn} = 0.7$.

Inattention. The previous projections have been constructed under the assumption of perfect attention. As a result, the effect of monetary policy announcements can be incredibly large—a manifestation of the forward guidance puzzle. Here, we examine the (type I) approach to dampen the forward guidance puzzle based on inattention.

Figure 6 demonstrates the effect of reducing the attention parameter, α, from 1 (dashed blue line) to 0.6 (dot-dashed aqua line) to 0 (dotted green line). Under commitment (top-row), there is a non-monotonic relationship between the time at the ZLB and the inattention parameter. Relative to the full attention case, when $\alpha = 0.6$, the policymaker promises a longer time at the ZLB (14 vs. 8 quarters). However, as α falls further, the time at the ZLB shrinks again. This is because for low values of α the promise to stay at the ZLB for longer and create an overshooting of inflation brings little or no benefit in period 0. But the overshooting to be delivered in the future when the ZLB episode is over is costly. This unfavourable trade-off makes it optimal to shrink the time the policymaker promises to keep rates at the ZLB. In contrast, under discretion (bottom-row) the period at the ZLB is monotonically decreasing in α. In the extreme case when $\alpha = 0$, the projected paths under commitment and discretion coincide.[27]

Small deviations. Another attractive feature of the toolkit is the ability to constrain the optimal policy projection to remain within a corridor of the baseline. Figure 7 presents an example in which the maximum deviation allowed from the baseline policy rate path is one percentage point (annualized). This experiment is interesting because it shows that a strikingly different interest rate path can generate near identical economic outcomes in terms of inflation and the output gap. Under commitment, the policy rate remains at the

[27]When $\alpha < 1$, the expected path of the policy rate (i.e., the path expected by the private sector) in any period is higher than the path promised by the policymaker (i.e., the path that is plotted). This explains why the outcomes in terms of inflation and the output gap are worse despite a lower-for-longer policy rate path. When $\alpha = 0$, the outcome is equivalent to the case in which the policymaker is only able to use unanticipated monetary policy shocks to deviate from the baseline.

Figure 6: Type I Inattention

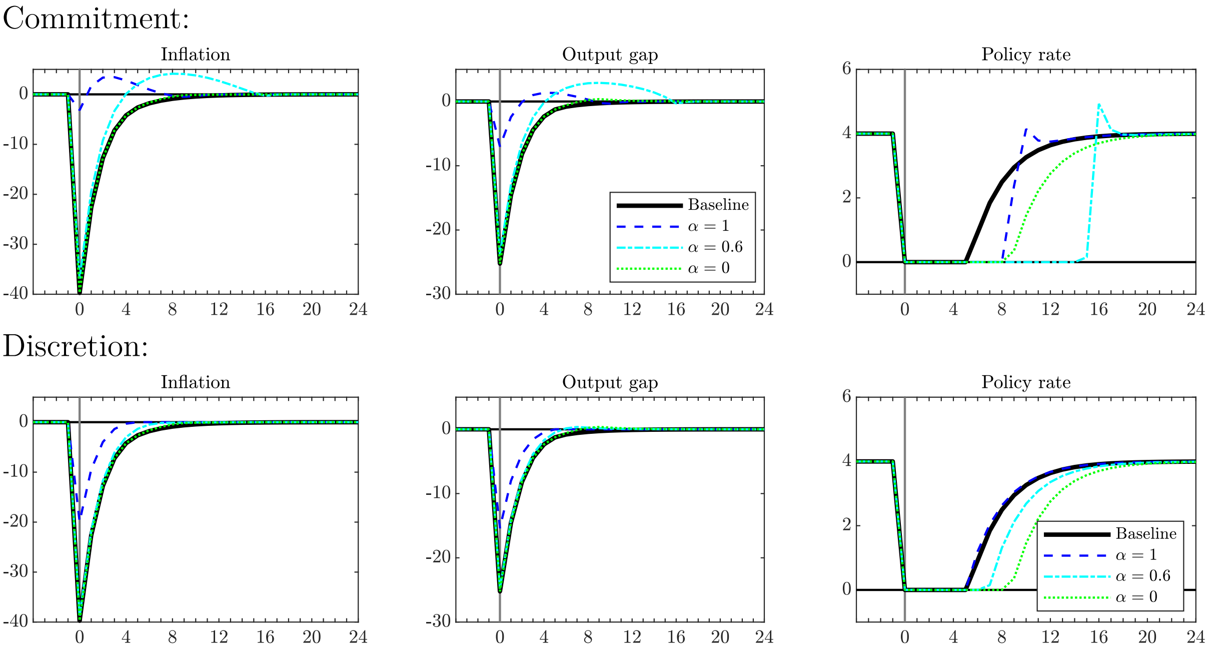

Note: All variables are given in percent. Inflation and the policy rate are annualized. The scenario is a 24% natural rate shock shock with $\rho_{rn} = 0.7$.

ZLB for an extended period of time before a sharp tightening. The alternative projection that represents a small deviation from the baseline is one in which the policymaker exits the ZLB earlier but tightens more gradually.[28] Thus, the toolkit reveals the flatness of social welfare across alternative policy paths, an insight not readily accessible from Lagrangian-based optimal policy simulations.

Figure 7: Small deviations

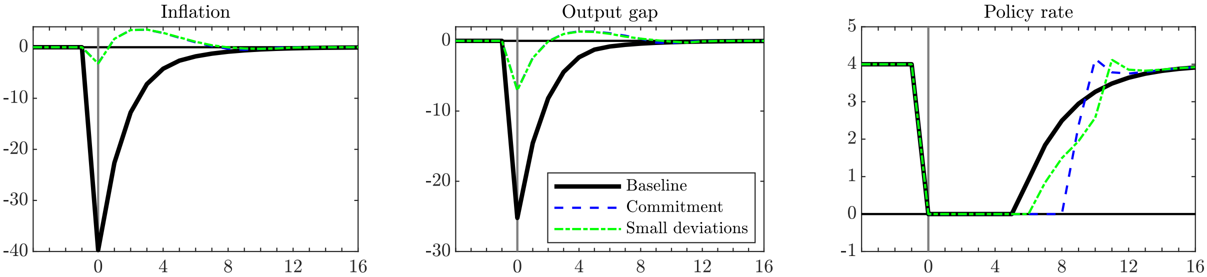

Note: All variables are given in percent. Inflation and the policy rate are annualized. The scenario is a 24% natural rate shock shock with $\rho_{rn} = 0.7$.

Fixed terms. Thus far, we have considered the case in which the policymaker either commits for H periods or else re-optimizes every period. Instead, we can imagine a

[28]The result follows from the IS equation (25). By solving (25) forward, one can derive that the output gap, x_t, is proportional to the undiscounted sum of future expected real interest rates. With inattention, this no longer holds (see Appendix Figure C.3). In terms of the losses, with the baseline loss normalized to 1, the loss under commitment is 0.0253 and the "small deviation" loss is 0.0256.

scenario in which a sequence of policymakers each serve a fixed term of D periods, where $H \pmod D = 0$. Each policymaker re-optimize at the start of their term but can fully commit for the remainder of the term.

Figure 8 displays the results for $D = \{6, 9\}$. $D = 6$ performs little better than the discretionary ($D = 1$) outcome since the ZLB binds in the baseline and the $D = 6$ policymaker is not able to commit to a longer period at the ZLB. In contrast, for $D = 9$ in this example, policy is able to closely match the full commitment paths for inflation and the output gap since it is able to replicate the additional periods at the ZLB.Thus, the toolkit helps pin down the exact location of the time-inconsistency in the projection under commitment.

Figure 8: Fixed term

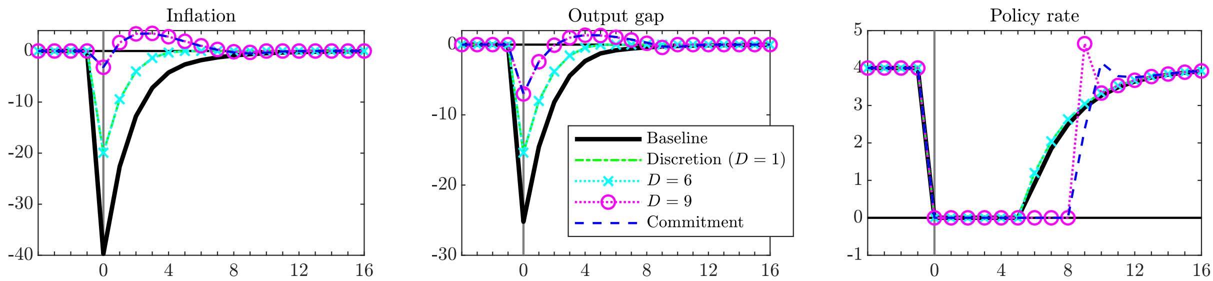

Note: All variables are given in percent. Inflation and the policy rate are annualized. The scenario is a 24% natural rate shock shock with $\rho_{rn} = 0.7$.

4 Optimal policy for the US: An illustration

This section illustrates the use of the toolkit in an application to the US economy. We focus on the historical episode of the Great Financial Crisis, which resulted in a severe economic downturn and the use of QE for the first time by the Fed.

4.1 Data and models

To implement optimal policy we need three ingredients. First, we need baseline projections. We put ourselves in the shoes of policymakers in early 2009 and assume a corresponding information set with macroeconomic projections and policy expectations from the Fed staff's long-term outlook published in the March 2009 Greenbook.[29] This scenario will form the baseline for our optimal policy projections.

Second, we need impulse responses to policy shocks. We consider two models. One is the seminal Smets and Wouters (2007) model (henceforth SW07) that has been estimated

[29] Available at www.federalreserve.gov/monetarypolicy/fomc_historical.htm.

for the US economy. Although this model captures key dynamics, the short-term policy rate is the only policy instrument meaning there is limited applicability to the episode in question. To capture the extended monetary toolkit, we also use the Sims and Wu (2020) model (henceforth SWu20) that includes a role for QE.

Third, the policymaker's loss function. We assume a quadratic loss function given by

$$\frac{1}{2}\mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \left\{ (\pi_t^a - \pi^*)^2 + \lambda_x x_t^2 + \lambda_{dr} \left(\Delta r_t^a\right)^2 + \lambda_q q_t^2 + \lambda_{dq} \left(\Delta q_t^2\right) \right\}, \tag{27}$$

where the loss function weights are given in Table 1. As is common in the literature (e.g., Svensson and Tetlow 2005; Debortoli et al. 2019) and in practice (e.g., Yellen 2012), the policymaker cares about keeping annualized inflation, π_t^a, at $\pi^* = 2\%$, closing the output gap, x_t, and making gradual changes to the policy rate, Δr_t^a.[30] Although the real side of the Fed's dual mandate is formulated in terms of maximum employment, we focus on the output gap in the absence of an unemployment gap in SW07. We can translate the unemployment gap into the output gap using Okun's law. We use an Okun's law coefficient of 2. Assuming an equal weighting between these two objectives gives the value for λ_x.[31]

In the SWu20 model that incorporates QE, we extend the loss function by a preference for keeping asset purchases, q_t, to a minimum and making gradual changes to the asset purchase portfolio, Δq_t. The weights on the level and change in asset purchases are based on Harrison et al. (2020). To ensure that policy announcements in the optimal policy

Table 1: Loss function weights

λ_x	λ_{dr}	λ_q	λ_{dq}
0.25	1	0.57	4.58

projections do not suffer from the forward guidance puzzle, we use the type I approach to mitigate the puzzle by assuming a constant share of inattentive agents. Based on empirical findings in Åhl (2017), de Groot and Mazelis (2020) calculate that a 30% share of inattentive agents approximates the data well. This value is in line with de Groot et al. (2020), who estimate the share of inattentive agents based on euro area data, and Christoffel et al. (2020), who calibrate the reaction of the New Area Wide Model to forward guidance shocks to be in line with empirical evidence.

[30] These preferences require a corresponding modification of steady state inflation from the estimated mode in SW07 and the calibration in SWu20.

[31] As is customary in the literature, policymakers are also assumed to care about limiting changes in the policy rate to avoid volatility in the federal funds rate.

4.2 Optimal unconstrained federal funds rate policy projections

The baseline projection is shown by the black line in Figure 9. The projections from 2009q2 to 2013q4 are taken from the Greenbook. Beyond 2013, we assume the inflation and output gaps will be closed by 2015. We then use the Kalman filter to solve for the transition paths for inflation and the output gap until the effects of shocks have dissipated and variables have converged to the steady state. In the staff projection, inflation is projected to remain well below 2% until the end of 2013. Notwithstanding the projected closing of the output gap, economic slack remains at the end of 2013. The (federal funds) policy interest rate stays at its effective lower bound (ELB) throughout the staff projection period. To illustrate the use of the COPPs toolkit, we allow policy to diverge from the baseline after 2009q1.

Figure 9 presents optimal policy projections using the SW07 model, under the assumption of commitment and without the ELB binding. The standard SW07 model (red-dash line) prescribes a reduction in the policy rate to −4%, resulting in inflation rising to over 4% in 2010, well above the inflation target.[32] Optimal policy allows this overshoot to lower the real interest rate and reduce the large negative output gap. Accommodation is withdrawn and the policy rate returns to positive values in 2011.

The standard SW07 model, however, exhibits elasticities that are significantly different from those in FRB/US (Laforte 2018), the Fed's workhorse model for policy simulations. A comparison of impulse responses to a standard monetary policy shock in both models confirms that the reaction of inflation in SW07 is several times larger than in FRB/US (see appendix Figure D.1). This is due to the relatively steep price and wage Phillips curves in the original SW07 estimation, which appears to be driven by the sample period extending to 1966.[33] To acknowledge the flattening of the price and wage Phillips curves, we calibrate the Calvo parameters governing the frequency of price and wage re-optimization to deliver results more in line with evidence from FRB/US.[34] This results in a muted reaction of inflation to monetary policy shocks and an increased sensitivity of the output gap.

The blue-dash line in Figure 9 shows the optimal policy projections of the adjusted SW07 with a flatter Phillips curve. This model version prescribes an even larger fall in the policy interest rate to −6%. The additional easing is necessary as inflation reacts more moderately while the output gap closes more quickly. The results are comparable to

[32] The inflation target is indicated by the dotted line at 2%. The dotted line displays the steady state value and also serves as the variables' target.

[33] This is also pointed out by Smets and Wouters (2007) in a robustness analysis that separately estimates the "Great Inflation" and "Great Moderation" sub-samples.

[34] This is achieved by increasing both parameters by 0.15.

Figure 9: Unconstrained optimal policy under commitment

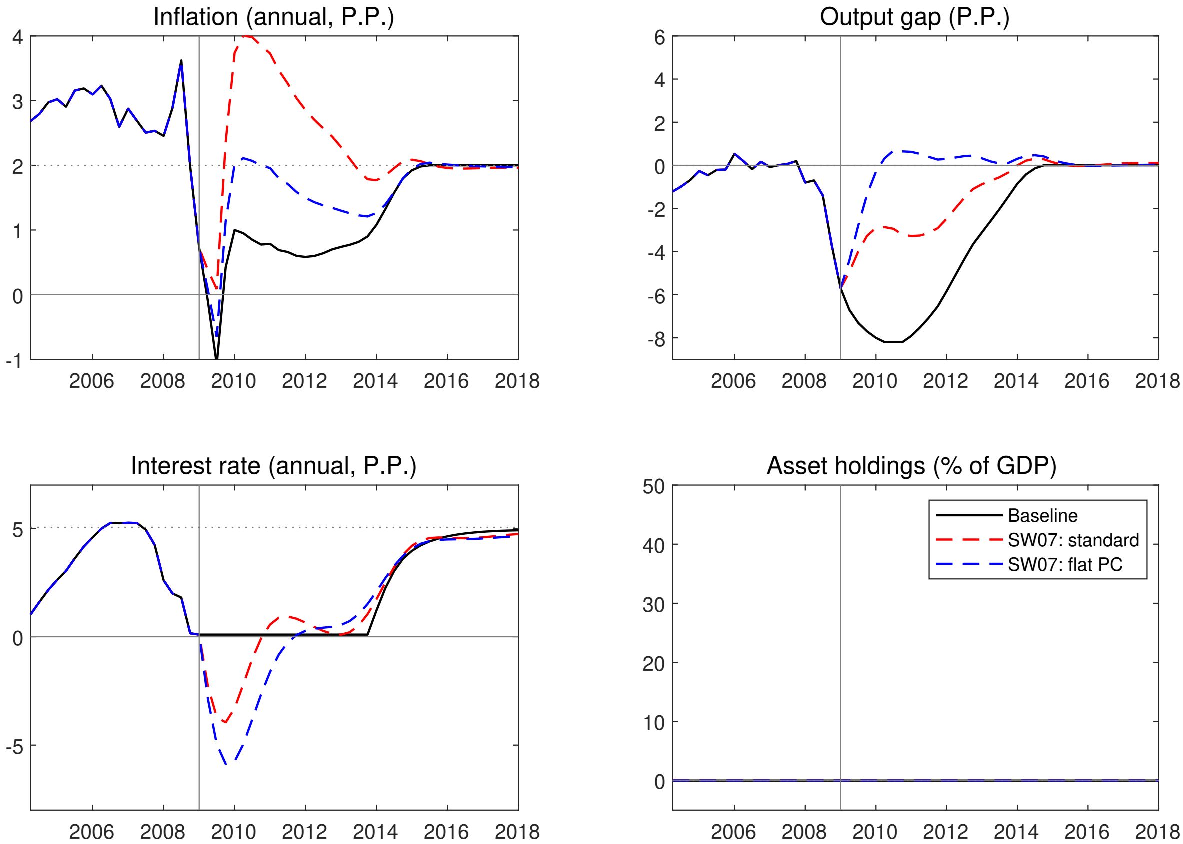

Note: Toolkit options: `Policy type = Commitment`. `Instruments = R`. `ELB constraints = Off`. `FG puzzle mitigation = On`. `Mitigation type = I`. `Mitigation parameter:` $\alpha = 0.7$.

the optimal policy simulations by the Fed staff in the March 2009 Greenbook, which also prescribes a fall in the policy rate to around -6% in the absence of the ELB.

4.3 Imposing the effective lower bound

Figure 10 presents constrained optimal policy projections using the adjusted SW07 model with the ELB imposed. The ELB prevents further rate cuts. Instead the policymaker can announce a longer stay at the ELB than projected in the baseline. The effectiveness of such an announcement depends on the degree of attention of agents to this guidance.

Given the fraction of inattentive agents, the blue-dash line shows the ELB constraint on policy results in lift-off from the ELB two quarters later than the baseline. There are limited effects on both inflation and the output gap relative to the baseline. This is due to a combination of limited attention to forward guidance, the weight assigned to output gap stabilization, and a flat forward curve in the baseline.

The red-dash line plots the optimal policy projection when all agents are attentive. Given that forward guidance is not dampened, the interest rate can be increased from the

Figure 10: Constrained optimal policy under commitment

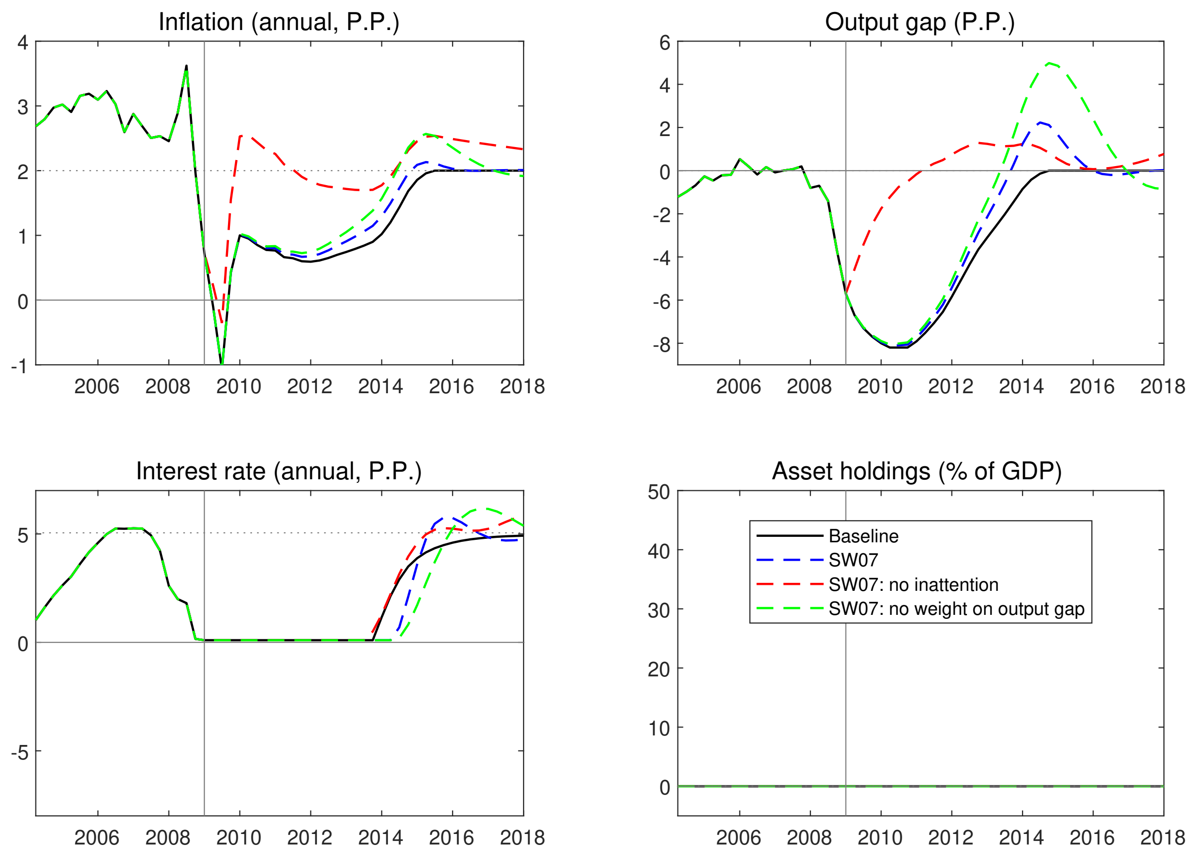

Note: Toolkit options: `Policy type = Commitment`. `Instruments = R`. `ELB constraints = On`. With the baseline loss normalized to 1, the loss of the blue, red, and green paths are 0.95, 0.14, and 0.86, respectively.

ELB one quarter earlier relative to the baseline. Inflation initially overshoots the target by half a percentage point and the output gap is closed 3.5 years earlier. These strong reactions are a powerful example of the forward guidance puzzle that befalls this class of models. The dampened attention to forward guidance emerges as an indispensable and practical tool to improve the performance of the simulations.

The green-dash line plots the effect of the policymaker placing no weight on closing the output gap. This change in loss function results in a later lift-off from the ELB and a slower convergence of the policy rate back to steady state. While inflation is unchanged in the near-term due to inattention, the additional forward guidance becomes effective in the medium-term, well before the prolonged stay at the ELB is actually implemented. The result is an increase in inflation of 50bp.

Had agents expected a steeper forward curve in the near-term, optimal policy would have had additional room for easing via forward guidance despite inattention to policy announcements. In this counterfactual, the short-term rate dynamics in the baseline are constructed via the Kalman filter without a forecast, which results in a gradual increase in the baseline policy path of about 75bps per year (see appendix Figure D.2). Optimal policy would have prescribed a stay at the ELB until late 2013. Such an announcement would have increased inflation by up to 25bps, and closed the output gap by early 2012, 2.5 years earlier than in the baseline.

4.4 Allowing QE

We now turn to the SWu20 model. This means we can analyze the effect of QE as an additional monetary policy instrument and to study the interaction of the policy rate instrument with QE.

We provide only a brief description of features that allow QE to have real effects in the model. First, firms in the model are required to issue long-term bonds to finance part of investment in new physical capital. Second, asset markets are segmented in that households can only indirectly access long-term bonds by holding short-term deposits with banks. Banks are introduced in a similar way to Gertler and Karadi (2011). A costly enforcement problem results in an endogenous leverage constraint on banks and a time-varying interest rate spread. QE is interpreted as central bank purchases of government bonds. QE has real effects in the model to the extent that banks are constrained by the costly enforcement problem. When banks are constrained, asset purchases ease this constraint in such a way that the total demand for bonds increases (and therefore purchases do not simply crowd out intermediary bond purchases). This results in higher bond prices, easing the loan-in-advance constraint facing the firms, and leads to higher investment and aggregate demand.

The toolkit allows us to restrict asset purchases to stay within limits to reflect operational constraints. We set two such constraints. First, we restrict the total stock of purchased assets as a percentage of GDP to be between 0% and 50%. Second, we restrict the flow of asset purchases to grow no faster than 3 percentage points per quarter from the initial period until the 50% maximum is reached.

Figure 11 presents the optimal policy projections from the SWu20 model. As with the SW07, we adjust the Calvo parameters on price and wage setting in SWu20 to be comparable with FRB/US.[35] This ensures that when limited to a single policy instrument (see the blue dash-dot line), the optimal policymaker in SWu20 prescribes a similar path to the blue-dash line in Figure 10. We consider this adjusted parameterisation as the baseline SWu20 in the subsequent simulations.

[35] This is achieved via an increase of both parameters by 0.18.

Figure 11: Two instrument optimal policy under commitment

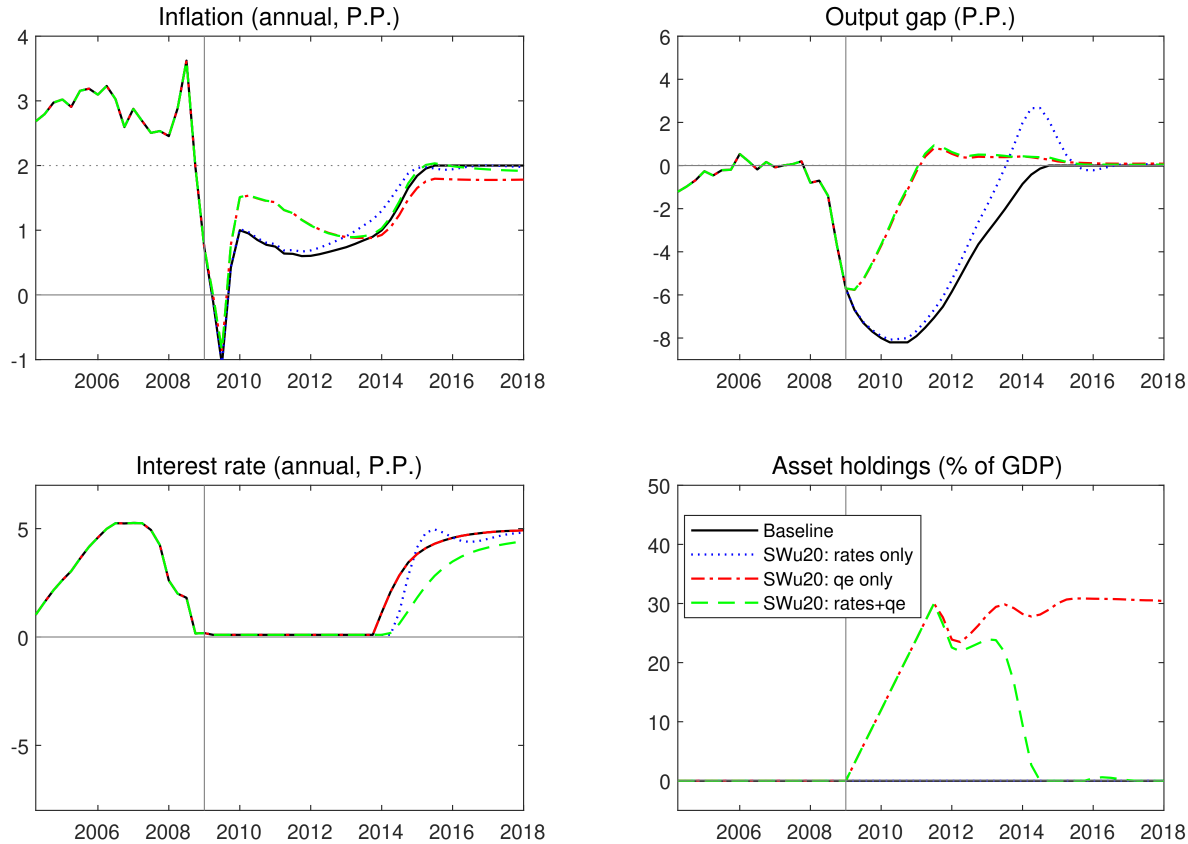

Note: Toolkit options: `Policy type = Commitment.` `Instruments = R & Q.` `ELB constraints = On.` `FG puzzle mitigation = On.` `Mitigation type = I.` `Mitigation parameter:` $\alpha = 0.7$. With the baseline loss normalized to 1, the loss of the blue and green paths are 0.93 and 0.25, respectively.

The red dash-dot line in Figure 11 plots the prescribed optimal policy path in which QE is set optimally but the policy rate is constrained to follow the baseline. In this scenario, QE is used aggressively, with total asset holdings announced to increase steadily for the first 10 quarters up to a total of more than 30% of GDP. QE is very effective. The additional stimulus allows the output gap to close 3.5 years earlier than the baseline. At this point, it is optimal for the policymaker to maintain a persistently large balance sheet. Inflation dynamics are not very responsive in this scenario, with near-term inflation receiving a small but transitory boost.

The green dash-dot line in Figure 11 plots the optimal policy path in which the policy rate and QE are set jointly. Although the macroeconomic consequences in terms of inflation and the output gap are similar in this scenario, the paths of the policy instruments are strikingly different. In this case, the policymaker reduces its asset holdings as soon as the output gap is closed. The tightening effect of a reduction in asset holdings is offset by the policymaker raising the policy interest rate more gradually.

This is a reflection of the relative properties of the two instruments. They have a different impact on inflation relative to GDP, and for this reason the instrument that is better at stabilising the variable that is further away from the target is selected. As seen in the impulse responses in Figure 12, the policy rate is better at stabilising inflation relative to GDP than QE. As a result, once the output gap is closed (green-dashed line in Figure 11), it is optimal to contract the Fed's balance sheet and use the policy rate to compensate and get inflation back to target.

Figure 12: Impulse responses to policy rate and asset purchase shocks

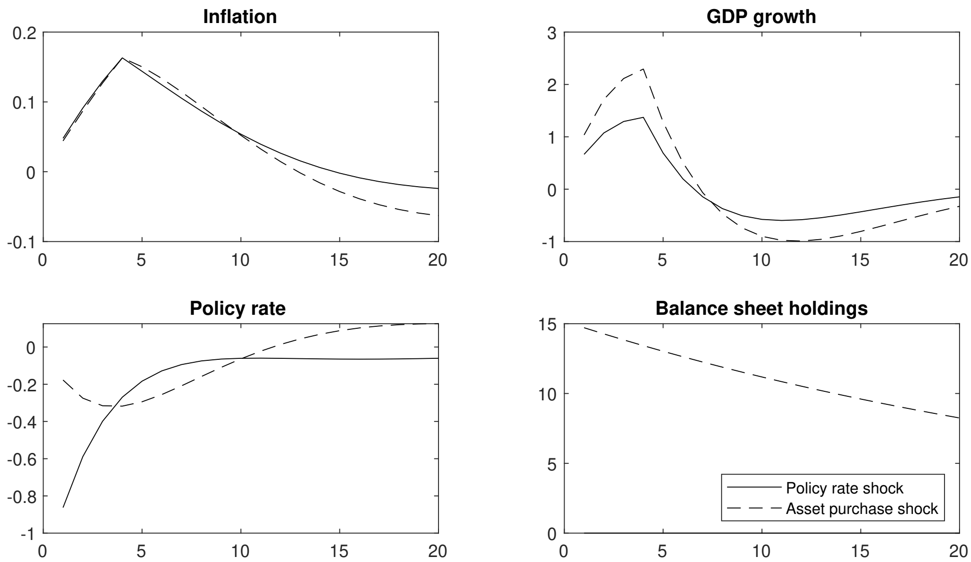

Note: SWu20 model. The IRFs are scaled to have the same maximum impact on inflation.

4.5 Optimal policy under discretion

The scenarios thus far have assumed that policymakers are able to fully commit to future policy actions. In this section we consider the case in which policymakers set policy optimally under discretion. If policymakers only have the standard policy interest rate available as an instrument, they will optimally choose not to deviate from the baseline (see blue dash-dot line in Figure 13). This is because i) the policymaker cannot credibly promise to overshoot inflation in the future to reduce the ex-ante real interest rates today, and ii) by the time interest rates are set to tighten in the baseline, both inflation and the output gap are already close to target.

Adding QE to the instrument mix shows that asset holdings are increased less aggressively and phased out more gradually than in the commitment case, resulting in more easing via QE (see green dash-dot line in Figure 13). The interest rate is increased earlier than in the baseline to limit the overshoot in the output gap, which comes at a cost in terms

Figure 13: Optimal policy under discretion

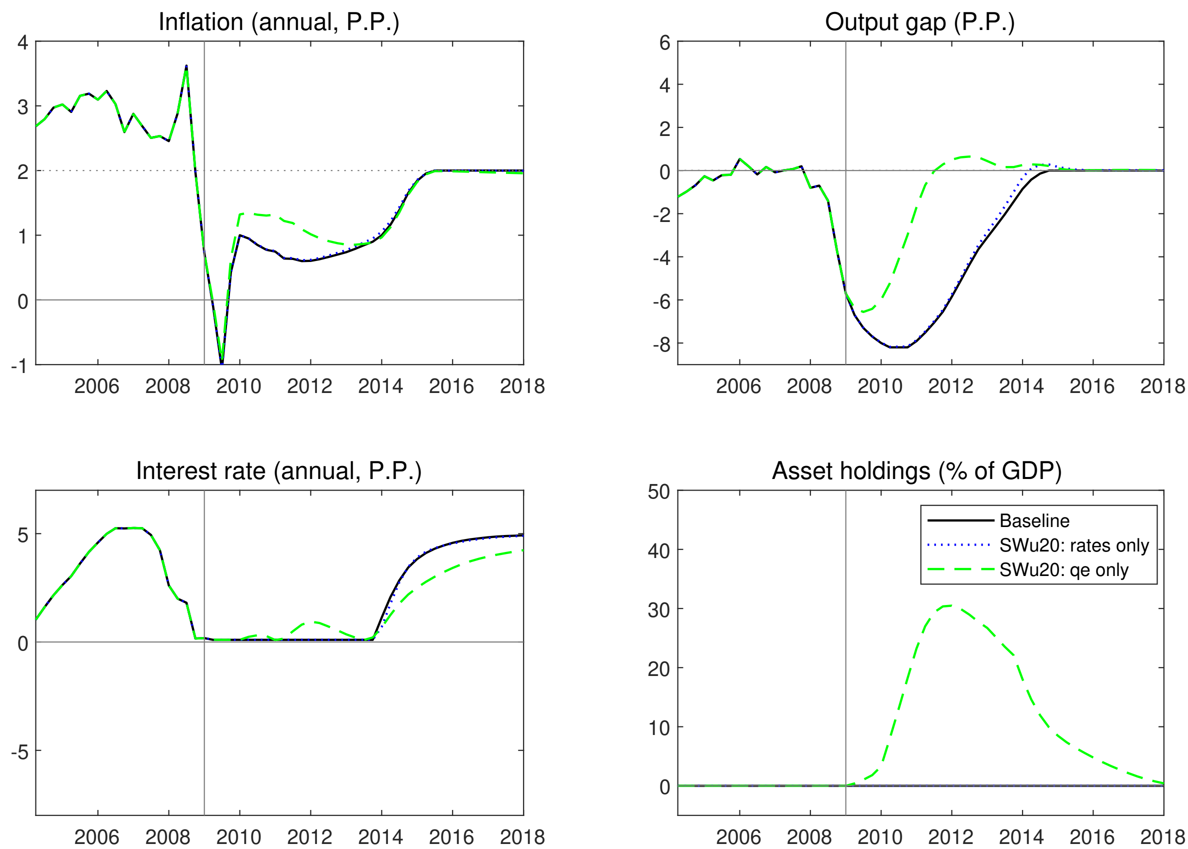

Note: Toolkit options selected: `Policy type = Discretion. Instruments = R & Q. ELB constraints = On. FG puzzle mitigation = On. Mitigation type = I. Mitigation parameter:` $\alpha = 0.7$. With the baseline loss normalized to 1, the loss of the blue and green paths are 0.98 and 0.43, respectively.

of the price stability objective. The greater persistence of the Fed's balance sheet under discretion vs. commitment is because the Fed cannot credibly commit to a lower path of interest rates to compensate for the contractionary effects of reducing the balance sheet.

4.6 Ex-post evaluation of Fed policy

Optimal policy projections are designed to find the optimal paths for the policy instruments conditional on a given baseline. If the baseline cannot be improved upon, optimal policy projections will be similar to the baseline projections. Alternatively, the deviation between the optimal and baseline projections gives a quantification of the failure of the baseline projections of the instruments to achieve optimality.

The March 2009 FOMC statement announced that the "Committee will maintain the target range for the federal funds rate at 0 to 1/4 percent and anticipates that economic conditions are likely to warrant exceptionally low levels of the federal funds rate for an extended period."[36] Our optimal policy projections call for lift-off of the policy rate around the same time as the path that was assumed in the Greenbook and that we use as the baseline projection. This means that the use of the policy rate instrument was consistent with optimality. This is because the baseline projection for the interest rate features four years at the ELB. Had the Fed announced in early 2009 an even longer period of low rates, this would not have provided additional support to the economy in the near-term. This is a consequence of our assumption, consistent with empirical evidence, that some agents are inattentive to forward guidance announcements. This degrades the value of policy commitments covering the distant future.

The activation of a further policy tool in the form of QE was therefore warranted at the time. But according to our optimal policy projections its quantitative deployment appears insufficient: while the FOMC statement committed to an increase of $300bn in treasuries and $750bn in mortgage-backed securities by the end of the year, this only amounts to around 11% of Fed holdings as a share of GDP. While our optimal projections arrive at similar figures for holdings by the end of 2009, they continue thereafter to increase to 30%. Optimal policy projections suggest that the economy would have recovered more quickly had a larger QE program been announced in 2009. The Fed eventually increased its balance sheet beyond the figures announced in March 2009 to reach around 25% in 2014, but this should have been announced earlier on.

5 Conclusion

In this paper, we propose an extension of the optimal policy projection methodology to incorporate constraints, limited-time commitment, discretionary policy, multiple instruments, and various methods for mitigating the forward guidance puzzle. A `Dynare`-compatible toolkit accompanies this paper providing an efficient and easy-to-use method for computing a large variety of optimal policy projections exercises.

The methodology and toolkit rely on using impulse responses to anticipated policy shocks to construct optimal policy projections around a baseline. As a result, this approach does not require a specific model's structural equations or to know the shocks underlying the baseline projection. The advantage is that it makes it easy to cross-check robustness of the optimal policy projections across different sets of impulse responses, representing uncertainty about the monetary policy transmissions. Using the toolkit to study robustly optimal policy is likely to be a fruitful direction for future research.

[36] Available at www.federalreserve.gov/newsevents/pressreleases/monetary20090318a.htm.

We present several examples and apply the toolkit to study the US Federal Reserve's policy at the height of the 2008-09 financial crisis, using two models and two policy instruments. Moreover, the toolkit lends itself to studying optimal policy in real-time both in policy settings and for research.

References

Adjemian, S. and M. Juillard (2011): "Accuracy of the extended path simulation method in a new Keynesian model with zero lower bound on the nominal interest rate," *Unpublished manuscript.*

Åhl, M. (2017): "How big is the toolbox of a central banker? Managing expectations with policy-rate forecasts: Evidence from Sweden," *Unpublished manuscript.*

Barnichon, R. and G. Mesters (2020): "Optimal policy perturbations," *Unpublished manuscript.*

Blake, A. P. and T. Kirsanova (2012): "Discretionary policy and multiple equilibria in LQ RE models," *Review of Economic Studies*, 79, 1309–1339.

Christoffel, K., O. de Groot, F. Mazelis, and C. Montes-Galdón (2020): "Using forecast-augmented VAR evidence to dampen the forward guidance puzzle," *ECB working paper 2495.*

Clymo, A. and A. Lanteri (2020): "Fiscal Policy with Limited-Time Commitment," *The Economic Journal*, 130, 623–652.

Cole, S. J. (2020): "The limits of central bank forward guidance under learning," *International Journal of Central Banking.*

D'Amico, S. and T. B. King (2015): "What does anticipated monetary policy do?" *Unpublished manuscript.*

de Groot, O. and F. Mazelis (2020): "Mitigating the forward guidance puzzle: inattention, credibility, finite planning horizons and learning," *ECB working paper 2426.*

de Groot, O., F. Mazelis, R. Motto, and A. Ristiniemi (2020): "Estimating the Strength of Forward Guidance Announcements," *Unpublished manuscript.*

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

Debortoli, D. and R. Nunes (2010): "Fiscal policy under loose commitment," *Journal of Economic Theory*, 145, 1005–1032.

Del Negro, M., M. P. Giannoni, and C. Patterson (2012): "The forward guidance puzzle," *FRB of New York Staff Report*.

Dennis, R. (2007): "Optimal Policy In Rational Expectations Models: New Solution Algorithms," *Macroeconomic Dynamics*, 11, 31–55.

Eggertsson, G. B., S. Egiev, A. Lin, J. Platzer, and L. Riva (2020): "A Toolkit for Solving Models with a Lower Bound on Interest Rates of Stochastic Duration," *forthcoming Review of Economic Dynamics*.

Fershtman, C. (1989): "Fixed rules and decision rules: Time consistency and subgame perfection," *Economics Letters*, 30, 191–194.

Gabaix, X. (2020): "A behavioral New Keynesian model," *American Economic Review*, 110, 2271–2327.

Galí, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton: Princeton University Press.

Galí, J. (2015): *Monetary policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, Princeton University Press.

Gertler, M. and P. Karadi (2011): "A model of unconventional monetary policy," *Journal of Monetary Economics*, 58, 17–34.

Gondzio, J. (2012): "Interior point methods 25 years later," *European Journal of Operational Research*, 218, 587–601.

Guerrieri, L. and M. Iacoviello (2015): "OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily," *Journal of Monetary Economics*, 70, 22–38.

Haberis, A., R. Harrison, and M. Waldron (2019): "Uncertain policy promises," *European Economic Review*, 111, 459–474.

Harrison, R., M. Seneca, and M. Waldron (2020): "Monetary policy options in a 'low for long' era," Presented at the 2020 Banca d'Italia and SUERF Workshop on the effectiveness of monetary policy in a low interest rate environment.

Harrison, R. and M. Waldron (2021): "Optimal policy with occasionally binding constraints: piecewise linear solution methods," *Bank of England working paper*.

Holden, T. and M. Paetz (2012): "Efficient Simulation of DSGE Models with Inequality Constraints," Quantitative Macroeconomics Working Papers 21207b, Hamburg University, Department of Economics.

Holden, T. D. (2016): "Computation of solutions to dynamic models with occasionally binding constraints," *Unpublished manuscript*.

——— (2019): "Existence and uniqueness of solutions to dynamic models with occasionally binding constraints," .

Laforte, J.-P. (2018): "Overview of the changes to the FRB/US model," FEDS Notes.

Laséen, S. and L. E. Svensson (2011): "Anticipated Alternative Policy Rate Paths in Policy Simulations," *International Journal of Central Banking*, 7, 1–35.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The power of forward guidance revisited," *American Economic Review*, 106, 3133–58.

Nakata, T., R. Ogaki, S. Schmidt, and P. Yoo (2019): "Attenuating the forward guidance puzzle: implications for optimal monetary policy," *Journal of Economic Dynamics and Control*, 105, 90–106.

Sims, E. R. and J. C. Wu (2020): "Evaluating Central Banks' Tool Kit: Past, Present, and Future," *forthcoming Journal of Monetary Economics*.

Smets, F. and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

Svensson, L. E. (2005): "Monetary Policy with Judgment: Forecast Targeting," *International Journal of Central Banking*, 1, 1–54.

——— (2010): "Inflation Targeting," in *Handbook of Monetary Economics*, ed. by B. M. Friedman and M. Woodford, Elsevier, vol. 3 of *Handbook of Monetary Economics*, chap. 22, 1237–1302.

Svensson, L. E. and R. J. Tetlow (2005): "Optimal Policy Projections," *International Journal of Central Banking*, 1, 177–207.

Svensson, L. E. and M. Woodford (2004): "Implementing Optimal Policy through Inflation-Forecast Targeting," *NBER working paper*.

Woodford, M. (2019): "Monetary policy analysis when planning horizons are finite," *NBER Macroeconomics Annual*, 33, 1–50.

Yellen, J. L. (2012): "Revolution and evolution in central bank communications," Speech given at the University of California Haas School of Business, Berkeley, California. Available at www.federalreserve.gov/newsevents/speech/yellen20121113a.htm.

A Derivations

A.1 Derivation of Equations (5)-(7)

Begin with the system of equations given by

$$AY_t = B\mathbb{E}_t Y_{t+1} + CY_{t-1} + \tilde{D}GM_t, \tag{A.1}$$

$$M_t = \mathcal{M}M_{t-1} + \mathcal{N}\tilde{V}_t. \tag{A.2}$$

This can be rewritten as follows

$$AY_t = B\mathbb{E}_t Y_{t+1} + CY_{t-1} + \tilde{D}\left(\boldsymbol{v}_{tt} + \boldsymbol{v}_{tt-1} + \cdots + \boldsymbol{v}_{tt-H}\right). \tag{A.3}$$

Postulate a solution of the form

$$\begin{aligned} Y_t = \mathcal{P}Y_{t-1} &+ \chi_{00}\boldsymbol{v}_{tt} + \chi_{10}\boldsymbol{v}_{t+1,t} + \cdots + \chi_{H0}\boldsymbol{v}_{t+H,t} \\ &+ \chi_{01}\boldsymbol{v}_{tt-1} + \cdots + \chi_{H-1,1}\boldsymbol{v}_{t+H-1,t-1} \\ &\vdots \\ &+ \chi_{0H}\boldsymbol{v}_{tt-H}. \end{aligned} \tag{A.4}$$

Substitute into (A.3) to give

$$\begin{aligned} 0 = &- A\begin{pmatrix} \mathcal{P}Y_{t-1} + \chi_{00}\boldsymbol{v}_{tt} + \chi_{10}\boldsymbol{v}_{t+1,t} + \cdots + \chi_{H0}\boldsymbol{v}_{t+H,t} \\ +\chi_{01}\boldsymbol{v}_{tt-1} + \chi_{11}\boldsymbol{v}_{t+1,t-1} + \cdots + \chi_{H-1,1}\boldsymbol{v}_{t+H-1,t-1} \\ \vdots \\ +\chi_{0H}\boldsymbol{v}_{tt-H} \end{pmatrix} \\ &+ B\mathcal{P}\begin{pmatrix} \mathcal{P}Y_{t-1} + \chi_{00}\boldsymbol{v}_{tt} + \chi_{10}\boldsymbol{v}_{t+1,t} + \cdots + \chi_{H0}\boldsymbol{v}_{t+H,t} \\ +\chi_{01}\boldsymbol{v}_{tt-1} + \chi_{11}\boldsymbol{v}_{t+1,t-1} + \cdots + \chi_{H-1,1}\boldsymbol{v}_{t+H-1,t-1} \\ \vdots \\ +\chi_{0H}\boldsymbol{v}_{tt-H} \end{pmatrix} \\ &+ B\begin{pmatrix} \chi_{01}\boldsymbol{v}_{t+1,t} + \chi_{11}\boldsymbol{v}_{t+2,t} + \cdots + \chi_{H-1,1}\boldsymbol{v}_{t+H,t} \\ \vdots \\ +\chi_{0H}\boldsymbol{v}_{t+1,t-H+1} \end{pmatrix} \\ &+ CY_{t-1} + \tilde{D}GM_{t-1} + \tilde{D}\left(\boldsymbol{v}_{tt} + \boldsymbol{v}_{tt-1} + \cdots + \boldsymbol{v}_{tt-H}\right), \end{aligned} \tag{A.5}$$

Collecting terms gives

$$\begin{aligned}
\boldsymbol{v}_{t,t-j}: &\quad 0 = -A\chi_{0j} + B\mathcal{P}\chi_{0j} + \tilde{D} \quad \rightarrow \quad \chi_{0j} = (A - B\mathcal{P})^{-1}\tilde{D}, \\
\boldsymbol{v}_{t+1,t-j}: &\quad 0 = -A\chi_{1j} + B\mathcal{P}\chi_{1j} + B\chi_{0j} \quad \rightarrow \quad \chi_{1j} = (A - B\mathcal{P})^{-1} B\chi_{0j}, \\
&\quad \vdots \\
\boldsymbol{v}_{t+H,t-j}: &\quad 0 = -A\chi_{Hj} + B\mathcal{P}\chi_{Hj} + B\chi_{H-1,j} \quad \rightarrow \quad \chi_{Hj} = (A - B\mathcal{P})^{-1} B\chi_{H-1,j},
\end{aligned}$$

for $j = \{0, H\}$. Denote $\Phi \equiv \chi_{00}$ and $\mathcal{F} \equiv (A - B\mathcal{P})^{-1} B$. Then all $\chi_{0j} = \Phi$. $\chi_{1j} = \mathcal{F}\Phi$, and, in general, $\chi_{Hj} = \mathcal{F}^H\Phi$. Thus, the solution is given by

$$\begin{aligned}
Y_t = \mathcal{P}Y_{t-1} &+ \Phi\boldsymbol{v}_{tt} + \mathcal{F}\Phi\boldsymbol{v}_{t+1,t} + \cdots + \mathcal{F}^H\Phi\boldsymbol{v}_{t+H,t} \\
&+ \Phi\boldsymbol{v}_{tt-1} + \cdots + \mathcal{F}^{H-1}\Phi\boldsymbol{v}_{t+H-1,t-1} \\
&\vdots \\
&+ \Phi\boldsymbol{v}_{tt-H}. \qquad \text{(A.6)}
\end{aligned}$$

A.2 Derivation of Equation (11)

We begin with (1), replace ∞ with T, drop the expectations operator, and rewrite in matrix form to give $L^{A\prime}\Omega L^A$, where $\Omega \equiv \text{diag}\left(1, \beta, \ldots, \beta^T\right) \otimes Q$. The problem is to choose $\hat{V}_0$ to minimize $L^{A\prime}\Omega L^A$ subject to the following constraints: $L^A = L^B + L^{IR}$ and $L^{IR} = \mathcal{A}_L\hat{V}_0$. Substituting the constraints into the objective and dropping terms that are independent of policy gives (11).

B Canonical New-Keynsian model

We use the canonical New-Keynesian model from (Galí 2008; Galí 2015) as the workhorse example. The private sector conditions are given by

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \kappa x_t + \epsilon_{p,t}, \tag{B.1}$$

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t\pi_{t+1} - r_{n,t}\right), \tag{B.2}$$

$$r_{n,t} = \rho_{rn} r_{n,t-1} + \epsilon_{r,t}. \tag{B.3}$$

Closing the model with a Taylor-type rule requires

$$i_t = \phi_\pi\pi_t + \phi_x x_t + \epsilon_{r,t}. \tag{B.4}$$

Closing model with optimal policy under commitment requires

$$\pi_t + \xi_{1,t} - \xi_{1,t-1} - \frac{1}{\beta\sigma}\xi_{2,t-1} = 0, \tag{B.5}$$

$$\vartheta x_t - \kappa\xi_{1,t} + \xi_{2,t} - \frac{1}{\beta}\xi_{2,t-1} = 0, \tag{B.6}$$

$$\{i_t > 0, \xi_2 = 0\} \quad \text{or} \quad \{i_t = 0, \xi_2 > 0\}. \tag{B.7}$$

Closing model with optimal policy under discretion requires

$$\vartheta x_t = -\kappa\pi_t - \xi_{2,t}, \tag{B.8}$$

$$\{i_t > 0, \xi_2 = 0\} \quad \text{or} \quad \{i_t = 0, \xi_2 > 0\}. \tag{B.9}$$

The parameter values used in the main text are given here: $\sigma = 1$, $\varphi = 5$, $\theta = 3/4$, $\beta = 0.99$, $\alpha = 1/4$, $\epsilon = 9$, $\phi_\pi = 1.5$, $\phi_x = 0.5/4$, $\rho_{rn} = 0.7$. $\omega = (1-\alpha)/(1-\alpha+\alpha\epsilon)$, $\lambda = (1-\theta)(1-\beta\theta)\,\omega/\theta$, $\kappa = \lambda\left(\sigma + (\varphi+\alpha)/(1-\alpha)\right)$, $\vartheta = \kappa/\epsilon$.

B.1 Writing the model in the form (2)-(3)

For clarity of exposition, we set $\rho_{rn} = \rho_p = 0$ in this example. The variables are $Y_t = [\pi_t, x_t]'$, $X_t = i_t$, $U_{1,t} = [\epsilon_{p,t}, r_{n,t}]$, and $U_{2,t} = \epsilon_{r,t}$; and the coefficient matrices are

$$A = \begin{bmatrix} 1 & -\kappa \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} \beta & 0 \\ 1/\sigma & 1 \end{bmatrix}, \quad C = 0, \quad D = \begin{bmatrix} 0 \\ -1/\sigma \end{bmatrix}, \quad E = \begin{bmatrix} 1 & 0 \\ 0 & 1/\sigma \end{bmatrix},$$

$$F = [\phi_\pi, \phi_x], \quad \tilde{G} = 1.$$

C Section 3 additional results

Figure C.1 replicates Galí (2008) Figure 5.1 using our optimal policy projection methodology, showing the optimal response to a cost-push shock under commitment and discretion, respectively. Figure C.2 replicates Galí (2008) Figure 6.4 using our optimal policy projection methodology, showing the optimal response to a technology shock in an New-Keynesian model with sticky wages. In the presence of sticky wages, the policymaker the divine coincidence no longer holds, and the policymaker is unable to close both the inflation and output gaps. Finally, Figure C.3 repeats the “small deviations” exercise from the main text but with Type I inattention. In the presence of inattention, the small deviations constraint is more detrimental to the policymaker's ability to stabilize inflation and the output gap.

Figure C.1:
Transitory cost-push shock

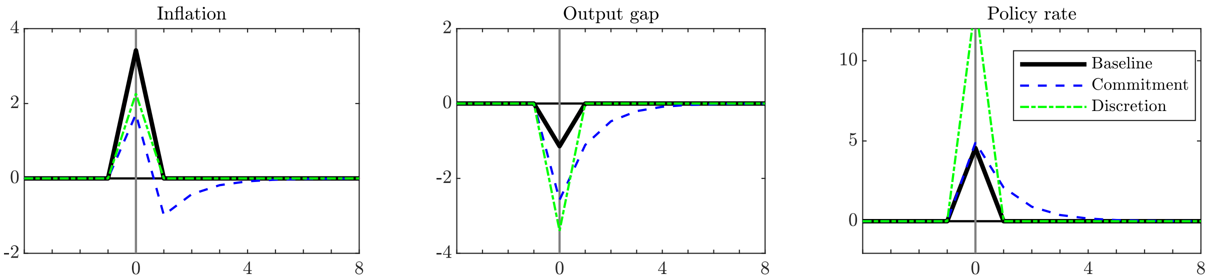

Note: Replicates Galí (2008) Figure 5.1

Figure C.2: New-Keynesian model with sticky wages:
Technology shock + Commitment

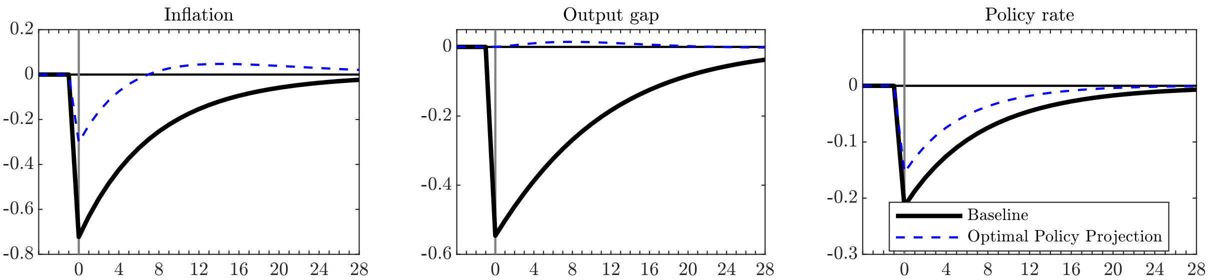

Note: Replicates Galí (2015) Figure 6.4. The baseline is the response with the Taylor rule given in Chapter 6.

Figure C.3:
Commitment + Small deviations + Type I ($\alpha = 0.6$)

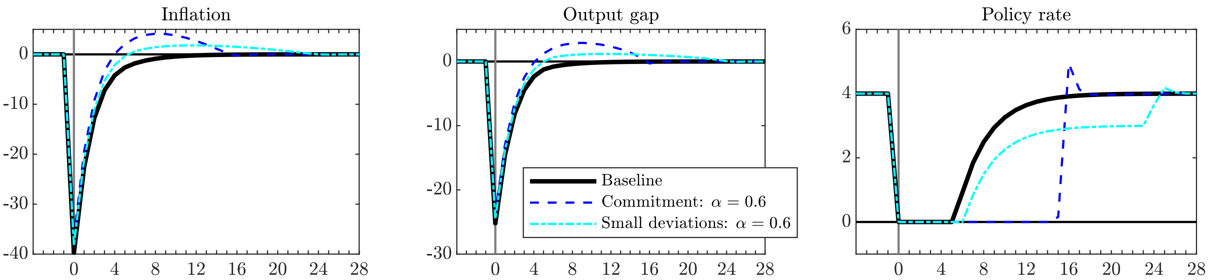

Note: This figure repeats the exercise in Figure 7 but with Type I inattention ($\alpha = 0.6$).

D Section 4 additional results

Figure D.1: Comparing SW07 with original and flat Phillips curves to FRB/US

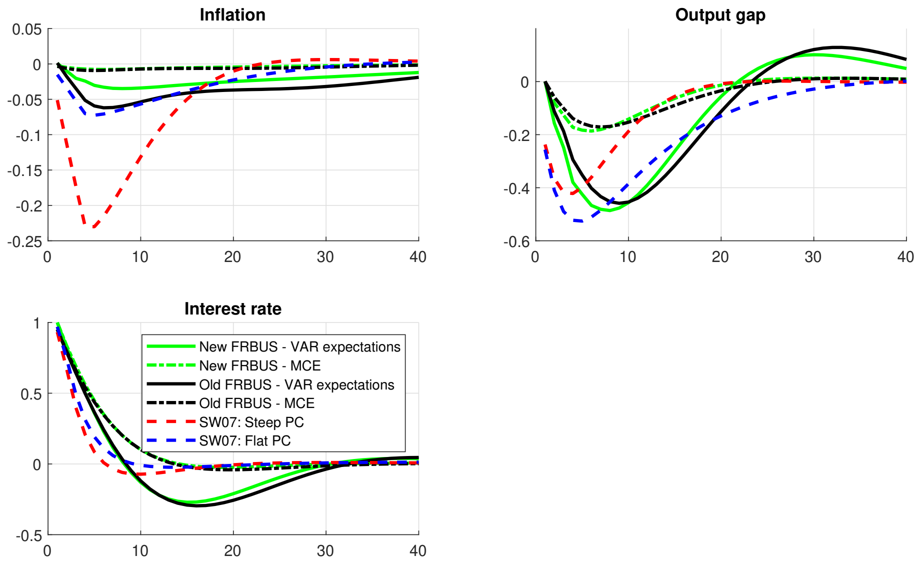

Note: FRB/US impulse responses are from Laforte (2018). MCE: Model consistent expectations.

Figure D.2: Optimal policy in 2009 with steeper forward curve

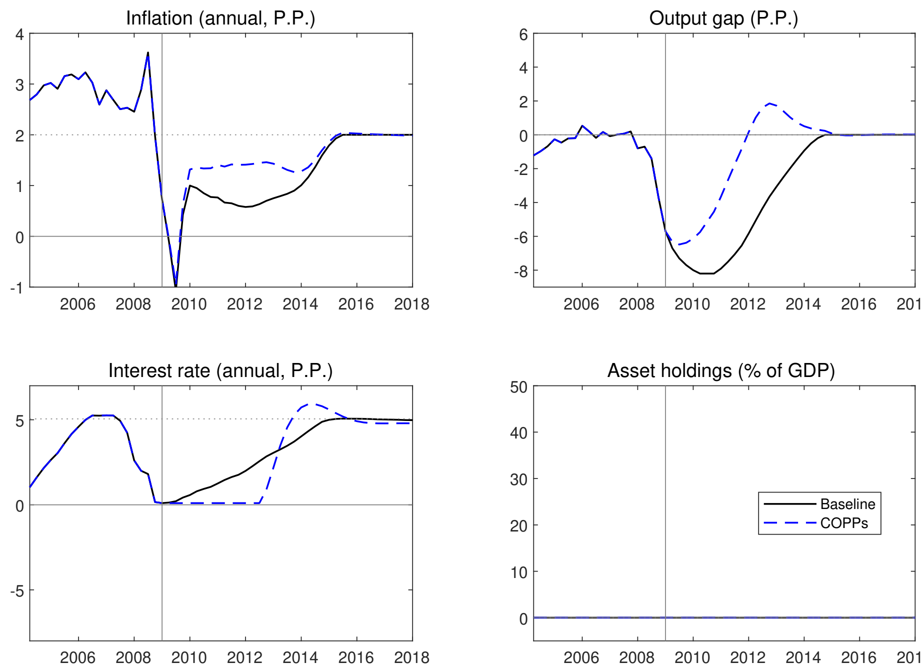

Note: This repeats the Figure 10 exercise but has a baseline interest rate path that exits the ELB earlier and rises steadily. Toolkit options: `Policy type = Commitment. Instruments = R. ELB constraint = On.`

Table D.1: Computational times

Figure	Model	No. of instruments	Policy	Constrained	H	Time (secs)
2	NK3	1	COM	UNC	60	2.08
2	NK3	1	DIS	UNC	60	3.53
4	NK3	1	COM	CON	20	2.20
4	NK3	1	COM	CON	40	1.90
4	NK3	1	COM	CON	60	2.11
4	NK3	1	DIS	CON	20	8.69
4	NK3	1	DIS	CON	40	20.29
4	NK3	1	DIS	CON	60	27.43
8	NK3	1	LC (D=6)	CON	60	8.73
8	NK3	1	LC (D=9)	CON	60	6.94
9	SW07	1	COM	UNC	40	3.43
10	SW07	1	COM	CON	40	3.77
11	SWu20	1	COM	CON	40	5.62
11	SWu20	2	COM	CON	40	7.44
13	SWu20	1	DIS	CON	40	45.60
13	SWu20	2	DIS	CON	40	200.51

Note: Timings based on Laptop with Processor: Intel(R) Core(TM) i7-8550U CPU @ 1.80GHz 1.99 GHz; System: 62-bit operating system, x64-based processor.