Threshold-based forward guidance: hedging the zero bound *

Richard Harrison[†] Lena Körber[‡] Matt Waldron[§]

June 5, 2015

Abstract

In this paper we study the efficacy of forward guidance as a policy that can be used to impart monetary stimulus at the zero lower bound (ZLB). Motivated by policies implemented by some central banks in response to the financial crisis, we use a simple New Keynesian model to study a particular form of forward guidance, whereby the policymaker makes a state-contingent commitment to hold the policy rate at the ZLB in a way that ensures that specific macroeconomic variables (e.g. inflation) do not breach particular 'thresholds'. In common with other similar policies, threshold-based forward guidance (TBFG) can be used to stimulate the economy at the ZLB via a commitment to hold the policy rate lower-for-longer than would otherwise have been the case. But TBFG also acts as a hedge against the asymmetric effects of shocks. That is because if further adverse shocks arise, prolonging the recession, the threshold will be expected to be breached at a later date and so the policy will provide additional stimulus. By contrast, if positive shocks arrive, so that the economy recovers more quickly than originally expected, the threshold will be expected to be breached sooner, thereby removing some of the policy stimulus. This hedging property of TBFG also means that there is a relatively low incentive for policymakers to renege on the policy, unlike in the case of calendar-based forward guidance. To maximise the benefits of TBFG, the thresholds must be chosen carefully and one contribution of this paper is to compute loss-minimising thresholds. Another contribution of the paper is to demonstrate that a unique equilibrium exists only if the policymaker specifies precisely what the conditions underpinning regime exit are (and the threshold conditions alone are not enough to determine that).

KEY WORDS: New Keynesian model; monetary policy; zero lower bound
JEL Classification: E12; E17; E20; E30; E42; E52

*We thank Toni Braun, Fabio Canova, Wouter Den Haan, Taisuke Nakata and seminar participants at the FRB Atlanta and Bank of England for their useful comments. Any views expressed in this paper are solely those of the authors and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee.

[†]Bank of England and Centre for Macroeconomics. Email: richard.harrison@bankofengland.co.uk.
[‡]Bank of England and London School of Economics, London, UK: lena.koerber@bankofengland.co.uk.
[§]Bank of England, London, UK. Email: matthew.waldron@bankofengland.co.uk

1 Introduction

The financial crisis of 2007/08 generated a severe and prolonged global contraction in output: the 'Great Recession'. In response, central banks around the world cut their policy rates to the zero lower bound (ZLB) and implemented a range of unconventional monetary policy measures, including an increased use of 'forward guidance' about the future path of the policy rate.

One interpretation of forward guidance is as the communication of a promise to hold the policy rate at the ZLB for long enough to reduce long-term real interest rates and provide near-term stimulus (Woodford, 2012). This type of behavior resembles the optimal commitment policy in New Keynesian models as first argued by Krugman (1998) and subsequently demonstrated by Eggertsson and Woodford (2003) among others. However, policymakers have tended to distance themselves from this interpretation, in part because they seem skeptical about their ability to credibly commit to behavior that is well known to be time inconsistent.[1]

In this paper we study a form of 'threshold-based' forward guidance (TBFG), in which the policymaker's commitment to hold the policy rate at the ZLB is state contingent in a particular way. The state-contingent commitment is designed to ensure that certain macroeconomic variables do not exceed pre-specified 'threshold' values in any state of the world in which the TBFG policy remains in effect.

We investigate whether this form of TBFG can be used as a temporary policy measure at the ZLB to improve outcomes, while limiting the extent to which the policymaker promises to behave in a time inconsistent manner. While our TBFG policies are clearly motivated by policies implemented by the FOMC and the Bank of England's MPC, our exercise falls well short of an evaluation of those policies as we abstract from many details that accompany real world central bank communications. Nevertheless, it is possible to daw some general lessons from our analysis.

The framework for our analysis is a simple New Keynesian model used in several other studies of policy at the ZLB (for example, Adam and Billi (2006) and Bodenstein et al. (2012)). The model consists of log-linearised equations describing the aggregate demand of optimising households (the 'IS' curve); and the pricing decisions of monopolistically competitive firms that adjust prices infrequently according to a Calvo (1983) scheme (the New Keynesian Phillips curve). The IS curve contains a stochastic 'demand shock' and the Phillips curve contains a stochastic 'cost push shock'.

The monetary policymaker sets the short-term nominal interest rate, subject to the ZLB. The policymaker minimizes the expected discounted value of a loss function derived from a second order approximation to household's utility. Our baseline assumption is that the policymaker acts with 'discretion'. That is, the policymaker does not have access to a commitment technology so that, each period, the policy rate is set to minimize expected discounted losses, taking the behavior of future policymakers as given. Under these assumptions, policy is time consistent. We solve the model using global methods to account for the nonlinearity introduced by the ZLB and by the form of the TBFG policies that we consider.

As is common in the literature on monetary policy at the ZLB, we examine what happens when a large negative demand shock causes the ZLB to bind. With our baseline assumption of time consistent monetary policy, we observe a deep recession. Because of the ZLB, the short-term nominal interest rate cannot be cut enough to reduce the *real* interest rate sufficiently to stabilize aggregate demand. This motivates us to consider experiments in which the policymaker announces a temporary deviation from time-consistent 'discretionary' monetary policy.

[1]For example, when describing the introduction of forward guidance by the Bank of England's Monetary Policy Committee, Bean (2013) argues that: "While such a time-inconsistent policy may be desirable in theory, in an individualistic committee like ours, with a regular turnover of members, it is not possible to implement a mechanism that would credibly bind future members in the manner required."

Our policy experiments examine 'one-off' temporary deviations from time consistent policy. Specifically, under TBFG the policymaker makes a state-contingent commitment to hold the policy rate at the ZLB for longer in expectation than would be the case if they continued to set policy in a time-consistent manner. Once the state of the economy is such that the TBFG regime has come to an end (i.e. once the economy has improved sufficiently), the policymaker reverts to setting policy in a time-consistent manner forever more.

One key contribution of our paper is to show that TBFG policy is incomplete in the absence of specific guidance about how the policymaker intends to interpret the threshold conditions. Put differently, in order for the private sector to understand the policy, it is not sufficient for the policymaker to announce a set of thresholds for macroeconomic variables because there are many policies that could be consistent with those thresholds. In order to overcome this indeterminacy (and uniquely define equilibrium), it is also necessary for the policymaker to announce precisely what the threshold conditions mean. One such announcement could be: "....commit to hold the policy rate at the ZLB for as long as possible subject to inflation not rising more than 1pp above target in any state of the world while the forward guidance regime remains in effect". This is the interpetation of the thresholds we use in this paper, but there are other equally valid intrepretations. For example, the policymaker could instead commit to hold rates at the ZLB for as long as possible subject to the thresholds being breached by the smallest possible amount prior to regime exit in all states of the world. It is worth noting that the macroeconomic effect of a TBFG with a given set of threshold conditions is dependent on the precise specification of the exit conditions.

Our baseline results compare the behavior of the model under time consistent policy and various forms of forward guidance with thresholds on both inflation and the output gap. We find that appropriately calibrated TBFG policies can improve welfare compared with fully time consistent behavior. The mechanism behind the result is straightforward. In line with the 'textbook' remedy to mitigating the ZLB constraint, TBFG can be used to stimulate activity and inflation today by promising higher inflation in the future.[2] As well as improving outcomes in expectation, TBFG can also be used to manage the variance of possible outcomes. Agents know that if further negative shocks arise, prolonging the recession, the policy rate will stay at the ZLB for longer because the threshold(s) will be breached at a later date. By contrast, if positive shocks arrive, so that the economy recovers more quickly from the recession than originally expected, then the threshold(s) will be breached sooner and the policy stimulus removed.

So TBFG policies can be viewed as a hedge against the asymmetric effects generated by the ZLB. This means that policymakers have relatively little incentive to deviate from them. This effect can be seen by comparing the distribution of outcomes under TBFG with those under calendar-based forward guidance (CBFG), in which the policymaker promises to hold the policy rate at the ZLB for a pre-specified length of time.[3] CBFG imparts stimulus regardless of the state of the economy. This can eliminate the negative skew in the distributions of the output gap and inflation that we observe under time-consistent policy, by raising expectations sufficiently to reduce the effect of the ZLB constraint. However, CBFG leads to worse outcomes for both positive and negative realisations of future demand shocks because it provides too much stimulus in 'good' states and insufficient stimulus in 'bad' states.[4] As a result, the variances of the distributions of the output gap and inflation are substantially larger than those for the

[2] A common theme in this work is that history dependent policies such as optimal commitment, price level targeting or the Reifschneider and Williams (2000) rule can significantly improve outcomes at the ZLB by using inflation expectations as a substitute for cutting the policy rate.

[3] Early incarnations of forward guidance by the FOMC and Bank of Canada had a calendar-based flavor, though also included (informal) threshold-based clauses.

[4] This result verifies the assertion of Campbell et al. (2012) that CBFG is likely to generate poor outcomes if the economy evolves differently to initial expectations as shocks arrive over time.

baseline time-consistent policy assumption.

Because our policy experiments are based on a temporary deviation from time-consistent behavior, they are (by definition) time inconsistent. As such, the experiments may be regarded as less than fully credible by agents in the model. We investigate this by computing a measure of the extent to which the policymaker could achieve better outcomes by reneging on the TBFG policy and reverting to the time-consistent policy at each point in time.[5] A corollary of the hedging property of TBFG is that the temptation to renege from TBFG is much smaller than for CBFG. For realisations of shocks in which the economy recovers more quickly than originally expected, CBFG generates too much stimulus and the policymaker has a strong incentive to revert to the time-consistent policy. By contrast, under TBFG, for realisations of the shocks in which the economy recovers more quickly, the exit thresholds are breached and policy automatically reverts to time-consistent behavior.

Of course, for TBFG to deliver better outcomes than fully time-consistent policy, the thresholds must be appropriately calibrated. In particular, the thresholds must be calibrated to generate an overshoot of goal variables from target. Otherwise, the policy is unable to increase expectations enough to impart any additional stimulus relative to the time-consistent policy. But there are infinitely many TBFG policies that satisfy this condition. One criterion for comparing alternative TBFG policies is the ex-ante loss. We use this criterion to compute optimal thresholds for both inflation and output gap threshold designs. Optimal TBFG policies achieve ex-ante losses that are close to the optimal commitment policy. In contrast, the losses associated with CBFG or time-consistent policies are significantly larger.

To our knowledge, this is the first paper to analyse TBFG policies similar to those actually implemented in response to the financial crisis in a fully stochastic setting. The closest paper to ours is Florez-Jimenez and Parra-Polania (2014), who also study TBFG in a small model. But their analysis is limited to a two-period model with a threshold defined in terms of an exogenous shock process. By contrast, we analyse TBFG policies of indefinite duration and model thresholds defined in terms of endogenous variables. Coenen and Warne (2013) consider a more realistic model and policy experiment, examining how a form of inflation forecast threshold can alter the performance of calendar-based forward guidance in the ECB's DSGE model. However, given the size of the model they use, they are restricted to perfect foresight approximations of expectations, whereas we compute a fully stochastic equilibrium.

This paper also relates to the growing literature on the responses of New Keynesian models to anticipated paths for the nominal interest rate, as the policies that we study require an ability to stimulate the economy in the near term by making promises about the future policy rate. For example, del Negro et al. (2012) document that imposing anticipated paths for the policy rate often creates unrealistically large responses of inflation and activity, a result that they label "the forward guidance puzzle". Since then, several papers have explored possible explanations. For example, Haberis et al. (2014) document that the puzzle may be reduced if the announcements made by the policymaker are not perfectly credible. McKay et al. (2015) illustrate the sensitivity to the assumption of forward-looking households. In their model, including a small fraction of households that face occasionally binding credit constraints eliminates the puzzle. Our work relates to this literature in two ways. First, as discussed above, CBFG embodies a commitment to hold rates at the ZLB for a particular amount of time *regardless* of the state of the economy. This kind of commitment is unlikely to be credible. As such, TBFG could be viewed as a 'solution' to the forward guidance puzzle because it embodies a more reasonable description of how the policymaker might behave in the event that the economy evolves differently to what had been expected. Second, the observation that forward guidance is particularly powerful in the sort of New-Keynesian model we use in this paper suggests that our results may be regarded

[5]The policymaker has an incentive to deviate if the welfare losses associated with reverting to time-consistent policy are less than the losses associated with continuing with the announced guidance.

as an upper bound on the effects of TBFG policies.

Finally, TBFG policy in our setting can be regarded as a constraint on the discretionary monetary policymaker in the sense that the policymaker is not permitted to increase the policy rate from the ZLB in certain states of the world. Under that interpretation, our results are consistent with existing work. For example, Nakata (2014a) finds that, in a model with an inefficient steady state allocation, welfare is larger without access to fiscal instruments if the policy maker cannot make commitments about future policy. Similarly but in a different context, Chari and Kehoe (2007) present a model in which restricting the amount of debt that can be issued is welfare-improving in a currency union.

The rest of the paper is organized as follows: Section 2 describes the policy experiments and the assumptions underpinning them. Section 3 details the model and the baseline description of policy and Section 4 defines equilibrium for both TBFG and CBFG policies. Section 5 describes the methods we use to solve our model. Section 6 outlines the parameterisation of the model and the calibration of the state of the economy prior to the implementation of forward guidance. Section 7 uses a TBFG policy explains the mechanism at work and reports optimal inflation and output gap thresholds. Section 8 compares TBFG polices to optimal commitment and Section 9 concludes.

2 The nature of the policy experiments

The policy experiments are ones in which a policymaker temporarily deviates from setting policy optimally but in the absence of a commitment device (optimal discretion). The temporary deviation is a one-off and fully credible forward guidance policy with the objective of achieving better outcomes, given an economic environment in which the policy rate has become constrained by the ZLB. As detailed in Section 4, the forward guidance policies can be characterised as a commitment by the policymaker to hold the policy rate at the ZLB in certain states of the world, in the case of threshold-based forward guidance (TBFG), or for a particular number of periods, in the case of calendar-based forward guidance (CBFG).

The precise sequence of events in all of our policy experiments is summarised in Figure 1. In period $t = 0$, a negative demand shock arrives that is sufficiently large to drive the policy rate to the ZLB. Having observed this shock and the subsequent outcomes, the policymaker announces a forward guidance policy that becomes effective in period $t = 1$ and remains in effect until the regime termination conditions have been met. Once the regime has ended, the policymaker reverts to setting policy by optimal discretion forever more.

Figure 1: Timeline of events for policy experiments

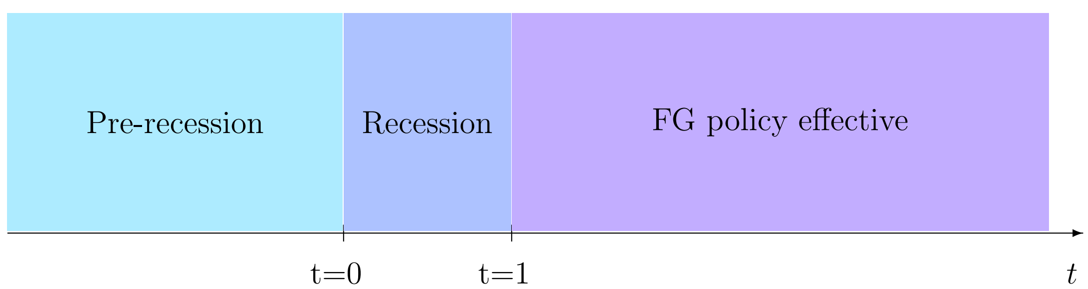

There are two overarching assumptions governing the nature of our experiments. First, the forward guidance policy is assumed to be transitory: before implementation, the policy is entirely unanticipated by agents in the model and, once the conditions for ending the regime have been met, the regime comes to an end with agents attaching no probability to the policy being implemented again in the future. This assumption is common to several other papers in the literature that study temporary deviations of policy from a rule governing the timeless behaviour of the policymaker (e.g. del Negro et al. (2012), Coenen and Warne (2013), Haberis et al. (2014)). This means that these policy experiments are not conducted under rational expectations and so are subject to the problem studied by Leeper and Zha (2003) and Harrison (2014) (in the context of CBFG). Specifically, one may obtain misleading results from implementing temporary policy regime change under the assumption that agents attach a zero *ex ante* probability to that regime change. In the context of our experiments with the policies implemented by some central banks in the wake of the financial crisis, it is arguably reasonable to believe that the forward guidance policy may not have been anticipated, but is perhaps less reasonable to believe that agents would not expect the policymaker to adopt a similar policy in the future, should the ZLB become a binding constraint on policy again. It is worth noting that the results of our policy experiments are likely to be sensitive to this assumption.[6]

Our second overarching assumption is that the forward guidance policy is fully credible. This assumption is seemingly at odds with a baseline description of policy being conducted in a fully time-consistent manner. Indeed, the mechanism by which the forward guidance policies we study are effective is through the manipulation of agents' expectations. In the absence of at least some credibility, the policymaker would be unable to affect agents' expectations and forward guidance of this sort would have no effect. Given the importance of this assumption, we pay particular attention to its likely validity by computing a measure of the incentive that the policymaker has to renege on the announced forward guidance policy. As argued in Nakata (2014b), if the incentive to renege is sufficiently small, then the assumption of full credibility seems more reasonable if the policymaker is concerned at least to some extent about their reputation.

3 The model

The model is identical to that used by Adam and Billi (2006, 2007) and Bodenstein et al. (2012) to study monetary policy at the zero lower bound (ZLB) under optimal commitment, optimal discretion and 'loose commitment' respectively.[7] It is a prototypical New Keynesian model in which a representative household supplies labour to firms and consumes a bundle of goods to maximize expected lifetime utility, and in which monopolistically competitive firms maximize the discounted sum of expected future profits subject to Calvo (1983) pricing rigidities. The first-order conditions for the household and firms, together with standard market clearing and aggregation conditions give rise to an Euler equation for output and an optimal pricing decision.[8] Following previous studies of monetary policy at the ZLB (e.g. Adam and Billi (2006), Adam and Billi (2007), Nakov (2008) and Bodenstein et al. (2012)), we use a partially log-linearized version of the model where the only nonlinearity is due to the ZLB and the remaining optimality conditions are log-linearised around the non-stochastic steady state.[9]

[6]Incorporating forward guidance at the ZLB into a full rational expectations model is the subject of our ongoing research.

[7]Under the loose commitment framework there is an exogenous, constant probability that the policymaker will renege on past commitments and re-optimise their policy.

[8]See Woodford (2003) for a detailed derivation and discussion.

[9]This is not an innocuous assumption. For example, Fernández-Villaverde et al. (2012) and Braun et al. (2013) have shown that non-linearities in the competitive equilibrium conditions can play an important role in the dynamics of New Keynesian models in the presence of an occasionally-binding ZLB. Investigating TBFG policies in the non-linear version of our model is the subject of our ongoing research.

The equilibrium conditions are:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + u_t \tag{1}$$
$$y_t = \mathbb{E}_t y_{t+1} - \sigma \left(r_t - \mathbb{E}_t \pi_{t+1}\right) + g_t \tag{2}$$
$$r_t \geq 1 - \frac{1}{\beta} \tag{3}$$

where: π is inflation, y is the output gap, and r is the policy rate (all expressed in deviations from steady state); $\beta < 1$ is the discount factor; $\kappa = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\frac{\sigma^{-1}+\omega}{1+\omega\theta}$ is the slope of the Phillips curve, where α is the probability that a firm cannot adjust its price, ω is the elasticity of a firm's real marginal cost with respect to its own output level and θ is the price elasticity of demand for the goods supplied by the monopolistic firms; σ is the intertemporal elasticity of substitution; u and g are exogenous disturbances to inflation and demand, often labeled cost push and demand shocks[10], both of which are assumed to follow AR(1) processes:

$$u_t = \rho_u u_{t-1} + \sigma_u \varepsilon_t^u \tag{4}$$
$$g_t = \rho_g g_{t-1} + \sigma_g \varepsilon_t^g \tag{5}$$

where $\varepsilon_t^u \sim$ *iid* N(0,1), $\varepsilon_t^g \sim$ *iid* N(0,1), ρ_u and ρ_g are the persistence parameters, σ_u and σ_g are the standard deviations, and with u_0 and g_0 as given.

The model is closed by the specification of monetary policy. Throughout our analysis, our baseline assumption is that the monetary policymaker follows optimal discretion. Specifically, we assume that the policymaker minimises the per-period loss (derived as a quadratic approximation to the representative agent's utility function[11]) subject to the equilibrium conditions from above, taking agents' expectations as given (Adam and Billi, 2007):

$$\begin{aligned}
&\min_{\{y_t,\pi_t,r_t\}} \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i (\pi_{t+i}^2 + \lambda y_{t+i}^2) \\
&s.t \;\; r_t \geq 1 - \frac{1}{\beta} \\
&\quad\;\; \pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + u_t \\
&\quad\;\; y_t = \mathbb{E}_t y_{t+1} - \sigma \left(r_t - \mathbb{E}_t \pi_{t+1}\right) + g_t \\
&\quad\;\; u_t = \rho_u u_{t-1} + \sigma_u \varepsilon_t^u \\
&\quad\;\; g_t = \rho_g g_{t-1} + \sigma_g \varepsilon_t^g \\
&\quad\;\; \mathbb{E}_t \{y_{t+i}, \pi_{t+i}, r_{t+i}\}_{i=1}^{\infty} \text{ given} \\
&\quad\;\; \{u_t, g_t\} \text{ given}
\end{aligned}$$

where $\lambda = \kappa/\theta$ is the relative weight on the output gap in the loss function.

In any period where the ZLB is not binding, the solution to this problem is the well-known targeting rule (e.g. Gertler et al. (1999)):

$$y_t = -\frac{\kappa}{\lambda}\pi_t \tag{6}$$

In the absence of an occasionally-binding ZLB, this rule describes the optimal policy response to shocks. In response to demand shocks, there is no trade-off between output and inflation

[10] The natural rate is related to the stochastic process, g, as follows: $g_t = \sigma r_t^*$, where r_t^* is the natural rate. The microfoundation of this shock is typically as a stochastic process for government spending (along with an assumption that government spending is entirely wasteful) or household's rate of time preference.

[11] See Woodford (2003) for a derivation and discussion.

stabilisation and the policymaker is able to achieve the first-best allocation[12] of inflation and the output gap at zero (i.e. the above rule delivers $y_t = 0$ and $\pi_t = 0$). In response to cost-push shocks, the policymaker is unable to stabilise the economy perfectly and the above targeting rule governs the policymaker's response to the trade-off that is created.

In the presence of an occasionally-binding ZLB, this result no longer holds (Adam and Billi (2007)). In particular, the policymaker is unable to perfectly stabilise the economy in the face of negative demand shocks if the policy rate becomes constrained by the ZLB. This creates a non-linearity, whereby outcomes for inflation and the output gap have a negative skew. Moreover, the possibility that the ZLB may become binding in the future affects outcomes today via agents' expectations. This requires us to use numerical methods to solve for the model's equilibrium, as described in Section 5.

4 Equilibrium in the policy experiments

The section defines equilibrium for threshold-based and calendar-based forward guidance policy given the environment described in Section 2 and given the model described in Section 3.

4.1 TBFG equilibrium definition

A pre-requisite for the existence of an equilibrium is that the status of the policy regime (i.e. whether or not the TBFG regime applies) must be known and must be forecastable by the private sector. As such, we characterise TBFG policy as a state-contingent indicator function defining the states of the world in which the TBFG policy regime is terminated, consistent with a set of thresholds on the endogenous variables (either inflation or output or a combination of the two). As explained in Appendix A using a simple deterministic example, it is not sufficient for the policymaker to announce a set of thresholds for the private sector to be able to determine how the policymaker intends to behave. It is also necessary for the policymaker to announce precisely how they intend to treat the threshold conditions. Put differently, the existence of a unique equilibrium (in which the status of the policy regime can be uniquely determined) requires the policymaker to reveal more information than just a set of threshold values. The precise interpretation we use in the definition of equilibrium below and in the remainder of this paper is that the indicator function determining regime exit should be such that the TBFG regime has the longest expected duration among all exit indicator functions subject to the thresholds not being breached in any state of the world while the TBFG policy regime remains in effect.[13]

We also restrict the feasible set of policies that the policymaker can implement to the announcement of a time-invariant exit indicator function. An implication of this restriction is that the distribution of expected outcomes (including the status of the policy regime) is identical regardless of the time period from which those expectations are taken, conditional on a particular state of the economy. While seemingly at odds with a baseline description of a policymaker who can re-optimse every period, this assumption is consistent with the motivation for forward guidance in this setting as a *temporary* commitment device to improve economic outcomes at the ZLB.

Formally, equilibrium in a one-off TBFG policy regime with inflation threshold, π^*, and output gap threshold, y^*, is defined by a regime exit indicator, $\mathbb{I}^{EXIT}(u, g) \in \{0, 1\}$, together

[12] Assuming that the steady-state distortions caused by the monopolistic competition are eliminated using a lump-sum transfer.

[13] It should be noted that there are alternative, equally valid, equilibrium definitions. For example, the policymaker could announce an exit indicator function consistent with the thresholds being breached by the smallest possible amount prior to regime exit in all states of the world.

with associated policy functions, $\pi^{FG}(u,g)$ and $y^{FG}(u,g)$, that satisfy:[14]

1. The competitive equilibrium conditions:
$y^{FG}(u,g) = \mathbb{E}^{FG}(u,g)\, y' - \sigma\left(1 - \frac{1}{\beta} - \mathbb{E}^{FG}(u,g)\,\pi'\right) + g$
$\pi^{FG}(u,g) = \beta\mathbb{E}^{FG}(u,g)\,\pi' + \kappa y^{FG}(u,g) + u$, where:
$\mathbb{E}^{FG}(u,g)\, y' = \int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\mathbb{I}^{EXIT}(u',g')\, y^{OD}(u',g')$
$\quad + \int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\left(1 - \mathbb{I}^{EXIT}(u',g')\right) y^{FG}(u',g')$
$\mathbb{E}^{FG}(u,g)\,\pi' = \int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\mathbb{I}^{EXIT}(u',g')\,\pi^{OD}(u',g')$
$\quad + \int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\left(1 - \mathbb{I}^{EXIT}(u',g')\right)\pi^{FG}(u',g')$
$u' = \rho_u u + \epsilon^{u'}$
$g' = \rho_g g + \epsilon^{g'}$
$\epsilon^{u'} \sim \mathbb{N}(0,\sigma_u)$
$\epsilon^{g'} \sim \mathbb{N}(0,\sigma_g)$.

2. The criterion for exit:
$\max \sum_{t=1}^{\infty}(t-1)\int_u\int_g \psi_t^{FG}(u,g)\,\mathbb{I}^{EXIT}(u,g)$
subject to: $\max\left(\pi^{FG}(.)\right) \leq \pi^*$ and/or $\max\left(y^{FG}(.)\right) \leq y^*$.

where $\psi_t^{FG}(u,g)$ measures the density across the state in the TBFG regime in period t and can be defined recursively as:

$$\psi_{t+1}^{FG}(u',g') = \int_u p(u'|u)\int_g p(g'|g)\,\psi_t^{FG}(u,g)\left(1 - \mathbb{I}^{EXIT}(u',g')\right)$$

where: $p(u'|u)$ is the probability of drawing u' conditional on u with $p(g'|g)$ defined analgously; $\psi_0^{FG}(u,g) = 1$ for $u = u_0$, $g = g_0$ and $\psi_0^{FG}(u,g) = 0\ \forall\ u \neq u_0$ and $g \neq g_0$ (i.e. there is a deterministic initial condition); $\int_u\int_g \psi_\infty^{FG}(u,g) = 0$ (i.e. policy will have reverted back to optimal discretion for sure in the limit).

Note that, by definition, the policy functions, $\pi^{FG}(u,g)$ and $y^{FG}(u,g)$, are only defined over states of the world where exit does not occur (where $\mathbb{I}^{EXIT}(u,g) = 0$). This reflects the nature of the policy: once exit has occurred, the regime is terminated and the equilibrium is determined by the policymaker acting under optimal discretion.

There are several features of this definition that are worth noting. First, expectations are defined as the probability weighted integral over all possible realisations of the shocks, taking into account the two different policy regimes: the case in which the forward guidance regime is still in effect, denoted with superscript ^{FG}, and the case in which policy has reverted back to optimal discretion, denoted with superscript ^{OD}. It is clear from this that the transmission of forward guidance policies in this model is via agents' expectations and that the macroeconomic effect of the policy depends on the precise exit conditions that the policymaker specifies. Second, and relatedly, TBFG can only affect outcomes to the extent that there are some states of the world in which the TBFG regime still applies *and* those are states of the world in which the policy rate would exceed the ZLB value (of $1 - \frac{1}{\beta}$) if policy were set under optimal discretion. From this it is clear that TBFG policy in this framework is a state-contingent form of 'lower-for-longer' policy. Third, the 'one-off' nature of the policy is embodied in the equilibrium definition because state-contingent outcomes under optimal discretion are taken as given (and are not a function of outcomes in the TBFG regime). Fourth, the initial condition for the economy in period $t = 0$ affects the equilibrium because it affects the probability density across the state and hence the expected duration of the policy. Different initial conditions would result in different equilibria (given a particular set of thresholds).

[14] For notational convenience we have dropped the time subscript. Variables without a $'$ superscript are measured at time t and those with a $'$ superscript are measured at time $t+1$.

This framework allows us to study a broad range of different policies. For example, one such policy would be a commitment by the policymaker to hold rates at the ZLB for as long as possible in expectation subject to inflation not rising above the target in any state of the world in which the regime could apply.[15] It is worth noting that although this setup allows us to study several different types of TBFG policy, it is much cruder than the real-world TBFG policies that have been implemented (e.g. by the Monetary Policy Committee of the Bank of England). These real-world policies have typically involved consideration of a broader range of factors, like emerging financial stability risks, as well as nuances in the interpretation of the thresholds as, for example, conditions that would trigger a re-assessment of the policy rather than as conditions that would automatically lead the policy to come to an end.[16] As such, our analysis is intended to draw some general conclusions about the efficacy and design of TBFG policies, rather than as direct commentary on policies that central banks have actually implemented.

4.2 CBFG equilibrium definition

CBFG policy is characterised as a scalar number of time periods, K, for which the policymaker commits to hold rates at the ZLB regardless of the state of the economy. Equilibrium is defined by a set of policy functions, $\{\pi_t^{FG}(u,g)\}_{t=1}^{K}$ and $\{y_t^{FG}(u,g)\}_{t=1}^{K}$, that satisfy:

1. The competitive equilibrium conditions:
$$y_t^{FG}(u,g) = \mathbb{E}_t^{FG}(u,g)\, y_{t+1} - \sigma\left(1 - \tfrac{1}{\beta} - \mathbb{E}_t^{FG}(u,g)\,\pi_{t+1}\right) + g$$
$$\pi_t^{FG}(u,g) = \beta\mathbb{E}_t^{FG}(u,g)\,\pi_{t+1} + \kappa y_t^{FG}(u,g) + u, \text{ where:}$$
$$\begin{aligned}\mathbb{E}_t^{FG}(u,g)\, y_{t+1} = &\int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\mathbb{I}_{t+1}^{EXIT} y^{OD}(u',g') \\ &+ \int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\left(1-\mathbb{I}_{t+1}^{EXIT}\right) y_{t+1}^{FG}(u',g')\end{aligned}$$
$$\begin{aligned}\mathbb{E}_t^{FG}(u,g)\, \pi_{t+1} = &\int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\mathbb{I}_{t+1}^{EXIT} \pi^{OD}(u',g') \\ &+ \int_{\epsilon^{u'}} p\left(\epsilon^{u'}\right)\int_{\epsilon^{g'}} p\left(\epsilon^{g'}\right)\left(1-\mathbb{I}_{t+1}^{EXIT}\right) \pi_{t+1}^{FG}(u',g')\end{aligned}$$
$$u' = \rho_u u + \epsilon^{u'}$$
$$g' = \rho_g g + \epsilon^{g'}$$
$$\epsilon^{u'} \sim \mathbb{N}(0,\sigma_u)$$
$$\epsilon^{g'} \sim \mathbb{N}(0,\sigma_g).$$

2. The criterion for exit:
$\mathbb{I}_t^{EXIT} = 0 \;\forall\; t <= K$ and $\mathbb{I}_{K+1}^{EXIT} = 1$.

As in the case of TBFG, it is clear from the above that CBFG affects economic outcomes in this setting via the manipulation of agents' expectations. The key distinction between the two policies is that regime exit is determined purely as a function of time under CBFG, while regime exit is determined purely as a function of the state of the economy under TBFG.

5 Solution method

5.1 Optimal discretion with a zero lower bound

The objective is to solve the model described in Section 3. That amounts to finding time-invariant policies for inflation, $\pi^{OD}(u,g)$, and the output gap, $y^{OD}(u,g)$, as functions of the state of the economy (outcomes for the cost-push and demand processes) that satisfy the

[15] This can be implemented as a special case in the definition of equilibrium in the text with $y^* = \infty$.

[16] One informal way to try to capture this kind of nuance would be to replace the state-contingent exit indicator with a state-contingent probability of exit. This is an extension we may explore in future work.

equilibrium conditions (i.e. the Phillips and IS curves) and that solve the policymaker's optimal discretion problem, subject to the ZLB constraint and the stochastic cost-push and demand processes.

There is no analytical solution to this problem (because of the ZLB constraint), so it is necessary to use numerical methods to approximate the solution. In doing so, we follow the approach described in Adam and Billi (2007). The approach is a time iteration implementation of policy function approximation using linear interpolation and quadrature to approximate expectations. The algorithm is initialised with a guess for the solution defined on a pre-specified grid of values for the state variables (cost-push and demand process outturns), for which we use the solution to a version of the model in which the ZLB contraint is ignored (which can be solved analytically). The output of each successive time iteration is a new guess at the solution on the state grid, using the previous guess to approximate agents' expectations for inflation and the output gap at each node in the state grid (which represents a particular combination of cost-push and demand process outturns). Conditional on expectations, the problem of the policymaker can be solved analytically as a sequence of independent static problems (for each node in the state grid) even with an occasionally-binding ZLB because the model has no endogenous state variables.[17] The time iteration is terminated when the difference between the latest guess for the solution (the output of the time iteration) and the previous guess (the input of the time iteration used to approximate expectations) is sufficiently small.

We implement the algorithm using a $20,000$ state grid formed of the tensor product of 100 and 200 node uni-dimensional grids of values for the cost-push and demand states respectively. These nodes are uniformly spaced between lower and upper bounds for each state, set to ensure that the policy experiment simulations do not require us to extrapolate the policy functions. This means that the lower and upper bounds for both states in the grid are functions of the particular parameterisation of the model we use. In the case of the baseline parametrisation outlined in Section 6, the bounds for the cost-push and demand state are set to ± 0.66 and ± 22 respectively (reflecting that the demand process is more persistent and has a higher variance than the cost-push process). In approximating expectations at each node in the state grid, we use a 25 node quadrature scheme formed of the tensor product of two separate 5 node Gauss-Hermite schemes for the cost-push and demand shocks. We terminate the time iteration when the largest absolute difference between the latest and previous guesses for the policy functions is less than $1e^{-6}$.[18]

5.2 TBFG policy experiments

The objective is to find policy functions for inflation, $\pi^{FG}(u, g)$, and the output gap, $y_t^{FG}(u, g)$, and an exit indicator function, $\mathbb{I}^{EXIT}(u, g)$, that satisfy the equilbrium conditions (i.e. the Phillips and IS curves) and the exit conditions of the regime, as defined in Section 4.1. In

[17] This works as follows: (i) use the first-order condition for the policymaker in equation (6) to substitute the output gap out of the Phillips curve (equation (1)) and rearrange to compute inflation as a function of expected inflation and the cost-push state; (ii) compute the output gap using the policymaker's first-order condition; (iii) rearrange the IS curve (equation (2)) to compute the interest rate as a function of the output gap, the expected output gap, expected inflation and the demand state; (iv) if the interest rate is greater than or equal the ZLB, then the solution (conditional on expectations) has been found and stop; (v) if the interest rate violates the ZLB constraint, then set it equal to the ZLB and; (vi) compute the output gap conditional on that, expectations and the demand state using the IS curve; (vii) compute inflation conditional on the output gap, expectations and the cost-push state using the Phillips curve.

[18] The algorithm takes 151 iterations to converge in 67 seconds in 64-bit MATLAB 2012b using a single Intel i7 CPU @ 2.90GHz. Key to that performance is the pre-computation of the state index numbers and weights for linear interpolaton in the approximation of expectations (noting that all the state variables are exogenous and so each possible realisation of next period's state given the quadrature scheme and this period's state is known in advance and does not vary across the iterations).

solving for those functions, we split the problem into two parts: a policy function iteration conditional on a guess for the exit indicator; an optimisation to solve the problem of maximising the expected duration of the policy subject to the threshold conditions not being breached.

5.2.1 Policy function iteration

Given a guess for the equilibrium indicator function, $\mathbb{I}^{EXIT}(u,g)$, we solve for the approximate policy functions for inflation and the output gap in a similar way to that described for the model under optimal discretion. However, the approximation of expectations in the forward guidance regime is more challenging than in the timeless solution under optimal discretion, reflecting that when the forward guidance policy regime is active, expectations are a weighted average of outcomes in two different policy regimes. Standard quadrature schemes like the 5 node Gauss-Hermite quadrature scheme used for the discretisation of the cost-push and demand shocks in approximating the timeless policy functions under optimal discretion do not perform well when applied to the approximation of expectations within the forward guidance regime because they do not provide a sufficiently accurate approximation of the probability of the exit conditions being met. Reflecting that, we employ a quadrature scheme designed to reflect better the probability density functions of the two shocks. More specifically, we adapt the Adda and Cooper (2003) methodology for Markov chain approximations of AR(1) proccesses to a quadrature scheme for the two shocks (which we combine together by tensor product in the standard way). This discretisation method works by dividing the distribution into N segments, each containing equal cumulative probability mass, and then setting the nodes in the quadrature scheme equal to the probabilistic mid-points of those intervals. This scheme performs better in approximating the probability of exit and, therefore, in approximating expectations than equivalent (in terms of number of nodes) Gauss-Hermite quadrature or naive Monte Carlo approaches.

We implement this algorithm using the same $20,000$ node state grid and linear interpolation scheme described in Section 5.1, initialising with the optimal discretion policy functions as our guess at the solution. For the quadrature, we use 20 nodes for both the cost-push and demand shocks selected using the discretisation scheme described above.[19]

5.2.2 Optimisation over the exit indicator function

An important pre-requisite to solving this sub-problem is to specify an approximate functional form for the exit indicator function (which is an unknown infinite-dimensional function). In order to make this problem tractable, we model the threshold at which exit occurs as a simple function of the cost-push and demand states to define exit as follows:[20]

$$\mathbb{I}^{EXIT}(u,g) = (a_u u + g > c) \tag{7}$$

Figure 2 graphically illustrates what the exit indicator function looks like for alternative patameterisations of a_u. In each case, a higher value for the constant, c, would shift the schedule

[19]For a given exit indicator function, computing the policy functions under this implementation takes 268 iterations to converge in 213 seconds in 64-bit MATLAB 2012b using a single Intel i7 CPU @ 2.90GHz.

[20]In future versions of the paper, we intend to generalise this functional form. In particular, the region of the state in which the policy rate is unconstrained by the ZLB in the model solved under optimal discretion can be well approximated with the addition of a cross-product term. The resulting exit indicator function would look as follows: $\mathbb{I}^{EXIT}(u,g) = (a_u u + g + a_{ug} ug > c)$. More generally, it is not clear what form the exit indicator function should take (since it is an endogenous object in our problem). In ongoing work, we are exploring a Markov chain approximation to the model in which exit can be directly modelled with 0-1 indicators in an attempt to provide some evidence for the functional form we will use (in addition to the a-priori knowledge of the shape of the threshold at which the ZLB binds in the optimal discretion model).

to the right (exit does not occur unless the demand state is higher) and a lower value would shift the schedule to the left (exit occurs at lower demand states).

Figure 2: Exit threshold function under alternative parameterisations

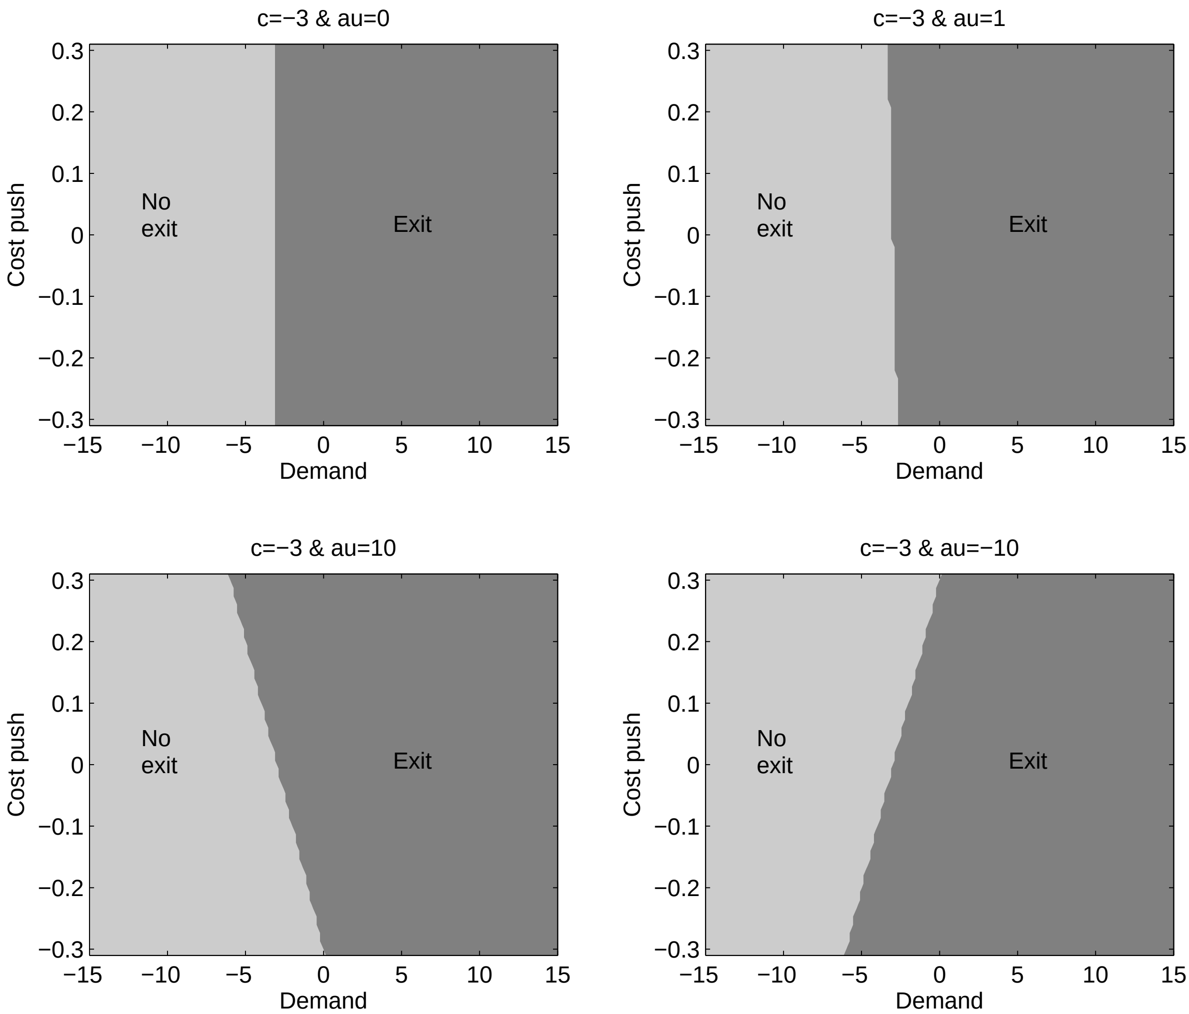

Notes: Each panel shows regions of exit (dark grey) and no exit (light grey) for different parameterisations of the exit indicator function in equation (7).

In general, the parameters of this function that deliver the maximum expected duration subject to the threshold conditions will depend on the particular threshold condition being studied and the initial conditions for the state. Given the calibration of the initial conditions described in Section 6 in the baseline calibration of the model, where the initial condition for the demand state is negative and the initial condition for the cost-push state is zero, then, for a given constant, c, the unconstrained maximum expected duration occurs at $a_u = 0$. This reflects that the distribution of the state converges along the demand axis towards zero (and so sloped exit thresholds 'cut-off' more of that distribution earlier in time). However, a key part of the equilibrium definition is that the thresholds should not be breached in any state of the world, so the constrained maximum expected duration could occur at non-zero values of a_u. Most obviously, inflation is a positive function of cost-push outturns and so one might expect $a_u > 0$ for inflation threshold policies (and that is what we find).

The optimisation approach we use to find the a_u and c paramaters of the exit threshold function that satisfies the definition of equilibrium is to split the problem into two parts. In the outer layer, we search over a_u parameterisations using the Brent optimisation algorithm (which seeks to take steps using a parabola approximation to the minimum, accepting those steps under certain conditions and using a golden search step on rejection). In the inner layer, we find the value for the constant c (given a particular a_u parameter) that maximises the expected duration of the policy regime without violating the threshold conditions (which requires us to solve for the policy functions given the exit indicator function characterised by the valus for a_u and c on that iteration). We approximate the expected duration of the policy using stochastic simulation with $200,000$ draws for the two shocks over 24 periods.[21]

5.3 CBFG policy experiments

Solving for the approximate policy functions that characterise a one-off CBFG policy is relatively straightforward via backward induction. In period K, the final period of the CBFG policy regime, the policy functions can be computed under the assumptions that the policy rate is pegged at the ZLB regardless of the state and that expectations are determined by outcomes in the optimal discretion regime. With the period K policy functions in hand, it is straightforward to work backwards from period $K-1$ to period 1 imposing that the policy rate is pegged at the ZLB and using the policy functions already computed for the period ahead to approximate expectations. We use the same state grid, linear interpolation and quadrature schemes as detailed above.

6 Parameterisation and experiment scenario

For the baseline, which we use to conduct the majority of the analysis in Section 7, we parameterise the model in exactly the same way as Adam and Billi (2006), Adam and Billi (2007) and Bodenstein et al. (2012).[22] The baseline parameter values we use are outlined in Table 1 (where the model is interpreted as a quarterly model). Table 1 also outlines two alternative parameterisations of the model, based on the "RBC" calibration of the model from Adam and Billi (2006). Sensitivity of our policy experiments to these alternative parameterisations is discussed in Appendix B and referred to in Section 7.

Table 1: Baseline model calibration and alternative variants

Parameter	Description	Baseline	Process sensitivity	RBC
α	Calvo parameter	0.6600	0.6600	0.6600
β	Discount factor	0.9913	0.9913	0.9913
σ	Interest elasticity of output gap	6.2500	6.2500	1.0000
θ	Price elasticity of demand	7.6600	7.6600	7.6600
ω	Elasticity of marginal cost	0.4700	0.4700	0.4700
ρ_u	Persistence of cost-push process	0.0000	0.3600	0.3600
σ_u	Standard deviation of cost-push shocks	0.1540	0.1710	0.1710
ρ_g	Persistence of demand process	0.8000	0.8000	0.8000
σ_g	Standard deviation of demand shocks	1.5240	0.2940	0.2940
κ	Slope of the Phillips curve	0.0240	0.0240	0.0569
λ	Weight on output in loss function	0.0031	0.0031	0.0074

[21] For a typical exit threshold function, exit occurs well before the 24^{th} period in almost all of the alternative paths.

[22] The parameters κ, σ and λ originate from Woodford (2003). The parameters of the stochastic processes and the discount factor were estimated by Adam and Billi (2006) on US data using the approach of Rotemberg and Woodford (1998).

As described in Section 2, the policy experiments are ones in which a large negative demand shock drives the policy rate to the ZLB, prompting the policymaker to implement a one-off forward guidance policy. We calibrate the size of the demand shock to deliver a fall in the output gap of 7.5pp, on average, in period one of our simulations for a policymaker who continues to follow optimal discretion. This is approximately equal to the amount by which quarterly GDP fell in the United States during the Great Depression.[23] Note that the demand shock required to deliver this will vary according to the particular parameterisation of the model being considered.

7 Inflation and output gap thresholds

7.1 Inflation thresholds

To inspect the mechanism at work, we start by discussing modal outcomes of three TBFG policies with alternative inflation thresholds compared to calendar-based forward guidance (CBFG) and optimal discretion. The first three panels of Figure 3 plot the modal responses of the endogenous variables (measured in quarterly deviations from steady state) given the alternative policy strategies. The bottom right panel shows the loss in each period associated with the paths for the output gap and inflation generated by each of the alternative policies. When calibrated appropriately, TBFG policy delivers better outcomes and smaller losses relative to the baseline case of optimal discretionary policy: in the initial period, inflation falls by 0.4pp under the discretionary policy but only by 0.2pp under a TBFG policy with an inflation threshold of 0.25 (the deviation of quarterly inflation from target).

The intuition for this result is quite simple. Under the fully discretionary policy, the distribution of future outcomes for the output gap and inflation is negatively skewed close to the ZLB. The negative skew arises because of the asymmetry in the policymaker's response to shocks: at the ZLB, a discretionary policymaker cannot cut rates in response to negative shocks that may arise in the future and cannot commit to any policy plans that would be inconsistent with loss minimization in the future. Other things equal, this reduces expectations of future inflation and activity, which in turn reduces current spending and inflation. In contrast, threshold-based guidance ensures that, for states of the world in which more negative demand shocks arrive, there will be additional stimulus because the policymaker has committed to continue to hold rates at the ZLB for all histories of shocks in which inflation does not breach the announced threshold.

The mechanism to stimulate inflation and activity today by promising inflation in the future is not unique to TBFG policies. A common theme of related work is that history dependent policies such as optimal commitment, price level targeting or the Reifschneider-Williams rule (e.g.) can significantly improve outcomes at the ZLB by using inflation expectations as a substitute for cutting the policy rate.[24]

We can also see from Figure 3 that the policymaker must engineer some overshoot of inflation from target to improve outcomes substantially: if the threshold is set to the inflation target, the equilibrium paths are very close to the discretionary policy. But under a TBFG policy with an inflation threshold of 0.75, outcomes are much improved. The same point is also illustrated by a CBFG policy in which the policymaker holds rates at the ZLB for 4 quarters regardless of

[23]Due to the parametrisation of the model, it is not possible to hit the observed falls in inflation and GDP during the Great Depression at the same time. In a future version of this paper, we intend to explore the sensitivity of the results to a scenario in which we calibrate the demand shock to match the fall in GDP during the Great Recession, rather than the Great Depression.

[24]See also Adam and Billi (2006), Adam and Billi (2007), Nakov (2008), Hills and Nakata (2014), Bundick (2014) and Chattopadhyay and Daniel (2014).

outcomes. This observation demonstrates that the amount by which the policymaker promises to ease future policy is the key driver of the extent to which forward guidance policy boosts activity and inflation in New-Keynesian models. On average, the policy rate stays at the ZLB for longer under TBFG policies when compared to optimal discretion, and the expected duration is increasing in the threshold value (Figure 4).

The distinction between expected outcomes and ex-post outcomes is important for understanding our results. We can examine these effects more directly by plotting the distribution of outcomes for inflation when alternative policies are followed (Figure 5). The top left panel shows the distribution of inflation for our baseline policy assumption of optimal discretion. As discussed above, the negative skew arises because in states of the world when demand is low, the policymaker has no ability to stimulate the economy by reducing the policy rate or by manipulating expectations. This has implications for policy even if the ZLB does not bind. As discussed in e.g. Nakov (2008), the optimal discretionary policy features a "deflationary bias", whereby the average rate of inflation falls short of its target. On the other hand, the output gap is above target on average: in the presence of an occasionally binding ZLB, demand shocks induce a policy trade-off (Adam and Billi (2006), Nakov (2008)).

The bottom right panel shows that if the inflation threshold is set equal to the inflation target (zero), then the distribution of inflation outcomes is almost identical to the baseline case of optimal discretionary policy (top left panel). That result is consistent with the finding that the policymaker needs to set an inflation threshold consistent with a temporary overshoot of the inflation target in order to influence outcomes.

The bottom left panel of Figure 5 illustrates that when the inflation threshold is set at 0.75, the distribution for inflation is narrowed dramatically. The TBFG policy provides stimulus in "bad" states, substantially reducing the negative skew in the distribution. But in "good" states, when positive demand shocks arrive, exit occurs earlier and and the additional stimulus associated with the TBFG policy is removed. The contrast with CBFG (shown in the top right panel) is stark. CBFG imparts stimulus regardless of the state of the economy. While it reduces the negative skew because is raises expectations sufficiently to reduce the impact of the ZLB constraint (Figure 3), it leads to worse outcomes in both good and bad states. That is because this policy provides too much stimulus in good states and insufficient stimulus in bad states. As a result, the variance of the distribution for inflation increases substantially.

Engineering an overshoot of inflation and the output gap is time inconsistent because once inflation and the output gap exceed their targets, the policymaker can reduce welfare losses by reneging on the policy, reverting back to discretion and increasing the policy rate. A measure of the size of the policymaker's incentive to renege in any given period can be computed as the sum of the welfare gains from reneging on the forward guidance policy and reverting to the time-consistent policy (weighted by the probability of being in each state and ignoring the welfare losses from reneging). More formally, denote the measure of time inconsistency of a particular policy, P, in period t, as $\mathbb{T}_t^P$:

$$\mathbb{T}_t^P = \int_{u_t} \int_{g_t} \psi_t\left(u_t, g_t\right)\left(\mathbb{L}_t^P\left(u_t, g_t\right) - \mathbb{L}_t^{OD}\left(u_t, g_t\right)\right) \mathbb{I}\left(\mathbb{L}_t^P\left(u_t, g_t\right) - \mathbb{L}_t^{OD}\left(u_t, g_t\right) > 0\right) \qquad (8)$$

where $\psi_t\left(u_t, g_t\right)$ is a measure of the density over the state in period t[25] and:

$$\mathbb{L}_t^J\left(u_t, g_t\right) = \sum_{s=t}^{\infty} \beta^{s-t} \mathbb{E}_t\left(u_t, g_t\right)\left(\pi_s^J\left(u_s, g_s\right)^2 + \lambda y_s^J\left(u_s, g_s\right)^2\right) \qquad (9)$$

[25]This is independent of the policy regime because the state variables are exogenous in this model.

is the welfare loss associated with policy $J \in \{P, OD\}$ given the state, $\{u_t, g_t\}$. And where $\mathbb{I}(.)$ is an indicator function taking a value of 1 if the loss associated with following the policy concerned exceeds that associated with optimal discretion and 0 otherwise.[26]

Figure 6 illustrates that for TBFG policies, the incentive to renege is quantitatively very small relative to CBFG. Even if the inflation threshold is set to 0.75, which delivered higher modal outcomes than under the CBFG policy, the incentive to renege from the TBFG policy is smaller than for the CBFG policy of four quarters (until the point that the CBFG comes to an end). This demonstrates that TBFG policies can be less time-inconsistent than CBFG, even when they impart more stimulus in expectation.

The fact that TBFG policies are time inconsistent does not necessarily make these policies uninteresting from the perspective of an applied policymaker because there may exist alternative mechanisms to overcome the time-inconsistency problem. For example, Nakata (2014b) demonstrates that policies of this sort can be made time consistent if the policymaker is concerned about their reputation and ZLB episodes are sufficiently frequent and persistent. In that context, TBFG policies are more likely to be supportable by a concern for reputation than CBFG policies because they embody less time inconsistency in the absence of reputational mechanisms.[27]

7.2 Welfare-maximizing inflation and output gap thresholds

In implementing real-world TBFG policy, central banks have experimented with alternative thresholds including inflation (Bank of Japan) and unemployment (Federal Reserve). Some central banks even specified a combination of thresholds rather than a single threshold to reflect their multiple objectives. For example, the Bank of England's Monetary Policy Committee announced in August 2013 that the policy rate would not be increased at least until unemployment had fallen to 7%, subject to a set of "knockouts" on inflation expectations and financial stability.

This section compares the inflation TBFG policies analysed in the previous section to TBFG policies with output gap thresholds.[28] One criterion for comparing alternative TBFG policies is the ex-ante loss. Figure 7 reports ex-ante losses associated with alternative inflation and output gap TBFG policies. The loss-minimizing threshold values are 0.75 and 3 for inflation and output gap TBFG policies respectively. Under our baseline parameterisation, the loss-minimising thresholds are associated with almost identical welfare losses.[29] To understand this result, Figures 9 and 10 document that the modal and mean paths are very similar for the loss-minimizing TBFG policies. Not surprisingly, this pattern can be also observed in time inconsistency measures (Figure 11).

While these alternative optimal TBFG policies share many common features, they differ in important aspects: Figure 8 reports how liftoff from the loss-minimizing TBFG regimes depends on the demand and cost-push state. If liftoff is governed by an inflation threshold, it can be triggered by either a demand or cost-push shock (left panel of Figure 8). The right

[26] In computing the time inconsistency measure, we replace the density measure with a discrete approximation on the state grid from which we can replace the integral with a finite sum – see, for example, Chapter 5 of Heer and Maussner (2005).

[27] Theoretically, it would be possible to make TBFG policies time consistent by allowing the central bank to issue option contracts where the buyer has the right (but is not obliged) to borrow at $\underline{r}$ and lend at: $\underline{r} + (r_t - \underline{r})(T_\pi - \pi_t)$, where $\underline{r}$ is the effective lower bound of the policy rate and T_π is the inflation threshold. The option expires when inflation exceeds its threshold for the first time. If the central bank honours its promise and keeps the policy rate at the ZLB until the threshold is reached, the option is out of the money. In contrast, if the central bank reneges on its promise and increases the policy rate before the threshold is met, then the option is in the money. We would like to thank our colleague Andrew Meldrum for suggesting this idea to us.

[28] Evaluating dual threshold designs is the subject of our ongoing work.

[29] The optimal threshold values and losses associated with them are model-specific, see Appendix B.1.

panel of Figure 8 reports the shape of the exit function for the optimal output gap threshold. In contrast to inflation thresholds, liftoff is independent of the cost-push shock. This reflects that cost-push shocks do not affect output directly in the baseline parameterisation of the model in which they are assumed to be *iid*.[30]

This result has important consequences. Inspection of Figure 9 reveals that the modal path of inflation under the optimal TBFG policy peaks at a value of 0.15 which is well below its threshold of 0.75. In contrast, the path for the output gap almost reaches its threshold value of 3. To understand this result, recall that our equilibrium definition requires that the thresholds are not breached in any state of the world. If the policymaker announces an output gap threshold, cost-push shocks do not determine exit. For inflation thresholds, in contrast, exit can be triggered by a positive cost-push shock and a valid equilibrium has to ensure that the inflation threshold is never breached for all histories of both demand and cost-push shocks.[31]

Unsurprisingly, making cost-push shocks autocorrelated overturns these results (Appendix B.1). In that situation, exit is not independent of cost-push shocks under output gap threshold designs, and mean and modal paths for the output gap do not exceed 1 while the output gap threshold is 3.

The potentially large differences between inflation threshold values and realized paths imply that the way in which the threshold conditions are interpreted can significantly alter the macroeconomic effects of TBFG policies. We adopt an interpretation that requires the thresholds not being breached in any state while the TBFG policy regime remains in effect. One alternative interpretation of the threshold conditions would be that the thresholds should be breached by the smallest possible amount prior to regime exit in all states of the world. For a given threshold value, the latter interpretation would injectmore stimulus: announcing an inflation threshold of 0.75, for example, would ensure that annual inflation rates were at least 3 percentage points above target prior to regime exit. In that case, it follows that the loss-minimising inflation threshold would likely be somewhat lower.

8 Comparison to alternative policies

Section 7 assesses TBFG against optimal discretionary policy. TBFG peforms better at the ZLB than optimal discretion because it embodies a commitment by the policymaker to set the policy rate in the future to improve outomes today. This feature is shared by other history dependent policies where the policymaker responds to a predetermined variable in the model. This modeling choice makes the evolution of policy rates in the future endogenous to outcomes during the recession. Examples of history dependent policies include optimal commitment, a Taylor rule with inertia and price-level targeting.[32] This section compares TBFG to optimal commitment. In future versions of the paper, we intend to braoden the discussion to other history-dependent policies.

[30]The presence of cost-push shocks in the model does affect output via the optimal behaviour of the policymaker under optimal discretion. It is also true that inflation expectations vary as a function of the cost-push parameter in the exit indicator function so that, indirectly, this parameter can affect output. This means that, in the general case, the cost-push parameter need not be zero.

[31]We therefore expect that this result is sensitive to how the exit function is specified. For example, if we introduce some curvature as planned, then exit determination will differ (most obviously in the case of inflation thresholds), which would affect the outomes for the endogenous variables at which exit occurs (which might mean that inflation tends to be closer to the threshold value on exit).

[32]The performance of these (and more) rules relative to optimal commitment and discretion in an economy with an occasionally binding ZLB has been assessed in Nakov (2008).

8.1 Monetary policy under optimal commitment

In the case of optimal commitment, the policymaker is able to commit to an interest rate plan that minimizes the entire discounted sum of future losses subject to the zero lower bound constraint on interest rates and the equilibrium conditions:

$$
\begin{aligned}
&\min_{\{y_t,\pi_t,r_t\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t(\pi_t^2+\lambda y_t^2) \\
&s.t \;\; r_t \geq 1-\frac{1}{\beta} \\
&\quad \pi_t = \beta\mathbb{E}_t\pi_{t+1}+\kappa y_t+u_t \\
&\quad y_t = \mathbb{E}_t y_{t+1}-\sigma\left(r_t-\mathbb{E}_t\pi_{t+1}\right)+g_t \\
&\quad u_t = \rho_u u_{t-1}+\sigma_u\varepsilon_t^u \\
&\quad g_t = \rho_g g_{t-1}+\sigma_g\varepsilon_t^g \\
&\quad \{u_0, g_0\} \;\text{ given}
\end{aligned}
$$

Under the assumption that the zero lower bound has not been binding in any period up to period t, the solution to this problem is the well-known targeting rule:[33]

$$y_t - y_{t-1} = -\frac{\kappa}{\lambda}\pi_t \tag{10}$$

The unconstrained solution to the optimal commitment problem implies that the policymaker trades-off the change in the output gap against inflation, rather than the level of the output gap as in the case of optimal policy under discretion. The presence of the lagged output gap in the rule arises as a consequence of the policymaker's ability to manipulate agents' expectations. The optimal commitment policy response to demand shocks (when the ZLB is not a binding constraint and never has been) is the same as in the optimal discretion case – the policymaker stabilities both inflation and the output gap at target. In response to a cost-push shock, however, the prescription is different. To see that, suppose that there is a positive cost-push shock at time 0. Under both optimal discretion and optimal commitment, the optimal response at time 0 is to allow inflation to rise above target and to allow a negative output gap to open up. After time 0, however, the two policies differ in their prescriptions. Under optimal discretion, the policymaker allows inflation and the output gap to gradually drift back to target as the cost push process dies away. Under optimal commitment, the policymaker commits to *continue* to reduce the output gap as long as inflation is above target. That credible commitment to act in the future reduces the impact of the shock today via agents' expectations. The result is that a policymaker who can credibly commit is better able to stabilize the economy in response to trade-off inducing shocks than one who cannot. This logic also extends to policy at the ZLB (Adam and Billi (2006)).

As in the case of optimal discretion, in the presence of an occasionally-binding ZLB constraint, it is not possible to solve for the equilibrium of the economy analytically. Furthermore, unlike in the case of optimal discretion, the above optimal targeting rule is invalid even if the ZLB is not binding in the current period, provided that it has bound at some point in history.[34] This is a direct consequence of the history dependence of policy which must be taken account of when the model is solved.

[33] See Gertler et al. (1999) for a derivation and discussion.

[34] There are analytical expressions that characterize the solution – see Adam and Billi (2006) – but they also include Lagrange multipliers from the first-order conditions to the Lagrangian representation of the constrained minimization problem.

Figure 12 documents that modal paths for optimal TBFG policies are close to the optimal commitment benchmark.[35] As documented in Adam and Billi (2006) and Nakov (2008), the optimal commitment policy stabilizes the economy by promising inflation above target and positive output gaps in the future. Also, the maximal incentive to renege is similar for optimal commitment and loss-minimizing TBFG policies (Figure 13).

Figure 14 compares the distribution of outcomes under optimal commitment with TBFG policies. Relative to optimal commitment, TBFG policies produce wider distributions for e.g. inflation but not dramatically so. Table 2 summarizes ex-ante losses for a variety of alternative policies considered in this paper. As discussed in Section 7.2, the loss-minimizing TBFG policies are both associated with an ex-ante loss of around 0.32. By comparison, the loss under optimal commitment is 0.21. If the policymaker follows a CBFG policy for 4 quarters, the welfare loss is almost equal to that under optimal discretion.

9 Conclusions

At least since the Bank of Japan's announcement in 1999 to "continue its [unprecedented accommodative monetary] policy until deflationary concerns subside" (Shirai (2013)), state-contingent forward guidance policies have been adopted by several central banks around the world when policy rates became constrained by the zero lower bound (ZLB). Motivated by this observation, we studied a stylised 'threshold-based' forward guidance (TBFG) policy experiment in a fully stochastic New Keynesian model that is solved using global methods. We model the TBFG as a temporary deviation from setting policy optimally on a period-by-period (or 'discretionary') basis. The temporary deviation from optimal discretion is intended to improve outcomes in response to a shock that has caused the ZLB to bind.

In that regard, TBFG can be thought of as a state-contingent form of 'lower-for-longer' policy, whereby the policymaker commits to hold rates at the ZLB for longer than would have been the case (under optimal discretionary policy) in at least some states of the world. By doing so the policymaker can gain leverage over inflation expectations and reduce the real interest rate in the same way as argued by Krugman (1998). But the state-contingency of the commitment also means that TBFG can act as a hedge against the asymmetric effects generated by the ZLB: if further negative shocks arise, prolonging the recession, the threshold will be breached at a later date, providing additional stimulus. In contrast, if positive shocks arrive, the threshold will be breached sooner and the policy stimulus removed. This allows the policymaker to manage the variance of possible outcomes, as well as to improve outcomes in expectation. This intuition is borne out in a quantitative analysis, where we find that TBFG policies are welfare improving relative to the optimal time-consistent policy and are associated with lower mean losses and a lower incentive to renege when compared to forward guidance based purely on calendar time.

Crucially, in order for TBFG policy to be effective, it is necessary for the private sector to understand precisely how the policymaker intends to behave. We demonstrate that that requires the policymaker to specify how they intend to intrepret the threshold conditions. For example, we adopt an interpretation that requires the thresholds not being breached in any state of the world while the TBFG policy regime remains in effect. In the absence of a specific intrepretation of this nature, there is a form of indeterminacy in which there are many policies and macroeconomic outcomes that could be consistent with a particular set of threshold conditions.

[35] The finding that the optimal commitment policy does not keep the policy rate at the ZLB longer than the optimal discretionary policy if the state evolves in line with expectations is just a coincidence in our particular experiment. If the initial condition for the demand state is set to -10 instead of -9.4, the modal ZLB duration under the optimal commitment policy is one period longer than under optimal discretion.

In addition to the requirement that the policymaker fully specify the threshold conditions, our analysis also has other practical implications. In particular, we show that the benefits of TBFG can vary substantially ex-ante depending on the thresholds used (as well as ex-post on how the economy evolves). Moreover, the best (loss-minimising) thresholds vary depending on the parameters of the economy (and the nature of the shocks that buffer it). These results suggest that thresholds need to be carefully designed to maximise the benefits of TBFG.

Table 2: Ex-ante loss and other statistics for alternative policies

	Ex-ante loss
Optimal discretion	0.845
Optimal commitment	0.207
Optimal inflation threshold	0.319
Optimal output gap threshold	0.331
CBFG (4 periods)	0.830

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. Mean ZLB durations are rounded to the next integer.

Figure 3: Modal responses to a large negative demand shock for TBFG policies with inflation thresholds, optimal discretion and calender-based forward guidance

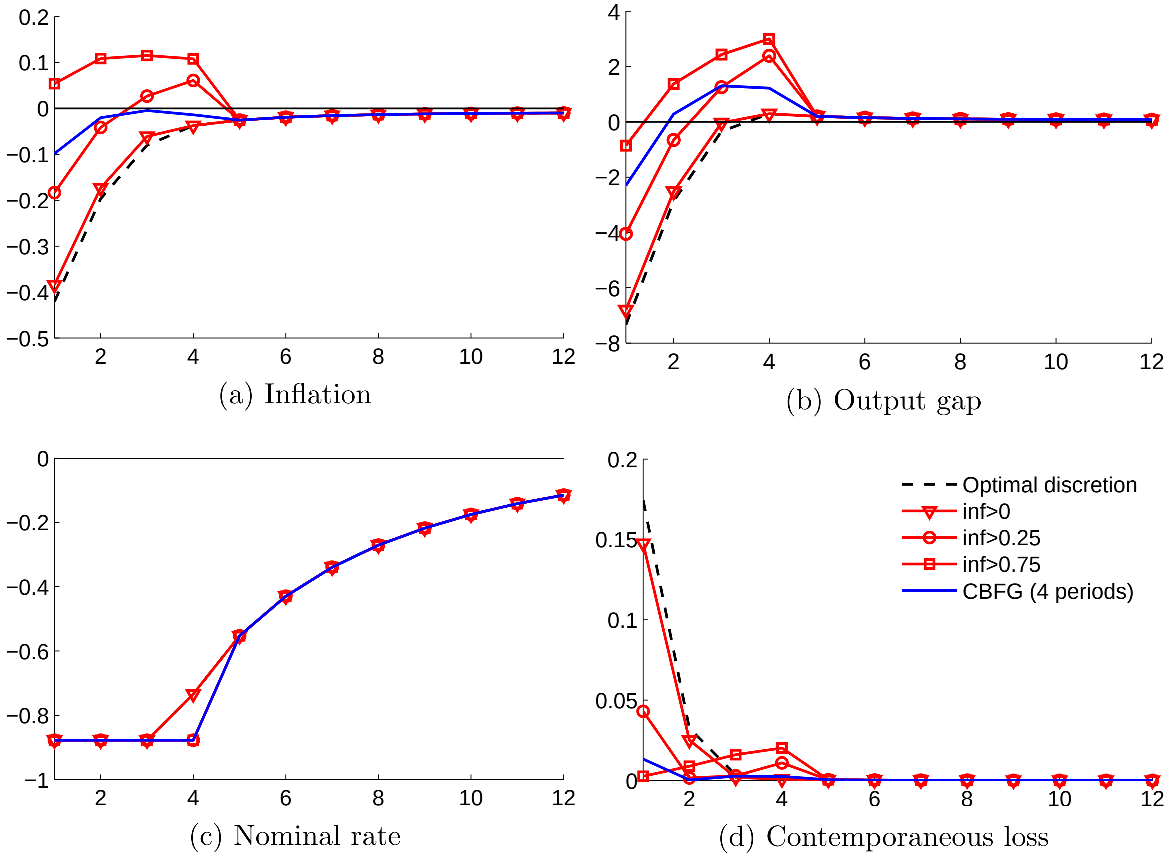

(a) Inflation

(b) Output gap

(c) Nominal rate

(d) Contemporaneous loss

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. Modal responses are computed under the assumption that no shocks arrive after the initial period 0 shock.

Figure 4: Probability of a binding ZLB after a large negative demand shock for TBFG policies with inflation thresholds, optimal discretion and calender-based forward guidance

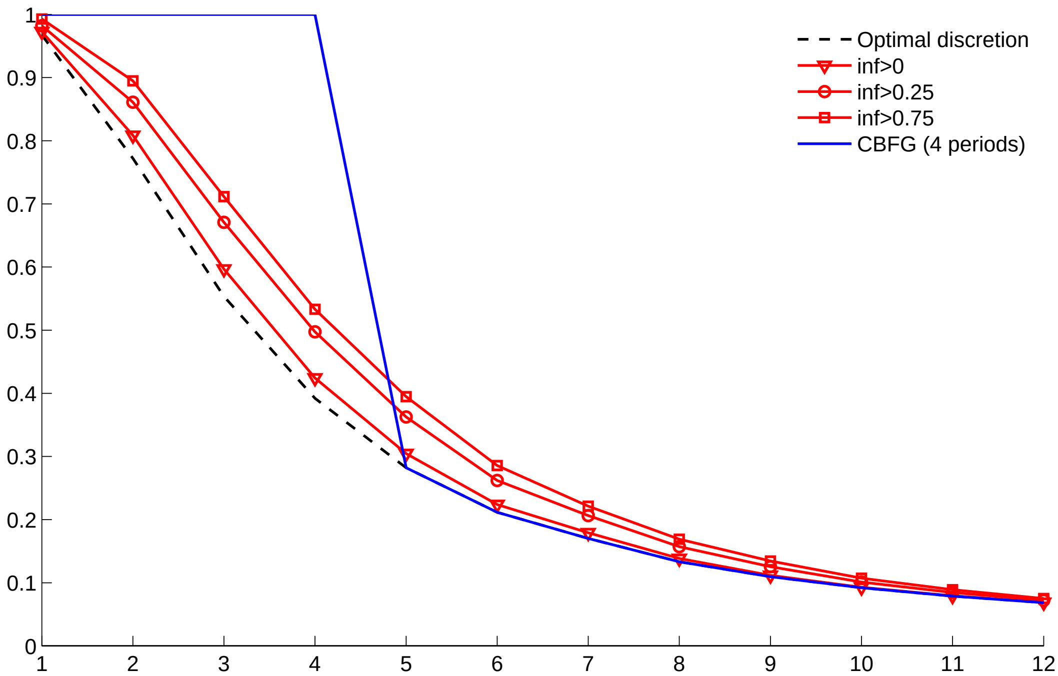

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$.

Figure 5: Distribution of inflation in responses to a large negative demand shock for alternative policies

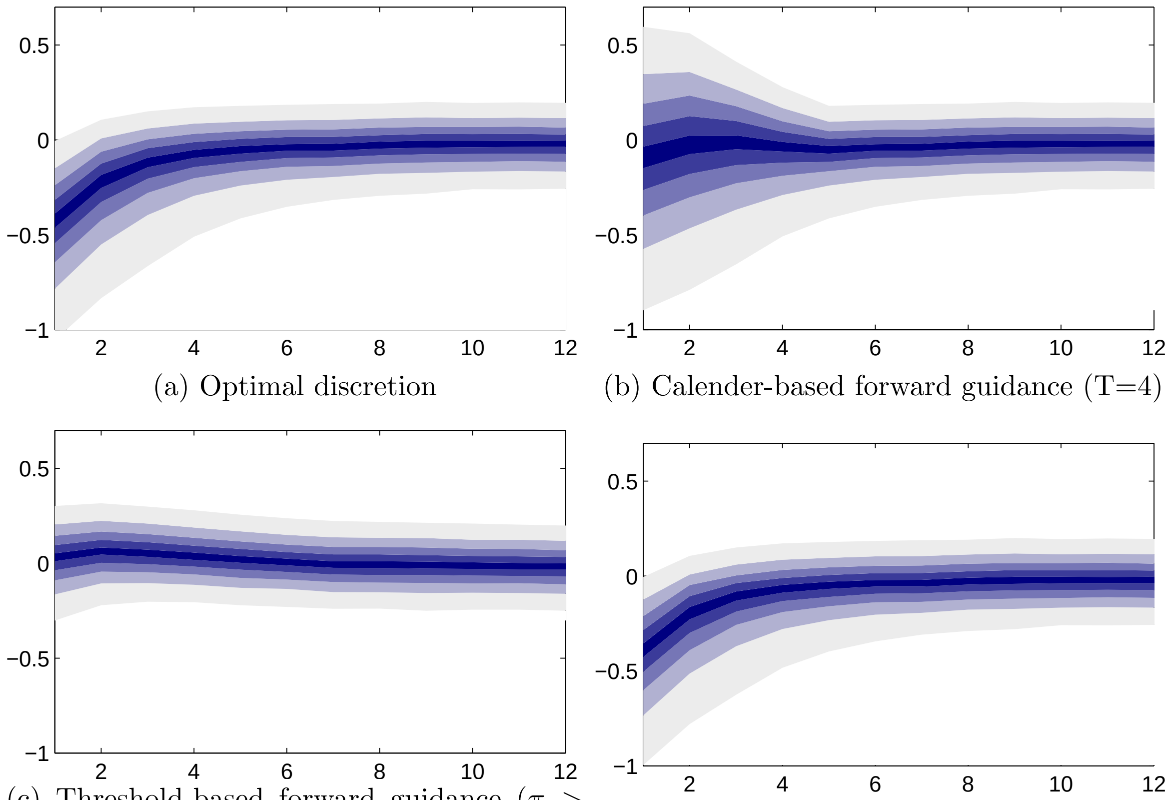

(a) Optimal discretion

(b) Calender-based forward guidance (T=4)

(c) Threshold-based forward guidance ($\pi > 0.75$)

(d) Threshold-based forward guidance ($\pi > 0$)

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. The distributions were computed by simulating the model solutions conditional on each alternative policy given 20,000 draws from the distributions of the shocks per period.

Figure 6: Measures of time inconsistency for TBFG policies with inflation thresholds, optimal discretion and calender-based forward guidance

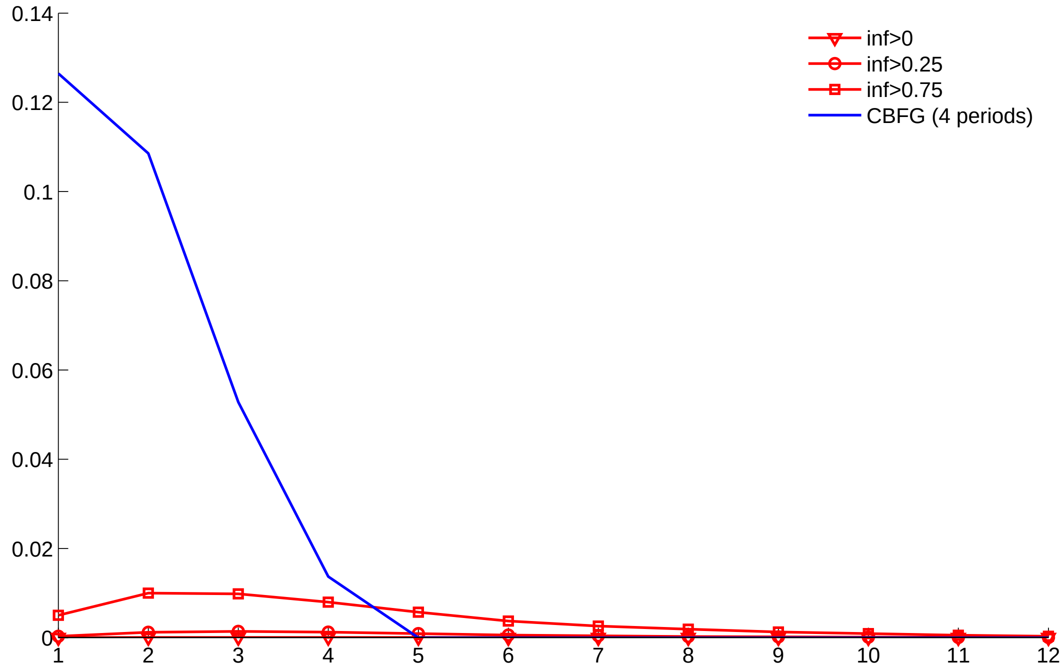

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. The time-inconsistency measure is computed as the difference between the discounted sum of losses given each policy (and taking into account all states of the world) relative to that which could be achieved by reneging and following optimal discretion. A higher number means that the policymaker can achieve better outcomes (lower losses) *on average* by reneging.

Figure 7: Ex-ante losses for inflation and output gap TBFG policies

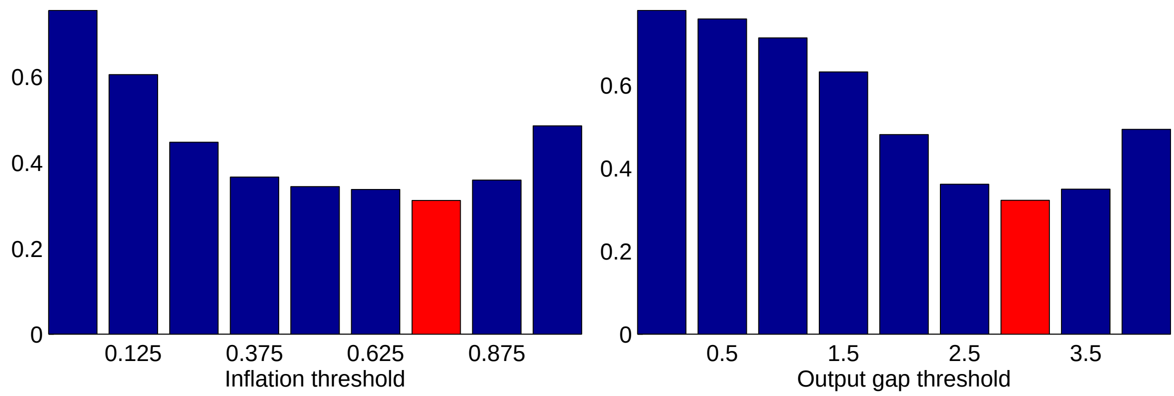

(a) Inflation thresholds

(b) Output gap thresholds

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. Ex-ante losses are computed as the expected discounted value of period losses. For comparison, the loss associated with the optimal discretionary policy is 0.84.

Figure 8: Liftoff across states for optimal l inflation and output gap TBFG policies

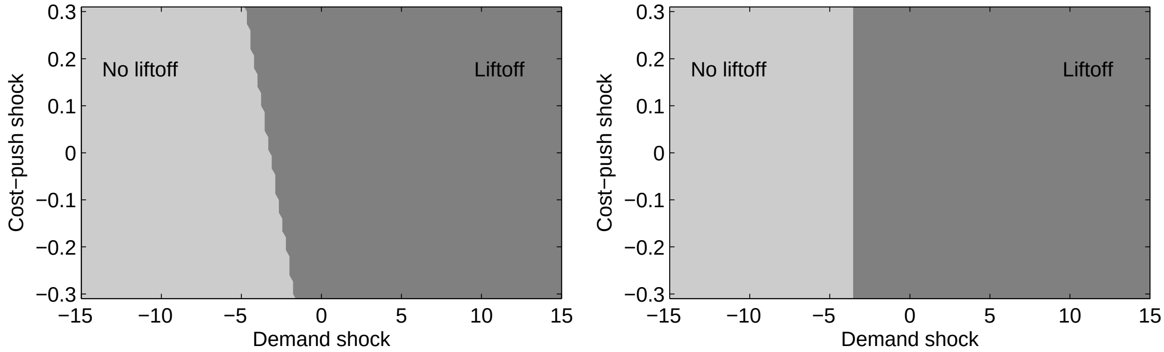

(a) Inflation thresholds (b) Output gap thresholds

Figure 9: Modal responses to a large negative demand shock for inflation and output gap TBFG policies

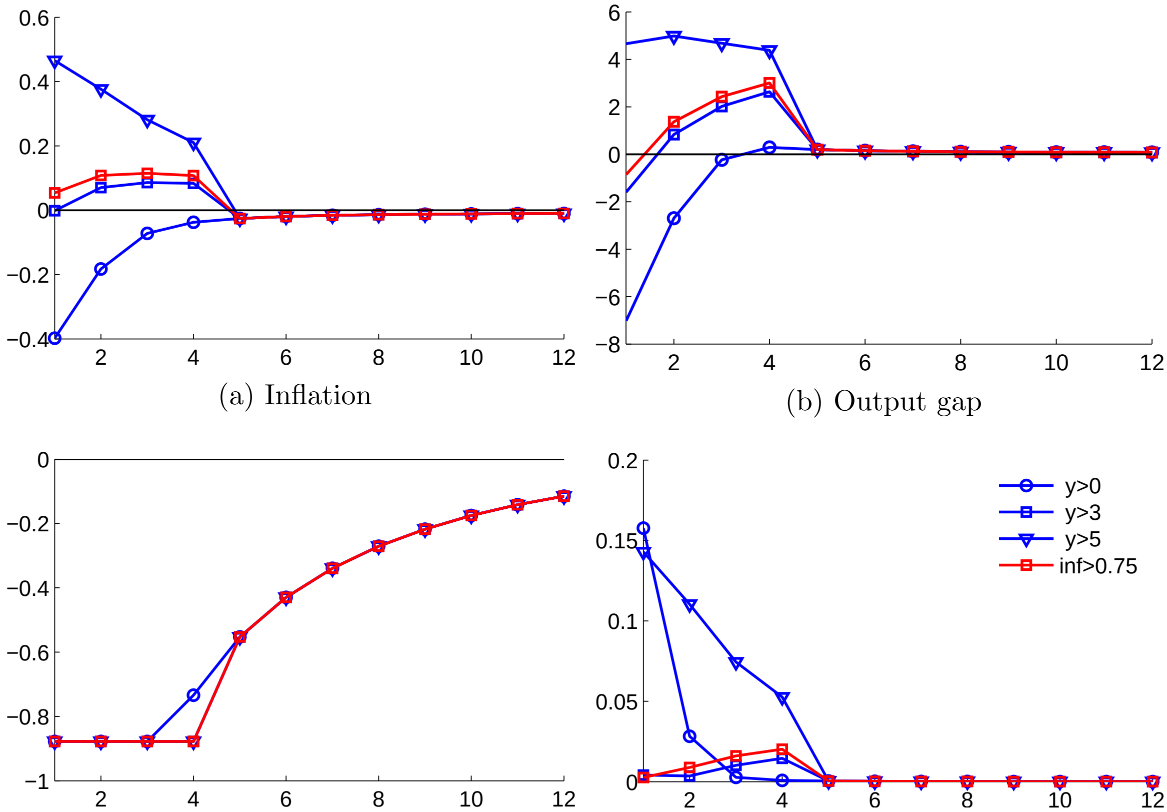

(a) Inflation (b) Output gap

(c) Nominal rate (d) Contemporaneous loss

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. Modal responses are computed under the assumption that no shocks arrive after the initial period 0 shock.

Figure 10: Mean responses to a large negative demand shock for inflation and output gap TBFG policies

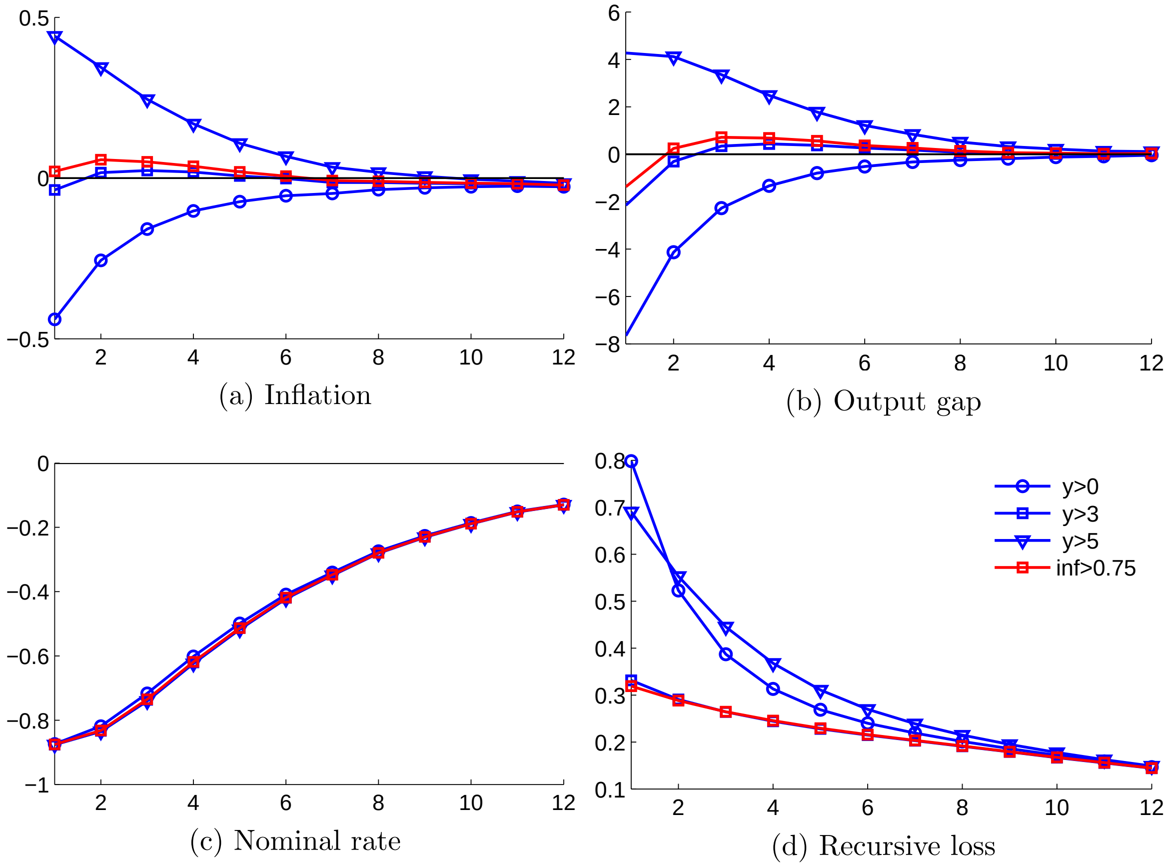

(a) Inflation

(b) Output gap

(c) Nominal rate

(d) Recursive loss

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. Modal responses are computed under the assumption that no shocks arrive after the initial period 0 shock.

Figure 11: Measures of time inconsistency for inflation and output gap TBFG policies

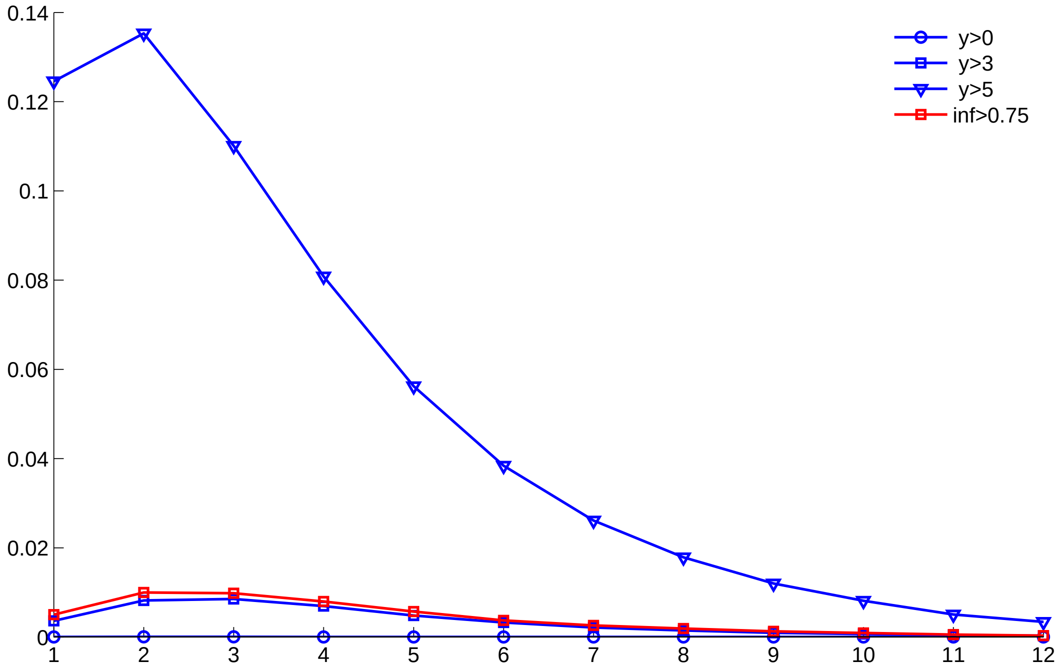

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. The time-inconsistency measure is computed as the difference between the discounted sum of losses given each policy (and taking into account all states of the world) relative to that which could be achieved by reneging and following optimal discretion. A higher number means that the policymaker can achieve better outcomes (lower losses) *on average* by reneging.

Figure 12: Modal responses to a large negative demand shock for optimal TBFG policies and optimal commitment

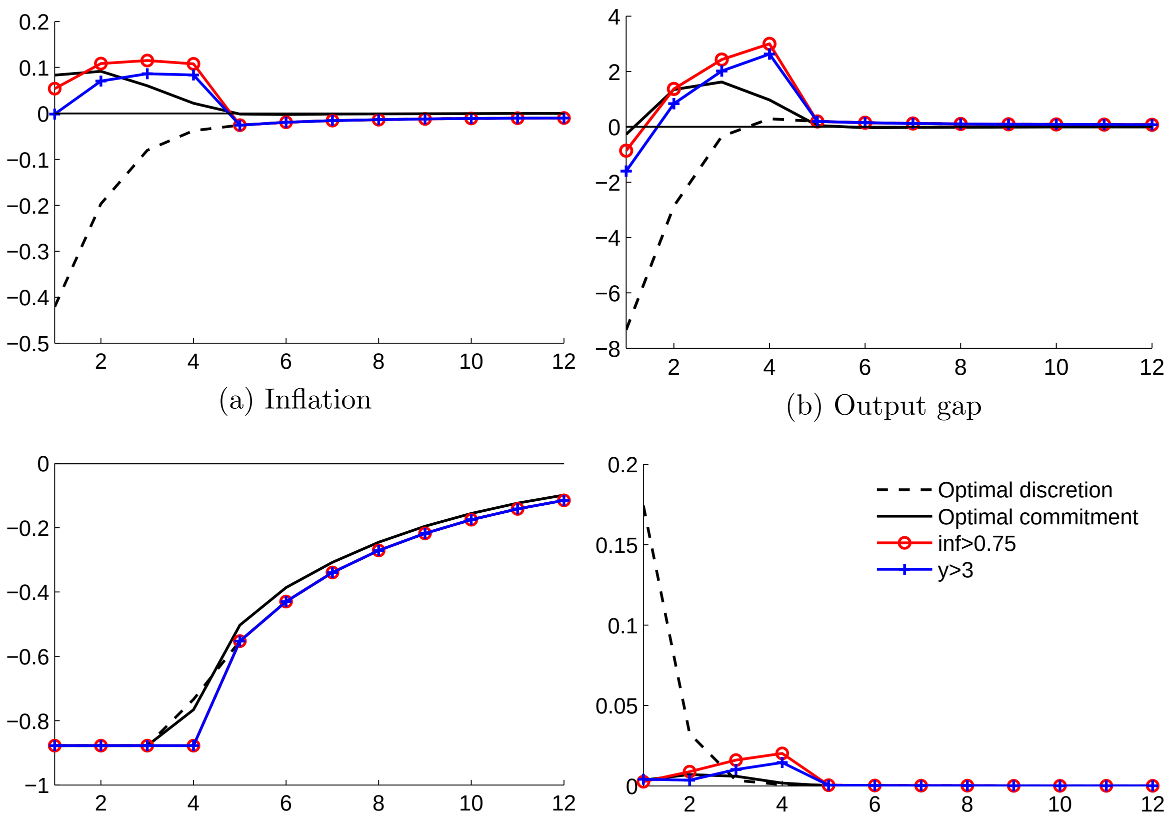

(a) Inflation

(b) Output gap

(c) Nominal rate

(d) Contemporaneous loss

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. Modal responses are computed under the assumption that no shocks arrive after the initial period 0 shock.

Figure 13: Measures of time inconsistency for optimal TBFG policies and optimal commitment

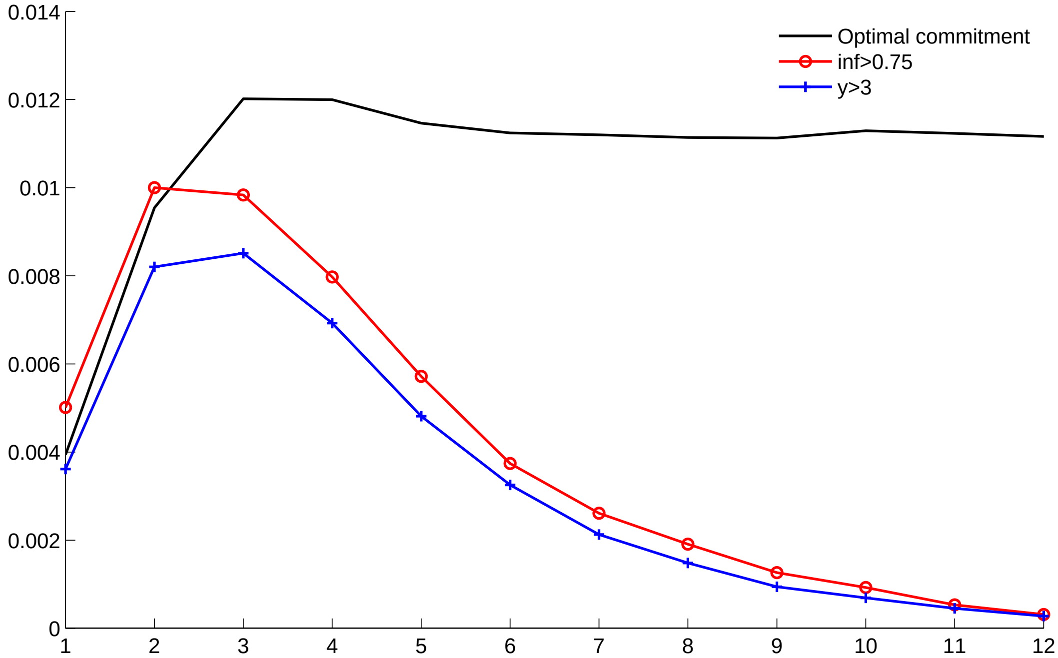

(a) TBFG with constant thresholds

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. The time-inconsistency measure is computed as the difference between the discounted sum of losses given each policy (and taking into account all states of the world) relative to that which could be achieved by reneging and following optimal discretion. A higher number means that the policymaker can achieve better outcomes (lower losses) *on average* by reneging.

Figure 14: Distribution of inflation in responses to a large negative demand shock for an optimal output gap TBFG policy and optimal commitment

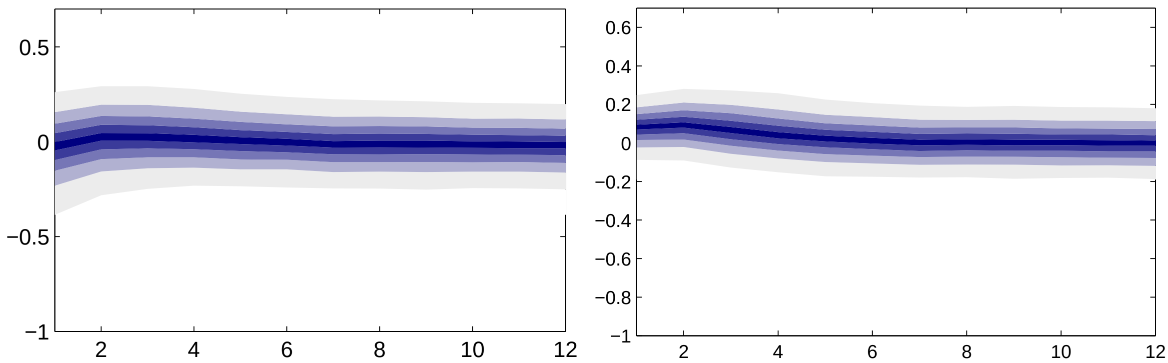

(a) State-dependent forward guidance ($\hat{y} > 3$) (b) Optimal commitment

Notes: The initial condition is $g_0 = -9.4, u_0 = 0$. The distributions were computed by simulating the model solutions conditional on each alternative policy given 20,000 draws from the distributions of the shocks per period.

Bibliography

Adam, K. and R. M. Billi (2007): "Discretionary monetary policy and the zero lower bound on nominal interest rates," *Journal of Monetary Economics*, 54, 728–752.

Adam, K. and R. M. R. M. Billi (2006): "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates," *Journal of Money, Credit, and Banking*, 38, 1877–1905.

Adda, J. and R. Cooper (2003): *Dynamic Economics*, The MIT press.

Bean, C. (2013): "Global aspects of unconventional monetary policies," Remarks at the Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole, Wyoming, Bank of England.

Bodenstein, M., J. Hebden, and R. Nunes (2012): "Imperfect credibility and the zero lower bound," *Journal of Monetary Economics*, 59, 135–149.

Braun, R. A., L. M. Körber, and Y. Waki (2013): "Small and orthodox fiscal multipliers at the zero lower bound," Tech. rep., Working Paper, Federal Reserve Bank of Atlanta.

Bundick, B. (2014): "Real Fluctuations at the Zero Lower Bound," Tech. rep., Mimeo.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–398.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012): "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 44, 1–80.

Carlstrom, C., T. Fuerst, and M. Paustian (2012): "Inflation and output in New Keynesian models with a transient interest rate peg," Bank of England working papers 459, Bank of England.

Chari, V. V. and P. J. Kehoe (2007): "On the need for fiscal constraints in a monetary union," *Journal of Monetary Economics*, 54, 2399–2408.

Chattopadhyay, S. and B. C. Daniel (2014): "The Inflation Target at the Zero Lower Bound," .

Coenen, G. and A. Warne (2013): "Risks to price stability, the zero lower bound and forward guidance: a real-time assessment," Working Paper Series 1582, European Central Bank.

del Negro, M., M. Giannoni, and C. Patterson (2012): "The forward guidance puzzle," Staff Reports 574, Federal Reserve Bank of New York.

Eggertsson, G. B. and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34, 139–235.

Fernández-Villaverde, J., G. Gordon, P. A. Guerrón-Quintana, and J. Rubio-Ramírez (2012): "Nonlinear Adventures at the Zero Lower Bound," Working Paper 18058, National Bureau of Economic Research.

Florez-Jimenez, M. L. and J. A. Parra-Polania (2014): "Forward guidance with an escape clause: When half a promise is better than a full one," Borradores de Economia 811, Banco de la Republica de Colombia.

Gertler, M., J. Gali, and R. Clarida (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

Haberis, A., R. Harrison, and M. Waldron (2014): "Transitory interest-rate pegs under imperfect credibility," Tech. Rep. 1422, Centre for Macroeconomics (CFM).

Harrison, R. (2014): "Estimating the effects of forward guidance in rational expectations models," Discussion Papers 1429, Centre for Macroeconomics (CFM).

Heer, B. and A. Maussner (2005): *Dynamic General Equilibrium Modelling*, Springer.

Hills, T. S. and T. Nakata (2014): "Fiscal Multipliers at the Zero Lower Bound: The Role of Policy Inertia," Finance and Economics Discussion Series 2014-107, Board of Governors of the Federal Reserve System (U.S.).

Krugman, P. R. (1998): "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 29, 137–206.

Leeper, E. M. and T. Zha (2003): "Modest policy interventions," *Journal of Monetary Economics*, 50, 1673–1700.

McKay, A., E. Nakamura, and J. Steinsson (2015): "The Power of Forward Guidance Revisited," Working Paper 20882, National Bureau of Economic Research.

Nakata, T. (2014a): "Optimal Fiscal and Monetary Policy With Occasionally Binding Zero Bound Constraints," Working paper, mimeo.

——— (2014b): "Reputation and Liquidity Traps," Finance and Economics Discussion Series 2014-50, Board of Governors of the Federal Reserve System (U.S.).

Nakov, A. (2008): "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking*, 4, 73–127.

Reifschneider, D. and J. C. Williams (2000): "Three lessons for monetary policy in a low-inflation era," *Conference Series ; [Proceedings]*, 936–978.

Rotemberg, J. J. and M. Woodford (1998): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy: Expanded Version," NBER Technical Working Papers 0233, National Bureau of Economic Research, Inc.

Shirai, S. (2013): "Monetary Policy and Forward Guidance in Japan," Speech, Bank of Japan.

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

——— (2012): "Methods of policy accommodation at the interest-rate lower bound," *Proceedings - Economic Policy Symposium - Jackson Hole*, 185–288.

A Equilibrium selection in a deterministic setting

This appendix builds intuition for the threshold-based forward guidance (TBFG) equilibrium concept outlined in Section 4.1 by use of an example in a deterministic setting. The basic environment is the same as that outlined at the beginning of Section 2: the economy begins with a very low state of demand in period $t = 0$; the policymaker (who ordinarily sets policy on a period-by-period basis to maximize welfare) announces a one-off fully credible forward guidance policy which takes effect in period $t = 1$. However, in this case we assume that the environment is deterministic in the sense that the probability of future shocks arriving is understood to be zero by all agents. The deterministic setting is instructive because any TBFG policy is equivalent to a particular calendar-based forward guidance (CBFG) policy in which the policymaker commits to hold rates at the zero lower bound (ZLB) for a specific number of periods.

Figure 15 shows alternative paths for the output gap in a deterministic version of the model outlined in Section 3 given alternative policies implemented from period 1 and given an initial condition for the state of demand, $g_0 = -11.95$.[36] In the case in which the policymaker continues to set policy following optimal discretion (the black line), the policy rate "lifts off" the ZLB in period 4 (after which the output gap is closed and inflation is at target at all times by virtue of the deterministic setting). The figure also shows paths for the output gap in cases where the policymaker credibly commits to hold rates at the ZLB for one, two and three additional periods in the blue line with circle markers, the red line with cross markers and the green line with plus markers respectively. The differences in outcomes are large and are a non-linear function of the duration of the CBFG policy as has been documented in, for example, Carlstrom et al. (2012).

Suppose that instead of announcing a CBFG policy, the policymaker instead announces a TBFG with an output gap threshold of 1.75, as indicated by the horizontal dashed black line in Figure 15. Which of the four alternative paths shown in Figure 15 is the equilibrium given this TBFG policy? In the absence of additional information, *any* of these paths could be equilibrium paths. To see that, consider the policies with liftoff in periods 5 (the blue line with circles) and 6 (the red line with crosses), which are arguably the most intuitive candidates. The policy with liftoff in period 6 would result in a path for the output gap along which the threshold was breached, but by the smallest amount of all such policies. While the policy with liftoff in period 5 delivers an output gap path that does not cross the threshold in any period, but which comes closest to doing so among all such policies.[37] But the policy with liftoff in period 7 could also be an equilibrium if the policymaker intended that the threshold be breached in every period (but by the smallest amount among all such policies) prior to liftoff. This demonstrates that even in a simple deterministic setting, the announcement of a threshold as part of a TBFG policy is not sufficient to pin down the equilibrium outcome, it is also necessary for the policymaker to specify precisely how they will determine regime exit. And, as an example of the necessity for precision in the policy announcement, suppose that the policymaker announces the output gap threshold along with a statement that the threshold should not be breached at any point prior

[36] Note that the model has been resolved for the deterministic case (with the standard deviations of the shocks set to 0) and the initial condition was set to deliver roughly the same fall in output in period 1 if the policymaker continues to set policy according to the optimal discretion prescription as the mean outcome for output in the stochastic version of the model used for the policy experiments in the main text.

[37] It should be noted at this point that New Keynesian models of the type used with endogenous state variables (e.g. indexation) exhibit 'sign-flipping' behavior, whereby outcomes for output and inflation are an increasing function of the duration for which rates are pegged at the ZLB until that duration crosses a certain threshold when the responses flip sign (see, e.g., Carlstrom et al. (2012)). In that context, the above statements should be interpreted as 'local' statements applying to policies with liftoff periods in the vicinity of that which occurs when the policymaker follows optimal discretion.

Figure 15: Output gap under alternative policies in a deterministic setting

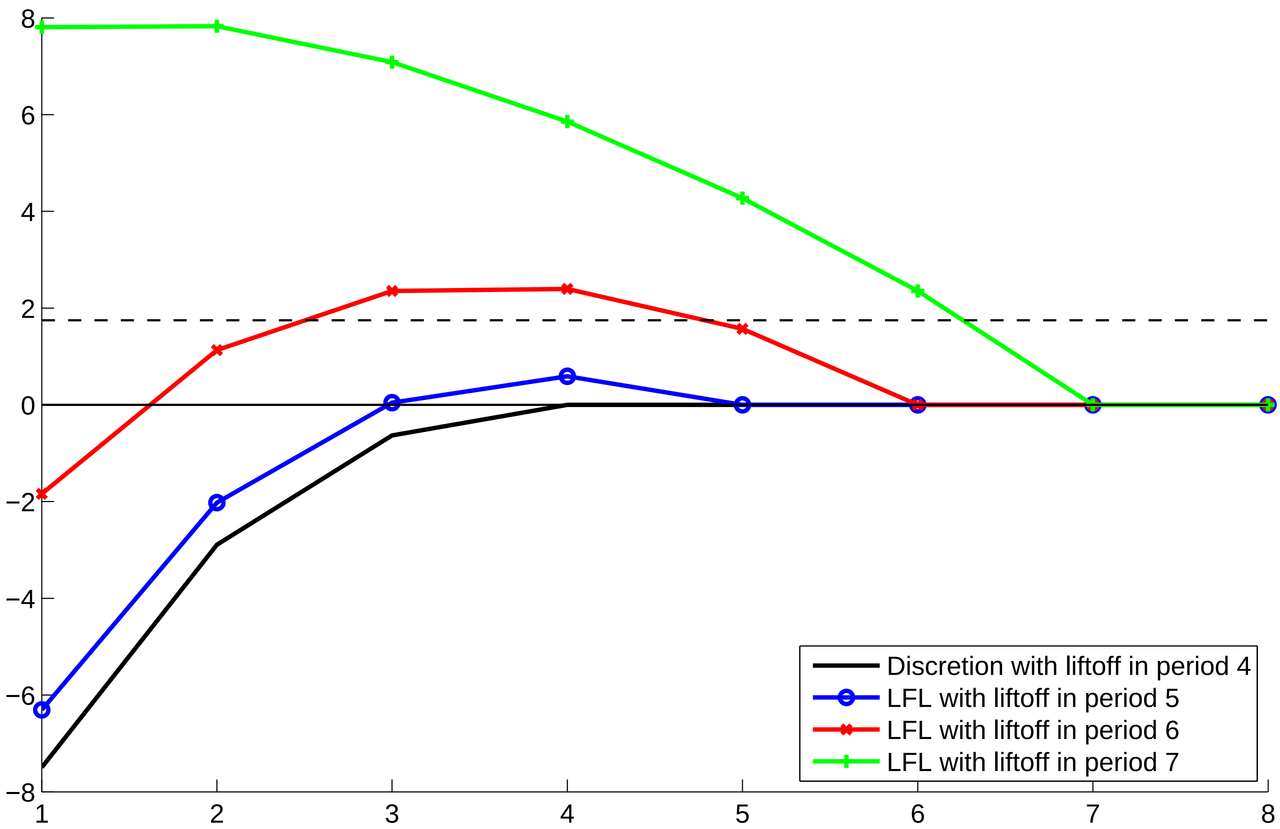

Notes: Computed from an initial condition of $g_0 = -11.95$ and $u_0 = 0$. No shocks arrive or are expected to arrive thereafter. Otherwise, the model is identical to that described in Section 3 of the main text with the baseline calibration outlined in Section 6.

to regime exit. This policy announcement would rule out the policies with liftoff in periods 6 and 7 as equilibria, but would leave open policies with liftoff in any period up to 5 because none of these would result in the threshold being breached in any period. In the main text, we select a unique equilibrium consistent with the threshold conditions as the longest expected duration for the policy subject to the condition that the threshold is not breached in any state of the world (which would select the blue line with circles in the above example).

Finally, we note that this analysis also reveals the dynamic aspect of the policy. The policy with liftoff in period 5 (the blue line with circles) did not breach the threshold in any period. However, in extending the period for which rates are held at the ZLB by 1, the policymaker delivers an equilibrium in which the threshold is breached, but not in the period prior to liftoff. This means that it is not sufficient to determine the validity of a candidate equilibrium by assessment of outcomes in the period prior to liftoff, it is necessary that the entire path (or distribution of outcomes in the stochastic case) is consistent with the specific interpretation of the threshold conditions that the policymaker has announced. Notice also that if the policymaker announced a TBFG policy that selected the policy with liftoff in period 5, then it would be possible to extend the liftoff period to period 6 conditional on agents having previously believed that liftoff would occur in period 5 without the threshold being breached. Our equilibrium concept rules out this type of time inconsistent behavior.

Figure 16: Liftoff across states for inflation and output gap TBFG policies for the autocorrelated cost-push shock calibration

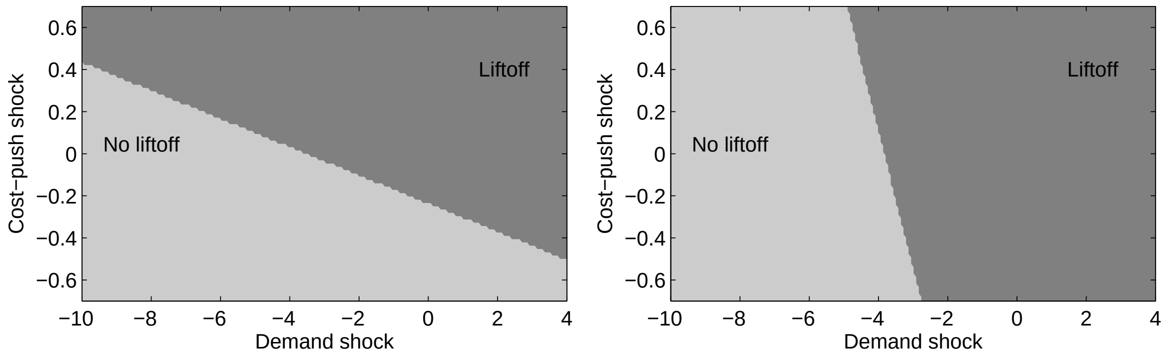

(a) Inflation thresholds (b) Output gap thresholds

B Sensitivity analysis

B.1 Autocorrelated cost-push shocks

In this section, we document how (and why) the optimal threshold values change in a version of the model in which cost-push shocks are not entirely transient and in which their variance relative to that of demand shocks is higher than in the baseline. The complete parametrisation of the model is outlined in Table 1 under "process sensitivity". Relative to the baseline parameterisation: the persistence of cost-push shocks in increased from 0 to 0.36; the standard deviation of cost-push shocks is increased from 0.154 to 0.171; and the standard deviation of demand shocks is lowered from 1.525 to 0.294.

Table 3 documents that the optimal threshold values depend on the specific parametrization of the stochastic processes. Compared to our baseline parametrization with *iid* cost-push shocks, the optimal inflation threshold is lower, and the distribution of liftoff across states is more sensitive to the cost-push shock (Figure 16).

Table 3: Optimal thresholds for the autocorrelated cost-push shock calibration

TBFG policy	Optimal inflation threshold	Optimal output gap threshold	Ex-ante loss
Inflation	0.5	–	0.64
Output gap	–	3	0.68

B.2 RBC calibration

[TO BE COMPLETED]

Figure 17: Modal responses to a large negative demand shock for optimal inflation and output gap TBFG policies for the autocorrelated cost-push shock calibration

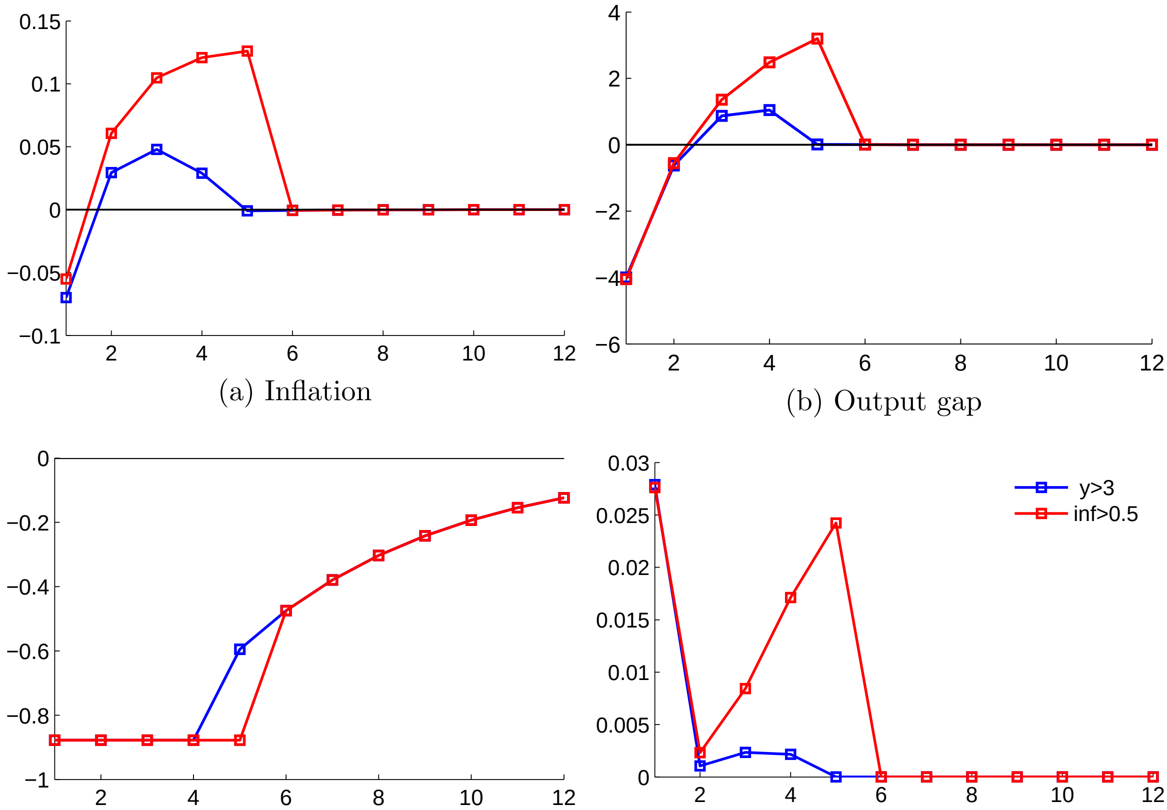

(a) Inflation

(b) Output gap

(c) Nominal rate

(d) Contemporaneous loss

Notes: The initial condition is $g_0 = -11.3, u_0 = 0$. Modal responses are computed under the assumption that no shocks arrive after the initial period 0 shock.

Figure 18: Mean responses to a large negative demand shock for optimal inflation and output gap TBFG policies for the autocorrelated cost-push shock calibration

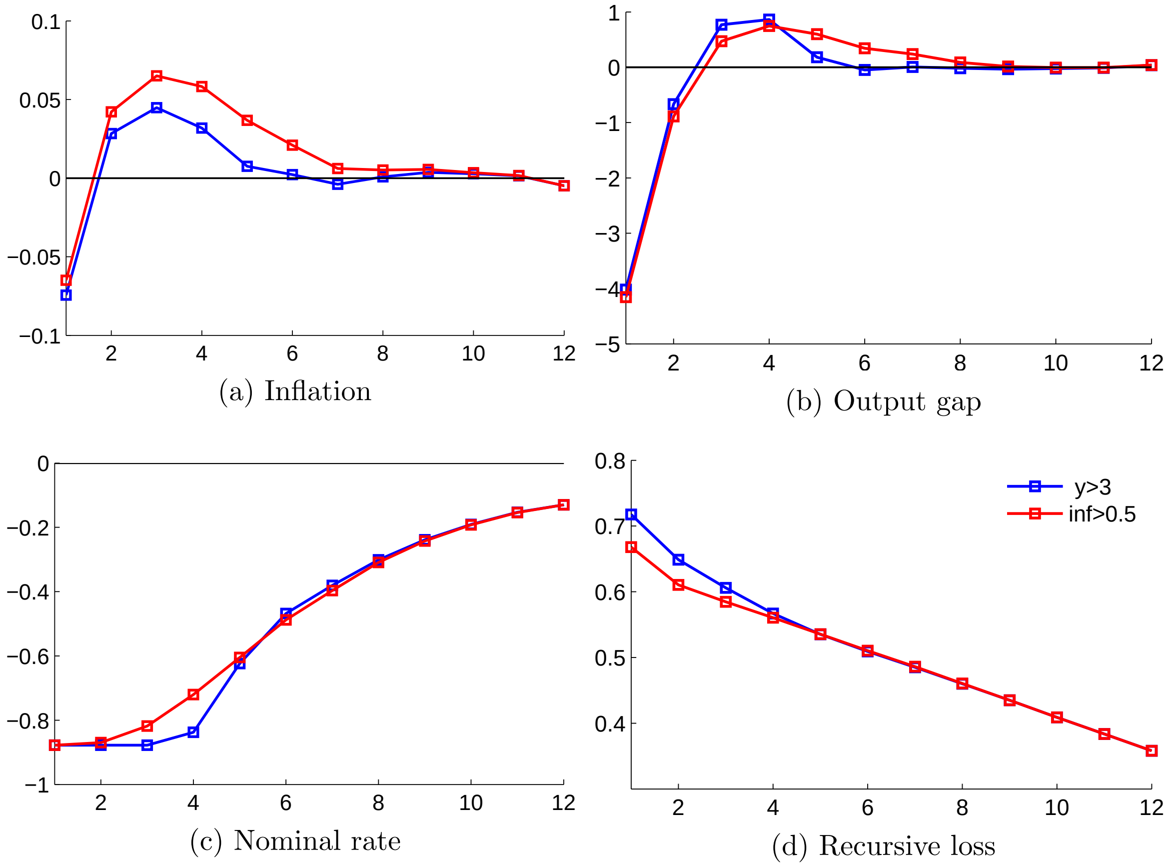

(a) Inflation (b) Output gap (c) Nominal rate (d) Recursive loss

Notes: The initial condition is $g_0 = -11.3, u_0 = 0$. Modal responses are computed under the assumption that no shocks arrive after the initial period 0 shock.