BANK OF ENGLAND

Staff Working Paper No. 561
Threshold-based forward guidance: hedging the zero bound

Lena Boneva, Richard Harrison and Matt Waldron

October 2015

BANK OF ENGLAND

Staff Working Paper No. 561

Threshold-based forward guidance: hedging the zero bound

Lena Boneva,[(1)] Richard Harrison[(2)] and Matt Waldron[(3)]

Abstract

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. We assume that the policy maker makes a state-contingent commitment to hold the policy rate at the zero lower bound (ZLB) in a way that ensures that specific macroeconomic variables (eg 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, exit would be expected to occur later and the policy would provide additional stimulus. In contrast, if positive shocks arrive, so that the economy recovers more quickly than originally expected, exit would be expected to occur sooner, thereby removing some of the policy stimulus. This hedging property of TBFG also means that there is a relatively low incentive for policy makers to renege on the policy, unlike lower-for-longer policies that depend purely on calendar time.

Key words: New Keynesian model, monetary policy, zero lower bound, forward guidance, thresholds.

JEL classification: E17, E31, E52.

---

(1) Bank of England and London School of Economics. Email: lena.koerber@bankofengland.co.uk
(2) Bank of England and Centre for Macroeconomics. Email: richard.harrison@bankofengland.co.uk
(3) Bank of England. Email: matthew.waldron@bankofengland.co.uk

We thank Toni Braun, Fabio Canova, Wouter Den Haan, Taisuke Nakata, attendees at the Sveriges Riksbank workshop on 'Deflation', the Econometric Society 2015 World Congress, and seminar participants at the FRB Atlanta, Bank of England and UCL for their useful comments.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

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 towards 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 motivation for 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 behaviour resembles optimal commitment policy at the ZLB in New Keynesian models as first argued by Krugman (1998) and subsequently demonstrated by Eggertsson and Woodford (2003).[1] However, policymakers have tended to distance themselves from this interpretation, in part because they seem skeptical about their ability to commit credibly to behaviour that is well known to be time inconsistent.[2]

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 way that ensures that selected macroeconomic variables do not exceed pre-specified 'threshold' values while 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.

Our analysis is motivated by policies implemented by the FOMC and MPC, both of whom stated that policy rates would not be increased at least until (among other conditions) the unemployment rate fell below particular threshold values. However, it falls well short of an evaluation of those real-world policies for two main reasons. First, we abstract from many of the details of those policies (e.g. consideration for financial stability concerns). Second, the communications that accompanied those policies tended to emphasise their role in clarifying central bank behaviour rather than in providing stimulus. With that in mind, our exercise could be regarded as an evaluation of 'what if' a central bank did employ TBFG to impart stimulus at the ZLB.

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 aggregate demand (the 'IS' curve) and the pricing decisions of firms (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 to minimise the expected discounted value of a loss function derived from a second order approximation to household's utility, subject to the ZLB constraint. Our baseline assumption is that the policymaker acts with 'discretion', taking the behaviour 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 stabilise aggregate demand. This motivates our experiments in which the policymaker attempts

[1] There are several other policy prescriptions (like price level targeting or the Reifschneider and Williams (2000) rule) that can also deliver better outcomes at the ZLB via the same mechanism.

[2] 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."

to improve outcomes by temporarily deviating from time-consistent policy. Specifically, under TBFG the policymaker makes a state-contingent commitment to hold the policy rate at the ZLB for longer than agents were expecting under the time-consistent policy. 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 back to setting the optimal discretionary policy 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 be able to understand the policy, it is not sufficient for the policymaker to announce a set of thresholds for macroeconomic variables. It is also necessary for the policymaker to announce precisely what the threshold conditions mean. There are many different interpretations of the threshold conditions. The approach we take in this paper is to define the set of feasible state-contingent commitments as those in which the threshold conditions are not breached in any state of the world in which the forward guidance regime remains in effect. And then to select a unique equilibrium from that set as that which maximises the expected duration of the regime.[3]

Our baseline results compare the behaviour 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 substantially improve welfare compared with fully time consistent behavior. Part of 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. But, 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 be held at the ZLB for longer. By contrast, if positive shocks arrive, so that the economy recovers more quickly from the recession than originally expected, then exit from the ZLB will occur sooner and the policy stimulus will be removed.

So TBFG can be viewed as a hedge against the asymmetric effects generated by the ZLB constraint. The magnitude of the effect can be seen by comparing losses 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 regardless of the state of the economy.[4] As in the case of TBFG, this can improve outcomes in expectation and eliminate the negative skew in outcomes induced by the ZLB constraint. However, CBFG leads to worse outcomes for both positive and negative realisations of future demand shocks than appropriately calibrated TBFG because it provides too much stimulus in 'good' states and insufficient stimulus in 'bad' states.[5] As a result, the variances of the distributions of the output gap and inflation are substantially larger.

Because our policy experiments are based on a temporary deviation from time-consistent behaviour, 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. 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

[3]The macroeconomic effects of a TBFG with a given set of threshold conditions is dependent on the precise specification of the exit conditions.

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

[5]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.

policy. By contrast, under TBFG, for realisations of the shocks in which the economy recovers more quickly, the exit thresholds are breached sooner and policy automatically reverts to time-consistent behavior.

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 approximate optimal values for both inflation and output gap thresholds. Unlike CBFG, optimal TBFG policies achieve ex-ante losses that are close to the optimal commitment policy.

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 specify thresholds 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 that model, they are restricted to perfect foresight approximations of expectations, whereas we compute a fully stochastic equilibrium.

The rest of the paper is organised as follows. Section 2 describes the policy experiments and the assumptions underpinning them. Section 3 details the model and the baseline description of policy. Section 4 defines equilibrium for both threshold-based and calendar-based forward guidance policies. Section 5 describes the methods we use to solve for equilibrium. 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 describes the simulation results, including comparisons of TBFG to CBFG and optimal commitment policy. Section 8 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 some arbitrary 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.

There are two overarching assumptions governing the nature of our experiments. First, the forward guidance policy is assumed to be transitory or 'one off': before implementation, the policy is entirely unanticipated by agents in the model and, once the the regime has ended, agents attach 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

Figure 1: Timeline of events for policy experiments

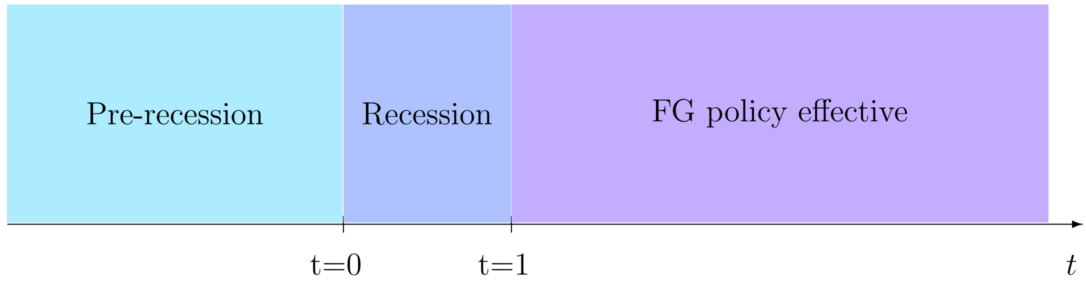

of policy from a rule governing the timeless behaviour of the policymaker (e.g. del Negro et al. (2012), Coenen and Warne (2013), Haberis et al. (2014)). This means that these policy experiments are not conducted under rational expectations and so are subject to the issues studied by Cooley et al. (1984) among others. Specifically, one may obtain misleading results from implementing a 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 policymakers to adopt a similar policy in the future, should the ZLB become a binding constraint on policy again. 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 by Nakata (2014), the assumption of full credibility may be reasonable if reneging on a policy has reputational costs for the policymaker. In that setting, the likelihood of the policymaker sticking to their policy plan (and hence the credibility of the announcement) depends on the costs and benefits of reneging: other things equal, a policy with a smaller incentive to renege is more likely to be viewed as credible than one with a larger incentive to renege.

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

[6]Modelling forward guidance at the ZLB with rational regime switching 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.

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 optimality conditions are log-linearised around the non-stochastic steady state.[9]

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[10]), taking agents' expectations as given. As in Adam and Billi (2007), the policymaker solves the following constrained minimisation problem:

$$\min_{\{y_t,\pi_t,r_t\}} \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i (\pi_{t+i}^2 + \lambda y_{t+i}^2)$$

$$s.t \;\; r_t \geq 1 - \frac{1}{\beta} \tag{1}$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + u_t \tag{2}$$

$$y_t = \mathbb{E}_t y_{t+1} - \sigma\left(r_t - \mathbb{E}_t \pi_{t+1}\right) + g_t \tag{3}$$

$$u_t = \rho_u u_{t-1} + \sigma_u \varepsilon_t^u \tag{4}$$

$$g_t = \rho_g g_{t-1} + \sigma_g \varepsilon_t^g \tag{5}$$

$$\mathbb{E}_t\{y_{t+i}, \pi_{t+i}, r_{t+i}\}_{i=1}^{\infty} \text{ given}$$

$$\{u_t, g_t\} \text{ given}$$

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; $\lambda = \kappa/\theta$ is the relative weight on output in the loss function; u and g are exogenous disturbances to inflation and demand, often called cost push and demand shocks[11], both of which are assumed to follow AR(1) processes with $\varepsilon_t^u \sim iid$ N(0,1), $\varepsilon_t^g \sim iid$ N(0,1), ρ_u and ρ_g the persistence parameters, and σ_u and σ_g the standard deviations.

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 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

[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.

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

[11]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.

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

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 requires us to use numerical methods to solve for the model's equilibrium, as described in Section 5.

4 Equilibrium in the forward guidance regime

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

4.1 TBFG equilibrium definition

We restrict the feasible set of policies that the policymaker can implement to the announcement of a time-invariant policy, whereby the distribution of expected outcomes (including the status of the policy regime) is a function of the state of the economy, not of the time period in which those expectations are taken. 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.

Conditional on that restriction, equilibrium is characterised by a time-invariant state-contingent indicator function defining the states of the world in which the TBFG policy regime is terminated, consistent with a pre-announced threshold on inflation or the output gap. As explained in Appendix A using a simple deterministic example, the threshold condition alone is not sufficient to determine uniquely the state-contingent exit indicator function. It is also necessary for the policymaker to announce precisely how they intend to interpret the threshold conditions. For example, the real-world TBFG policies implemented by the FOMC and the Bank of England's MPC drew a distinction between 'thresholds' and 'triggers'. Thresholds were conditions that, if breached, would prompt a reassessment of the policy, whereas if a trigger condition were breached the regime would automatically come to an end and the policy rate raised. Under this taxonomy, the equilibrium definition we use in this paper is a kind of state-contingent trigger. The state-contingent exit indicator function is such that were the policy regime to be maintained in at least some additional states, the threshold condition would be breached in at least some (possibly different) states. An alternative way of thinking about this is that we restrict the set of feasible exit indicator functions to those that ensure that the threshold condition is not breached in any states of the world in which the forward guidance regime remains in effect. From that set, we select as the unique equilibrium the one that maximises the expected duration of the policy (equivalent to maintaining the regime in as many states as possible subject to the threshold condition not being violated).[13]

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

[13] It should be noted that there are alternative, equally valid, equilibrium definitions. In future work, we intend to explore the distinction between thresholds and triggers by incorporating probabalistic exit into the analysis, whereby the breach of a threshold condition would trigger regime exit with a non-zero, but non-unitary probability.

[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 $t+1$.

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:
$\max \sum_{t=1}^{\infty} t \int_u \int_g \psi_t^{FG}(u,g)$
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 likelihood that the TBFG regime is still in effect in period t, 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 analogously; $\psi_0^{FG}(u,g) = 1$ 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).

There are three 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}. So the transmission of forward guidance policies in this model is via agents' expectations and the macroeconomic effect of the policy depends on the precise exit conditions that the policymaker specifies. It follows that 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. In this framework, TBFG is a state-contingent form of 'lower-for-longer' policy. Second, 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). Third, the initial condition for the economy in period $t = 0$ affects the equilibrium because it affects the likelihood of the TBFG policy being in place at each date (ψ) and hence the expected duration of the policy. Different initial conditions would result in different equilibria (given a particular set of thresholds).

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

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 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 our initial guess, we use the solution to a version of the model in which the ZLB constraint is ignored (which can be solved analytically). The algorithm is then comprised of an outer layer and an inner layer. In the outer layer, 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). In the inner layer, outcomes for the endogenous variables are solved analytically as a sequence of independent static problems (for each node in the state grid) conditional on the approximation of expectations.[15] 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 $1e^{-6}$.[16]

5.2 TBFG policy experiments

The objective is to find policy functions for inflation, $\pi^{FG}(u, g)$, and the output gap, $y^{FG}(u, g)$, 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 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 in Section 5.1. 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

[15] First, solve for outcomes on the assumption that the ZLB is not binding in the following way: (i) use the first-order condition for the policymaker in equation (6) to substitute the output gap out of the Phillips curve (equation (2)) 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 (3)) to compute the interest rate as a function of the output gap, the expected output gap, expected inflation and the demand state. If the interest rate is greater than or equal the ZLB, then the solution (conditional on expectations) has been found and stop. If the interest rate violates the ZLB constraint then: (i) set the interest rate equal to the ZLB; (ii) compute the output gap conditional on the interest rate, expectations and the demand state using the IS curve; (iii) compute inflation conditional on the output gap, expectations and the cost-push state using the Phillips curve.

[16] 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).

forward guidance policy regime is active, expectations are a weighted average of outcomes in two different policy regimes. Standard quadrature like the 5 node Gauss-Hermite 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.[17]

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:[18]

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

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

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 baseline calibration described in Section 6, 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 a mean of 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

[17] As an indicitive guide to computational effort (which varies depending on the 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.

[18] 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 is 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 stochastic processes, which would permit exit to be directly modelled with 0-1 indicators.

Figure 2: Exit threshold function under alternative parameterisations

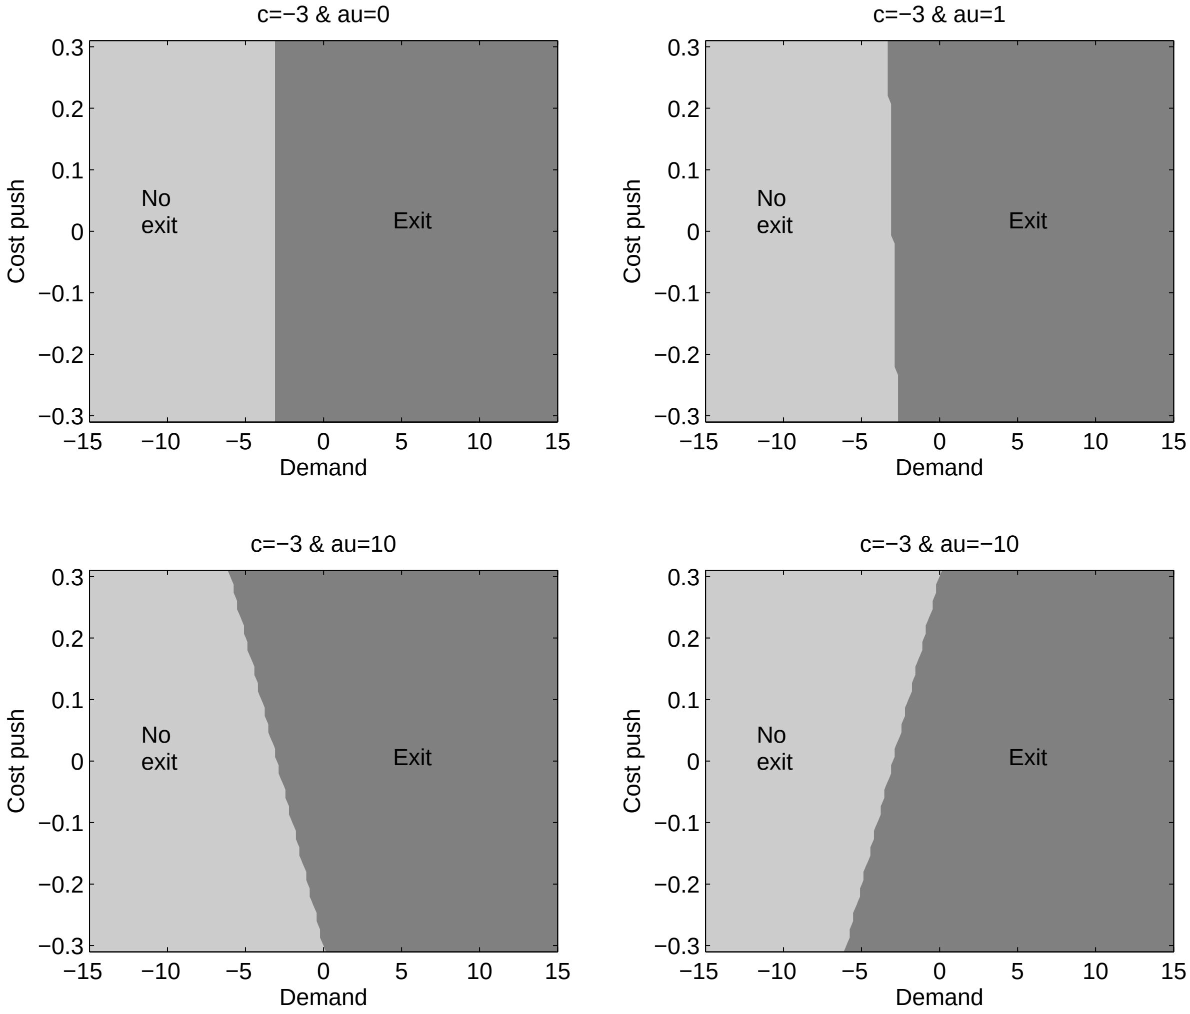

Notes: Each panel shows regions of exit (dark grey) and no exit (light grey) for different parameterisations of the exit indicator function in equation (7).

expected duration could occur at non-zero values of a_u. For example, inflation is a positive function of the cost-push state 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.[19]

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

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).[20] The baseline parameter values we use are outlined in Table 1 (where the model is interpreted as a quarterly model). Sensitivity of our policy experiments to an alternative parameterisation for the stochastic processes is discussed in Appendix B and referred to in Section 7.

Table 1: Baseline model calibration

Parameter	Description	Value
α	Calvo parameter	0.6600
β	Discount factor	0.9913
σ	Intertemporal elasticity of substitution	6.2500
θ	Price elasticity of demand	7.6600
ω	Elasticity of marginal cost	0.4700
ρ_u	Persistence of cost-push process	0.0000
σ_u	Standard deviation of cost-push shocks	0.1540
ρ_g	Persistence of demand process	0.8000
σ_g	Standard deviation of demand shocks	1.5240
κ	Slope of the Phillips curve	0.0240
λ	Weight on output in loss function	0.0031

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.[21]

[20]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).

[21]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.

7 Results

As a summary of the results, Table 2 compares ex-ante losses of the welfare-maximising inflation and output gap TBFG policies with those of the welfare-maximising CBFG policy and optimal discretion. It shows that the policymaker cannot achieve quite the same level of welfare as under optimal commitment, but can substantially improve outcomes relative to optimal discretion and CBFG.

Table 2: Ex-ante losses as a ratio to loss under the optimal commitment policy

Policy	Ratio
Optimal discretion	4.1
Optimal inflation TBFG	1.5
Optimal output gap TBFG	1.6
Optimal CBFG	4.0

Notes: Ex-ante losses computed from a stochastic simulation of $20,000$ draws over 24 periods from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

The rest of this section discusses these results in more detail by comparing TBFG to CBFG and optimal discretion in Section 7.1, discussing loss-minimisng inflation and output gap thresholds in Section 7.2, and comparing TBFG to optimal commitment in Section 7.3.

7.1 TBFG compared to CBFG and optimal discretion

To inspect the mechanism at work, we start with the modal outcomes of three TBFG policies with alternative inflation thresholds compared to 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 per-period outcomes for the output gap and inflation generated by each of the alternative policies. Under the fully discretionary policy, inflation and the output gap are negative at the ZLB because the policymaker cannot cut rates below the ZLB 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. By contrast, when calibrated appropriately, CBFG and TBFG policy deliver better outcomes and smaller losses relative to the baseline case of optimal discretionary policy.

The mechanism behind the improvement in outcomes delivered by (appropriately-calibrated) TBFG and CBFG is evident from panel (c). The policy rate is held at the ZLB for one additional period relative to optimal discretion under both of the positive inflation TBFG policies and the CBFG policy. By promising looser policy in the future, the policymaker can boost inflation and activity today via the effect of the commitment on expectations. This mechanism is not unique to TBFG or CBFG policies. A common theme of related work is that history dependent policies such as optimal commitment, price level targeting or the Reifschneider-Williams rule can significantly improve outcomes at the ZLB by using inflation expectations as a substitute for cutting the policy rate.[22]

We can also see from Figure 3 that the policymaker must commit to a TBFG with an above-target inflation threshold in order to improve outcomes relative to the optimal discretionary baseline: if the threshold is set equal to the inflation target (zero), the equilibrium paths are very close to those under discretionary policy. But under a TBFG policy with an inflation threshold

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

Figure 3: Modal responses under alternative inflation TBFG policies, optimal discretion and CBFG

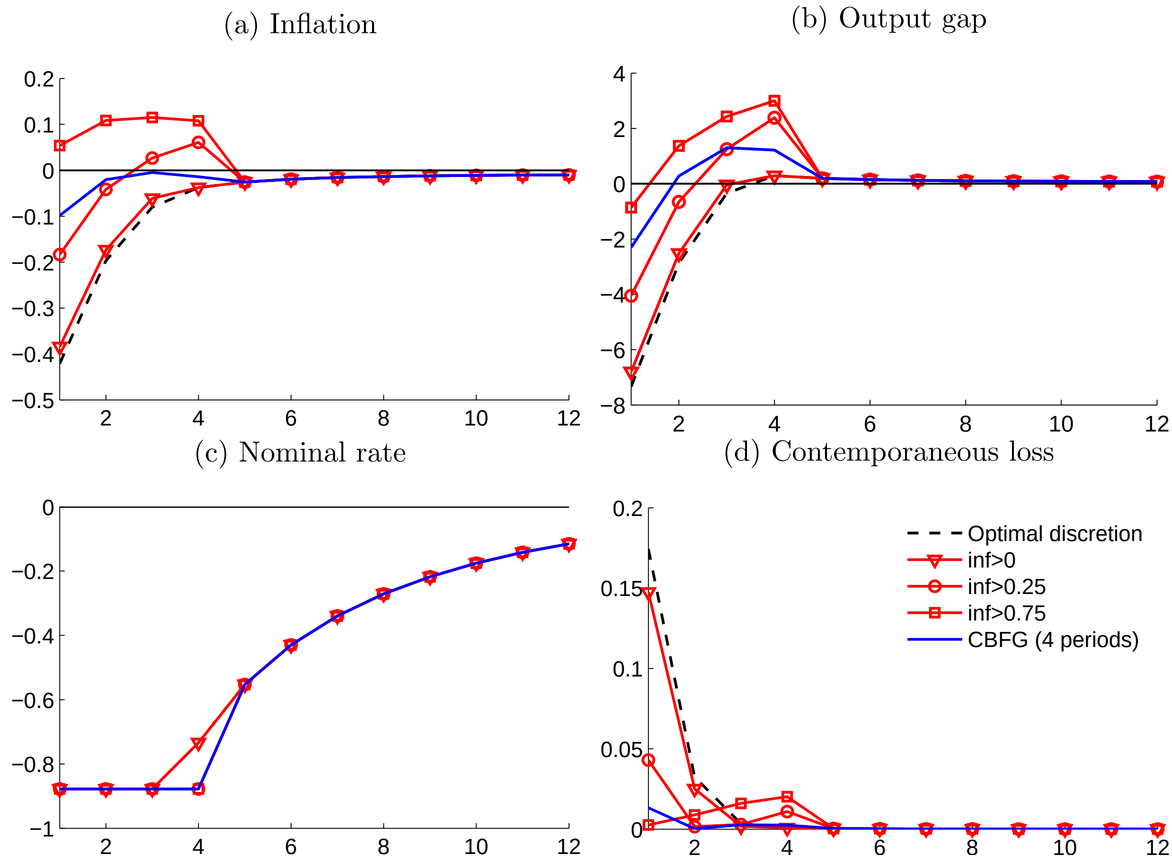

Notes: Computed under the assumption that no shocks arrive after the initial period 0 and from an initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

of 0.75, outcomes are much improved. 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 the differences between these alternative policies. Figure 5 shows the distribution of outcomes for inflation. Panel (a) shows that the distribution of inflation is negatively skewed under optimal discretion because in states of the world where 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. Accordingly, 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)).

Panel (d) 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 (panel (a)). That is because there are very few states of the world in which the policymaker is committing to hold rates at the ZLB for longer than they would have done had they continued to follow the optimal discretionary policy. Panel (c) of Figure 5 illustrates that when the inflation threshold is set at 0.75, the distribution of inflation is narrowed dramatically. The TBFG policy provides stimulus in "bad" states, substantially reducing the negative skew in

Figure 4: Probability of rates being at the ZLB under alternative inflation TBFG policies, optimal discretion and CBFG

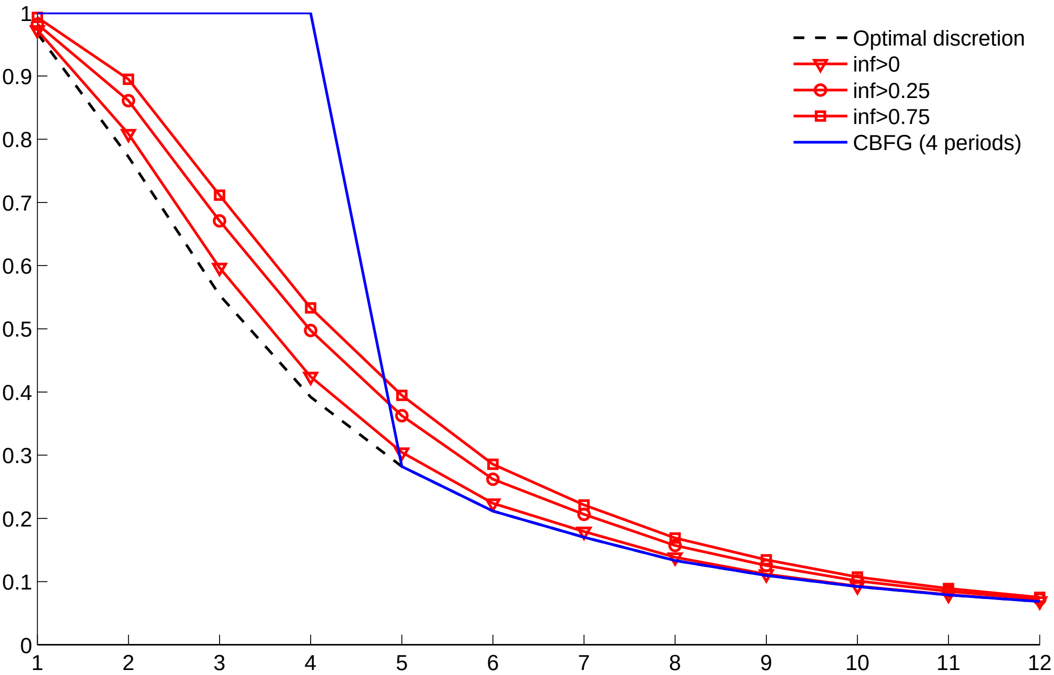

Notes: Computed from a stochastic simulation of 20,000 draws over 24 periods from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

the distribution. But in "good" states, when positive demand shocks arrive, exit occurs earlier and the stimulus is removed. The contrast with CBFG (panel (b)) is stark. CBFG imparts stimulus regardless of the state of the economy. While it reduces the negative skew because it raises expectations sufficiently to reduce the impact of the ZLB constraint (Figure 3), it leads to worse outcomes in both good and bad states than appropriately-calibrated TBFG. That is because CBFG 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.

The results above demonstrate that a TBFG policy with an above-target inflation threshold can achieve substantially better outcomes at the ZLB than the optimal discretionary policy. But 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 improve welfare by reneging on the policy and reverting to discretion (with an increase in the policy rate). A measure of the size of the policymaker's incentive to renege in any given period can be computed as the probability-weighted integral of the welfare gains from reneging on the forward guidance policy and reverting to the time-consistent policy (ignoring states in which welfare is higher if policy remains in the forward guidance regime). More formally, denote the measure of time inconsistency of a particular policy, P, in period t, as $\mathbb{T}_t^P$:

$$\mathbb{T}_t^P = \int_u \int_g \psi_t^P(u,g)\left(\mathbb{L}_t^P(u,g) - \mathbb{L}_t^{OD}(u,g)\right)\mathbb{I}\left(\mathbb{L}_t^P(u,g) - \mathbb{L}_t^{OD}(u,g) > 0\right) \quad (8)$$

where $\psi_t^P(u,g)$ is a measure of the likelihood that policy P is in effect in period t, $\mathbb{I}(.)$ 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 and:

$$\mathbb{L}_t^J(u,g) = \sum_{s=t}^{\infty} \beta^{s-t}\mathbb{E}_t(u,g)\left(\pi_s^J(u,g)^2 + \lambda y_s^J(u,g)^2\right) \quad (9)$$

is the welfare loss associated with policy $J \in \{P, OD\}$ given the state, $\{u,g\}$.

Figure 5: Distribution of inflation for alternative inflation TBFG policies, optimal discretion and CBFG

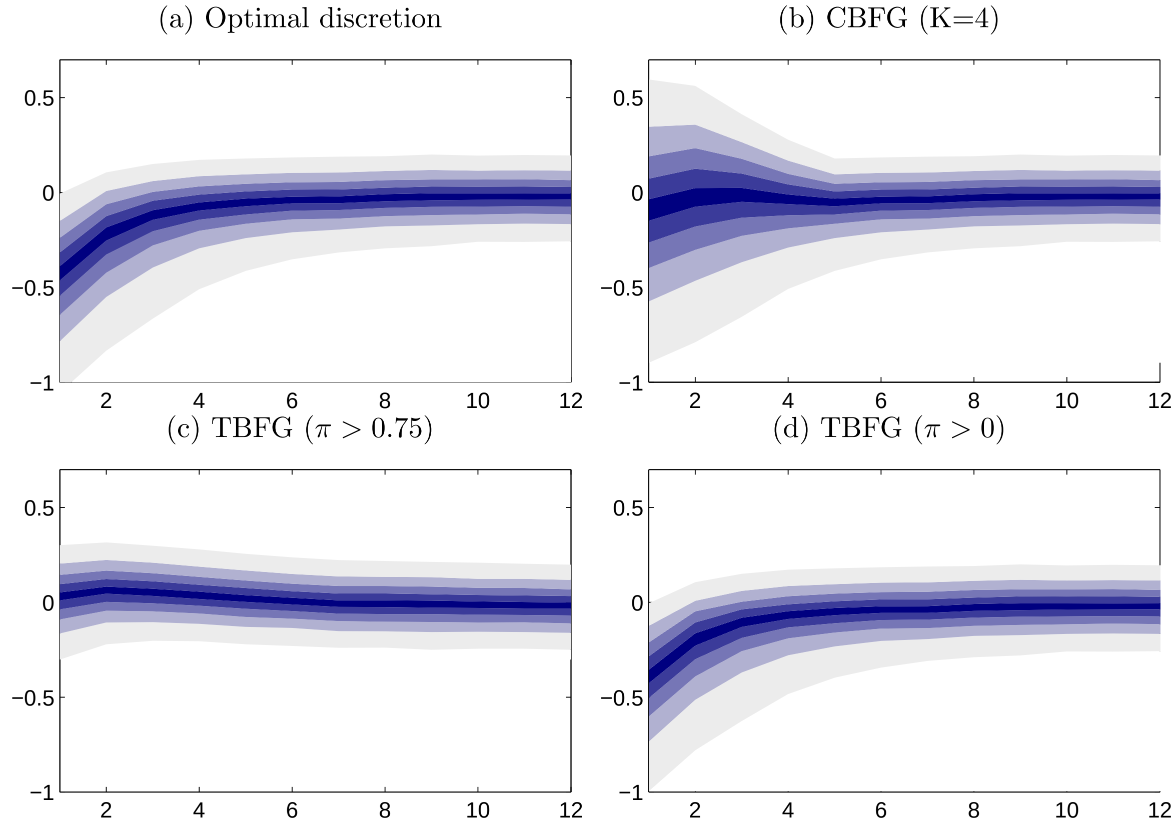

Notes: Computed from a stochastic simulation of 20,000 draws over 24 periods from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

Figure 6 illustrates that the incentive to renege on TBFG 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.

Although TBFG is less time inconsistent than CBFG, it is nevertheless (by definition) still time inconsistent. But that 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 (2014) 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.[23]

[23]Theoretically, it could 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. See Tinsley (1999) for an early variant of this type of idea.

Figure 6: Time-inconsistency measures for alternative inflation TBFG policies and CBFG

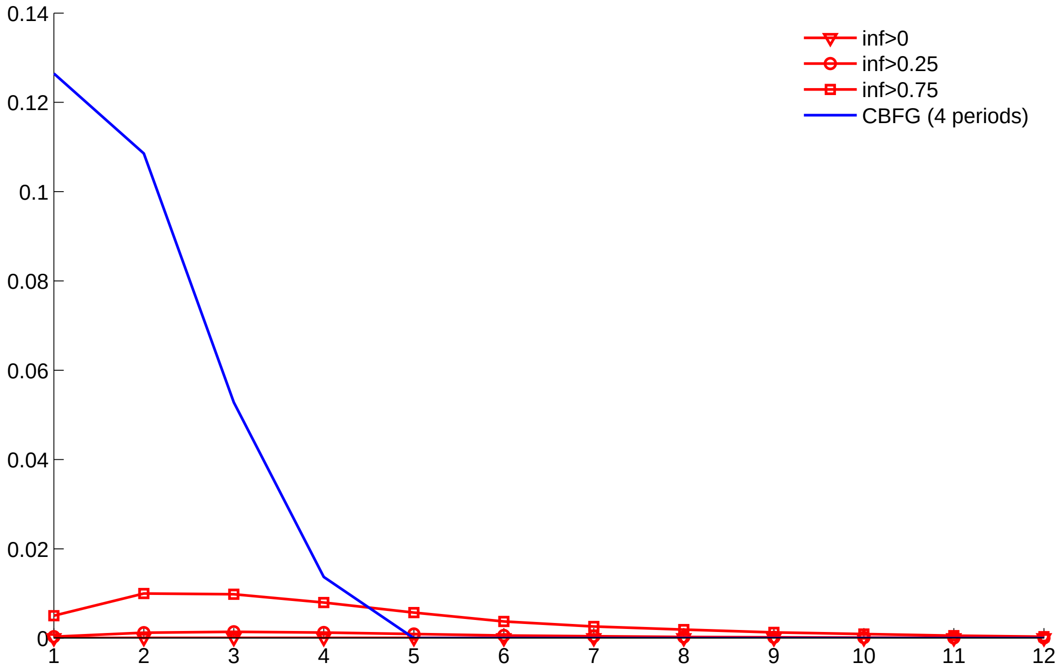

Notes: Computed by replacing the density measure in equation (8) with a discrete approximation on the state grid from from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$, which we use to replace the integral with a finite sum – see, for example, Chapter 5 of Heer and Maussner (2005).

7.2 Loss-minimising inflation and output gap thresholds

This section compares inflation TBFG policies to output gap TBFG policies. 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 for inflation TBFG and 3 for output gap TBFG. Under our baseline parameterisation, the loss-minimising thresholds are associated with almost identical welfare losses.[24]

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

(a) Inflation thresholds (b) Output gap thresholds

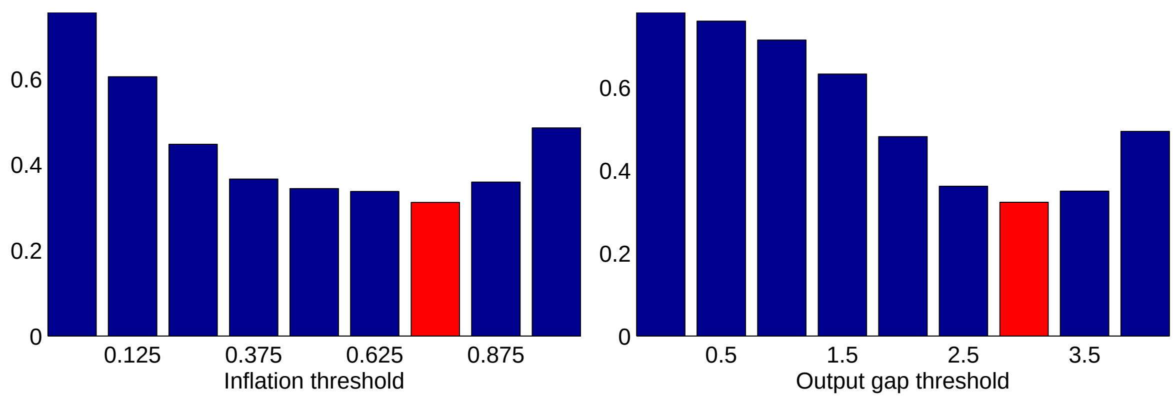

Notes: Computed from a stochastic simulation of 20,000 draws over 24 periods from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

While these alternative optimal TBFG policies have similar losses, they do not deliver the

[24]The optimal threshold values and losses associated with them are model-specific. See Appendix B for loss-minimising thresholds in a version of the model in which the cost-push process is more important relative to the demand process than in the baseline model.

same outcomes in all circumstances. Figure 8 reports how regime exit in the two alternative loss-minimising TBFG policies depends on the demand and cost-push state. For inflation TBFG, exit can be triggered by either a demand or cost-push shock (left panel of Figure 8). By contrast, exit is independent of the cost-push state for the loss-minimising output gap TBFG policy. 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*.[25] Unsurprisingly, making cost-push shocks autocorrelated overturns these results (Appendix B). In that situation, exit is not independent of cost-push shocks under output gap threshold designs.

Figure 8: Regime exit indicator function for optimal inflation and output gap TBFG policies

(a) Optimal inflation TBFG policy (b) Optimal output gap TBFG policy

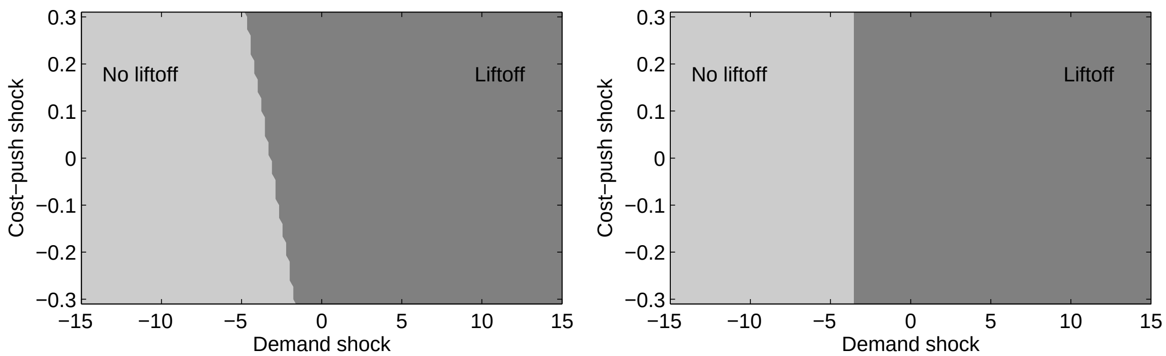

There is an interaction between the model (structure of the economy), the exit indicator function and the precise interpretation of the threshold conditions being applied by the policymaker. Inspection of Figure 9 reveals that the modal path of inflation under the optimal inflation TBFG policy peaks at a value of 0.15 which is well below its threshold of 0.75. By 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. But under an inflation TBFG policy, 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.[26] Although the exit threshold could be pushed outwards (rightwards and/or upwards) in some states without the thresholds being breached, that would mean that the thresholds would be breached in some other states via the boost in expectations that would result (given that expectations are the probability weighted integral of ouctomes in all states of the world). This observation reveals 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. Alternative interpretations which required that the thresholds be breached prior to exit would clearly yield different ouctomes and would imply more stimulus than the one we have used (for a given threshold value).

[25]The presence of cost-push shocks in the model does affect output via the optimal behaviour of the policymaker under optimal discretion, but that reflects an active response by the policymaker to act on the inflationary consequences of cost-push shocks, which is not present when rates are pegged at the ZLB.

[26]We therefore expect that this result is sensitive to the precise functional form of the exit indicator. For example, if we were to introduce some curvature, then exit determination would differ and so would the outomes for the endogenous variables prior to exit (which could mean that inflation tends to be closer to the threshold value on exit).

Figure 9: Modal responses under alternative inflation and output gap TBFG policies

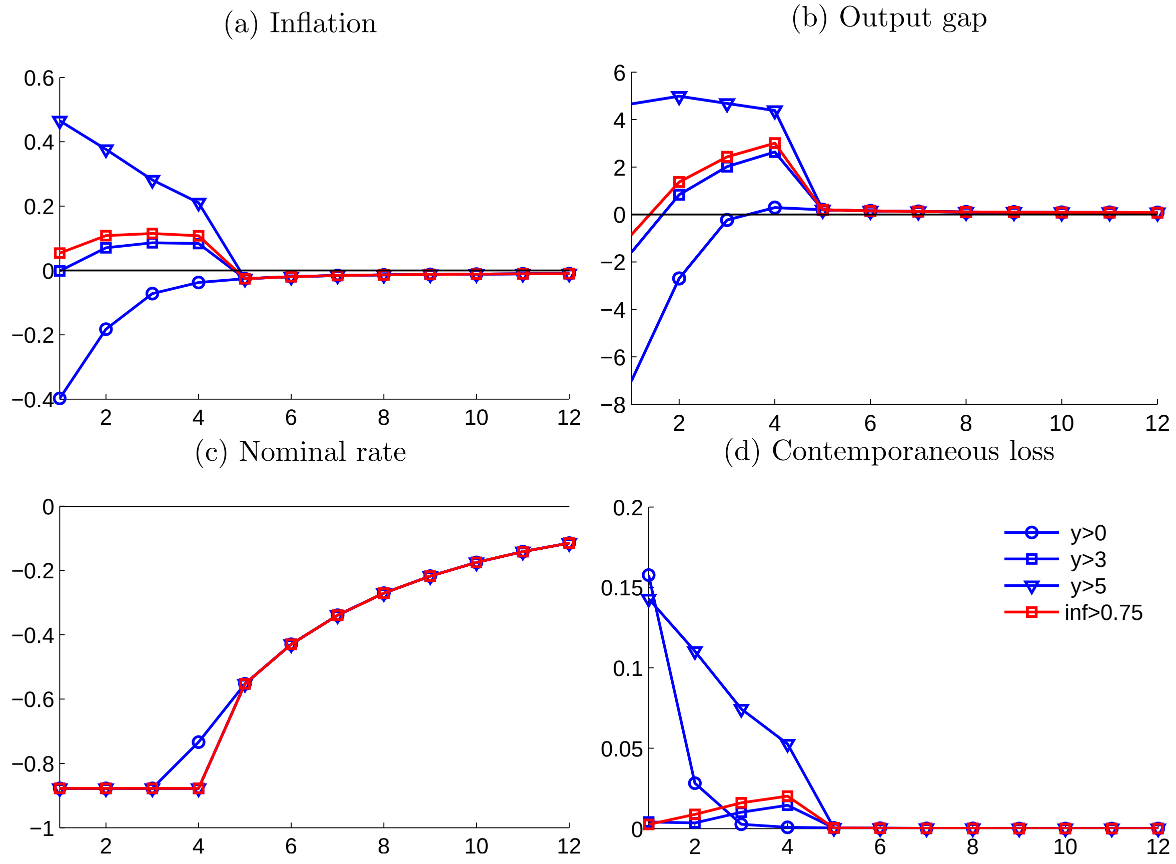

Notes: Computed under the assumption that no shocks arrive after the initial period 0 and from an initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

7.3 TBFG as an approximation to optimal commitment at the ZLB

Section 7.1 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 outcomes today. This section compares TBFG to optimal commitment. In the case of optimal commitment, the policymaker is able to commit to an interest rate plan that minimises the entire discounted sum of future losses subject to the zero lower bound constraint on interest rates and the equilibrium conditions:

$$\min_{\{y_t,\pi_t,r_t\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t(\pi_t^2 + \lambda y_t^2)$$

$$s.t \;\; r_t \geq 1 - \frac{1}{\beta}$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + u_t$$

$$y_t = \mathbb{E}_t y_{t+1} - \sigma\left(r_t - \mathbb{E}_t \pi_{t+1}\right) + g_t$$

$$u_t = \rho_u u_{t-1} + \sigma_u \varepsilon_t^u$$

$$g_t = \rho_g g_{t-1} + \sigma_g \varepsilon_t^g$$

$$\{u_0, g_0\} \;\; \text{given}$$

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:[27]

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

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

Figure 10: Modal responses under optimal TBFG policies, optimal commitment and discretion

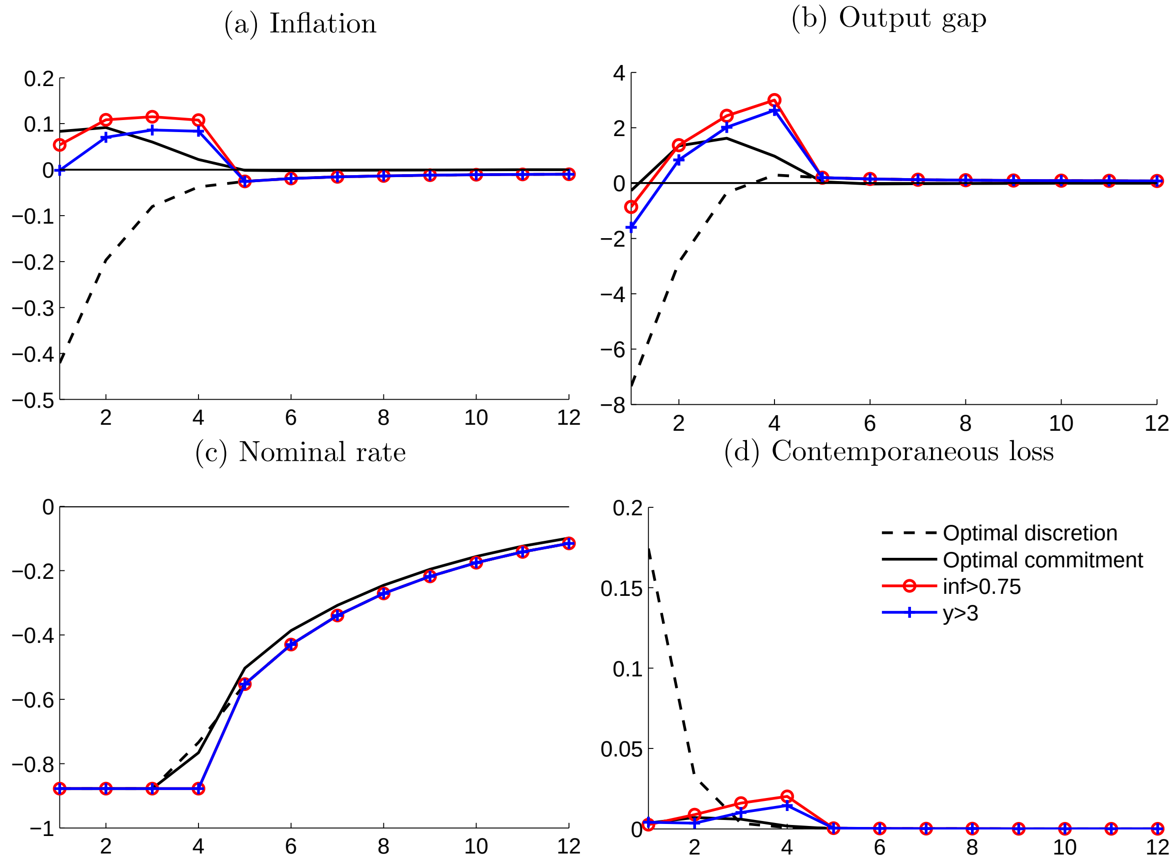

Notes: Computed under the assumption that no shocks arrive after the initial period 0 and from an initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

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 stabilises 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.[28] This is a direct consequence of the history dependence of policy which must be taken account of when the model is solved.

Figure 10 shows that modal paths for the optimal TBFG policies are close to the optimal commitment benchmark.[29] 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.

Figure 11 compares the distribution of inflation outcomes under optimal commitment with the optimal output gap TBFG policy. Relative to optimal commitment, the inflation distribution is slightly wider under the TBFG policy, but not dramatically so. One potential reason for that is that optimal commitment policy includes concern for both output and inflation. It follows that a dual threshold policy might be able to achieve better outcomes than a single threshold forward guidance policy. This is a question we leave for future research.

Figure 11: Distribution of inflation for optimal TBFG policy and optimal commitment

(a) Optimal output gap TBFG (b) Optimal commitment

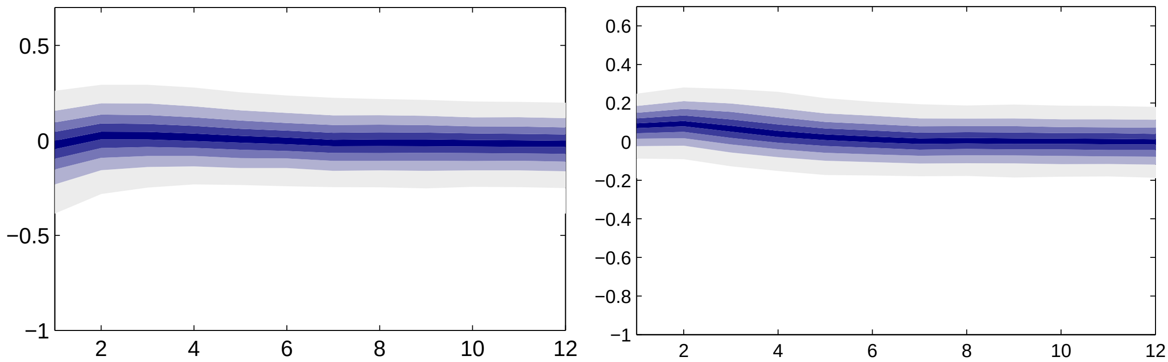

Notes: Computed from a stochastic simulation of 20,000 draws over 24 periods from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$.

That TBFG policies come close to replicating outcomes under optimal commitment at the ZLB is certainly of relevance to policymakers. TBFG has some practical advantages over optimal commitment. In particular, it may be much easier for the public to understand than a fully state-contingent optimal commitment policy, particularly as thresholds can be specified directly on goal variables that are used to frame a lot of central bank communications. In this way, TBFG could be thought of as an approximate implementation of optimal commitment policy at the ZLB.

[28]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.

[29]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. This is a consequence of the discrete-time setting used here. Werning (2011) uses a continuous-time setting to show that optimal commitment always involves setting the policy rate at the ZLB for longer than under optimal discretion.

8 Conclusions

Motivated by forward guidance policies implemented by the FOMC and the MPC of the Bank of England, this paper has studied the efficacy of stylised 'threshold-based' forward guidance (TBFG) as a temporary policy tool that can be used to impart stimulus at the zero lower bound (ZLB). We have shown that TBFG can improve outcomes at the ZLB 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, reduce the real interest rate and improve outcomes in the same way as first 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 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 interpret 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 interpretation 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.

In future versions of this paper we intend to extend the analysis in this paper in some or all of the following ways: a generalisation of the state-contingent forward guidance regime exit function to non-linear functional forms; further sensitivity analysis around the relationship between the structure of the economy and optimal thresholds; analysis of dual thresholds on both activity and inflation; a comparison of TBFG with other policies that have been proposed at the ZLB (like temporrary price-level targeting); an analysis of the distinction between 'thresholds' beyond which regime exit may occur but not with certainty and 'triggers' beyond which regime exit occurs with certainty (as analysed in this version of the paper). In addition, it is worth noting a limitation of the above analysis that it is conducted under the assumption of non-rational and ransitory regime switching. An avenue for future research would be to explore a generalisation of the setting to one in which there may be multiple TBFG regimes as and when the ZLB is a binding constraint.

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.

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.

Cooley, T. F., S. F. LeRoy, and N. Raymon (1984): "Econometric policy evaluation: Note," *The American Economic Review*, 467–470.

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).

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.

Nakata, T. (2014): "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.

Tinsley, P. A. (1999): "Short rate expectations, term premiums, and central bank use of derivatives to reduce policy uncertainty," *Finance and Economic Discussion Series Working Paper No. 1999-14*.

Werning, I. (2011): "Managing a Liquidity Trap: Monetary and Fiscal Policy," NBER Working Papers 17344, National Bureau of Economic Research, Inc.

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 an optimal discretionary basis) 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, by definition, there is a single state in each period that is perfectly forecastable by agents. This means that there is no uncertainty about when the policy rate will liftoff and so there is an equivalence between TBFG and 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 12 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, $g_0 = -11.95$.[30] 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 12. Which of the four alternative paths shown in Figure 12 is the equilibrium given this TBFG policy? In the absence of additional information, *any* of these paths could be an equilibrium. 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.[31] 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

[30]The model has been resolved for the deterministic case (with the standard deviations of the shocks set to 0). The initial condition for the state 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.

[31]It should be noted at this point that New Keynesian models 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 ZLB durations that do not result in sign-flipping.

Figure 12: Output gap under alternative policies in a deterministic setting

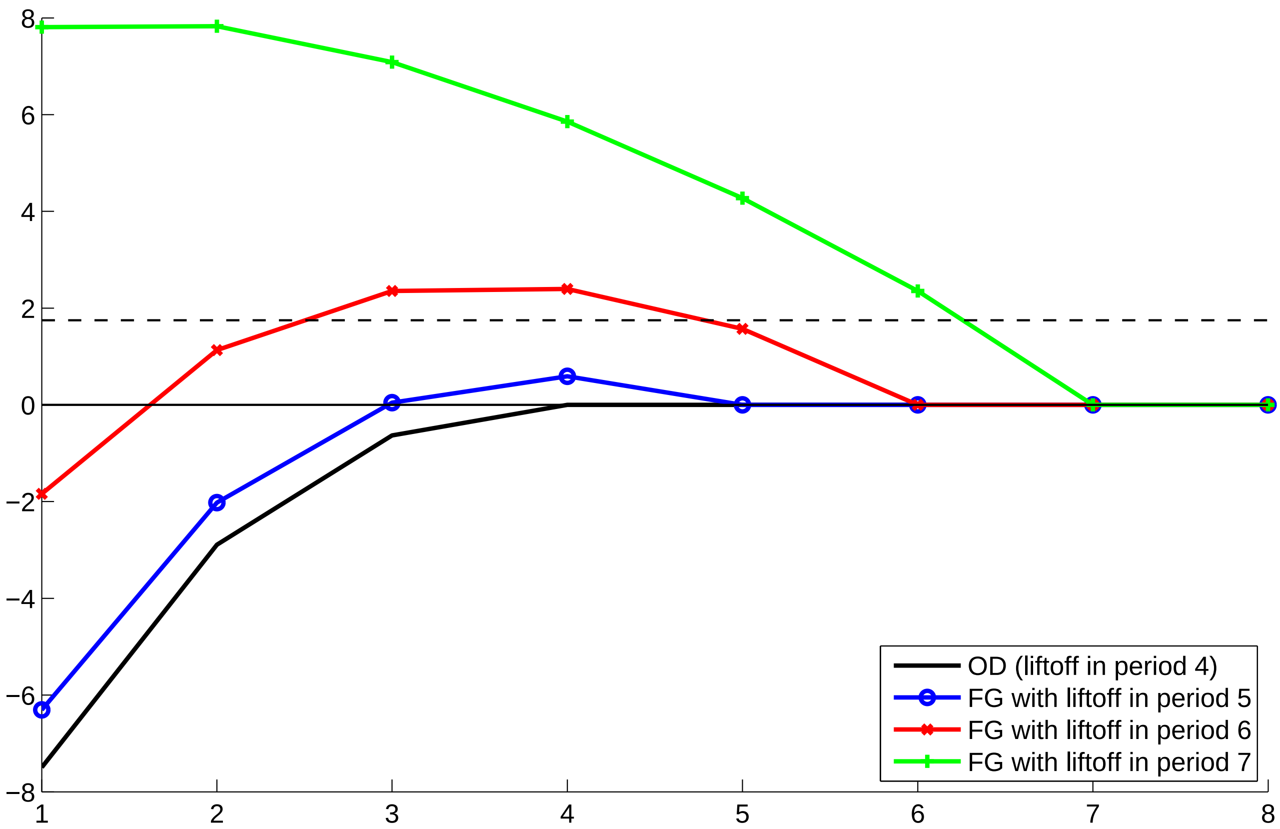

Notes: Computed from an initial condition of $g_0 = -11.95$ and $u_0 = 0$. No shocks arrive or are expected to arrive thereafter. Otherwise, the model is identical to that described in Section 3 of the main text with the baseline calibration outlined in Section 6.

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).[32]

The above example also reveals that equilibrium selection concerns the entire path of outcomes. It is not sufficient to determine exit on a period-by-period basis because the entire expected path for rates matters for outcomes in the preceeding periods. To see that, suppose that the economy is on the path determined by the policy with liftoff in period 5 (the blue line with circles) along which the threshold is not breached in any period. Notice that, on arrival in period 5, it would be possible for the policymaker to extend the period for which rates are held at the ZLB by 1 without the threshold conditon being breached. However, note that if agents had known that the policymaker would behave in this way prior to period 5, then

[32]As explained in the main text, the equilibrium object for a given TBFG policy is a state-contingent exit indicator function. The blue line with circles in Figure 12 would be suppported by an exit indicator function of the following form: $\mathbb{I}^{EXIT}(g,u) = g > c$, where $g_0^5 \geq c < g_0^6$. In words, exit does not occur for any demand state less than that prevailing in period 5 (which is equal to g_0^5 by virtue of the deterministic assumption), but must occur for all demand states that equal or exceed that prevailing in period 6 (g_0^6). This reveals a non-unique mapping from equilibrium selection to the exit indicator function, reflecting the combination of the deterministic and discrete time settings. This non-uniqueness is resolved in the stochastic case, where there is a distribution across the state with liftoff occurring probabalistically.

the equilibrium path would be governed by the policy with liftoff in period 6 along which the threshold is violated in periods 3 and 4. Our equilibrium concept, in which exit is a function of the state and not time, rules out outcomes in which the policymaker behaves *ex post* in a way that agents had not expected *ex ante*.

B Sensitivity to stochastic process calibration

In this appendix, 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. 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. We recalibrate the initial condition for the demand state to deliver roughly the same average fall in output (under the optimal discretionary policy) as in the model with the baseline parameterisation.

Table 3 documents that the optimal threshold values depend on the specific parameterisation of the stochastic processes. Compared to our baseline parametrisation 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 13). The intuition for both results is that cost-push induced inflation is more costly because it persists for longer and because cost-push shocks explain a higher proportion of the variance of outcomes than in the baseline calibration of the model. As a result, the optimal TBFG policies take a greater concern for cost-push shocks. Notice also that the optimal inflation TBFG policy delivers a lower ex-ante loss than the optimal output gap policy. This was not the case under the baseline parameterisation and reflects the additional cost of cost-push induced inflation in this variant of the model.

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

Figure 13: Regime exit indicator function for optimal inflation and output gap TBFG policies in autocorrelated cost-push shocks variant of the model

(a) Optimal inflation TBFG policy (b) Optimal output gap TBFG policy