DISCUSSION PAPER SERIES

DP11749

THRESHOLD-BASED FORWARD GUIDANCE: HEDGING THE ZERO BOUND

Lena Boneva, Richard Harrison and Matt Waldron

MONETARY ECONOMICS AND FLUCTUATIONS

ISSN 0265-8003

THRESHOLD-BASED FORWARD GUIDANCE: HEDGING THE ZERO BOUND

Lena Boneva, Richard Harrison and Matt Waldron

Discussion Paper DP11749
Published 05 January 2017
Submitted 05 January 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

THRESHOLD-BASED FORWARD GUIDANCE: HEDGING THE ZERO BOUND

Abstract

"Threshold-based forward guidance" (TBFG) is a state-contingent commitment to hold the policy rate at the zero lower bound until macroeconomic variables breach particular "thresholds". Though such guidance has been implemented in practice, little is known about how this policy works in theory. This paper fills that gap by studying threshold-based guidance as a tool to improve outcomes at the zero bound within a general equilibrium framework. Policymakers have rarely advocated using TBFG to provide stimulus, in part reflecting skepticism about their ability to commit credibly to time inconsistent behavior. We show that TBFG can be used to provide temporary stimulus, while also limiting the time inconsistency of policy promises. We also show that existence of a unique equilibrium requires the policymaker to specify how the thresholds should be interpreted, as well as their values. The optimal design of the threshold conditions depends on the relative importance of those shocks that induce a trade-off between stabilizing output and inflation and those that do not. With an appropriate choice of thresholds, TBFG outperforms forward guidance based purely on calendar time and comes close to mimicking outcomes under optimal commitment policy.

JEL Classification: E17, E31, E52

Keywords: monetary policy, zero lower bound, forward guidance, thresholds

Lena Boneva - lena.boneva@bankofengland.co.uk
Bank of England and CEPR

Richard Harrison - richard.harrison@bankofengland.co.uk
Bank of England and Centre for Macroeconomics

Matt Waldron - matthew.waldron@bankofengland.co.uk
Bank of England

Acknowledgements
We thank Guido Ascari, Toni Braun, Fabio Canova, Alex Clymo, Wouter Den Haan, Taisuke Nakata, attendees at the Riksbank workshop on "Deflation", the Econometric Society 2015 World Congress, the ECB workshop on "Non-standard monetary policy measures" and seminar participants at the FRB Atlanta, Bank of England and UCL for their useful comments. The views expressed in this paper are our own and not those of the Bank of England.

Threshold-based forward guidance: hedging the zero bound *

Lena Boneva† Richard Harrison‡ Matt Waldron§

December 19, 2016

Abstract

'Threshold-based forward guidance' (TBFG) is a state-contingent commitment to hold the policy rate at the zero lower bound until macroeconomic variables breach particular 'thresholds'. Though such guidance has been implemented in practice, little is known about how this policy works in theory. This paper fills that gap by studying threshold-based guidance as a tool to improve outcomes at the zero bound within a general equilibrium framework. Policymakers have rarely advocated using TBFG to provide stimulus, in part reflecting skepticism about their ability to commit credibly to time inconsistent behavior. We show that TBFG *can* be used to provide temporary stimulus, while also limiting the time inconsistency of policy promises. We also show that existence of a unique equilibrium requires the policymaker to specify how the thresholds should be interpreted, as well as their values. The optimal design of the threshold conditions depends on the relative importance of those shocks that induce a trade-off between stabilizing output and inflation and those that do not. With an appropriate choice of thresholds, TBFG outperforms forward guidance based purely on calendar time and comes close to mimicking outcomes under optimal commitment policy.

Key words: monetary policy; zero lower bound; forward guidance; thresholds
JEL Classification: E17; E31; E52

*We thank Guido Ascari, Toni Braun, Fabio Canova, Alex Clymo, Wouter Den Haan, Taisuke Nakata, attendees at the Riksbank workshop on "Deflation", the Econometric Society 2015 World Congress, the ECB workshop on "Non-standard monetary policy measures" and seminar participants at the FRB Atlanta, Bank of England and UCL for their useful comments. The views expressed in this paper are our own and not those of the Bank of England.

†Bank of England and CEPR: lena.boneva@bankofengland.co.uk.
‡Bank of England and Centre for Macroeconomics. Email: richard.harrison@bankofengland.co.uk.
§Bank of England, London, UK. Email: matthew.waldron@bankofengland.co.uk

1 Introduction

In response to the global financial crisis, central banks cut policy rates towards the zero lower bound 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 such form of guidance is 'threshold-based forward guidance' (TBFG), in which the policymaker ties liftoff from the zero bound to outcomes for certain macroeconomic variables, promising not to increase the policy rate (at least) until those variables breach pre-specified 'threshold' values.

Though TBFG has been implemented in practice by both the FOMC and the MPC of the Bank of England, little is known about how this policy works in theory. This paper fills this gap by studying TBFG as a tool to improve outcomes at the zero lower bound within a general equilibrium framework.

Intriguingly, neither the FOMC, nor the MPC motivated their TBFG policies as tools to impart additional stimulus.[1] Instead, both emphasized an intent to provide greater clarity about their objectives in order to reduce uncertainty about the policy outlook. That policymakers have not attempted to use TBFG policy to inject stimulus appears, in part, to reflect skepticism about their ability to commit credibly to behavior that is well known to be time inconsistent (Nakata, 2015).

In this paper we show that TBFG *can* be used as a temporary commitment device to improve outcomes at the zero bound, while also limiting the extent to which the policymaker promises to behave in a time inconsistent manner. Thus, our paper advocates using TBFG policy for a different purpose to those which motivated real-world implementations.

We use a simple New-Keynesian model to study a scenario in which a large negative demand shock causes the zero bound to bind. This generates a deep recession, given a policymaker who optimizes on a period-by-period basis (optimal discretion). This motivates our TBFG experiments in which the policymaker attempts to improve outcomes by temporarily deviating from time-consistent policy.

The model we use is also the workhorse for several other studies of monetary policy at the zero bound (for example, Adam and Billi (2006) and Bodenstein et al. (2012)). The policy instrument is the short-term nominal interest rate (subject to a zero lower bound constraint) and the objective is to minimize a welfare-based loss function. The policymaker acts under discretion, so that policy is time consistent. This assumption incorporates the skepticism of real-world policymakers about their ability to commit to future policy actions. We solve the model using projection methods to account for the nonlinearity induced by the zero bound.

We model TBFG as a one-off commitment not to increase the policy rate from the zero bound at least until particular macroeconomic variables have breached pre-announced threshold values. Once the economy has recovered sufficiently that the conditions have been met, the policymaker reverts to setting policy under optimal discretion forever more.

Exit from the TBFG policy is assumed to be probabilistic. The probability of exiting the TBFG policy is an increasing function of the amount by which the thresholds have been breached. This makes the expected exit date (a key determinant of the stimulus imparted by the policy) a continuous random variable and is necessary to generate a unique equilibrium. More generally, we show that TBFG policies are incomplete in the absence of sufficient detail about how the policymaker will act when the thresholds are breached.

We find that TBFG can substantially improve welfare relative to the time consistent policy. In line with the 'textbook' remedy to mitigating the zero bound constraint, TBFG can stimulate

[1] Both policies were designed around threshold values for unemployment with 'knock-outs' that would render the policy inactive if inflationary pressure were sufficiently strong. In both cases, the threshold value for unemployment was set above prevailing FOMC and MPC estimates for the medium-term equilibrium unemployment rate, thereby implying no 'overshooting' commitment. See Yellen (2013) and Bank of England (2013) for discussions of the respective TBFG policies.

activity and inflation today by a promise of higher inflation in the future. In addition, TBFG can reduce the variance of the distribution of possible outcomes. If further negative shocks arise, prolonging the recession, the policy rate will be held at the zero bound for longer. By contrast, if positive shocks arrive, so that the economy recovers more quickly than originally expected, then liftoff from the zero bound will occur sooner and the stimulus will be removed.

In this way, TBFG can be viewed as a hedge against the asymmetric effects generated by the zero lower bound. The magnitude of the effect can be seen by comparing losses under threshold-based guidance with those under 'calendar-based forward guidance' (CBFG), whereby the policymaker promises to hold the policy rate at the zero bound for a pre-specified length of time regardless of the state of the economy.[2] While CBFG can improve expected outcomes and eliminate the negative skew in outcomes induced by the zero lower bound, it leads to worse outcomes for both positive and negative realizations of future demand shocks as it provides too much stimulus in 'good' states and insufficient stimulus in 'bad' states.

Our TBFG policy experiments represent a temporary deviation from time-consistent behavior and are by definition time inconsistent. As such, the experiments may be regarded as less than fully credible. We assess this by computing a measure of the extent to which the policymaker could achieve better outcomes by reneging on the threshold-based policy and reverting to the time-consistent policy. A corollary of the hedging property of TBFG is that the temptation to renege is much smaller than for calendar-based guidance. For realizations of the shocks in which the economy recovers more quickly than originally expected, the exit conditions are more likely to be met and policy automatically reverts to time-consistent behavior.

We also find that threshold-based policies can achieve ex-ante losses that are close to the optimal commitment policy. To deliver this result, the values of the thresholds must be chosen carefully. A general requirement is that the thresholds should generate an overshoot of the inflation target and/or a positive output gap. However, our sensitivity analysis demonstrates that optimized threshold values depend on the structure of the economy, the nature of the disturbances, and the interpretation of the threshold conditions. In particular, in the baseline calibration of the model in which demand shocks are dominant in driving the model's dynamics, optimized inflation and output gap threshold policies deliver similar losses. By contrast, an optimized inflation threshold performs better than an optimized output gap threshold in a version of the model in which cost-push shocks are more important. Nevertheless, there is a range of (inflation and output gap) threshold values for which welfare can be improved in all of the model specifications we consider.

To our knowledge, this is the first paper to analyze threshold-based policies in a fully stochastic setting. Campbell et al. (2012) demonstrate the potential for CBFG to perform poorly when unexpected shocks hit the economy and conjecture that the use of TBFG policies would improve performance. We verify that conjecture. English et al. (2015) point out that the public announcement of thresholds can serve as a commitment device, allowing the policymaker to provide additional stimulus as a state-contingent form of lower-for-longer policy. They also show that TBFG policies that deliver better performance on average can imply a larger number of states of the world in which the policymaker could achieve better outcomes by reneging on the implied commitment. Our baseline assumption of time-consistent optimal policy allows us to formally model TBFG as a commitment device and to quantify the extent to which such policies are time inconsistent. Chattopadhyay and Daniel (2014) explore the effects of temporary increases in the inflation target in a Taylor rule when the monetary policymaker hits the zero bound. While this policy proposal shares some similarities with our approach it embeds the assumption that policymakers can commit to future actions (via alternative Taylor rules). As described above, we view our assumption of time-consistent policy as an approximation to

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

the concerns of real-world policymakers about their ability to commit credibly to future actions.

Walsh (2016) asks whether policymakers without access to a commitment technology can make credible CBFG announcements. Using a 'sustainable plans' approach, he finds that such announcements may be credible when shocks may drive the economy to the zero bound in future. Nakata (2014) uses a similar approach to show that reputational effects may be sufficient to enforce the Ramsay policy if the frequency of zero bound episodes is sufficiently high. While the credibility of deviations from time-consistent policies is also of interest to us, we focus on *temporary* deviations in the context of threshold-based policies. One limitation is that the TBFG policies we study are understood to be one off deviations from time-consistent policy that will never be repeated. Incorporating TBFG into systematic policy behavior is beyond the scope of the present paper and remains an interesting topic for future work.

The closest paper to ours is Florez-Jimenez and Parra-Polania (2016), who also study threshold-based guidance in a small model. However, their analysis is limited to a two-period model with a threshold defined in terms of an exogenous shock process. We analyze threshold-based policies of stochastic duration and specify thresholds in terms of endogenous variables. Both Coenen and Warne (2013) and English et al. (2015) study TBFG in more realistic multi-period settings with perfect foresight approximations of expectations. We extend the analysis in those papers by defining and computing a fully stochastic equilibrium, formally modeling TBFG as a state-contingent commitment, quantifying the size of the associated time-inconsistency problem, and by computing loss-minimizing thresholds.

The rest of the paper is organized as follows. Section 2 details the model and the baseline description of policy. Section 3 describes the policy experiments and the assumptions underpinning them. Section 4 defines equilibrium given the forward guidance policies. Section 5 outlines the parameterization of the model and the calibration of the state of the economy prior to the implementation of forward guidance. Section 6 describes the simulation results, including comparisons of threshold-based guidance to calendar-based guidance and optimal commitment policy. Section 7 examines the sensitivity of the results to alternative calibrations of the model and interpretations of the threshold conditions. Section 8 concludes.

2 The model

We use a prototypical New Keynesian model in which a representative household supplies labor 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 model is identical to that used by Adam and Billi (2006, 2007) and Bodenstein et al. (2012) to study monetary policy at the zero bound under optimal commitment, optimal discretion and 'loose commitment' respectively.

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.[3] Following previous studies of monetary policy at the zero bound (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 zero bound and the optimality conditions are log-linearized around the non-stochastic steady state.

Throughout our analysis, our baseline assumption is that the monetary policymaker sets policy under optimal discretion. Specifically, we assume that the policymaker chooses the policy instrument each period to minimize a loss function derived from a quadratic approximation to

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

the representative agent's utility function, taking agents' expectations as given:

$$\min_{\{y_t,\pi_t,r_t\}} \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i (\pi_{t+i}^2 + \lambda y_{t+i}^2) \tag{1}$$

$$s.t \;\; r_t \geq 1 - \frac{1}{\beta} \tag{2}$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + u_t \tag{3}$$

$$y_t = \mathbb{E}_t y_{t+1} - \sigma \left(r_t - \mathbb{E}_t \pi_{t+1} \right) + g_t \tag{4}$$

$$u_t = \rho_u u_{t-1} + \sigma_u \varepsilon_t^u \tag{5}$$

$$g_t = \rho_g g_{t-1} + \sigma_g \varepsilon_t^g \tag{6}$$

$$\mathbb{E}_t \{ y_{t+i}, \pi_{t+i}, r_{t+i} \}_{i=1}^{\infty} \;\; \text{given} \tag{7}$$

$$\{u_t, g_t\} \;\; \text{given} \tag{8}$$

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; $\kappa = \frac{(1-\xi)(1-\xi\beta)}{\xi} \frac{\sigma^{-1}+\omega}{1+\omega\theta}$ is the slope of the Phillips curve, where ξ 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, 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.

A Markov perfect equilibrium comprises policy functions $y^{OD}(u_t, g_t)$, $\pi^{OD}(u_t, g_t)$ and $r^{OD}(u_t, g_t)$ that solve the problem (1)-(8). Given the existence of a ZLB constraint (2), there is no analytical solution to the problem. Instead, the solution can only be approximated. Our approach to doing so follows Adam and Billi (2007) and is described in Appendix C.

3 The nature of the policy experiments

The policy experiments are ones in which a policymaker temporarily deviates from the optimal discretionary policy. The temporary deviation is a one-off and fully credible forward guidance policy with the objective of stimulating output and inflation, given an economic environment in which the policy rate has become constrained by the zero bound.

The precise sequence of events in all of our policy experiments is shown in Figure 1. In period $t = 0$ a negative demand shock arrives that is sufficiently large that the policy rate consistent with optimal discretionary policy is negative and hence constrained by the zero bound. 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

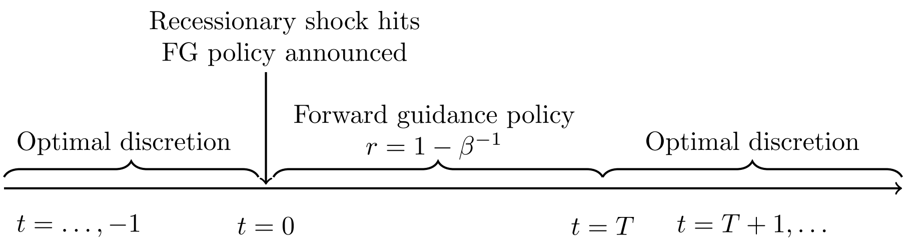

In this paper we analyze two alternative assumptions about the date T at which the policymaker reverts to setting policy under discretion. In the case of calendar-based guidance, the exit date is fixed $T = K\ (\geq 1)$ at $t = 0$. In contrast, threshold-based guidance links exit from the zero bound to macroeconomic conditions, namely whether certain variables have crossed particular threshold values. The values of the thresholds are announced at date $t = 0$. This implies that the exit date T is a random variable because the dates at which the thresholds may be breached in the future depend on the sequence of shocks that might hit the economy in the meantime.

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 and, once the policy has ended, agents attach no probability to it being implemented again. This assumption is common to other papers that study temporary deviations of policy from 'usual' policy behavior (e.g. del Negro et al. (2012), Coenen and Warne (2013), Haberis et al. (2014), English et al. (2015)). Such policy experiments are not conducted under fully rational expectations and so are subject to the issues discussed 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. Relaxing this assumption would substantially increase the complexity of the problem and we leave it as a question for future research. We do note, though, that with the policies implemented in the wake of the financial crisis in mind, it is perhaps reasonable to believe that explicit forward guidance policy may not have been fully anticipated.

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.

4 Equilibrium in the forward guidance regime

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

4.1 Threshold-based guidance equilibrium

Threshold-based forward guidance policy is characterized by threshold values for the output gap, $\bar{y}$, or inflation, $\bar{\pi}$, or both, as well as precise information about how the threshold conditions should be interpreted. A key aspect of our approach is the assumption that exit from the forward guidance policy is probabilistic. That is, in the event that the threshold conditions are breached, exit from the forward guidance policy will occur with some probability strictly less than unity. Our motivation for this is to ensure that there is a unique equilibrium under threshold-based guidance, but we also note that this approach is more consistent with the real-world policies enacted by policymakers at the Federal Reserve and Bank of England, which specified that the policy rate would remain at the zero bound "at least until" thresholds were breached.[4]

Appendix A motivates the probablistic approach in more detail. Using a simple deterministic example, we show that merely announcing threshold-based conditions for exit from the forward guidance policy may be insufficient for either existence or uniqueness of equilibrium. In the context of this example we show that there is an inherent inconsistency between an assumption that exit occurs with certainty in the period that the threshold conditions are breached and agents expecting the policy to remain in place for long enough to generate that threshold breach. This type of result is a feature of the 'overshooting' generated by policy stimulus at the zero bound. In such cases, the equilibrium associated with a promise to exit with certainty once the threshold conditions have been breached may not be unique or may not exist at all. We also demonstrate that our probabilistic exit assumption alleviates this problem by making the expected exit date (a key determinant of the stimulus imparted by the policy) a continuous random variable.

We assume that the probability of exit is increasing in the distance between the threshold variable and the threshold value according to an exponential distribution function. For example, for a TBFG policy defined in terms of an output gap threshold, $\bar{y}$, the probability mapping function is:

$$f\left(y-\bar{y}\right)=\begin{cases}0 & \text{if } \ y\leq\bar{y}\\ 1-\exp\left(-\alpha_y^{-1}\left(y-\bar{y}\right)\right) & \text{if } \ y>\bar{y}\end{cases} \tag{9}$$

where the parameter $\alpha_y > 0$. The rationale for our choice of this function and its parameterization are discussed in more detail in Section 5. At this point, we highlight that equation (9) embodies the assumption that the threshold is fully credible, in the sense that the probability that the policymaker exits is zero if the threshold is not breached ($y \leq \bar{y}$).

In order to simplify the computation of equilibrium and to facilitate an approximation of the case in which exit occurs instantaneously following a threshold breach (a 'trigger'), we employ a within-period timing assumption. Specifically, we assume that the sequence of events in each period t is as follows. First, shocks $\{\epsilon_t^u, \epsilon_t^g\}$ are realized and observed by all agents. Next, the private sector chooses $\{y_t, \pi_t\}$. Finally, the policy rate, r_t, is announced consistent with threshold-based guidance. For example, in the case of an output gap threshold, $\bar{y}$, the policymaker sets $r_t = 1-\beta^{-1}$ for certain if $y_t \leq \bar{y}$, otherwise they set r_t according to the optimal discretionary policy with probability $f\left(y_t - \bar{y}\right)$ and $r_t = 1-\beta^{-1}$ with probability $1-f\left(y_t-\bar{y}\right)$. This timing assumption implies a modified version of the IS curve (4) in which agents form expectations about the policy rate within period t:

$$\begin{aligned} y_t &= \mathbb{E}_t y_{t+1} - \sigma\left(\mathbb{E}_t r_t - \mathbb{E}_t \pi_{t+1}\right) + g_t \\ \mathbb{E}_t r_t &= f\left(y_t-\bar{y}\right) r_t^{OD} + \left(1-f\left(y_t-\bar{y}\right)\right)\left(1-\beta^{-1}\right) \end{aligned} \tag{10}$$

where r^{OD} denotes the interest rate under optimal discretion.[5]

Our timing assumption simplifies the computation of equilibrium because exit from the forward guidance policy depends on variables that have already been determined by the decisions

[4]For example, see the extract from the FOMC statement of December 2012 in Appendix B.

[5]The timing assumption can be motivated by assuming that, at the start of each period, perfectly competitive financial intermediaries lend to households for the duration of the period. The financial intermediaries borrow using short-term government bonds. Perfect competition ensures that the rate offered to households equals the expected rate paid on bonds.

of private agents earlier in the period. By an appropriate choice of a function that maps outcomes for the endogenous variables to the probability of exit, this also allows us to approximate the case in which exit from the forward guidance policy occurs almost instantaneously when the threshold has been breached without introducing simultaneity between private sector and policymaker decisions.[6] Combined with a suitable calibration of the exit probability function, this allows us to examine threshold-based policies that approximate 'trigger' policies.

To define the equilibrium under the threshold-based guidance regime, we introduce some notation. The vector of endogenous variables is $x \equiv [\pi, y]$ and the state vector is $s \equiv [u, g]$. We drop time subscripts for period t variables and denote values in the following period using primes (e.g. s' is the state vector in the following period). We use the following notation for conditional expectations:

$$\mathbb{E}_{s'|s} h(x') = \int h(x(s'))\,\mathrm{d}W(s'|s) \tag{11}$$

for any function h, where $W(s'|s)$ is the distribution function of s' given s.

Formally, equilibrium in a threshold-based policy regime with inflation threshold, $\bar{\pi}$, and output gap threshold, $\bar{y}$ (so that $\bar{x} \equiv [\bar{\pi}, \bar{y}]$), is defined by policy functions, $\pi^{FG}(s)$ and $y^{FG}(s)$, that satisfy:

1. The competitive equilibrium conditions:

$$\begin{aligned} y^{FG}(s) &= p(s)\,\mathbb{E}_{s'|s} y^{OD}(s') + (1 - p(s))\,\mathbb{E}_{s'|s} y^{FG}(s') \\ &\quad - \sigma \left\{ \begin{array}{l} p(s)\, r^{OD}(s) + (1 - p(s))(1 - \beta^{-1}) \\ - \left[p(s)\,\mathbb{E}_{s'|s}\pi^{OD}(s') + (1 - p(s))\,\mathbb{E}_{s'|s}\pi^{FG}(s')\right] \end{array} \right\} + g \end{aligned} \tag{12}$$

$$\pi^{FG}(s) = \kappa y(s) + \beta\left[p(s)\,\mathbb{E}_{s'|s}\pi^{OD}(s') + (1 - p(s))\,\mathbb{E}_{s'|s}\pi^{FG}(s')\right] + u \tag{13}$$

2. The mapping from outcomes to the probabilities of exiting:

$$p(s) = f\left(x^{FG}(s) - \bar{x}\right) \tag{14}$$

Notice that expectations are defined as the probability weighted integral over all possible realizations of the shocks, accounting for the two different policy regimes: the case in which the forward guidance regime is still in effect (superscript ^{FG}), and the case in which policy has reverted back to optimal discretion (superscript ^{OD}). This makes clear that the transmission of forward guidance is via agents' expectations. Outcomes can only be affected to the extent that there are some states of the world in which the forward guidance regime still applies *and* in which the policy rate would be set above the ZLB under optimal discretion. This discussion makes clear that that threshold-based guidance is a state-contingent form of 'lower-for-longer' policy in this setting. The effect of a given policy hinges on the state-contingent probability of exit, $p(s)$, which in turn depends on the threshold values and the interpretation of the threshold conditions via the function that maps outcomes to the probability of exit.

4.2 Calendar-based forward guidance equilibrium

Calendar-based forward guidance policy is characterized 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, $\{\pi_t^{FG}(s),\ \ y_t^{FG}(s)\}_{t=1}^{K}$, that satisfy:

[6] For similar reasons to those discussed in Appendix A, simultaneous decision making may lead to situations in which an equilibrium does not exist. For example, conditional on the policymaker staying at the zero bound in the current period the private sector may choose an output gap or inflation level that breaches the threshold, but conditional on the policymaker exiting in the current period optimal private sector decisions may not breach the thresholds.

1. The competitive equilibrium conditions:
$y_t^{FG}(s) = \mathbb{E}_t^{FG}(s)\, y_{t+1} - \sigma\left(1 - \frac{1}{\beta} - \mathbb{E}_t^{FG}(s)\, \pi_{t+1}\right) + g$
$\pi_t^{FG}(s) = \beta \mathbb{E}_t^{FG}(s)\, \pi_{t+1} + \kappa y_t^{FG}(s) + u$, where:
$\mathbb{E}_t^{FG}(s)\, y_{t+1} = \mathbb{E}_{s'|s}\left[\mathbb{I}_{t+1}^{EXIT} y^{OD}(s') + \left(1 - \mathbb{I}_{t+1}^{EXIT}\right) y_{t+1}^{FG}(s')\right]$
$\mathbb{E}_t^{FG}(s)\, \pi_{t+1} = \mathbb{E}_{s'|s}\left[\mathbb{I}_{t+1}^{EXIT} \pi^{OD}(s') + \left(1 - \mathbb{I}_{t+1}^{EXIT}\right) \pi_{t+1}^{FG}(s')\right]$
$u' = \rho_u u + \epsilon^{u'}$
$g' = \rho_g g + \epsilon^{g'}$
$\epsilon^{u'} \sim \mathbb{N}(0, \sigma_u)$
$\epsilon^{g'} \sim \mathbb{N}(0, \sigma_g)$.

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 threshold-based guidance, it is evident from the equilibrium definition that calendar-based guidance 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 as a function of time under calendar-based guidance, while regime exit is determined as a function of the state of the economy under threshold-based guidance.

4.3 Approximating the solutions

Solving for the equlibria defined above requires approximating the relevant policy functions. In both cases, reflecting that the forward guidance policies are transitory, a pre-requisite for the approximation is a solution to the model described in Section 2, in which policy is conducted under optimal discretion. The solution method is described in Appendix C.1. It uses time iteration and function approximation to find a fixed point in the space of policy functions $y^{OD}(s)$, $\pi^{OD}(s)$ and $r^{OD}(s)$. Having solved for these policy functions, solving for equilibrium for a given TBFG policy requires finding a fixed point of a fixed point: a fixed point for the policy functions $y^{FG}(s)$ and $\pi^{FG}(s)$ conditional on a guess for the exit probabilities $p(s)$; and a fixed point where all three objects are consistent. The approach to solving this problem uses the same fundamental building blocks as for the solution to the model with optimal discretion and is explained in Appendix C.2. Solving for a CBFG policy is much more straightforward. As described in Appendix C.3, this is not a fixed point problem and can be solved via backward induction from the terminal period of the policy, $t = K$.

5 Parameterization and experiment scenario

We conduct the majority of the analysis in Section 6 using exactly the same parameterization of the model as Adam and Billi (2006), Adam and Billi (2007) and Bodenstein et al. (2012).[7] Those baseline parameter values are outlined in Table 1. Sensitivity of our policy experiments to an alternative parameter values is discussed in Section 7.

A key element of our approach is the function that determines exit probabilities p from the extent to which the threshold $\bar{x}$ is breached. As noted in Section 4.1, we choose an exponential distribution function (9). This distribution is attractive for two reasons. First, it has just one parameter, which minimizes the number of degrees of freedom when specifying the distribution.

[7]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). Model variables are measured in quarterly deviations from steady state scaled by 100 so that, for example, the interest rate lower bound is $100 \times \left(1 - \beta^{-1}\right)$.

Table 1: Baseline model calibration

Parameter	Description	Value
ξ	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

The second, related, reason is that the α parameter acts as an index of the extent to which the function approaches a 'trigger': as $\alpha \to 0$, the density function gets steeper (see Figure 2).

Figure 2: Alternative calibrations of the mapping $f(x - \bar{x}) = 1 - \exp(-\alpha^{-1}(x - \bar{x}))$

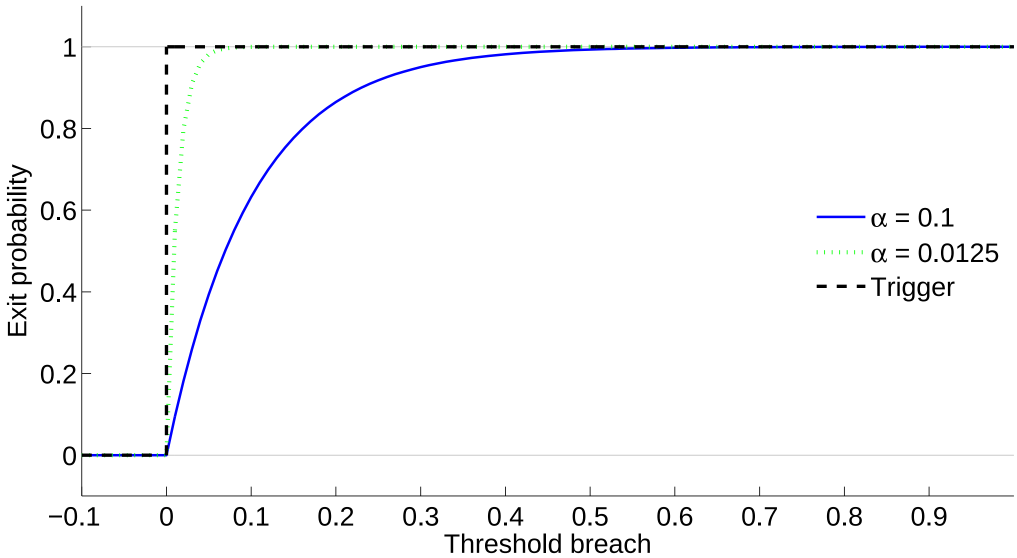

We use different baseline values of α_y and α_π (for output gap thresholds and inflation thresholds respectively) because the slope of the Phillips curve implies that the size of output gap and inflation responses to shocks are markedly different. In Appendix B we use survey evidence to find baseline values of these parameters that imply f functions for output gap and inflation thresholds that are similarly 'steep'. This delivers baseline values of $\alpha_y = 0.1$ and $\alpha_\pi = 0.0125$.

As explained in Section 4.1, allowing for the possibility that exit from the forward guidance regime will not occur with certainty once the threshold has been breached is important to deliver an equilibrium. However, we choose our baseline parameterizations of α_y and α_π to be close to a 'trigger'. This decision is motivated by a desire to ensure that the results are not excessively influenced by the choice of the f function. In other words, we wish to minimize the extent to which the results rely on the assumed link between threshold breaches and exit probabilities. Although the approximation to a trigger is a useful benchmark, Appendix B presents evidence that these thresholds are tighter than implied by survey evidence of financial market participants when the FOMC announced threshold-based guidance in December 2012. Reflecting that evidence, we examine sensitivity of our results to a looser calibration for α_y and

α_π in Section 7.1.

We also examine the behavior of dual thresholds, applied to both the output gap and inflation. Such policies are arguably better approximations to real-world policies. Indeed, the threshold-based forward guidance policies implemented by the FOMC and the Bank of England's MPC applied conditions to more than one macroeconomic variable.[8] We consider a simple approximation to a 'dual threshold' applying to both the output gap and inflation, where the probability of exit is only positive if both variables breach the thresholds (an 'AND' specification).[9] The probability of exit for this dual threshold specification is defined as:

$$\begin{aligned} f(\tilde{x}) = & \left[1 - \exp\left(-\alpha_y^{-1}\tilde{y}\right)\right]\left[1 - \exp\left(-\alpha_\pi^{-1}\tilde{\pi}\right)\right] \\ & \times \left[1 - \exp\left(-\min\left\{\alpha_y^{-1}\tilde{y}, \alpha_\pi^{-1}\tilde{\pi}\right\}\right)\right]\mathcal{I}(\tilde{y})\,\mathcal{I}(\tilde{\pi}) \end{aligned} \tag{15}$$

where $\tilde{y} \equiv \max(0, y - \bar{y})$, $\tilde{\pi} \equiv \max(0, \pi - \bar{\pi})$, and $\mathcal{I}(z)$ is an indicator function taking a value of 1 if $z > 0$ and zero otherwise. This is the joint CDF of a bivariate exponential distribution.

As described in Section 3, the policy experiments are ones in which a large negative demand shock drives the policy rate to the zero bound, prompting the policymaker to implement a one-off forward guidance policy. We calibrate the size of the demand shock to deliver a modal fall in the output gap of 7.5pp 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. More specifically, we find the value for g_0 that delivers $y_1 = -7.5$ given $u_0 = 0$, $\varepsilon_1^u = \varepsilon_1^g = 0$ and a policymaker acting under optimal discretion. The result is $g_0 = -9.4$. Section 7.1 examines sensitivity to an alternative initial condition calibrated to match the fall in output in the United States during the Great Recession.

6 Results

Here we consider results for cases in which the threshold values ($\bar{\pi}$ and $\bar{y}$) have been optimized to deliver the minimum *ex ante* welfare loss (as measured by equation (1)). To do this, we solve for the policy functions for pairs $\{\bar{y}, \bar{\pi}\}$ on a grid.[10] For each pair, we then simulate 100,000 paths[11] of length 24 periods[12] from which we compute the mean discounted loss.

6.1 Headline results

We first consider the performance of alternative specifications of threshold-based guidance, with reference to the baseline policy of optimal discretion. Specifically, we consider policies based

[8] In both cases, a threshold was applied to the unemployment rate but the policy was also contingent on inflation expectations remaining well anchored. In the case of the Bank of England, the conditions applied to inflation expectations (and also financial stability) were termed 'knockouts' indicating a lexicographic dominance over the unemployment threshold: see Monetary Policy Committee (2013) for a comprehensive discussion.

[9] An alternative is an 'OR' specification, where the probability of exit from the forward guidance regime is zero if neither variable satisfies the threshold condition but positive if at least one variable breaches its threshold. For a given set of threshold values, this specification makes exit more likely. As a result, the optimized threshold values are lower than they are for the 'AND' specification. When the threshold values are optimized the 'AND' and 'OR' specifications deliver similar expected outcomes (and losses). Results are available on request.

[10] The $\bar{\pi}$ grid runs from -0.4 to 0.4 in increments of 0.05 and the $\bar{y}$ grid runs from -1 to 6 in increments of 0.25.

[11] A reasonably large number of draws is required to compute summary statistics for the model given the probabilistic nature of exit from the TBFG regime. The results we report are robust to recomputing the simulations with a larger number of draws. For example, the maximum absolute difference in losses computed in simulations using 500,000 draws is less than 0.3% and the loss-minimizing thresholds reported in Table 2 were identical.

[12] The probability of any of the forward guidance policies remaining active in period 24 is less than 0.002%.

on a single inflation threshold ($\bar{\pi}$), a single output gap threshold ($\bar{y}$) and the 'dual' threshold specification, as described by equation (15).

Table 2 records the optimized threshold values for the alternative threshold specifications and the average losses achieved. All threshold specifications considered reduce losses relative to the baseline optimal discretion policy by more than 50% (final column).

Table 2: Results for baseline calibration of model

Threshold type	$\bar{\pi}^*$	$\bar{y}^*$	Loss	$\frac{\text{Loss}}{\text{Loss(OD)}}$
Inflation threshold	0.15	–	0.377	0.444
Output gap threshold	–	2	0.334	0.394
Dual threshold	-0.1	1.75	0.332	0.391

Although the expected loss associated with the optimized inflation and output gap thresholds is similar, the alternative policies do not deliver the same outcomes in all circumstances. Under an inflation threshold, exit from forward guidance can be triggered by either a demand or cost-push shock. By contrast, exit is much less dependent on cost-push shocks for an output gap threshold-based policy. This reflects that cost-push shocks do not affect output directly in the baseline parameterization of the model in which they are assumed to be *iid*.[13]

This logic also explains why an output gap threshold can deliver comparable results to an inflation threshold despite the relatively small weight on output in the loss function (Table 1). The benefit of an inflation threshold is that it can directly mitigate losses arising from high inflation. The cost is that exit from the forward guidance regime can be triggered by transitory cost-push shocks in states of world where underlying inflationary pressure is weak (because demand is weak). Unsurprisingly, this result is overturned when cost-push shocks are autocorrelated, as shown in Section 7.1. In that case, the benefit associated with avoiding high inflationary losses exceeds the cost of cost-push driven exit in weak demand states of the world and an optimized inflation threshold performs better than an optimized output gap threshold.

For the dual threshold in which exit can only occur if both inflation and the output gap exceed the threshold values, the optimized values for the inflation and output gap thresholds ($\bar{\pi}^*$ and $\bar{y}^*$) are both lower than the optimized values for individual inflation and output gap thresholds. The intuition for the result is relatively straightforward. For given values of $\bar{y}$ and $\bar{\pi}$, there are fewer states of the world in which both of the thresholds are simultaneously breached than in which either threshold is breached individually. As a result, the optimal state contingent stimulus can be achieved with lower threshold values because a positive exit probability requires both thresholds to be simultaneously satisfied. Overall, although the dual threshold specification permits more control over the state-contingency of the guidance, it delivers a similar loss to an output gap only threshold, reflecting that the cost-push shock in relatively unimportant in the baseline calibration of the model.

6.2 Inspecting the mechanism

To inspect the mechanism at work, we examine the most likely (that is, modal) outcomes from the optimized dual threshold-based policy and compare them to the modal outcomes under the baseline optimal discretion policy. These paths are constructed by assuming that the shocks for period 1 onwards are equal to their modal value of zero: $\epsilon_t^g = \epsilon_t^u = 0, t = 1, \ldots$.

Panels (a) and (b) of Figure 3 plot the modal responses of the endogenous variables (measured in quarterly deviations from steady state) given the alternative policy strategies. Panels

[13]The presence of cost-push shocks does affect output via the behavior of the policymaker under optimal discretion. However, that effect reflects an active response to the inflationary consequences of cost-push shocks, which is not present when rates are constrained by the zero bound.

(c) and (d) show the modal paths for the policy rate, together with policy rate expectations along the modal path prior to liftoff in period 4 for optimal discretion and TBFG with an optimized dual threshold respectively.

Figure 3: Modal responses under optimized dual threshold policy and optimal discretion

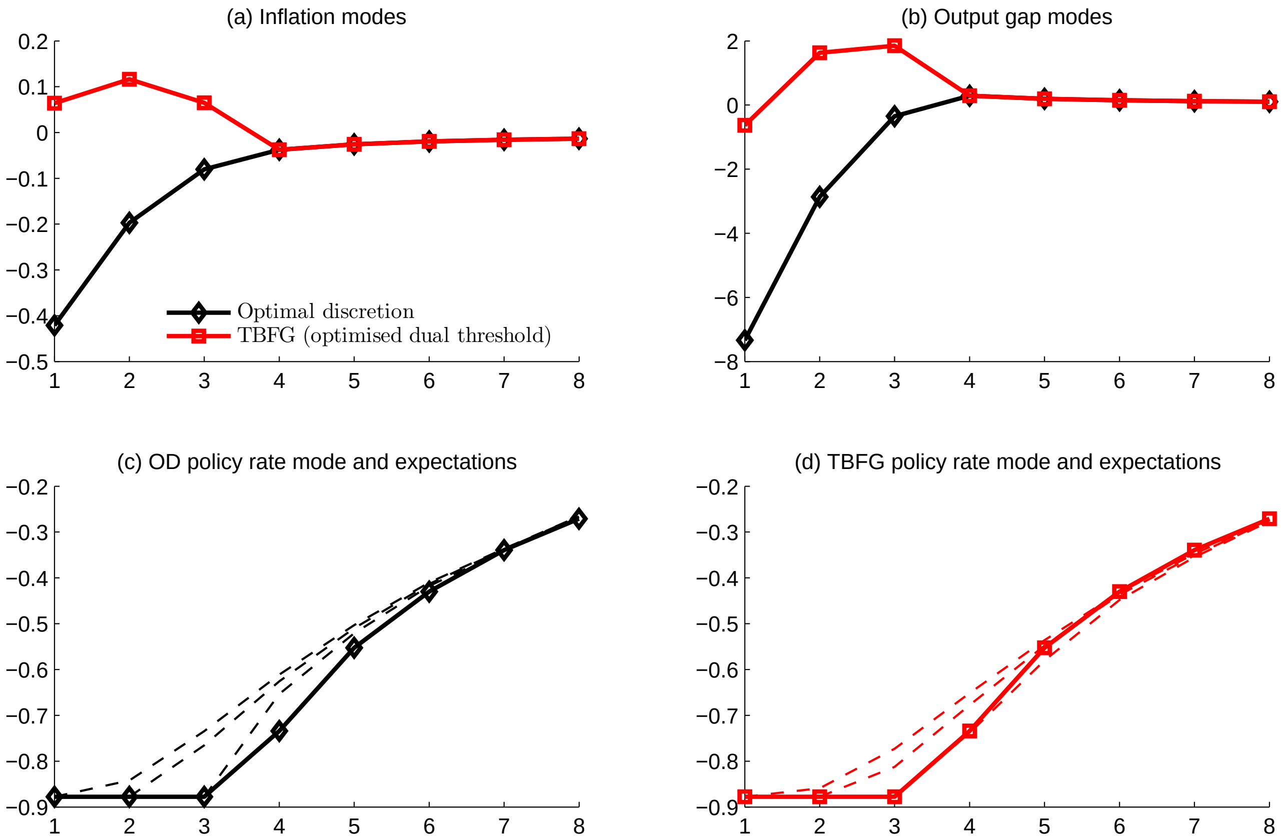

Notes: The baseline policy assumption of optimal discretion (OD) is represented by the black lines with diamond markers. The loss-minimizing threshold-based forward guidance (TBFG) policy is shown in the red lines with square markers. This corresponds to a dual threshold 'AND' specification with $\{\bar{y}^*, \bar{\pi}^*\} = \{1.75, -0.1\}$. The modal responses are computed from the initial condition $g_0 = -9.4, u_0 = 0$ under the assumption that no shocks arrive after period 0. Given that exit is stochastic in the TBFG case, the modal exit period is taken to be the one with the highest ex-ante probability of occurring along the modal path which, in this example, is period 3 with a 0.43 probability. The dashed lines in panels (c) and (d) show expectations for the policy rate prior to liftoff along the modal paths under optimal discretion (c) and the optimized dual TBFG policy (d). All variables are expressed in quarterly percentage point deviations from steady state.

Under the discretionary policy, inflation and the output gap are negative while the policy rate is at the zero bound because the policymaker cannot cut rates below zero and cannot commit to any policy plans that would be inconsistent with loss minimization in the future. Indeed, panel (c) shows that the expected outcomes for the policy rate lie above the modal outcomes at all points prior to liftoff, consistent with a policymaker who cannot commit to hold the policy rate below the rate implied by period-by-period optimization.

Despite the fact that the modal path for the policy rate is identical under threshold-based forward guidance, outcomes for inflation and the output gap are materially different. That is because TBFG embodies a temporary commitment mechanism to hold the policy rate lower than under optimal discretion in certain states of the world. The result is that expectations for the policy rate in future periods are lower than under optimal discretion at all points along the modal path prior to liftoff in period 4 (which can be seen by comparing the dashed lines in panels (c) and (d)). Relative to optimal discretion, the lower expected path for the policy rate increases the expected amount of policy stimulus, which raises inflation and the output gap in periods 1, 2 and 3.[14]

6.3 Threshold-based guidance versus other policies

In this section we compare the optimized dual threshold policy to two alternatives: an optimized calendar-based forward guidance policy (CBFG) and the optimal commitment policy. For each alternative policy, we compute the equilibrium policy functions as described in Sections C.3 and C.4.

Figure 4 shows that modal paths for the optimized dual TBFG policy are close to the optimal commitment benchmark.[15] 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. An (appropriately calibrated) calendar based guidance has similar effects. A CBFG policy that holds the policy rate at the zero bound for an additional period (that is, setting $K = 4$) stimulates inflation and activity today via the effect of a commitment to looser policy in the future. This mechanism is not unique to the policies we consider here. A common theme of related work is that history dependent policies such as optimal commitment, price level targeting or the Reifschneider and Willams (2000) modification of a simple rule can substantially improve outcomes at the zero bound by manipulating inflation expectations to reduce the ex-ante real interest rate as a substitute for cutting the policy rate.[16]

The modal paths under optimal commitment, CBFG and TBFG in panels (a)–(c) of Figure 4 are relatively similar. Although, relative to the other policies, CBFG fails to stabilize the output gap as effectively in period 1, the overshoot in periods 2 and 3 is smaller. As a result, inflation is better stabilized in periods 2 to 4 and per-period losses are smaller in these periods. Despite these differences, all three alternative policies appear to substantially outperform the baseline assumption of optimal discretion. However, a proper assessment of the policies requires consideration of the distribution of outcomes, rather than the outcomes corresponding to a particular realization of the shocks.

Figure 5 plots the frequency of simulations in which the policy rate is at the zero bound for the alternative policies considered. Under TBFG, the probability of remaining at the zero bound is persistently higher than the baseline assumption of optimal discretion. In contrast, under CBFG, the policy rate is held at the zero bound with certainty for four periods (and is thereafter set according to optimal discretion). This observation suggests that the overall performance of CBFG may be compromised by the lack of state contingency in the policy setting for the first four periods. In the first two periods, optimal commitment generates a slightly higher probability of rates at the zero bound compared to TBFG. Thereafter, the probability that rates are at the zero bound is lower, gradually falling below the probability generated by optimal discretion. This latter result reflects the fact that optimal commitment policy operates

[14] Note also that inflation and the output gap deviate from the optimal discretion equilibrium in period 3 even though the policymaker exits the threshold-based policy in this period. This reflects our within-period timing assumption: output and inflation are chosen before the policymaker chooses the policy rate. These choices are influenced by the possibility that the forward guidance regime will remain in effect. The effect of this on expectations for the policy rate can be seen in panel (d), which shows that, conditional on the modal state in period 3, the expected policy rate lies below the modal path in all future periods reflecting that there is always a non-zero probability that the forward guidance policy will still be active.

[15] 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 zero bound for longer than under optimal discretion.

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

Figure 4: Modal responses under optimized forward guidance and optimal policy

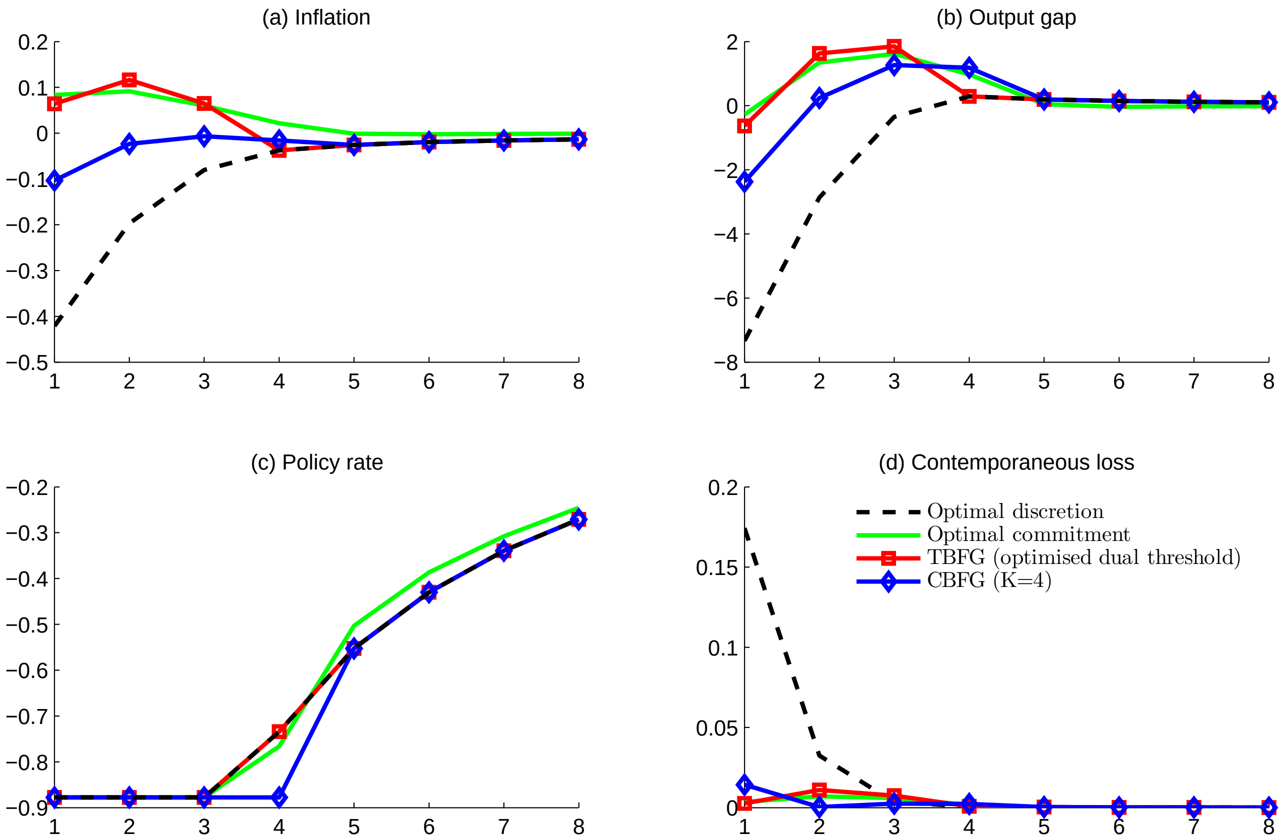

Notes: The baseline policy assumption of optimal discretion is represented by the dashed black lines (no markers). The solid green lines (no markers) correspond to optimal commitment. The blue lines with diamond markers show the responses under a calendar-based forward guidance (CBFG) in which the policymaker commits to hold the policy rate at the zero bound for $K = 4$ periods. The loss-minimizing threshold-based forward guidance (TBFG) policy is shown in the red lines with square markers. This corresponds to a dual threshold 'AND' specification with $\{\bar{y}^*, \bar{\pi}^*\} = \{1.75, -0.1\}$. The responses are computed from the initial condition $g_0 = -9.4, u_0 = 0$ under the assumption that all subsequent shock realizations are equal to zero. The responses in panels (a)-(c) are plotted as quarterly percentage point deviations from steady state. The contemporaneous loss in panel (d) is computed on a period-by-period basis using the loss function in equation (1).

through the policymaker's ability to manage expectations. Since the private sector internalizes these effects when forming expectations of future outcomes, the probability of hitting the zero bound in future periods is reduced.

Figure 6 plots the distributions of outcomes under the baseline policy of optimal discretion, optimized calendar-based guidance, optimal commitment and optimized threshold-based guidance. The distributions of inflation and the output gap are negatively skewed under optimal discretion (first column) 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.[17]

The contrast between CBFG (second column) and TBFG (fourth column) is stark. Calendar-based guidance 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 zero bound constraint (Figure 3), it leads to worse outcomes in both good and bad states than

[17] This has implications for policy even if the zero bound 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 zero bound, demand shocks induce a policy trade-off (see, for example, Adam and Billi, 2006; Nakov, 2008).

Figure 5: Probability that the policy rate is at the zero bound under alternative policies

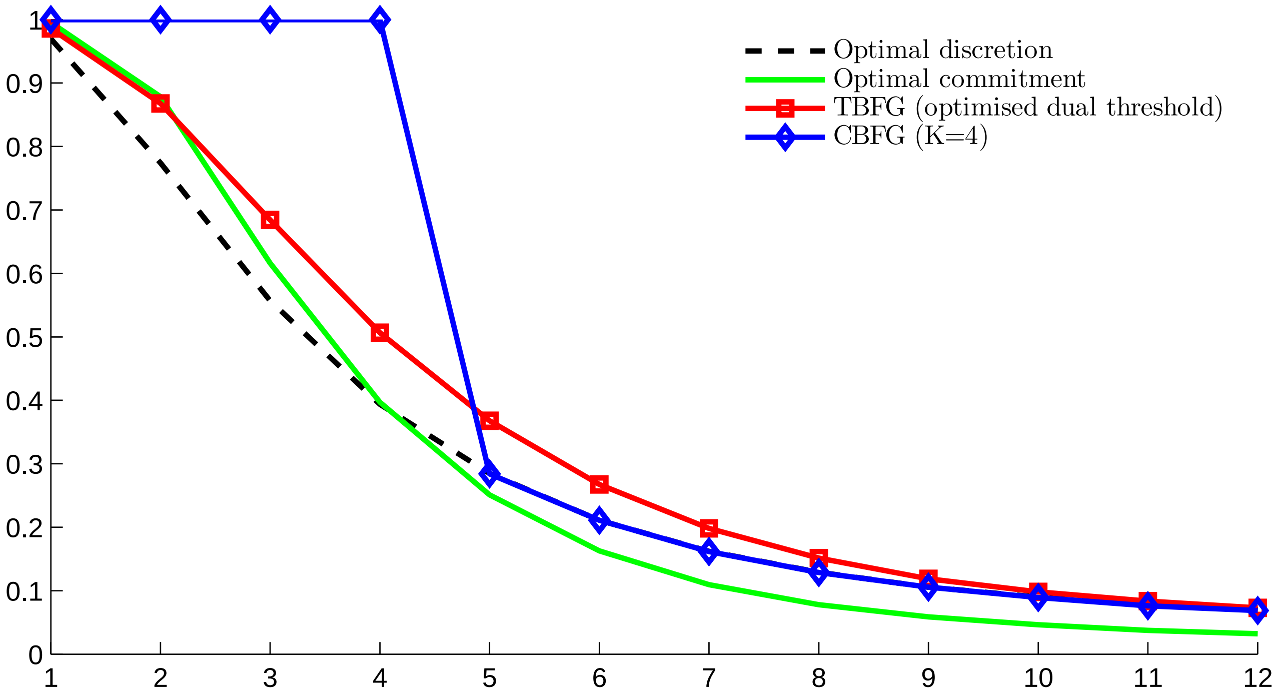

Notes: Computed from a stochastic simulation of $100,000$ draws over 24 periods from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$. The baseline policy assumption is optimal discretion (dashed lines). The results from optimal commitment are shown in the solid green line (no markers) and are computed with the Lagrange multipliers initialized at zero. CBFG ($K = 4$) refers to a calendar-based forward guidance policy in which rates are held at the zero bound for 4 periods (blue lines with diamond markers). Results from the loss-minimizing threshold-based forward guidance (TBFG) policy is shown in the red lines with square markers. This corresponds to a dual threshold 'AND' specification with $\{\bar{y}^*, \bar{\pi}^*\} = \{1.75, -0.1\}$.

appropriately-calibrated threshold-based policies. That is because CBFG provides too much stimulus in good states and insufficient stimulus in bad states. As a result, the variance of the distributions of inflation and the output gap increases substantially.

TBFG cannot stabilize the distributions of inflation and the output gap as well as the optimal commitment policy (third column). Nevertheless, the improvement over the distribution generated by the baseline policy of optimal discretion (first column) is substantial. The state-contingent promise to hold rates at the zero bound until the thresholds are breached provides additional stimulus in bad states of the world, hedging against some of negative skew associated with the zero bound.

That TBFG comes close to replicating outcomes under optimal commitment at the zero bound is certainly of relevance to policymakers. Threshold-based guidance 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 often 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 zero bound.

The results above demonstrate that an appropriately calibrated TBFG policy can achieve substantially better outcomes at the zero bound than the optimal discretionary policy. But engineering an overshoot of inflation and/or 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

Figure 6: Distributions of outcomes under alternative policies

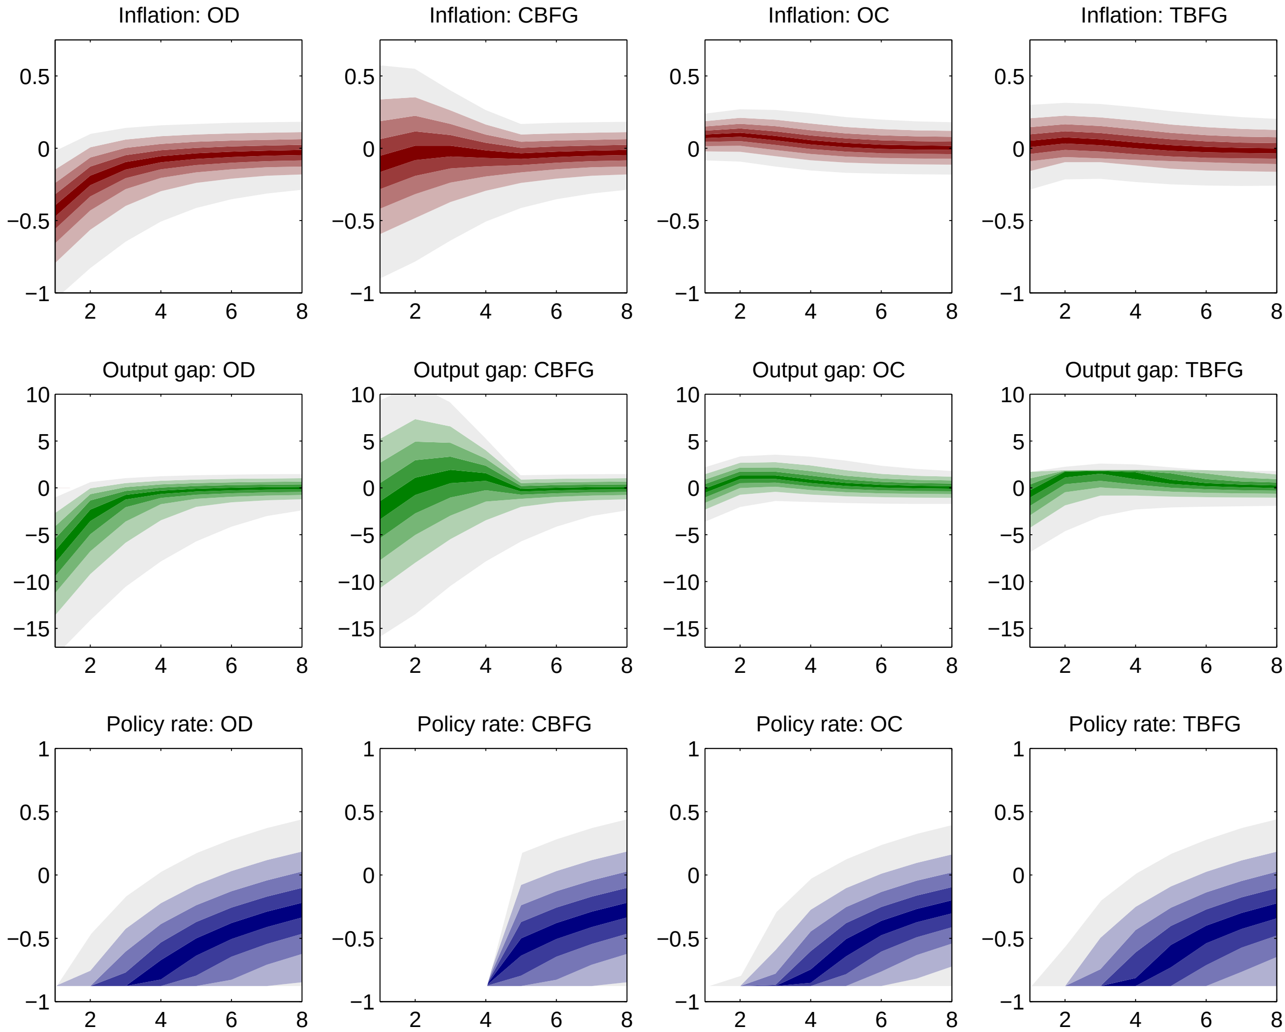

Notes: Distributions of endogenous variables when policy is set according to: optimal discretion (OD); calendar-based forward guidance (CBFG); optimal commitment (OC) and threshold-based forward guidance (TBFG). The CBFG policy is one in which the policy rate is held at the zero bound for $K = 4$ periods. The TBFG policy is the loss-minimizing threshold-based forward guidance policy: a dual threshold 'AND' specification with $\{\bar{y}^*, \bar{\pi}^*\} = \{1.75, -0.1\}$. Distributions for all cases are computed from a stochastic simulation of 100,000 draws over 24 periods from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$. The distributions are represented in 'fan charts' with each shaded band containing 10% of the cumulative probability mass.

policy remains in the forward guidance regime). More formally, denote the measure of time inconsistency of a particular policy, P, in period t, as $\mathbb{T}_t^P$:

$$\mathbb{T}_t^P = \int_u \int_g \psi_t^P (u, g) \left(\mathbb{L}_t^P (u, g) - \mathbb{L}_t^{OD} (u, g)\right) \mathbb{I} \left(\mathbb{L}_t^P (u, g) - \mathbb{L}_t^{OD} (u, g) > 0\right) \mathrm{d}g \, \mathrm{d}u \tag{16}$$

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:

$$\mathbb{L}_t^J (u, g) = \sum_{s=t}^{\infty} \beta^{s-t} \mathbb{E}_t (u, g) \left(\pi_s^J (u, g)^2 + \lambda y_s^J (u, g)^2\right) \tag{17}$$

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

Figure 7 illustrates that the incentive to renege on threshold-based guidance is small relative to the incentive to renege on calendar-based guidance. Even though the TBFG policy generates higher modal inflation and the output gap than the CBFG policy (see Figure 3), the incentive to renege on the threshold-based guidance is smaller than for calendar-based guidance (until the calendar-based policy reverts to the optimal discretionary policy from period 5 onwards). This demonstrates that threshold-based guidance can be less time inconsistent than calendar-based guidance, even when it imparts more stimulus in expectation.

Figure 7: Time-inconsistency measures for optimized forward guidance policies

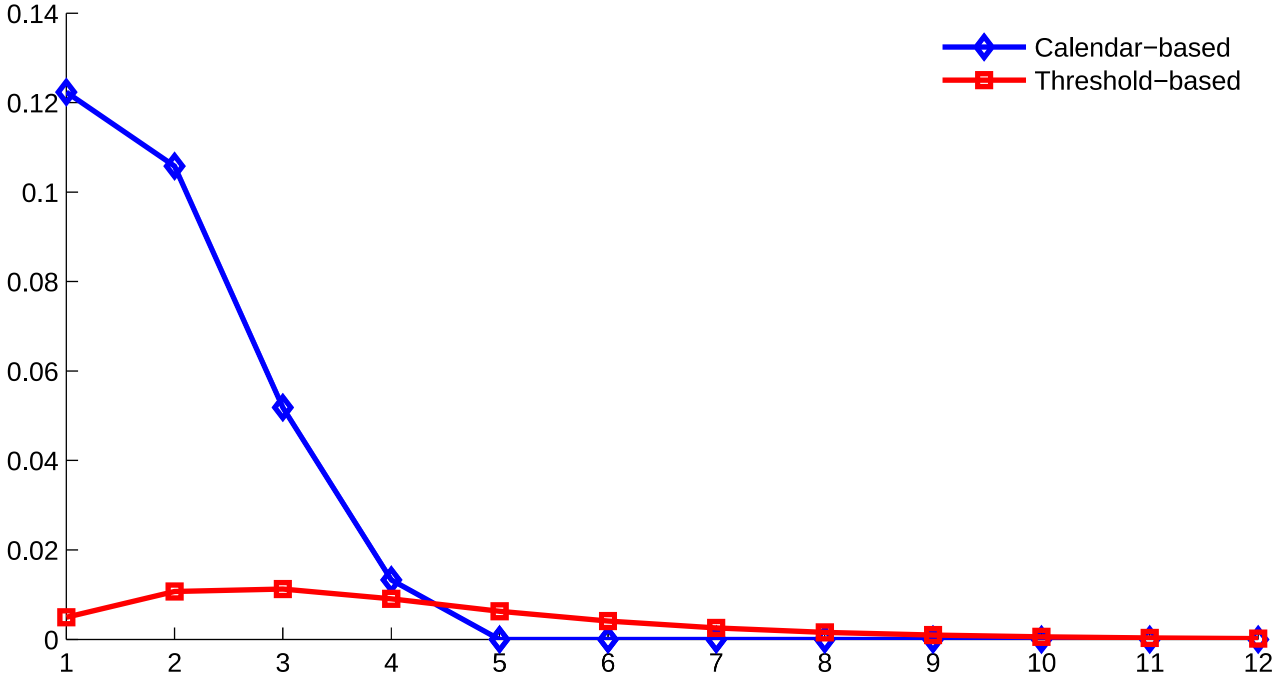

Notes: Computed by replacing the density measure in equation (16) with a discrete approximation on the state grid from from the initial condition for the state of $g_0 = -9.4$ and $u_0 = 0$, which we use to replace the integral with a finite sum – see, for example, Chapter 5 of Heer and Maussner (2005). The calendar-based forward guidance (CBFG) policy is one in which the policy rate is held at the zero bound for $K = 4$ periods. The threshold-based forward guidance (TBFG) policy is the loss-minimizing dual threshold 'AND' specification with $\{\bar{y}^*, \bar{\pi}^*\} = \{1.75, -0.1\}$.

Although there may be less temptation to renege from TBFG it is nevertheless time inconsistent to some degree. But that does not necessarily make such policies uninteresting from the perspective of a 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 supported as reputational equilibria if zero bound episodes are sufficiently frequent and persistent. Similarly, Walsh (2016) uses a 'sustainable plans' approach to show that calendar based forward guidance may be a sustainable plan when shocks may drive the economy to the zero bound in future. In that context, TBFG is more likely to be supportable by a concern for reputation than CBFG because it embodies less time inconsistency in the absence of reputational mechanisms.[18]

[18]In principle, it could be possible to make threshold-based 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 an inflation threshold. The option expires when inflation exceeds its threshold for the first time. If the central bank honors its promise and keeps the policy rate at the zero bound 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.

7 Sensitivity and robustness analysis

In this section, we consider the sensitivity of our results along several dimensions. In Section 7.1, we consider the sensitivity of our results to alternative parameterizations of the model, an alternative calibration of the mapping from threshold breaches to the probability of exiting the policy and an alternative specification of the initial demand state (so that the recession is less severe than the baseline case). In Section 7.2, we consider whether our results for alternative specifications can provide any simple rules of thumb for calibrating threshold-based policies designed to stimulate the economy in the face of a recessionary shock.

7.1 Sensitivity analysis

Before presenting the results of our sensitivity analysis, we first briefly summarize the alternative specifications that we consider.

Model parameterization

The results obtained in Section 6 will naturally depend on the parameter values used. To test the sensitivity of our results, we and consider three alternative parameterizations of the model, based on the sensitivity analysis presented in Adam and Billi (2007).

Table 3: Alternative model parameterizations

	Baseline	'RBC calibration'	'Larger cost shock'	'Lower IES'
σ	6.2500	1.0000	6.2500	1.0000
ρ_u	0.0000	0.3600	0.3600	0.0000
σ_u	0.1540	0.1710	0.1710	0.1540
σ_g	1.5240	0.2940	1.5240	0.2940
κ	0.0240	0.0569	0.0240	0.0569
λ	0.0031	0.0074	0.0031	0.0074
g_0	-9.4000	-2.8750	-11.2500	-3.2500
g_0/σ	-1.5040	-2.8750	-1.8000	-3.2500

Table 3 presents the alternative parameterizations of the model that we consider. The table documents the baseline values of the relevant parameters and the values used in the alternative variants. Since some of the parameterizations change the values of 'composite' parameters (namely κ and λ) the values of these parameters are also reported. The changes in model calibration also change the value of the initial demand state, g_0, that is consistent with an initial fall in output of 7.5% (in the 'modal' case). The values of g_0 used to generate this outcome (to ensure that the simulations across model variants are more easily comparable) is also reported. All other parameters remain at the baseline values shown in Table 1.

Following Adam and Billi (2007) we consider an 'RBC calibration' which, relative to the baseline model, features a smaller intertemporal elasticity of substitution ($\sigma = 1$), which increases the slope of the Phillips curve (κ) and the weight on the output gap in the loss function (λ). This variant of the model also features cost push shocks with greater persistence ($\rho_u = 0.36$) and standard deviation ($\sigma_u = 0.171$). The standard deviation of the demand shock ($\sigma_g = 0.294$) is smaller than the baseline value.[19] The RBC calibration involves changes in both the responsiveness of demand to real interest rates and the properties of the disturbances to the model, in particular the relative importance of cost push shocks (u) and demand shocks

[19]This implies that the standard deviation of the disturbance to the IS curve measured in 'real interest rate units' (that is, $\sigma^{-1}\sigma_g$) is virtually unchanged from the baseline parameterization.

(g). To isolate the relative importance of these factors, we also consider parameterizations in which these groups of parameters are changed in turn. Specifically, the 'larger cost shock' variant holds σ and σ_g at their baseline values and increases only ρ_u and σ_u in line with the RBC calibration. The a 'lower IES' calibration that reduces σ and σ_g in line with the RBC calibration, while holding ρ_u and σ_u at their baseline values.

Threshold specification

As discussed in Section 5, our baseline specification for the mapping from threshold breaches to the probability that the policymaker will exit from the forward guidance regime is intended to mimic a 'trigger' (see Figure 2). Survey evidence, presented in Appendix B, suggests that financial market participants may have viewed the FOMC's December 2012 threshold-based guidance as less strict than a trigger. That is, market participants seem to have assigned a non-negligible probability to exit from the forward guidance regime occurring only once macroeconomic variables had moved substantially beyond the stated threshold values. We mimic these types of beliefs by recomputing the threshold-based forward guidance experiments using a 'flatter' f function (9). To implement this, we double the values of the α parameters.[20]

Scale of the recession

Our baseline specification of the recession scenario was chosen to deliver a modal fall in output in the first quarter of the simulation equal to around 7.5% when policy is set under optimal discretion. The magnitude of the recession is therefore broadly in line with the Great Depression: Eggertsson (2011) asserts that US output fell by around 30% on an *annual* basis. In this section we consider a specification that is more in line with the Great Recession. Christiano et al. (2011) argue that the Great Recession generated a decline in annual output of around 7%. We therefore set the initial condition for the demand state to $g_0 = -6.925$ which generates a modal fall in quarterly output of around 1.75% in the first quarter of the simulation under optimal discretion.

Sensitivity results

Table 4 confirms that the optimal threshold values depend on the specific parameterization of the model. However, some broad themes emerge from the results. First, TBFG reduces losses relative to optimal discretion by around 30% or more in all cases considered. The relative improvement over the baseline assumption of optimal discretion does depend on the parameterization of the model, as discussed below. Second, there are regular patterns between the optimized values for single and dual threshold specifications. In all cases, the dual specification (in which exit cannot occur until both thresholds have been breached) features values for the inflation and output gap thresholds that are smaller than the optimized values for the respective single threshold policies. Furthermore, in all cases the ex-ante loss associated with the dual threshold policy is lower than for single threshold policies, reflecting that a dual threshold policy gives the policymaker more flexibility in managing the state-contingency of the policy.

Turning to the differences from the baseline specification, we first consider the RBC calibration. Relative to the baseline specification, the optimized values for inflation thresholds are larger and the optimized values for the output gap thresholds are smaller. Compared with the baseline parameterization, the benefits of threshold-based guidance are larger: optimally chosen thresholds achieve losses of less than 30% of the level delivered by optimal discretion (compared with around 40% for the baseline variant of the model).

[20] So for inflation and output gap thresholds we use $\alpha_\pi = 0.025$ and $\alpha_y = 0.2$ respectively.

Table 4: Results for alternative calibrations

Variant	Threshold type	$\bar{\pi}^*$	$\bar{y}^*$	$\frac{\text{Loss}}{\text{Loss(OD)}}$
Baseline calibration	Inflation threshold	0.15	–	0.444
	Output gap threshold	–	2	0.394
	Dual threshold	-0.1	1.75	0.391
'RBC' calibration	Inflation threshold	0.25	–	0.264
	Output gap threshold	–	1	0.284
	Dual threshold	0.15	0	0.254
Larger cost shock	Inflation threshold	0.1	–	0.556
	Output gap threshold	–	2.25	0.548
	Dual threshold	0	0.75	0.528
Lower IES	Inflation threshold	0.25	–	0.215
	Output gap threshold	–	1	0.198
	Dual threshold	0	0.75	0.188
$\alpha_y = 0.2; \alpha_\pi = 0.025$	Inflation threshold	0.1	–	0.447
	Output gap threshold	–	2	0.394
	Dual threshold	-0.1	1.5	0.393
'Great Recession' calibration	Inflation threshold	0.05	–	0.754
	Output gap threshold	–	1.5	0.704
	Dual threshold	-0.15	1.25	0.7

We also observe that, for the RBC variant, the loss from the optimized inflation threshold is smaller than loss from the optimized output gap threshold. This finding is consistent with the increased role for cost push shocks under this calibration. Importantly, when cost push shocks are persistent, they affect the output gap at the zero bound through the effects on inflation expectations and hence the real interest rate. So the output gap becomes a less reliable indicator of the demand state (g) if a sizable fraction of its variability is explained by cost push shocks and a threshold-based policy based solely on the output gap performs worse than a threshold-based policy based on the inflation rate.

This logic is confirmed from the results for the variant of the model with a lower intertemporal elasticity of substitution. That case is consistent with the baseline parameterization: we observe smaller losses with an optimized output gap threshold than with an inflation threshold. This suggests that it is the increased role of cost push shocks in the RBC calibration that makes the output gap a less useful threshold variable relative to the baseline parameterization.

For both the 'lower IES' and 'RBC' variants, the demand state g_0 consistent with an initial 7.5% fall in output implies a much longer average duration at the zero bound. This can be seen from the final row of Table 3 which reports g_0/σ, mapping the size of the initial demand state into interest rate units. The demand state needs to be substantially weaker for the 'lower IES' and 'RBC' variants, which generates a more prolonged period at the zero bound under the baseline policy of optimal discretion. This implies that the benefits from additional stimulus are substantial so that losses relative to the baseline case can be reduced by more than 80% in some cases.

However, the 'larger cost shock' variant also exhibits some contrasting results. Compared to the baseline parameterization (with *iid* cost-push shocks), the optimal inflation threshold is *lower*. One reason is that the modal path for inflation under optimal discretion is higher than the modal path generated by the baseline parameterization of the model. This is despite the fact that a larger negative demand state is required to deliver an initial 7.5% fall in output ($g_0 = -11.25$ versus -9.4 in the baseline). Because the distribution for inflation is higher (than the baseline parameterization) under optimal discretion, a given level of additional stimulus can be achieved with a lower inflation threshold. The output gap threshold is higher than the baseline parameterization, reflecting the possibility that a positive output gap could be generated by a negative cost-push shock (which was absent in the *iid* case).

Table 4 reveals that our results are relatively insensitive to the choice of α coefficients. As we would expect, the optimal threshold values under the higher α calibrations are (very slightly) lower than for the baseline case. This reflects the fact that a given threshold breach is associated with a lower exit probability for a higher value of α. That is, a given output gap threshold policy could be approximated by an alternative policy with a lower value of $\bar{y}$ and a higher α_y. These results suggest that our choice of baseline values for α are close to a 'trigger', since doubling those values has very little effect.

Finally, Table 4 shows that, for a scenario calibrated to match the Great Recession, the optimal threshold values are lower than the baseline scenario because the initial fall in output is smaller. As we would expect, a smaller recessionary shock means that the optimal amount of stimulus is also smaller. We also observe that the ability of threshold-based guidance to improve outcomes relative to optimal discretion is somewhat reduced: *ex ante* losses under optimal thresholds are around 70% of the losses achieved under optimal discretion, whereas losses are more than halved in the baseline case. Again, this reflects the fact that the recession scenario is less severe in this instance, so that the zero bound on the policy rate has a smaller effect on the distributions of outcomes.

7.2 Robust 'rules of thumb' for thresholds?

In this section we consider whether, despite the differences across the alternative parameterizations discussed above, our results suggest any general guidelines for the calibration of threshold-based guidance policies. Figure 8 plots the losses measured relative to those achieved under optimal discretion for alternative output gap thresholds (left panel) and inflation thresholds (right panel).

Figure 8: Losses (relative to optimal discretion) for alternative model variants

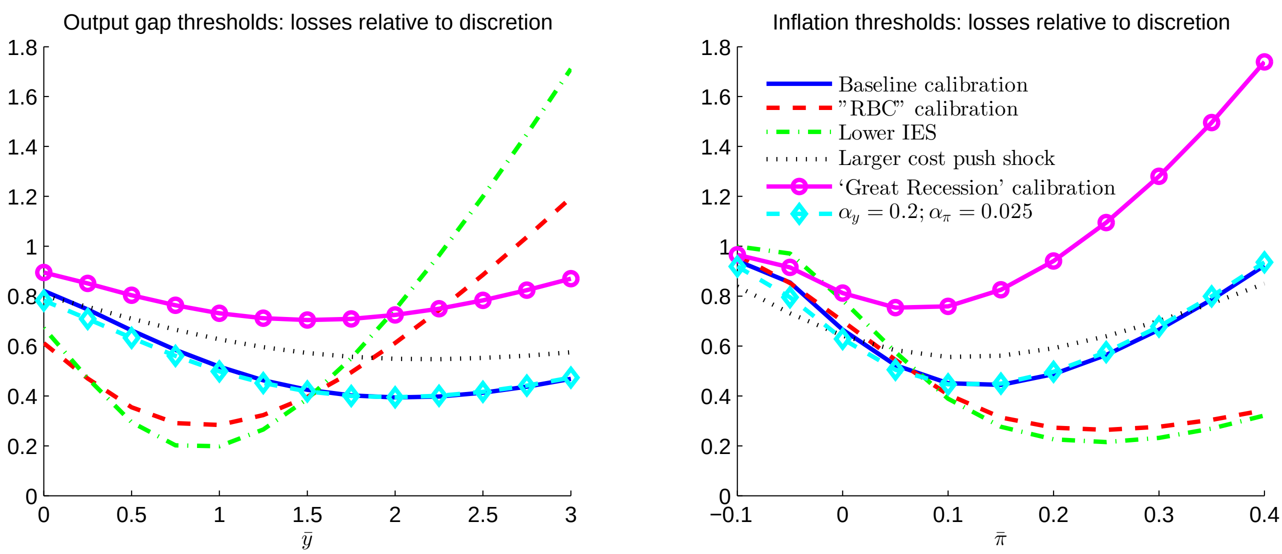

Notes: The left panel shows the losses, measured as a fraction of those attained under optimal discretion, for a range of values of the output gap threshold, $\bar{y}$. Each line shows the results for a distinct variant of the model/experiment. The right panel shows analogous results for the case of the inflation threshold $\bar{\pi}$.

These results suggest that for a policymaker who is uncertain about the structure of the economy, or about private agents' interpretation of the threshold-based policy (α), an output gap threshold between 1 and 1.5 per cent would generate reasonable welfare gains in all cases. Similarly, inflation thresholds between 0.1 and 0.2 per cent (measured as a quarterly rate, relative to steady-state inflation) will typically generate welfare gains of 40% or more. The exception to these rules of thumb is the case in which the initial recessionary shock is relatively small (the 'Great Recession' variant). In that case, the benefits of threshold-based guidance are relatively small and, intuitively, it pays to be more conservative in the calibration of the thresholds. Nevertheless, output gap and inflation thresholds at the lower end of the rule of thumb range (1% and 0.1% respectively) would be welfare-improving even in this case.

Finally, we note that the profiles of losses for the baseline calibration (blue solid lines) lie (very slightly) to the right of the profiles of losses generated by the experiment using higher values of α (dashed cyan lines with diamond markers). This suggests that, other things equal, policymakers should choose higher threshold values if they believe that private agents will interpret the threshold as a 'trigger' (so that exit from the lower bound is almost certain once the threshold has been breached).

8 Conclusion

In this paper we have examined threshold-based forward guidance (TBFG) as a tool to escape a liquidity trap within a general equilibrium framework. Though TBFG has been implemented in practice, monetary policymakers have rarely described these policies in terms of providing additional stimulus, in part because of their skepticism about their ability to commit credibly to time inconsistent behavior. We show that TBFG *can* be used as a temporary policy to stimulate the economy at the zero bound, while also limiting the extent to which a policymaker must promise to behave in a time inconsistent manner.

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.

Bank of England (2013): *Monetary Policy Trade-offs and Forward Guidance*, London: 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.

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,” .

Christiano, L., M. Eichenbaum, and S. Rebelo (2011): “When Is the Government Spending Multiplier Large?” *Journal of Political Economy*, 119, 78–121.

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. (2011): “What fiscal policy is effective at zero interest rates?” in *NBER Macroeconomics Annual 2010, Volume 25*, University of Chicago Press, 59–112.

English, W. B., J. D. Lpez-Salido, and R. J. Tetlow (2015): “The Federal Reserves Framework for Monetary Policy: Recent Changes and New Questions,” *IMF Economic Review*, 63, 22–70.

Florez-Jimenez, M. L. and J. A. Parra-Polania (2016): “Forward guidance with an escape clause: When half a promise is better than a full one,” *Applied Economics*, 48, 1372–1381.

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.).

Monetary Policy Committee (2013): "Monetary policy trade-offs and forward guidance," Bank of England.

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

——— (2015): "Credibility of optimal forward guidance at the interest rate lower bound," FEDs Notes, August, accessed: 15-08-2016.

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. Willams (2000): "Three Lessons for Monetary Policy in a Low-Inflation Era," *Journal of Money, Credit and Banking*, 32, 936–66.

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*.

Walsh, C. (2016): "Sustainable forward guidance at the ELB," *mimeo*.

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.

Yellen, J. L. (2013): "Communication in monetary policy," Speech at the Society of American Business Editors and Writers 50th Anniversary Conference, Board of Governors of the Federal Reserve System.

A Equilibrium selection with threshold-based policies

In this appendix, we provide some intuition for why a threshold-based guidance policy may fail to deliver a unique equilibrium using a simple deterministic example. We also explain how introducing probabilistic exit from the forward guidance policy helps to resolve this problem.

A.1 Intuition using a deterministic example

Here we build intuition for our approach to threshold-based forward guidance described in Section 4.1 by using a deterministic example. The basic environment is the same as that outlined at the beginning of Section 3: the economy begins with a very low state of demand in period $t = 0$; the policymaker (who ordinarily sets an optimal discretionary policy) 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 lift off from the zero bound and so there is an equivalence between threshold-based and calendar-based forward guidance in which the policymaker commits to hold rates at the zero lower bound for a specific number of periods.

Figure 9 shows alternative paths for the output gap in a deterministic version of the model outlined in Section 2 given alternative policies implemented from period 1 and given an initial condition for the state of demand, $g_0 = -11.95$.[21] In the case in which the policymaker continues to set policy following optimal discretion (the black line), the policy rate 'lifts off' 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 in cases where the policymaker credibly commits to hold rates at the zero bound 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 diamond markers respectively. The differences in outcomes are large and are a non-linear function of the duration of the calendar-based guidance policy as has been documented by, for example, Carlstrom et al. (2012).

Suppose that instead of announcing a calendar-based policy, the policymaker announces a threshold-based policy with an output gap threshold of 1.75, as indicated by the horizontal dashed black line in Figure 9. Which of the paths shown in Figure 9 is the equilibrium for the threshold-based 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 plausible candidates. The policy with liftoff in period 6 would result in a path for the output gap along which the threshold was breached in *some periods* prior to exit, but by the smallest amount of all such policies. The policy with liftoff in period 5 delivers an output gap path that does not breach the threshold in any period, but which comes closest to doing so among all such policies. 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 discussion demonstrates that even in a simple deterministic setting, the announcement of a threshold is not sufficient to determine an equilibrium outcome. It is also necessary for

[21]The model has been re-solved 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 Section 6.

Figure 9: Outcomes under alternative policies in a deterministic setting

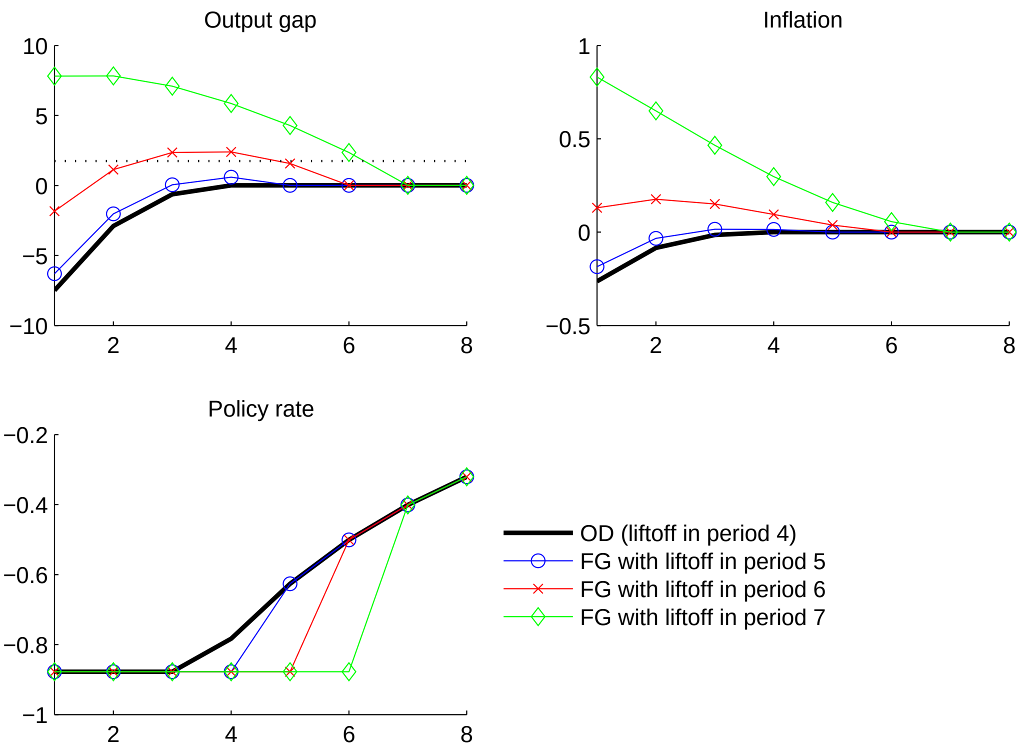

Notes: Computed from an initial condition of $g_0 = -11.95$ and $u_0 = 0$. No shocks arrive or are expected to arrive thereafter. Otherwise, the model is identical to that described in Section 2 with the baseline calibration outlined in Section 5.

the policymaker to specify precisely how they will determine regime exit. 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 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.[22]

The deterministic example considered here also reveals that equilibrium selection concerns the entire expected path of the policy rate. 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 preceding 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 zero bound by 1 period without the threshold condition being breached. However, note that if agents had known that the policymaker would behave in this way prior to period 5, then expectations would be governed by the policy with liftoff in period 6, in which case the threshold is expected to be breached in periods 3 and 4.

Conversely, suppose that the policymaker announces that exit will occur in the period immediately after the threshold has been breached. The path corresponding to liftoff in period 6 cannot be supported as an equilibrium in this case. To see this, note that this path first breaches the threshold in period 3, implying that exit would occur in period 4. However, if agents know that exit will occur with certainty in period 4 their expectations would be consistent with outcomes under the baseline policy (optimal discretion, OD) and the resulting

[22] In a previous version of this paper, we selected a unique equilibrium by maximizing the expected duration of the policy regime 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 deterministic example of Figure 9).

equilibrium would be as depicted in the black lines.

It is evident that one issue that complicates the selection of an equilibrium such as the red line in Figure 9 is the dynamic responses of the macroeconomic variables in response to delayed liftoff. That is, in order to generate near term stimulus the policymaker must typically promise a subsequent overshoot of the output gap (or indeed inflation). But such an overshoot implies that the desired trajectory for the output gap approaches steady state equilibrium from above. The fact that the output gap is falling in the period prior to exit from the forward guidance regime complicates the use of a threshold-based policy that ties exit to a condition that the output gap must *exceed* a particular threshold.

A.2 Probabilistic exit

We now extend our deterministic example to consider the case in which a breach of the threshold conditions is associated with *probabilistic* exit from the forward guidance regime.

To examine the outcomes in this case, we use the algorithm developed by Haberis et al. (2014). A full explanation and derivation of the algorithm is provided by Haberis, Harrison, and Waldron (henceforth HHW), so here we summarize the key elements as they apply to our specific model and policy experiment.

The HHW algorithm is applied to linear, deterministic models with the following timing structure:

- At the start of period t households and firms choose output and inflation $\{y_t^*, \pi_t^*\}$. They do so on the basis of the expected path of the policy rate set by the policymaker, *including the rate that will be set in period t.*

- With probability p_t the policymaker exits from the zero bound and sets the policy rate according to the usual policy rule (that is, optimal discretion). With probability $1 - p_t$ the policymaker continues to hold the policy rate at the zero bound.

- If the policymaker reverts to optimal discretion at date t then it continues to set policy under optimal discretion forever more. So the forward guidance policy is a 'one-off' policy and reversion to optimal discretion is an absorbing state.

- If the policymaker continues to set the policy rate at the zero bound in date t then we enter date $t+1$ with households and firms choosing $\left\{y_{t+1}^*, \pi_{t+1}^*\right\}$ as described above.

- If we reach the end of a particular period K and the policymaker has not reverted to setting policy according to optimal discretion in any of the periods $t = 1, \dots, K$, then the policymaker reverts to optimal discretion for sure in period $K+1$. At the start of the policy ($t = 0$) the date K is known to all agents.

HHW show that there are $K+1$ possible paths for the model, corresponding to reversion to optimal discretion in periods $t = 1, \dots, K+1$. They also show how to solve simultaneously for these paths, conditional on a set of exit probabilities $\{p_t\}_{t=1}^K$. The solution procedure is a 'stacked time' approach in which the $K+1$ possible paths are stacked together, weighted according to the probabilities of each path occurring. Those probabilities are straightforward to compute given the fact that reversion to optimal discretionary policy is an 'absorbing state'.

Solving for the case of threshold-dependent exit probabilities requires a straightforward iterative approach. Conditional on a guess for the exit probabilities $\{p_t\}_{t=1}^K$, the equilibrium paths can be computed using the algorithm described above. These paths can then be used to compute a new guess for the exit probabilities, using a pre-specified mapping from threshold breaches to exit probabilities. Iteration between updating equilibrium paths and exit probabilities continues in this way until a fixed point is reached.

Figure 10: Threshold-based probabilistic exit

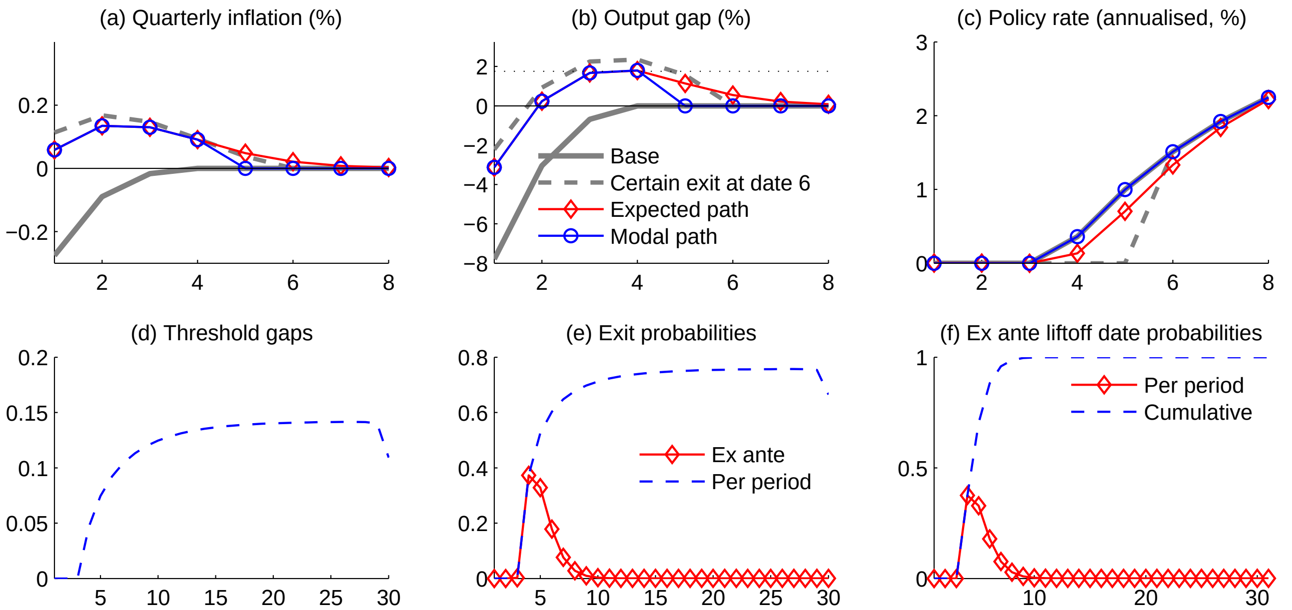

Notes: Computed from an initial condition of $g_0 = -11.95$ and $u_0 = 0$. No shocks arrive or are expected to arrive thereafter. The model is identical to that described in Section 2 with the baseline calibration outlined in Section 5. Exit from the zero bound is determined by a probabilistic exit criterion: the probability of exit in period t is $1 - \exp\left(-\alpha^{-1}\left(y_t - \bar{y}\right)\right)$ if the output gap y_t exceeds the threshold $\bar{y}$ and 0 otherwise. The parameters are $\alpha = 0.1$ and $\bar{y} = 1.75$ (the dotted horizontal line in panel (b)).

We apply the HHW algorithm to our deterministic example, using the baseline calibration for the mapping between (output gap) threshold breaches and exit probabilities in the main text: the mapping from threshold breach to exit probability is given by (9) with $\alpha = 0.1$ and $\bar{y} = 1.75$. We set K, the maximum number of periods of the forward guidance policy, to 30.

The results are shown in Figure 10. Panels (a)–(c) show the responses of inflation, the output gap and the policy rate. The solid gray lines are the baseline case of optimal discretion. The dashed gray lines correspond to a calendar-based forward guidance policy with exit *for certain* in period 6. The dotted horizontal line in the output gap chart is at 1.75% and we note that the dashed gray line (exit in period 6 with certainty) is the trajectory that is 'close' to the threshold but cannot be supported as an equilibrium in which exit triggered with certainty when the threshold is breached (as discussed in Appendix A.1 above).

Still focusing on the top row, the red lines with diamond markers are the *expected* paths from the simulation in which exit from the zero bound is probabilistic. The blue lines are the *modal* paths, corresponding to exit at $t = 4$.

The fact that the modal path is for exit in period 4 is confirmed by the second row of Figure 10. In panel (e), we plot the probability that the policymaker exits the forward guidance policy and reverts to the optimal discretion policy. The red line with diamond markers corresponds to the *ex ante* probabilities of reversion. These are computed using information at the beginning of the experiment (i.e. the start of period 1). The modal exit date is date 4, which has the highest *ex ante* probability of occurring. The blue dashed line in panel (e) shows the probabilities of reverting in each period, *conditional on the FG policy still being in effect at that date.* These probabilities rise over time because, as the economy recovers back to steady state, the stimulus associated with continuing to hold the policy rate at the zero bound increases. Indeed, panel (d) plots the 'threshold gaps' $(y_t - \bar{y})$, conditional on the policymaker having not already exited the forward guidance policy. The threshold gaps increase over time consistent with an increasing amount of stimulus associated with delayed exit from the zero bound.

Finally, panel (f) shows the *ex ante* distribution of liftoff dates. The blue dashed line shows the cumulated ex ante liftoff probability (i.e. the probability that liftoff occurs before a particular period). This shows that there is very small probability that the policy rate remains at the ZLB beyond period 10.

In terms of the economics underpinning these results, there are several points to note.

We first note that the modal path for the policy rate is identical to the baseline path corresponding to optimal discretionary policy. Here stimulus is being provided *in expectation* because there is a chance that the policymaker will exit later than the baseline case. A similar effect occurs in our the stochastic model studied in the main text because of the probability that additional stimulus (later exit) will be applied in future 'bad' states. But in the case considered here it is probabilistic exit alone that gives rise to the expected stimulus.

What matters for inflation and output is the *expected* path of the policy rate. Probabilistic exit affects the expected path for rates in a fairly smooth fashion (the expected path is a weighted average of paths in which exit occurs at a discrete date). We can see that the expected path for the policy rate in panel (c) lies below the baseline path for the policy rate, gradually converging to it. This reflects the fact that the sequence of *ex ante* exit probabilities imply a monotonically increasing probability of liftoff (dashed blue line in panel (f)) so that the *ex ante* probability that the policymaker has returned to setting policy according to optimal discretion approaches unity as the horizon increases.

This observation provides intuition for why probabilistic exit helps to pin down the equilibrium under a threshold-based policy. In a discrete time, deterministic model exit must occur at a particular date. The example in Appendix A.1 showed that there may be cases where applying a threshold-based exit criterion *with certainty* is incompatible with exit at date T because that implies insufficient stimulus for the threshold to be breached at that date, but also incompatible with exit at date $T+1$ because that would generate enough stimulus to breach the threshold at an earlier date. Probabilistic exit implies that the expected exit date is a continuous random variable. In the simulation presented in Figure 10, the *ex ante* expected number of periods at the zero bound is 5.09. This provides a substantial amount of stimulus, even though the most likely outcome is for exit to occur in period 4 (as in the baseline case in which policy is set using optimal discretion). The expected duration of the probabilistic forward guidance policy lies in between the dates $T(=5)$ and $T+1(=6)$ that are incompatible with equilibrium if the threshold-based exit criterion is applied with certainty.

A monotonic and convex mapping from the threshold breach to the exit probability creates a feedback between threshold breach and exit probability that supports an equilibrium. Other things equal, a set of beliefs about exit probabilities (that is, $\{p_t\}_{t=1}^{K}$) that imply a longer duration at the zero bound will impart more stimulus and increase the extent to which the threshold variable is expected to breach the threshold. A monotonic and convex mapping to exit probabilities will tend to map a large expected stimulus into higher probabilities of exit, which will reduce the expected stimulus.

Finally, we note that it does not appear that our results are sensitive to the assumption that policy reverts to optimal discretion with certainty by period 31, as the probabilities attached to late exit are extremely low.[23] There is a kink in the profiles for the 'per period' threshold gaps and exit probabilities (panels (d) and (e)) in the final period of the simulation. These are conditional on reversion to optimal discretion not having occurred already. Because reversion to optimal discretion happens with certainty in period 31, the stimulus of setting policy at the ZLB in period 30 is much reduced relative to stimulus in previous periods (because outcomes in previous periods incorporate some probability of remaining at the ZLB to period 30).

[23] Recomputing the equilibrium using $K = 15$ generates results that are almost identical.

B Calibrating the exit probability function

As noted in the main text, the behavior of the exponential distribution approximates a 'trigger' (exit occurs for sure when the threshold is breached) as $\alpha \to 0$. Our goal is to find positive values for α that allow us to approximate a trigger, while retaining the ability to compute an equilibrium in light of the discussion in Section 4.1. This reflects our desire to choose a conservative variant of the exit probability function, f.

To inform our calibration, we use information from the Primary Dealer Survey conducted by the Federal Reserve Bank of New York following the introduction of threshold-based forward guidance by the FOMC in December 2012.[24] Our main focus is the survey conducted in January 2013, following the FOMC's announcement in their statement on 12 December 2012 that they would pursue threshold-based guidance:

> the Committee decided to keep the target range for the federal funds rate at 0 to $\frac{1}{4}$ percent and currently anticipates that this exceptionally low range for the federal funds rate will be appropriate at least as long as the unemployment rate remains above $6\frac{1}{2}$ percent, inflation between one and two years ahead is projected to be no more than a half percentage point above the Committee's 2 percent longer-run goal, and longer-term inflation expectations continue to be well anchored.

The January 2013 Primary Dealer Survey investigated respondents' views about the conditions under which the FOMC would lift off from the zero bound, in light of the forward guidance issued in December 2012.

Question 12a of the survey asked respondents for their *modal* estimates of the joint outcomes for the unemployment rate and headline 12-month PCE inflation rate at the date of the first increase in the Federal funds rate. For the unemployment rate, the median and 75th percentile responses were 6.5% indicating that the majority of respondents viewed the threshold as a trigger. However, the 25th percentile response was 6.25% so that 25% of respondents thought it most likely that liftoff would occur with the unemployment rate 0.25 percentage points or more below the threshold.

Using a simple Okun's law relationship, this implies that 25% of respondents thought it most likely that liftoff would occur when the output gap was more than 0.5 percentage points above the threshold value.[25] Given our calibration of $\alpha = 0.1$, the probability of the policymaker exiting the forward guidance regime when the output gap is 0.5 percentage points above the threshold value is 0.67%. Although it is impossible to directly map from a survey distribution of heterogeneous beliefs about the modal rate of unemployment consistent with liftoff to our f function, our calibration implies that outcomes which a substantial proportion of market participants thought to be the *most* likely are, under our calibration, very unlikely. This suggests that our calibration is closer to a trigger than real world guidance was perceived to be by market participants.

In terms of beliefs about inflation at liftoff, question 12c concerns the conditions under which exit would occur for alternative assumptions about unemployment and the inflation projection. Specifically, respondents were asked at what maximal unemployment rate they thought the FOMC would lift off from the zero bound conditional on alternative assumptions about the outlook for inflation.

Table 5 contains the responses to this question. Most respondents regard the 2.5% inflation forecast threshold as a trigger: liftoff would occur at unemployment rates higher than 6.5% for

[24]The survey results are available here: https://www.newyorkfed.org/markets/primarydealer_survey_questions.html.

[25]Okun's law is a widely used rule of thumb that a one percentage point unemployment gap corresponds to a two per cent output gap.

Table 5: Maximum unemployment rates at liftoff conditional on inflation forecast

Inflation forecast :	2.50%	2.75%	3.00%
Percentile 25	6.50%	6.50%	6.70%
Percentile 50	6.50%	7.00%	7.25%
Percentile 75	6.50%	7.00%	7.80%

inflation forecasts higher than 2.5%. However, 25% of respondents thought that an inflation forecast of 2.75% would be compatible with the unemployment threshold (that is liftoff would occur at an unemployment rate equal to or below the threshold of 6.5%).

Mapping to our model is difficult, not least because in the model the threshold relates to the current inflation rate whereas the relevant variable in the FOMC's guidance was expected inflation. However, we can say that a sizable fraction of respondents (one quarter of them) thought that exceeding the inflation forecast threshold by 0.25 percentage points would be possible while remaining within the FOMC's threshold-based guidance regime. Mapping 0.25 percentages points to a quarterly inflation rate gives around 0.06 percentage points. Setting $\alpha = 0.0125$ in our f function implies that the probability of exceeding the threshold by 0.06 percentage points without triggering exit is 0.8%, which is comparable to the 'tightness' we applied to the output gap threshold.

C Solution method

C.1 Optimal discretion with a zero lower bound

To solve the model described in Section 2 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 initialized 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 zero bound 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.[26] 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 are unlikely to 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 parameterization of the model we use. In the case of the baseline parametrization outlined in Section 5, 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 $1e^{-6}$.[27]

C.2 Threshold-based guidance experiments

The objective is to find policy functions for inflation, $\pi^{FG}(s)$, and the output gap, $y^{FG}(s)$, that satisfy the equilibrium conditions (i.e. the Phillips and IS curves) and are consistent with the exit probabilities of the regime, as defined in Section 4.1. We use a time iteration

[26] First, solve for outcomes on the assumption that the zero bound is not binding in the following way: (i) use the first-order condition for the policymaker if unconstrained ($y_t = -\frac{\kappa}{\lambda}\pi_t$) to substitute the output gap out of the Phillips curve (equation (3)) 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 (4)) 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 to the zero bound $(1-\beta^{-1})$, then the solution (conditional on expectations) has been found and stop. If the interest rate violates the zero bound constraint then: (i) set the interest rate equal to $1-\beta^{-1}$; (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.

[27] 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 interpolation in the approximation of expectations (noting that all the state variables are exogenous and so each possible realization 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).

approach, solving for policy functions conditional on a guess for exit probabilities and updating exit probabilities according to the function f in an iterative fashion.

The structure of the algorithm is as follows, where the subscript $\langle i \rangle$ denotes iteration i:

0. Initialize policy functions $y^{FG}(s)$, $\pi^{FG}(s)$ and the probabilities $p(s)$. Policy functions are initialized using the policy functions under optimal discretion:

$$x^{FG}_{\langle 0 \rangle}(s) \equiv \begin{bmatrix} y^{FG}_{\langle 0 \rangle}(s) \\ \pi^{FG}_{\langle 0 \rangle}(s) \end{bmatrix} = \begin{bmatrix} y^{OD}(s) \\ \pi^{OD}(s) \end{bmatrix}$$

The probabilities are initialized as

$$p_{\langle 0 \rangle}(s) = \delta_0 f\left(x^{FG}_{\langle 0 \rangle}(s) - \bar{x}\right) + (1 - \delta_0)\,\mathbf{1}$$

where $\delta_0 \in (0, 1]$ is a damping factor and $\mathbf{1}$ is the unit vector (i.e., exit with certainty for every state s).

Then for each iteration, $i = 1, \ldots$:

1. Taking $p(s)$ as given, solve equations (12) and (13) using time iteration. This is done in an analogous fashion to optimal discretion.[28] Denoting policy functions based on time iteration j conditional on $p_{\langle i-1 \rangle}(s)$ as $x^{FG}_{\langle j|i-1 \rangle}$, iteration proceeds until $\left\| x^{FG}_{\langle j|i-1 \rangle} - x^{FG}_{\langle j-1|i-1 \rangle} \right\|_\infty < \tau$. The resulting policy functions are denoted $x^{FG}_{\langle *|i-1 \rangle}$.

2. The policy functions from step 1 are used to update the exit probabilities $p_{\langle i \rangle}(s)$. We first compute the probabilities consistent with the latest estimate of the policy functions using the mapping (18):
$$\widetilde{p}(s) = f\left(x^{FG}_{\langle *|i-1 \rangle}(s) - \bar{x}\right) \tag{18}$$
The probabilities are then updated for the next iteration i according to
$$p_{\langle i \rangle}(s) = (1 - \delta)\, p_{\langle i-1 \rangle}(s) + \delta \widetilde{p}(s) \tag{19}$$
where $\delta \in (0, 1]$ controls damping.

3. Check for convergence.

   (a) Convergence is achieved if:

      (i.) The policy functions have converged $\left\| x^{FG}_{\langle *|i \rangle} - x^{FG}_{\langle *|i-1 \rangle} \right\|_\infty < \varepsilon_x$; *and*

      (ii.) The exit probabilities have converged $\left\| \widetilde{p}_{\langle i \rangle}(s) - \widetilde{p}_{\langle i-1 \rangle}(s) \right\|_\infty < \varepsilon_p$.

   (b) If convergence is not achieved return to step 1.

We implement this algorithm using the same 20,000 node state grid and linear interpolation scheme described in Section C.1 and the same quadrature nodes for cost-push and demand shocks. The overall tolerance applied to policy function convergence is the same as for the optimal discretion solution (i.e., $\varepsilon_x = 1 \times 10^{-6}$) and the convergence for exit probabilities is set

[28] For example, in each iteration we solve the IS curve (12) for $y^{FG}(s)$ for each s by computing the terms on the right hand side as follows. The terms $p(s)$, $\left(1 - \beta^{-1}\right)$, $g(s)$ and $r^{OD}(s)$ are known. The expressions $\mathbb{E}_{s'|s} y^{OD}(s')$ and $\mathbb{E}_{s'|s} \pi^{OD}(s')$ are also known (we pre-compute them using the OD policy functions). The expectations $\mathbb{E}_{s'|s} y^{FG}(s')$ and $\mathbb{E}_{s'|s} \pi^{FG}(s')$ can be computed by numerically approximating the integral using the previous guess for the FG policy functions (just as we do in each iteration of the solution for the OD policy functions). The Phillips curve is solved in an analogous manner.

to $\varepsilon_p = 1 \times 10^{-4}$.[29] The 'within iteration' convergence criterion for the time iteration step 1 was set to a relative loose value, $\tau = 1 \times 10^{-1}$, in conjunction with substantial damping of the exit probability updates: $\delta = 0.025$ in most cases.[30]

This design reflects the very strong feedback from the exit probabilities p to the policy functions x, which motivates substantial damping in the updating of exit probabilities in Step 2. Given this damping, the algorithm takes longer to converge (around 10 times longer than the optimal discretion solution). For this reason, the tolerance τ is set loosely to avoid over-refining policy function estimates conditional on exit probabilities that are far away from the equilibrium probabilities. As the policy functions converge, updating of the policy functions and exit probabilities becomes sequential, speeding convergence.[31]

C.3 Calendar-based guidance experiments

Solving for the approximate policy functions that characterize a one-off calendar-based forward guidance policy is relatively straightforward via backward induction. In period K, the final period of the regime, the policy functions can be computed under the assumptions that the policy rate is pegged at the zero bound 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 zero bound 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.

C.4 Optimal commitment

In Section 6 we compare outcomes generated by threshold-based forward guidance with other policies. A natural benchmark is the optimal commitment policy as it delivers the best achievable outcomes. 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:

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

As in the case of optimal discretion, in the presence of an occasionally-binding zero bound constraint, it is not possible to solve for the equilibrium of the economy analytically.[32] Fur-

[29] This was rarely the binding constraint on convergence. The exit probabilities had typically converged to within 1×10^{-5} or less by the time that the policy functions had converged.

[30] More moderate damping was used for the initialization, $\delta_0 = 0.05$. For cases in which the threshold values $\bar{x}$ were further from zero, even stronger damping was required and $\delta = 0.005$ was used.

[31] That is Step 1 converges to the required tolerance, τ in a single iteration.

[32] 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.

thermore, unlike in the case of optimal discretion, the optimal targeting rule that would apply in the absence of the zero bound is invalid even if the zero bound is not binding in the current period, provided that it has bound at some point in the past. This is a direct consequence of the history dependence of policy which must be accounted for when the model is solved. We solve the model using the approach of Adam and Billi (2006).