EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Sandra Daudignon, Oreste Tristani

Monetary policy and the drifting natural rate of interest

Revised September 2024

No 2788 / February 2023

Non-Technical Summary

Empirical research suggests that the long-run natural interest rate – the real interest rate consistent with output at its long-run equilibrium and stable inflation – is not constant over time. It declined over the past few decades, probably reaching levels around zero in the 2010s, and it may now be increasing again. The future level of the long-run natural rate is uncertain.

These empirical results raise questions for the conduct of monetary policy, due to the effective lower bound constraint on nominal interest rates. One question concerns the most appropriate monetary policy response to a reduction (or increase) of the long-run natural interest rate. A broader question regards the implications for monetary policy of the risk that the long-run natural rate may change unpredictably in the future.

This paper provides answers to these questions based on a standard modelling framework modified to account for the possibility of random changes in the long-run natural rate of interest. The paper also takes explicitly into account the effective lower bound constraint on nominal interest rates.

The results of the analysis suggest that monetary policy ought to be over-expansionary, compared to an ideal situation in which the long-run natural rate were constant and the effective lower bound were not a constraint on short-term policy rates. The reason is that the risk of future reductions in the long-run natural rate tends to impart a downward bias on output and inflation expectations, because the central bank is constrained in its ability to provide sufficient monetary accommodation at the effective lower bound. To offset this bias in expectations, in the absence of shocks, the central bank should maintain a negative gap between the real interest rate and the natural rate; and the gap should increase following any reduction in the long-run natural rate. In other words, the paper finds that the neutral rate – i.e., the policy rate consistent with stable inflation and the natural rate at its long-run level – is lower than the long-run natural rate, and increasingly so, the further the long-run natural rate descends towards zero.

The paper also analyses the ability of simple rules to ensure good macroeconomic outcomes. It specifically focuses on price level targeting rules, because they have the advantage of not requiring knowledge of the long-run natural rate. Simple rules rely more heavily on a conventional policy stimulus in the face of the risk of future changes in the long-run policy rate. The neutral rate will thus be at the effective lower bound as soon as the long-run natural rate falls below 1%.

Abstract

Empirical analyses find that the *long-run natural rate*, or the real rate prevailing over a long-run equilibrium where nominal rigidities are absent, is subject to permanent shocks. How should monetary policy react to such shocks? Our paper answers this question in a variant of the new Keynesian model. Because of the zero lower bound (ZLB) on nominal interest rates, the mere possibility of future movements towards zero of the long-run natural rate imparts a downward bias on inflation expectations. To offset this bias, a central bank optimizing under commitment should not only rely on forward guidance at the ZLB, as recommended by the existing literature, but also adopt an expansionary bias away from the ZLB. The *neutral rate*, i.e. the real policy rate consistent with stable inflation in the long-run, should fall more than one-to-one with the long-run natural rate, as the latter approaches zero. This is the case both under optimal commitment policy, and if optimal policy is implemented through a price level targeting rule.
Keywords: Zero lower bound, Optimal monetary policy with commitment, Liquidity trap, New Keynesian model.
JEL Codes: C63, E31, E52.

1 Introduction

The real short-term interest rate that emerges once transitory economic shocks have been left behind, often called the *natural*, or *neutral*, rate of interest, or simply r-star, is a useful, albeit elusive, long-run guidepost for monetary policy. One of the properties that make r-star elusive is that it is not constant over time. For example, in 2018 Fed Chairman Powell talked about "shifting stars" (with reference to both r-star and the natural rate of unemployment u-star) and the difficulty of "guiding policy by the stars in practice".[1] The state of the economy in 2023 was a case in point. Estimates of various empirical notions of r-star placed it at levels around zero at the end of the 2010s. If one had assumed that it would remain constant at this level in the future, one would have worried that an aggressive monetary policy tightening following the 2022-23 inflation outburst may cause a recession and a return of interest rates to the zero lower bound (ZLB).[2] If instead one had believed r-star to be higher, the risk of hitting the ZLB again would not have been a major concern. Either assumption would have been difficult to defend, because r-star is known to be time-varying.

The aim of this paper is to study how monetary policy ought to take into account the observed time-variation in the long-run natural rate. In contrast to the existing theoretical literature, we recognise that the *long-run* natural rate of interest is subject to stochastic shifts. We then study what these shifts imply for the optimal conduct of monetary policy under commitment while taking into account the ZLB constraint on nominal rates. Our key finding is that, in spite of the availability of forward guidance at the ZLB, monetary policy should be characterized by a type of expansionary bias away from the ZLB. Such bias will be larger, the lower the current value of the long-run natural rate.

The starting point of our analysis is the available empirical evidence, which typically finds it appropriate to model the long-run natural rate as an integrated process. The long-run natural rate is instead constant in the standard new Keynesian theory, and determined by model parameters: the rate of time preference and the steady state productivity growth rate. An obvious theoretical option to make the long-run natural rate time-varying would be to assume that the long-run productivity growth rate follows a random walk. Alas, this assumptions is implausible. In the model, productivity growth would reach arbitrarily large, positive and negative, values in

[1] Powell (2018).

[2] The recent euro area experience has shown that the lower bound on nominal interest rates is not zero, but negative due to cash storage costs. In our theoretical model, cash storage costs are ignored, so the lower bound is equal to zero.

finite time, while over the post-WWII period average productivity growth fluctuated within a relatively narrow, positive range. To avert this shortcoming, in our model we do allow for permanent shocks to productivity growth, but *only within finite, upper and lower, boundaries.* Motivated by the historical evidence, we set the boundaries at 3% and 0, respectively. These values are consistent with upper and lower boundaries for the long-run natural rate.

We use this version of the new Keynesian model to ask three main questions. How should optimal monetary policy respond to shocks to the long-run natural rate? How frequently can the zero lower bound be expected to bind under optimal monetary policy, if the long-run natural rate can change over time? Can optimal policy be implemented, at least approximately, through price level targeting rules that have been shown to work well in models with a constant, long-run natural rate?

Before summarizing our results, we need to establish two definitions. We will distinguish between, on the one hand, the "real rate prevailing over a long-run equilibrium in the absence of nominal rigidities" and, on the other hand, the "real policy rate consistent with absence of inflationary or deflationary pressures in the long-run". Following Obstfeld (2023), we will refer to the former notion as long-run *natural* rate, or $\bar{r}$, and to the latter as (long-run) *neutral* rate, or r^*.

Regarding the optimal response of monetary policy to $\bar{r}$ shocks, our results suggest that it should be charaterized by an expansionary bias in (the risky) steady state. More specifically, the neutral rate should fall more than one-to-one with the long-run natural rate. In other words, due to the risk of future ZLB episodes, the central bank should maintain a negative gap between the neutral rate r^* and the long-run natural rate $\bar{r}$; and the gap should increase following any exogenous reduction in $\bar{r}$. This is a novel result in comparison to the previous literature on optimal commitment, which has so far emphasised the central bank's ability to promise higher inflation in the wake of a binding ZLB as a sufficiently strong tool to make any "pre-emptive" easing unnecessary. A degree of pre-emptive easing is known to be a feature of the adjustment path under optimal discretionary policy, that is a situation in which the central bank is unable to make credible promises and is therefore powerless at the ZLB – Adam and Billi (2007) and Nakov (2008).[3] In our model, the pre-emptive easing is necessary *in spite* of the central bank's ability to make credible promises, because permanent, downward shocks to $\bar{r}$ are more pernicious than temporary shocks. Any reduction in $\bar{r}$ makes a binding ZLB (in expectation) permanently more likely, so that the central bank will simply have less scope for fulfilling its promise to engineer more inflation in the future. From a quantitative perspective, we find that the neutral rate r^* should be zero as soon as $\bar{r}$ falls to 75 basis points – that is, a $\bar{r}$ level well

[3] See Nakata and Schmidt (2019) for a proposal on how to reduce the deflationary bias.

above empirical estimates prevailing at the end of the 2010s. This approach would allow the central bank to continue pursuing near price stability: in annualized terms, (unconditional) optimal inflation in the model is essentially zero.

How should monetary policy respond in the short run to permanent $\bar{r}$ shocks? One could conjecture that, as long as the policy rate is not constrained by the ZLB, there should be no transitional dynamics to the new long-run equilibrium: the real rate could be immediately adjusted to offset any inflationary/deflationary pressure caused by the new $\bar{r}$ level. We show that this is not the case under optimal policy. We have already emphasized that the neutral rate will eventually fall more than one-to-one in response to a negative $\bar{r}$ shock. We also demonstrate that the real policy rate ought to adjust gradually after the shock – that is, the real policy rate will be temporarily higher than the neutral rate. This implies that monetary policy will be contractionary along the adjustment process, and that a permanent reduction in $\bar{r}$ will be followed by a temporary disinflation. These results are consistent with the general principle of history-dependence of optimal policy under commitment.

The answer to our second question is a direct implication of the properties of optimal policy. Since the neutral rate is lower than $\bar{r}$, it will also reach the zero level earlier than $\bar{r}$, as the latter falls. We show that, for other parameter values identical to those in previous studies, the ZLB incidence under optimal policy is equal to one third once we allow $\bar{r}$ to fluctuate between 0 and 3%. We also show that one would obtain wildly different results regarding the ZLB frequency in a model with constant $\bar{r}$, if one solved the model for different calibrations of $\bar{r}$. Our model delivers results that are robust to uncertainty as to the future $\bar{r}$.

We finally show that variants of the price level targeting rule put forward in Eggertsson and Woodford (2003) continue delivering economic outcomes relatively close to those obtained under optimal commitment. In our model this is the case especially if the price level target is not fixed as in Eggertsson and Woodford (2003), but includes a small upward drift – optimally equal to 10 basis points.

Our paper contributes to the literature on the consequences of the ZLB for optimal monetary policy – see Krugman (1998), Eggertsson and Woodford (2003), Jung et al. (2005), Adam and Billi (2006), Nakov (2008), Levin et al. (2010), Billi (2011). All these papers assume a constant $\bar{r}$. They also rely on calibrations consistent with a relatively high level of the nominal interest rate (typically 3.5%). The promise to maintain policy rates low for longer after adverse shocks is therefore a sufficiently powerful tool of macroeconomic stabilisation when the ZLB constraint binds. More recent contributions to this literature have studied calibrations with different, steady-state values of the natural rate. Billi et al. (2023) studies optimal policy in a new Keynesian economy where the steady state natural rate is negative, hence optimal inflation must be positive to ensure existence of an equilibrium. It demonstrates that monetary policy can stabilise output and inflation around their steady states, even if the policy rate remains almost always at zero.

Other contributions to the literature on the ZLB have focused on the performance of simple policy rules. Many papers have analysed simple *instrument* rules – see amongst others Reifschneider and Williams (2000), Mertens and Williams (2019), Bianchi et al. (2021), Kiley and Roberts (2017). More recently, Andrade et al. (2019) and Andrade et al. (2021) adopt a richer and more realistic model specification and look for the inflation rate that should optimally be assigned to a central bank as the target of a Taylor rule. The papers find that the target should increase almost one-to-one with the steady-state natural rate, once the latter falls below 5% (in annualised terms). Fernández-Villaverde et al. (2021) studies the interaction of the ZLB with household inequality. In contrast to all these papers, we allow for time-variation in $\bar{r}$. We additionally focus on the performance of a simple *target* rule, notably price level targeting – see Eggertsson and Woodford (2003), Vestin (2006). Price level targeting has the advantage of not requiring explicit knowledge of the natural rate of interest. As demonstrated by Eggertsson and Woodford (2003), constant price level targeting is particularly effective against the ZLB in a model where the long-run natural rate is constant, because it induces positive inflation expectations after a deflationary period.

The paper is organised as follows. Section 2 briefly summarises the empirical evidence on the dynamics of the natural rate of interest. The model is presented in Section 3, where we also state the optimal policy problem. Section 4 describes the solution method and key features of our calibration. Our main results on optimal monetary policy under commitment and on price level targeting rules are illustrated in Sections 5 and 6, respectively. Section 7 offers some concluding remarks.

2 An overview of the available empirical evidence

This section briefly summarises the empirical evidence that motivates our theoretical model. Since different papers rely on different notions of "natural rate of interest", we start by establishing a few definitions.

2.1 A few definitions

The theoretical, new Keynesian literature defines the natural rate as the short-term real interest level that would be observed in the absence of nominal rigidities, or r^n (Woodford, 2003).[4] It is a summary statistic that captures many possible exogenous determinants of economic fluctuations, including preference and technology shocks, and is therefore subject to *transitory*, high-frequency variations (for example, Edge et al. (2008), Justiniano and Primiceri (2010)). The theoretical literature has so far assumed r^n to be constant in the long-run, i.e. in the steady state of the models.

Most of the empirical literature in reduced form has instead focused on notions with a long-run equilibrium flavor. By construction, these magnitudes will not be affected at all by transitory shocks, but only by shocks with permanent effects. For example, Laubach and Williams (2003) explicitly focuses on a horizon prevailing "once transitory shocks [...] have abated". Other analyses focus on the value of the real short-term interest rate expected to prevail in the distant future (for example Hamilton et al. (2016)). In the rest of our paper we will denote the empirical notion, or long-run natural rate, as $\bar{r}$. As stated above, this will be the "real rate prevailing over a long-run equilibrium in the absence of nominal rigidities". It will also represent the long-run value of r^n. When $\bar{r}$ is time-varying, r^n will thus be affected by both transitory and permanent shocks.

The long-run natural rate $\bar{r}$ will most of the time coincide with the neutral rate r^*, i.e. the "real policy rate consistent with absence of inflationary or deflationary pressures in the long-run". However, the two definitions are not always and necessarily identical. One of the main results of our analysis will indeed be to show that r^* can be different from $\bar{r}$ due to the ZLB constraint.

2.2 Empirical evidence of long-run trends in real rates

The very low inflation rates observed over the 2010s, while policy interest rates remained close to zero, were considered as suggestive evidence of a very low level of the natural rate. This generated renewed interest in measuring $\bar{r}$ empirically – see, for example, Hamilton et al. (2016), Holston et al. (2017), Fiorentini et al. (2018). In these econometric studies $\bar{r}$ is typically modelled as a random walk process, since it does not appear to converge to a constant value over time. By and large, these papers find that $\bar{r}$ has fallen in recent decades, but there is considerable uncertainty as to its future evolution. For example, Holston, Laubach and Williams (2017) finds that, in 2016, $\bar{r}$ was between 0 and 1% in the United States and possibly slightly negative in the euro area. Using an alternative approach, Fiorentini et al. (2018)

[4]This notion can only be computed in the context of a structural model, since it requires the evaluation of an economic equilibrium in which prices and wages are counterfactually assumed to be perfectly flexible.

estimates that the long-run natural rate in 2016 was slightly above 1% in the U.S. and as low as -1% in the euro area. Hamilton et al. (2016) forecasts the real rate in the U.S. to asymptote to a value slightly lower than 0.5% by 2021. Some very recent estimates for the U.S. (reported in Obstfeld (2023)) suggest that $\bar{r}$ may be again in safely positive territory, namely in a range between 0.6% and close to 2%.[5] Platzer and Peruffo (2022) forecasts it to reach a trough of 0.38% by 2030 and then rise again to 1% in the long run.

Empirical models disagree on the exact determinants of the time-variation in $\bar{r}$. The approach pioneered in Laubach and Williams (2003) emphasises time-variation in trend productivity growth. Hamilton et al. (2016) argues that the relationship between the long-run natural rate and trend GDP growth is tenuous. Other papers emphasize the demographic transition that is ongoing in many Western economies (Carvalho et al. (2016), Gagnon et al. (2016), Aksoy et al. (2019)), an increase in the required premium for safety and liquidity (e.g. Caballero and Farhi, 2017, Krishnamurthy and Vissing-Jorgensen, 2012, Del Negro et al. (2017)) and rising income inequality (Platzer and Peruffo (2022)).

We draw one main lesson from the review of the empirical literature. There is overwhelming evidence of time variation in $\bar{r}$, so that its future level is uncertain. This will be the key distinguishing feature of our model. Regarding the determinants of $\bar{r}$, their choice does not affect our optimal results as long as they are independent of monetary policy. We will follow Laubach and Williams (2006) and attribute its time variation to permanent shocks to productivity growth. This assumption has the advantage of being more easily calibrated based on TFP-data from Fernald (2014).

3 The model

The model we employ in our analysis is relatively standard. In this section we briefly summarize it to highlight the few modifications that we introduce in order to allow for variations in the long-run natural rate of interest. The model is described in more detail in appendix A.

As in the standard new Keynesian model, households consume a composite good C_t, which is the Dixit-Stiglitz aggregate of a continuuum of differentiated goods. Differently from the usual formulation of the model, we assume that bond holdings provide utility benefits on top of their pecuniary return. This assumption allows the model to produce possibly very low values of $\bar{r}$ without postulating unrealistically low, or even negative, growth rates of productivity in the long run.

[5] See also IMF (2022), p.21.

The representative household j demands an amount $C_{j,t}$ of the composite good in order to maximise intertemporal utility

$$\mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t U_t \left(C_{j,t}, H_{j,k,t}, \frac{M_{j,t}}{P_t} \right)$$

subject to a sequence of usual budget constraints. In the above equation, $H_{j,k,t}$ are hours worked in all firms in the economy $k \in [0,1]$ and $M_{j,t}$ are nominal non-state contingent bonds issued by the government and yielding a gross nominal return I_t^m. The assumption of bonds-in-the-utility has also been adopted in Fisher (2015) and Krishnamurthy and Vissing-Jorgensen (2012) – see also Sidrauski (1967). We will specifically follow Michaillat and Saez (2021) and postulate that households derive utility from their relative real bond holdings, so that temporary utility is

$$U_{j,t} = \bar{C}_t \cdot \left[\log C_{j,t} + \upsilon \left(\frac{M_{j,t}}{P_t} - \frac{M_t}{P_t} \right) - \frac{\gamma}{1+\upsilon} \int_0^1 H_{j,k,t}^{1+\upsilon} \mathrm{d}k \right]$$

where the function $\upsilon(\cdot)$ is increasing and concave. Since preference shocks are often used as triggers of ZLB episodes in the new Keynesian literature, we also introduce this type of shock, $\bar{C}_t$, in our utility specification. We demonstrate below that this shock is observationally equivalent to a temporary productivity growth shock. In our analysis we will assume that $\bar{C}_t = \Delta_t \bar{C}_{t-1}$, for $t > 1$ and $\bar{C}_0 = 1$, where $\delta_t = log(\Delta_t)$ will follow a stationary autoregressive process such that

$$\delta_t = \rho_\delta \delta_{t-1} + \sigma_\delta \varepsilon_{\delta,t}, \quad \varepsilon_{\delta,t} \sim \mathcal{N}(0,1).$$

where $\mathcal{N}(.)$ denotes the normal distribution.

On the production side, there is a continuum of firms, indexed by $k \in [0,1]$, producing differentiated goods under monopolistic competition and sticky prices, and using the production technology $Y_{kt} = \bar{A}_t (H_{kt})^\phi$, where $\bar{A}_t$ denotes economy-wide productivity. We allow for permanent shocks to the gross productivity growth rate $\Xi_t \equiv \frac{\bar{A}_t}{\bar{A}_{t-1}}$. In log terms, $\xi_t \equiv \log(\Xi_t)$ is assumed to follow a bounded unit root process between large but finite boundaries ξ^H and ξ^L

$$\xi_t = \xi_{t-1} + \sigma_\psi \varepsilon_t^\psi, \quad \varepsilon_t^\psi \sim \mathcal{TN}(0, 1, \tfrac{\xi^L - \xi_{t-1}}{\sigma_\psi}, \tfrac{\xi^H - \xi_{t-1}}{\sigma_\psi}),$$

where $\mathcal{TN}(.)$ denotes the truncated standard normal distribution.

The demand side of the model can be summarized by a variant of the standard Euler equation, which can be written as

$$\frac{1}{I_t^m} = \mathrm{E}_t \Big[\beta \Delta_{t+1} \frac{C_t}{C_{t+1}} \frac{1}{\Pi_{t+1}} \frac{1}{1 - \Delta_t^m} \Big]$$

where I_t is the gross interest rate on a complete portfolio of state-contingent bonds, the spread $\Delta_t^m \equiv \frac{I_t - I_t^m}{I_t}$ is in equilibrium given by the marginal rate of substitution between wealth (in the form of real bond holdings) and consumption, $\Delta_t^m = C_t/\bar{A}_t \upsilon' (0)$. Here we have not yet introduced any detrending and the definition remains correct, given that Δ_t^m does not change if one defines it in terms of detrended interest rates . Note that the Euler equation would become standard in the $\Delta_t^m = 0$ case in which $I_t^m = I_t$.

The Euler equation can be linearized around the non-stochastic steady state with zero inflation.[6] Due to productivity growth, it is well known that output and consumption need to be detrended by the *level* of productivity, $\bar{A}_t$, before linearization. If we denote detrended consumption as $\tilde{C}_t \equiv C_t/\bar{A}_t$, the Euler equation can be rewritten as

$$\frac{1}{I_t^m/\Xi_t} = \mathrm{E}_t\left[\beta\Delta_{t+1}\frac{\tilde{C}_t}{\tilde{C}_{t+1}}\frac{\Xi_t}{\Xi_{t+1}}\frac{1}{\Pi_{t+1}}\frac{1}{1-\Delta_t^m}\right].$$

This expression shows that the nominal interest rate on government bonds, I_t^m, will inherit the stochastic trend in the *growth rate* of productivity, Ξ_t. We therefore define a stationary variable $\breve{I}_t^m \equiv I_t^m/\Xi_t$ and end up with an equation where all variables are stationary in the deterministic steady state.

Note that we could alternatively have assumed no preference shocks in the utility function. The Euler equation would have simplified to $\frac{1}{I_t^m} = \mathrm{E}_t\left[\beta\frac{C_t}{C_{t+1}}\frac{1}{\Pi_{t+1}}\frac{1}{1-\Delta_t^m}\right]$. We could then have assumed that productivity growth is driven not only by permanent shocks Ξ_t, but also by transitory shocks Δ_t such that $\frac{\bar{A}_t}{\bar{A}_{t-1}} = \frac{\Xi_t}{\Delta_t}$. In detrended terms, the Euler equation would have been identical.

The Euler equation could be directly linearized in terms of detrended consumption or, using the aggregate resource constraint $C_t = Y_t$, in terms of detrended output. It is however customary to express the linearized Euler equation in terms of the output *gap*, that is output in deviations from the output level which would be observed in the natural equilibrium. This is the equilibrium which would prevail in the absence of price rigidities and whose equilibrium real rate, which is independent of monetary policy, is the natural rate.

In terms of the output gap x_t, the Euler equation in our model could be written as

$$x_t = (1-\Delta^m)\left[\mathrm{E}_t x_{t+1} - (\breve{\imath}_t^m - \mathrm{E}_t\pi_{t+1} - \breve{r}_t^n)\right] \tag{1}$$

[6] As is common in new Keynesian analyses of monetary policy at the ZLB, we linearize the model equations so that the only source of nonlinearity is represented by the ZLB itself. An alternative option would be to solve the fully nonlinear model. Our approach maximises the comparability of our results with the rest of the literature, as well as being computationally less demanding.

where $x_t \equiv \tilde{c}_t - \tilde{c}_t^n$, i.e. the logarithm of detrended consumption/output in deviation from its natural level, and $\breve{r}_t^n = \mathrm{E}_t\tilde{c}_{t+1}^n - \tilde{c}_t^n/(1-\Delta^m) - \mathrm{E}_t\delta_{t+1} + \sigma_\psi \mathrm{E}_t\varepsilon_{t+1}^\psi$ is the detrended short-run natural rate (in deviation from its steady state value $-\ln\beta + \ln(1-\Delta^m)$). Note that equation (1) is very similar to the standard linearised Euler equation of the new Keynesian model except for the "discount factor" $(1-\Delta^m)$. Δ^m is the liquidity spread in the non-stochastic steady state, so that $1-\Delta^m$ is a coefficient smaller than 1.[7]

The detrended natural rate in deviation from its steady state can be solved out explicitly in terms of the exogenous states:

$$\breve{r}_t^n = \bar{\delta}_t + \sigma_\psi \mathrm{E}_t\varepsilon_{t+1}^\psi \tag{2}$$

where $\bar{\delta}_t \equiv -\rho_\delta\delta_t$. In the rest of the paper, we do not use δ_t, but work directly with the derived process $\bar{\delta}_t = \rho_{\breve{r}^n}\bar{\delta}_{t-1} + \sigma_{\breve{r}^n}\varepsilon_t^{\bar{\delta}}$. Thus $\breve{r}_t^n$ is entirely driven by preference shocks but, as highlighted above, equation (2) would be identical if δ_t denoted transitory technology shocks.

We define the long-run natural rate $\bar{r}_t$ as

$$\bar{r}_t = -\ln\beta + \ln(1-\Delta^m) + \xi_t, \tag{3}$$

which varies over time with the long-run rate of productivity growth ξ_t.

In contrast to the standard new Keynesian model, in our framework the natural rate of interest is subject to both permanent and temporary shocks. Shifts in productivity growth, ξ_t, are permanent and cause variations in $\bar{r}_t$. For given $\bar{r}_t$, persistent, but stationary fluctuations will be induced by shocks $\breve{r}_t^n$.

Assuming that firms are subject to Calvo-style price rigidities, the supply side of the model, in linearized form, can be summarised through a standard Phillips curve

$$\pi_t = \kappa x_t + \beta \mathrm{E}_t\pi_{t+1} \tag{4}$$

provided that the output gap x_t is defined as above in terms of detrended output in deviation from detrended natural output.

[7]Michaillat and Saez (2021) also emphasizes that that the bonds-in-utility specification leads to a form of discounting in the linearized Euler equation of the model. Michaillat and Saez (2021) goes on to demonstrate that this specification solves a number of anomalies of the new Keynesian model at the ZLB, including the so-called *forward guidance puzzle* (Giannoni et al. 2015). While the Euler equation features discounting also in our version of the model, optimal policy results would remain practically unchanged if we dropped the bonds-in-the-utility assumption. This is the case, on the one hand, because optimal policy can react endogenously to changes in future output and inflation induced by interest rate promises, and, on the other hand, because the Euler equation discount factor in our calibration is very close to 1. As stated at the beginning of the section, we adopt the bonds-in-utility assumption because it leads to a plausible calibration of trend productivity growth.

Finally, note that, in deviation from the non-stochastic steady state, the ZLB constraint on nominal interest rates will be written as

$$\breve{i}_t^m \geq -\bar{r}_t. \tag{5}$$

In other words, feasible deviations of the policy rate from its non-stochastic steady state have a time-varying lower bound that depends on the prevailing $\bar{r}$. The lower bound will be higher – thus feasible, negative deviations of the policy rate from its non-stochastic steady state will be smaller – the closer to zero $\bar{r}$.

To summarize, the constraints of the optimal policy problems are equations (1), (4) and inequality (5). These constraints are characterized by two differences from the standard new Keynesian model. First, the Euler equation (1) includes a form of discounting. Second, the lower bound constraint in equation (5) is effectively time-varying, rather than constant.[8]

Appendix A shows that, up to a second order approximation, household temporary utility can be written as in the stationary case as $U_t^{CB} = -\pi_t^2 - \lambda x_t^2$ for $\lambda = \kappa/\theta$.

We derive and analyze optimal monetary policy under commitment. This has the benefit of providing an ideal, normative benchmark, since the central bank is able to exploit in full the gains from credibility. In reality, these gains are harder to reap. Since policy with commitment is time-inconsistent, it is not clear that central bank promises would always be credible. The literature has therefore considered also other notions of optimal policy, from the opposite extreme of absence of commitment, or discretion (for example Clarida et al. (1999)), to the intermediate cases of loose commitment (for example Debortoli et al. (2014)).

Optimal policy under commitment requires

$$\lambda_{x,t} = -2\lambda x_t + \beta^{-1}\left(1-\Delta^m\right)\lambda_{x,t-1} + \kappa\lambda_{p,t} \tag{6}$$

$$\lambda_{p,t} = -2\pi_t + \beta^{-1}\left(1-\Delta^m\right)\lambda_{x,t-1} + \lambda_{p,t-1} \tag{7}$$

plus $\lambda_{x,t} = 0$ when the nominal rate is unconstrained, $\breve{i}_t^m > -\bar{r}_t$, and $\lambda_{x,t} > 0$ when the nominal rate is at the ZLB, $\breve{i}_t^m = -\bar{r}_t$. In the above equations $\lambda_{x,t}$ is the lagrange multiplier associated to equation (1), which is proportional to the multiplier on the ZLB constraint (this explains why the complementary slackness condition is expressed in terms of $\lambda_{x,t}$); $\lambda_{p,t}$ is the multiplier associated to the Phillips curve.

The equilibrium is characterized by the solution of the system given by equations (1), (4), (5), (6) and (7).

[8] Masolo and Winant (2019) solves the new Keynesian model under a Taylor rule allowing for a lower bound that can vary stochastically between two discrete values.

4 Solution and calibration

We solve the model using projection methods. The details on the method are provided in appendix B.1 and C. The main difficulty that we face compared to standard models is that expectation terms do not take the form of unbounded integrals, as shocks to the rate of growth of productivity are bounded. These boundaries induce some reflecting behaviour. Once trend productivity growth reaches its lower boundary, it can only increase. Conversely, it can only fall once it reaches its upper boundary. We take this into account by combining both Gauss-Hermite quadrature and Gauss Legendre quadrature when computing expectations.[9]

The values of key parameters are reported in Table 1.

Table 1: Calibration of structural parameters

Parameter	Definition	Value	Source
Model specific			
σ_ψ	Standard deviation of shocks to ξ	0.108 (%, annualised)	Fiorentini et al. (2018)
$[\xi^L, \xi^H]$	Boundaries of ξ	[0,3] (%, annualised)	Fernald (2014)
Δ^m	Convenience yield	0.73 (%, annualised)	Del Negro et al. (2017)
β	Subjective discount factor	0.9982	Lindé et al. (2016)
Standard			
$\rho_{\breve{r}}$	Autoregressive parameter	0.8	
$\sigma_{\breve{r}}$	Standard deviation of shocks to $\breve{r}_t^n$	1.176 (%, annualised)	
α	Share of firms keeping prices fixed	0.66	
ω	Elasticity of firms' marginal cost	0.48	
θ	Price elasticity of demand	7.66	
κ	Slope of the Phillips curve	0.056	Implied
λ	Weight on output in the loss function	0.0073	Implied

We set the boundaries, ξ_L and ξ_H at 0 and 3%, respectively. These values are broadly consistent both with the long time-series for utilization-adjusted TFP growth

[9]We have not experienced numerical convergence problems in our analysis. We believe this to be due to our focus on optimal monetary policy with commitment. To the best of our knowledge, all of the studies documenting multiple solutions or non-existence of a solution at the ZLB rely on simple monetary policy rules (e.g., Davig and Leeper 2007, Richter and Throckmorton 2015). In a different setting, Roulleau-Pasdeloup (2020) shows that the solution is (locally) unique provided that the degree of commitment is large enough.

in the U.S. (see Fernald (2014)) and with the results for the euro area in Holston et al. (2017). More specifically, we take 20-year moving averages of the yearly data in Fernald (2014) to capture the low-frequency component of TFP growth. Starting from around 2.2% in the late 1960s, the 20-year moving average of utilization-adjusted TFP growth undergoes a slow, but persistent decline to 0.5% in the 1990s, before increasing again in the early 2000s. Over the whole 2010s, average TFP growth remains stable around the 1% mark.[10]

As shown in expression (3), the long-run natural rate will also be affected by the parameters Δ^m and β. For Δ^m, we follow Krishnamurthy and Vissing-Jorgensen (2012) and Del Negro et al. (2017) which finds that the convenience yield amounts to 0.73% on average (in annualised terms). In terms of our model specification, this implies $\Delta^m = 0.0018$. We finally set $\beta = 0.9982$, a value found in recently estimated models – see for example Lindé et al. (2016).

Figure 1: Unconditional distribution of the natural rate

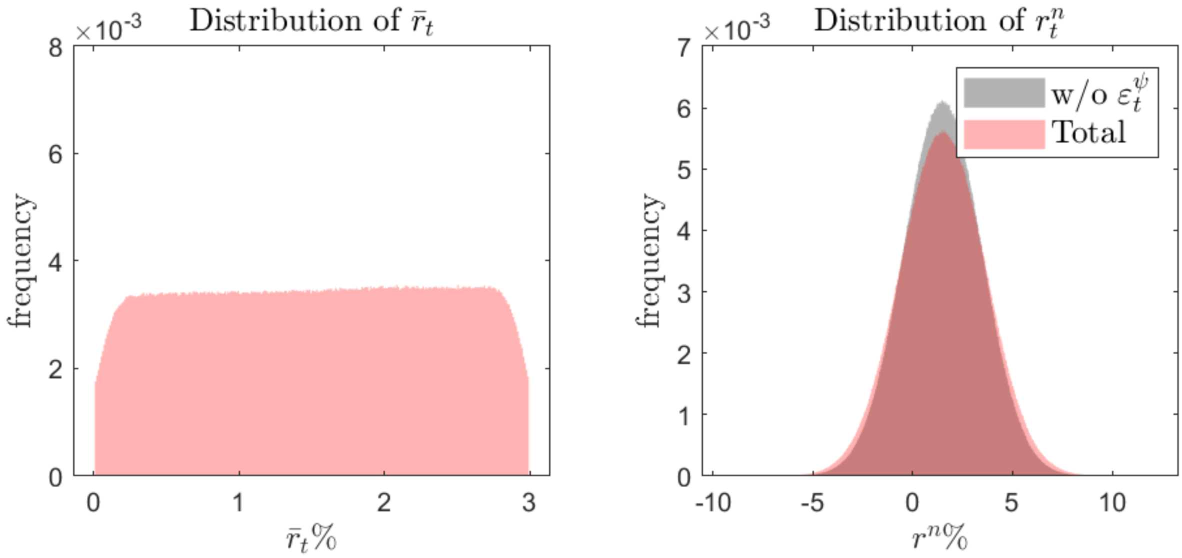

Note: both the natural rate r_t^n and the long-run natural rate $\bar{r}_t$ are expressed in annualised terms.

We finally need to calibrate the standard deviations of permanent productivity

[10]Based on an unobserved component model, Holston et al. (2017) estimates that the euro area productivity trend growth rate was about 3% in the 70s and declined to 1% over the period until 2015.

shocks, σ_ψ, which drive fluctuations in the long-run natural rate. We calibrate σ_ψ based on the estimates in Fiorentini et al. (2018), that is based on historical data at annual frequency over the period 1891-2016 for a set of 17 advanced economies. The conditional standard deviation σ_ψ takes an annualised value of 0.1%, which is about 10 times smaller than the conditional standard deviation of the stationary component $\sigma_{\breve{r}^n}$ (equal to 1.18% in annualised terms following Adam and Billi (2006)). This suggests that temporary shocks play a dominant role on the distribution of the natural rate.

The unconditional distribution of $\bar{r}_t$ consistent with our calibration is shown in the left-hand side panel of figure 1. The distribution is uniform over most of its support, and it has somewhat lower mass close to its boundaries, zero and 3%, that induce some reflecting behavior in $\bar{r}_t$.

The right hand side panel in figure 1 shows the ergodic distribution of the natural rate r_t^n (where by definition $r_t^n = \bar{r}_t + \breve{r}_t^n$). Transitory shocks on their own induce a normal unconditional distribution for the natural rate. The additional source of variability induced by the long-run natural rate is relatively small. It reduces the mass around the unconditional mean, but it does not dramatically change the overall shape of the distribution.

For all other parameter values, we simply follow Adam & Billi (2006), which in turn draws on Woodford (2003).[11]

5 Optimal policy

The existing literature based on models with constant $\bar{r}$ has already highlighted key features of optimal commitment policy. While the ZLB limits the central bank's ability to cut policy rates in reaction to large disinflationary shocks, the ability to commit gives monetary policy the option to respond to such shocks by credibly promising "to be irresponsible" (Krugman (1998)). When the shocks are transitory, the central bank can pledge to create excess inflation in the future through forward guidance, i.e. promise to keep interest rates low for longer than necessary, once the shocks have abated. The pledge will be reflected in private-sector expectations and feed back to higher current inflation – or lower current disinflation. In turn, higher inflation will support economic activity through the ensuing reduction of the real rate.

Forward guidance in response to transitory, disinflationary shocks is also a feature

[11] The main results in Adam and Billi (2006) are based on a high interest rate elasticity of output. Our calibration corresponds to that in section 6.2 of that paper.

of our model. We do not discuss these shocks in detail, because their effects are qualitatively comparable to those observed in a model with constant $\bar{r}$.[12]

We focus instead on the impact of permanent $\bar{r}_t$ shocks, because they produce a novel feature of optimal commitment policy. Clearly, as $\bar{r}_t$ falls closer to zero, the likelihood that policy rates hit the ZLB will increase. In turn, a higher likelihood of hitting the ZLB constraint will also strengthen the deflationary bias in private-sector expectations. Intuitively, the bias can become so large when $\bar{r}_t$ is close to zero, that forward guidance may be unable to offset it, even if it extends far into the future. If this is the case, a type of pre-emptive easing becomes necessary: even in the absence of shocks, it becomes optimal to keep the neutral rate r_t^* below $\bar{r}_t$ so as to produce a positive inflationary pressure that can offset the strong deflation bias. This type of pre-emptive easing has been previously found to be desirable for a central bank acting under discretion, but it was hithertho unexplored as a feature of optimal commitment policy. We will illustrate and quantify this feature in the realistic case in which $\bar{r}_t$ can randomly evolve between zero and 3%.

Before delving into the conditional dynamics of the equilibrium in response to shocks, we briefly summarize the unconditional properties of the model under optimal policy. We show in Table 2 the unconditional means of all endogenous variables. Consistently with the distribution displayed in Figure 1, the unconditional mean of the natural rate of interest is approximately 1.5%. Both inflation and the nominal interest rate are only about 5 basis points (annualised) higher than in the non-stochastic steady state, on average. In spite of the possibility that $\bar{r}_t$ reaches near-zero values due to permanent shocks, optimal inflation is essentially zero as in the standard new Keynesian model. In the rest of this section, we discuss the policy feature that makes this possible.

Table 2: Unconditional means

r^n	x	π	i^m	RR spread
1,51	0,00	0,05	1,57	0,00

Note: the output gap, x is expressed in percentage points; all other variables in annualised percentage points; "RR spread" denotes the difference between the real interest rate $i^m - \pi$ and the natural rate r^n.

[12]Impulse responses to transitory shocks in our model are available in the online appendix.

5.1 The long-run natural rate $\bar{r}_t$ and the neutral rate r_t^*

We analyze the effects of permanent shocks to $\bar{r}$ in terms of the risky steady state. For each variable, this is the (theoretical) point that would be reached once the effects of all current shocks have abated and if no more shocks happened in the future, but agents in the model continued factoring in the possibility of future shocks – see Coeurdacier et al. (2011). This is different from the non-stochastic steady state, which describes a situation in which no shocks can ever occur and agents in the model are aware that this is the case.

Figure 2 displays the evolution of the economy after an illustrative sequence of negative, permanent shocks to $\bar{r}_t$ which take it from 1.5% to 0.5%. No other shock is assumed to occur in this simulation. At the beginning of the simulation, all endogenous variables are at their risky steady state consistent with the initial $\bar{r}_t$. Following each permanent shock to $\bar{r}_t$, all endogenous variables reach a new risky steady state (after an adjustment period discussed in Section 5.2).

The key result that we wish to emphasize with this figure is that the neutral rate r_t^* falls more than one-to-one with the long-run natural rate $\bar{r}_t$, after a negative permanent shock. This can be read from the bottom right panel in the figure, which shows the difference between the real policy rate and the long-run natural rate. The neutral rate would be identical to the long-run natural rate in a version of the model in which the ZLB constraint were ignored. Compared to this benchmark, the ZLB induces optimal policy to adopt an expansionary bias at the risky steady state, i.e. after the adjustment process to the shock has been completed.

Figure 2 shows that a small expansionary bias is already present when $\bar{r}_t = 1.5\%$. At this level of $\bar{r}_t$, the neutral rate is $r_t^* = 1.3\%$. The spread between the two rates is negative and equal to -0.2%. The expansionary bias increases as the long-run natural rate falls to $\bar{r}_t = 1.18\%$ and $\bar{r}_t = 0.85\%$. The neutral rate declines to $r_t^* = 0.8\%$ and $r_t^* = 0.2\%$, respectively. Correspondingly, the spread between the two variables increases (in absolute value) to -0.4% and -0.6%. However, once the long-run natural rate falls further to $\bar{r}_t = 0.53\%$, the spread goes a bit down (in absolute value) to -0.54%.

Figure 2: Simulations of a sequence of permanent shocks under optimal policy

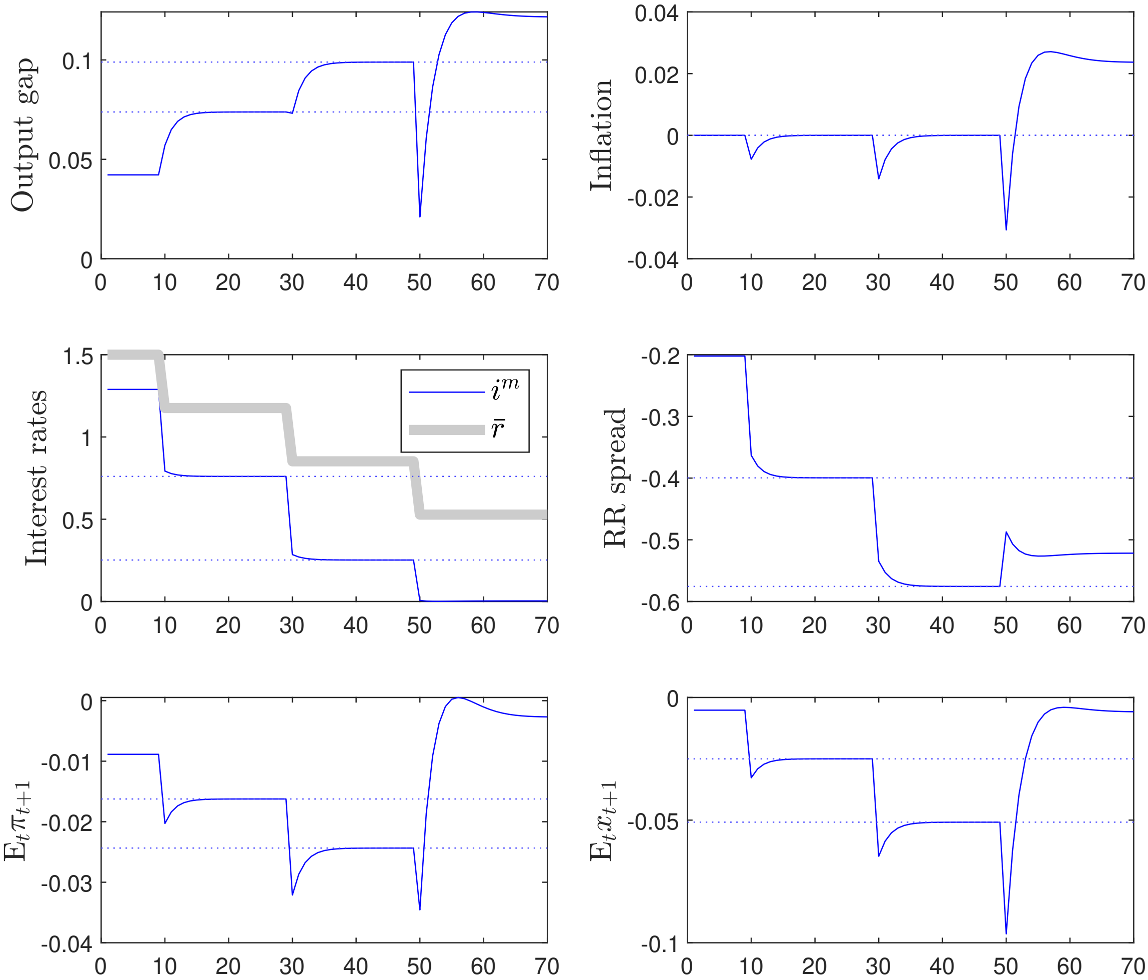

Note: the figure shows the results of the simulation of a sequence of three negative permanent shocks to the long run natural rate $\bar{r}_t$*, denoted by the thick, grey solid line in the middle left panel. Starting from a value equal to 1.5% at time 0, it falls successively by 0.324 p.p., i.e., three times the standard deviation of shocks* σ_ψ*, at time 10, 30 and 50. At the end of this sequence of shocks,* $\bar{r}_t = 0.528\%$*. The sequence of permanent shocks occurs in isolation, while all other shocks are set to zero. The output gap is expressed in percentage points; all other variables in annualised percentage points. "RR spread" denotes the difference between the real interest rate* $i_t^m - E_t\pi_{t+1}$ *and the natural rate* $\bar{r}_t$

To understand this result, note that the output gap is positive in all risky steady states, and that its values increase, as the long-run natural rate falls towards zero. Nevertheless, as long as the long-run natural rate does not fall below $\bar{r}_t = 0.85\%$, the risky steady state of inflation is zero. This suggests that the negative spread between $\bar{r}_t$ and r^*_t is necessary to stimulate the economy – i.e. produce a positive output gap – and thus offset the deflationary bias that characterizes inflation (and output gap) expectations. Inflation can then be stabilized at zero.

Obviously, the spread between $\bar{r}_t$ and r^*_t can only be set at the desired level as long as there is room for cutting the policy rate. Results change when $\bar{r}_t$ falls towards 0.5%. At that point, r^*_t reaches zero, the ZLB constraint binds, and monetary policy can no longer induce a sufficient reduction in the policy rate. The only remaining option for the central bank to provide monetary accommodation is to promise an overly expansionary reaction to future shocks. Inflation expectations will move upwards and contribute to a reduction in the real rate. However, this approach has side effects. Since inflation expectations rise, inflation can no longer be stabilized at zero, but it will become positive.

The expansionary bias that we have documented – as long as there is room for the policy rate to be cut – was hithertho unexplored as a feature of optimal commitment policy. In this case, the literature has emphasised forward guidance as the key policy recommendation at the ZLB. Given the focus of the literature on transitory shocks, there is always the option of keeping policy rates low for longer – that is, low once the natural rate has returned to its (higher) starting level. Permanent, downward shocks to $\bar{r}_t$ are more pernicious than transitory ones, because they are (in expectations) never followed by a return of $\bar{r}_t$ to its (higher) starting level. Hence, while making a binding ZLB more likely, downward permanent shocks to $\bar{r}_t$ reduce the effectiveness of forward guidance.

The expansionary bias that characterizes optimal commitment policy is akin to a pre-emptive monetary policy easing – that is, an easing in the face of the mere risk of future adverse shocks. A type of pre-emptive easing is known to characterize optimal discretionary policy – Adam and Billi (2007) and Nakov (2008). In this case, a central bank is unable to make credible promises at the ZLB, hence it tries to ease policy pre-emptively. In contrast to these models, we analyze optimal policy under commitment, hence the central bank does have the credibility to make promises at the ZLB. We find that it should act pre-emptively in spite of its credibility.

A form of pre-emptive easing under commitment is also highlighted in Eggertsson and Woodford (2003) as the optimal response to an *announced*, future deflationary shock. The paper analyzes this type of shock to discuss the recommendation to "keep the powder dry", that is the idea to avoid cutting policy rates all the way to zero too quickly in order to save ammunition for future emergencies. Contrary to this idea, Eggertsson and Woodford (2003) finds that policy should react pre-emptively to the shock, i.e. cut policy rates upon announcement, and not when the shock actually takes place. In contrast Eggertsson and Woodford (2003), we find that pre-emptive easing characterizes the risky steady state, rather than the adjustment process to a (transitory) shock. Indeed, section 5.2 demonstrates that the adjustment process following a permanent shock is characterized by gradualism, rather than aggressiveness.

5.2 The adjustment process after permanent shocks

Our model also allows us to study the adjustment process following a permanent shocks. We show the results in Figure 3, which focuses on the same shocks analyzed in Figure 2.

The key feature emerging from Figure 3 is that the adjustment of all endogenous variables to the new risky steady state is not instantaneous. This may be surprising, given the purely forward-looking nature of the standard new-Keynesian model. Based on equations (1) and (4), one may expect that, following a $\bar{r}_t$ shock, the economy would immediately jump to its new risky steady state. More specifically, consider for example the case in which the shock occurs when $\bar{r}_0 = 1.5\%$. Note that, since the policy rate is in positive territory, $\lambda_{x,t-1} = \lambda_{x,t} = 0$ in the optimal policy equations (6) and (7). This implies that, after the shock, the central bank could immediately set the real interest rate at its new, neutral level. The output gap would immediately jump to its new stochastic steady state. Equations (6) and (7) then show that the (in expectation) permanent increase in the output gap would be accompanied by one period of negative inflation. Thereafter, inflation would return to zero and the output gap would stay at its higher level.

Figure 3: Impulse responses to permanent shocks under optimal policy starting from different values of $\bar{r}_0$

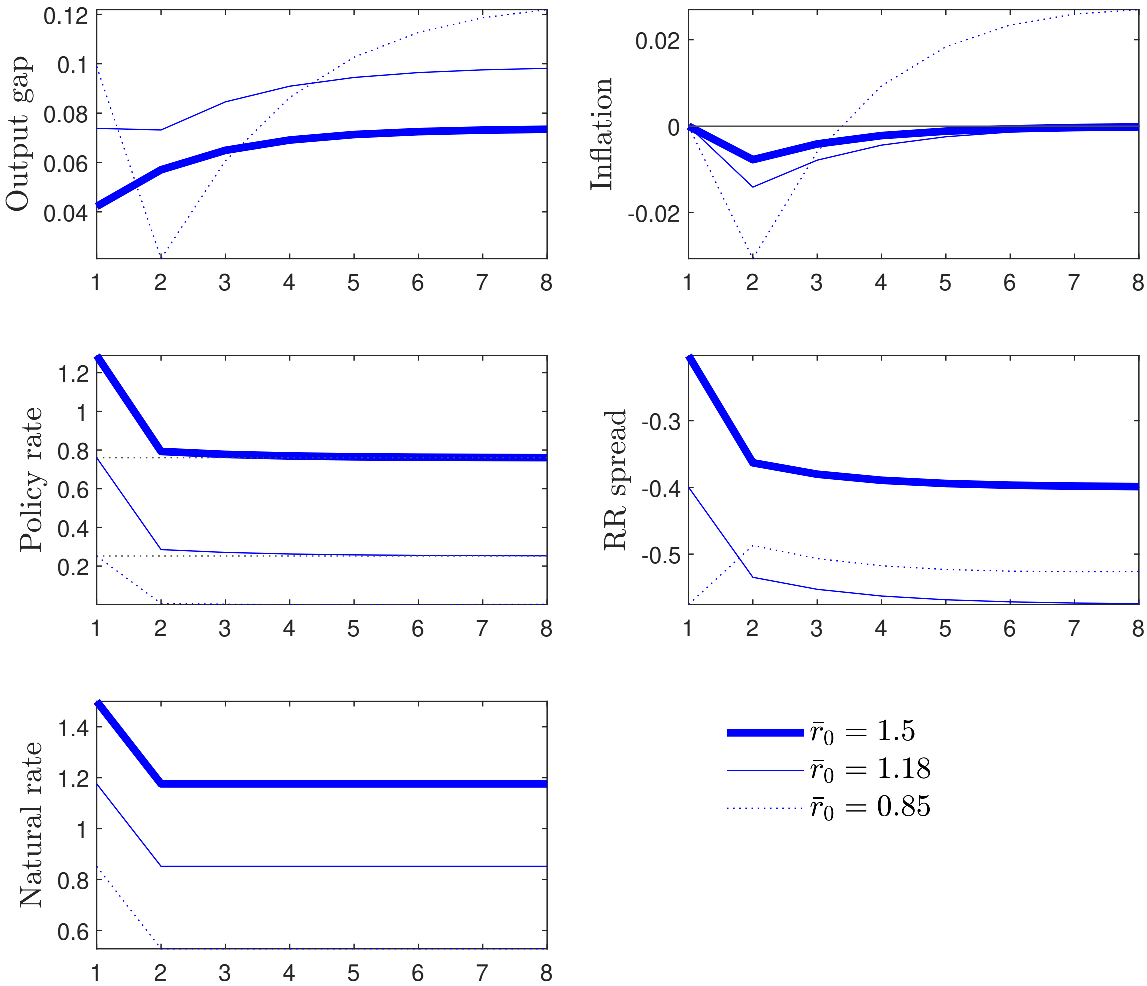

Note: the figure shows impulse responses to permanent shocks of a given size (0.324 percentage points, i.e., three times the standard deviation of the shock σ_ψ) starting from different initial values of the long-run natural rate. The response of the output gap is in percentage points; all other responses are expressed in annualised percentage points.

To understand why a different outcome is desirable, recall that a general feature of optimal policy under commitment is history-dependence – Woodford (2003). The central bank can use its ability to make credible promises to implement a different dynamic response to the shock. As shown in Figure 3, it can announce a gradual decline of the real interest rate towards its new, neutral level. As a result, the output gap will be less positive on impact, and it will slowly grow towards its new stochastic steady state. From the optimal policy equations (6) and (7), this will result in an initially smaller, if more persistent fall in inflation.

Hence, the key benefit of this dynamic policy response is to generate superior inflation stabilization in the period when the shock occurs. This comes at the cost of a more prolonged period of negative inflation. However, this cost is contained, since negative inflation will be reflected in expectations and will thus contribute to offset the inflationary pressure created by the persistently positive output gap.

The history-dependence of the optimal response to $\bar{r}_t$ shocks is consistent with those of the dynamic response to transitory cost-push shocks analyzed, for example, in Woodford (2010). In contrast to that case, however, the price in level in Figure 3 will not return to its initial value, but it will be permanently lower. This is the case both when the initial $\bar{r}_0 = 1.5\%$ and when $\bar{r}_0 = 1.18\%$, i.e. when policy rates are not yet at the ZLB and can be set to achieve the desired degree of policy stimulus. Results change when the initial long-run natural rate is at 0.85%. In this case the policy rate reaches the ZLB, so the dynamic response of the previous two cases becomes infeasible. The output gap can only be pushed up over time by the promise to create excess inflation in reaction to other shocks, but this promise will be reflected in expectations resulting in positive inflation in the new stochastic steady state.

Interestingly, a slow adjustment of endogenous variables after a permanent shock appears to be a feature of the data. Our impulse responses are qualitatively consistent with the results in Schmitt-Grohé and Uribe (2022), which characterises empirically the macroeconomic impact of permanent shocks to the natural rate when there is sufficient room for reductions in the policy rate. Schmitt-Grohé and Uribe (2022) also finds that a permanent real interest rate shock is deflationary in the short run and leads to a slow reduction in the real interest rate.[13] The adjustment process in Schmitt-Grohé and Uribe (2022) is however much more drawn out and lasts years. The differences may be due to our focus on optimal policy, while Schmitt-Grohé and Uribe (2022) does not make any specific assumptions regarding the behaviour of the central bank. Our model also abstracts from many sources of persistence which are likely to be a feature the data.

[13] Schmitt-Grohé and Uribe (2022) finds that the shock also leads to a reduction in output. This is again consistent with our theoretical results. Even if the output gap increases in the figure, output falls due to the permanent reduction in productivity caused by the shock.

5.3 The frequency of the ZLB under optimal policy

As shown in Table 1, we calibrate the standard parameters of our model as in the existing literature, except for allowing $\bar{r}_t$ to fall to zero. It is therefore interesting to check the implications of a lower $\bar{r}_t$ for the expected future frequency of the ZLB. These implications are explored in Table 3.

The first column in the table shows the ZLB frequency obtained in a model with constant $\bar{r}$ and calibrated so that $\bar{r} = 3.5\%$, as in Adam and Billi (2006). Consistently with the results of that paper, the incidence of the ZLB is 5%. Starting from this benchmark, Table 3 shows that the ZLB frequency would increase very quickly, if the model with constant $\bar{r}$ were solved for lower and lower values of $\bar{r}$ – shown in the second-to-sixth columns in the table. It would rise to 50% for $\bar{r} = 1\%$ and to 84% for $\bar{r} = 0$. Similarly, the average duration of a ZLB episode would progressively increase from just over 2 quarters when $\bar{r} = 3.5\%$ to 22 quarters for $\bar{r} = 0$.

Clearly, results on the ZLB frequency are highly sensitive to the calibration of $\bar{r}$. In section 2, we have reported recent estimates placing this notion anywhere between 1/2 and 2 percent. Which steady state calibration should be trusted at the current juncture?

Table 3: ZLB frequency

	w/o ε_t^{ψ}						with ε_t^{ψ}
$\bar{r}$	3.5 percent	2 percent	1.5 percent	1 percent	0.5 percent	0 percent	
ELB frequency (x100)	4,95	22,85	35,56	49,91	68,03	84,37	36,77
ELB duration (quarters)	2,40	3,99	5,21	7,18	11,22	21,93	6,24

Our model allows us to draw conclusions on the incidence of the ZLB without taking a stance on the value of $\bar{r}$ which will prevail in the future. Our results, shown in the last column of Table 3, suggest that the ZLB would bind approximately one third of the time. Note that this outcome is very similar to that of a model with constant $\bar{r} = 1.5\%$. This is not surprising, given that 1.5% is also the unconditional mean of $\bar{r}_t$ (Table 2). In other words, our model provides a possible rule of thumb for the calibration of models that ignore uncertainty as to the future value of the long-run natural rate. The rule of thumb is to fix the steady state natural rate in those models at the unconditional mean, rather than at the most recent value, of available empirical estimates of the long-run natural rate. Of course, this rule of thumb is only valid in a situation, such as the one we analyze, in which the distribution of $\bar{r}_t$ is symmetric around its mean.

A ZLB frequency equal to one third is strikingly high. It should of course be interpreted with care, since it is based on the basic new Keynesian model, which provides a very simple characterization of economic and inflation dynamics. Nevertheless, our results are of the same order of magnitude as those obtained in Kiley and Roberts (2017) using the much richer Federal Resere Board's FRB/US model and the assumption that monetary policy follows a simple interest rate rule.[14] Moreover, a higher incidence of the ZLB under optimal commitment policy than under simple rules is to be expected, since the neutral rate will more quickly reach zero as $\bar{r}$ falls, due to the expansionary bias that characterizes optimal commitment policy.

Our model also suggests that the average ZLB duration with time-varying $\bar{r}_t$ would be roughly 1.5 years. This is not a particularly long period in the wake of the experience following the Great recession.

6 Implementation: price level targeting

So far we have analyzed the properties of an equilibrium in which monetary policy is set optimally according to a Ramsey plan. Amongst the many requirements of such policy approach for the central bank is the observation of $\bar{r}_t$. This requirement is clearly unrealistic, given the existing range of possible estimates of $\bar{r}_t$.

In this section we investigate how the optimal equilibrium described in Section 5 can be implemented in practice. We specifically focus on implementation through a simple policy rule which does not require knowledge of the long-run natural rate of interest.

An obvious candidate is the price level targeting rule put forward in Eggertsson and Woodford (2003). This is a rule of the form

$$p_t + \frac{\lambda}{\kappa} x_t = P^* \tag{8}$$

where P^* is the target gap-adjusted (log-)price level (a given constant). Equation (8) commits the central bank to counteract shocks destabilizing the price level. When perfect stabilization is not feasible because of the ELB, the rule commits the central bank to undo actual deflation with future inflation so as to bring the gap adjusted (log-) price level (henceforth GAPL) back to target. Over long time periods, therefore, inflation will be zero on average.

[14] Kiley and Roberts (2017) finds that the ZLB would bind two-fifths of the time if the long-run nominal interest rate were 3% and the inflation target 2%, so that $\bar{r}$ would equal 1%.

Figure 4: Welfare loss under price level targeting with drift

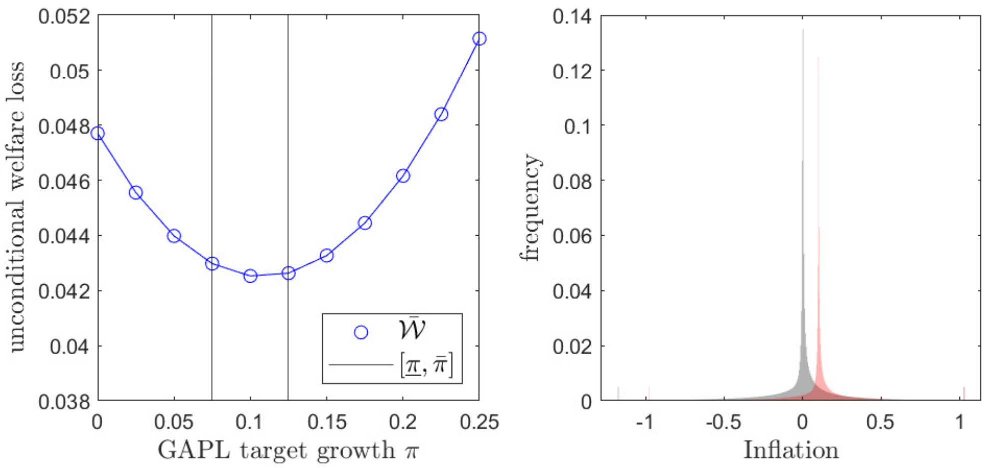

Note: the left panel shows the welfare-theoretic central bank loss function for different values of the GAPL target growth π^* in equation (10). The right panel shows the unconditional distribution of inflation winsorized at 1% under price level targeting when π^* is equal to zero (grey) and 10 basis points (red). Both the price level target π^* and inflation are expressed in annualised percentage terms.

Eggertsson and Woodford (2003) study the properties of the price level targeting rule (8) within a stationary model, in which optimal inflation is zero on average. By contrast, as shown in Figure 2 for very low values of $\bar{r}$, optimal policy in our model can call for a positive rate of inflation. One would therefore expect the Eggertsson and Woodford (2003) rule to have a poorer performance in our model.

For this reason, we also consider a possible refinement of rule (8), which is tailored to our environment. The refinement is to make the price level target no longer constant, but evolving along an exogenous trend. More specifically, we consider the rule

$$p_t + \frac{\lambda}{\kappa} x_t = P_t^* \tag{9}$$

where the time-varying P_t^* follows an exogenous, deterministic trend π^* such that

$$P_t^* = P_{t-1}^* + \pi^* \tag{10}$$

Clearly, the adjusted rule (9)-(10) boils down to rule (8) when $\pi^* = 0$.

Figure 5: Impulse responses to a sequence of permanent shocks under various rules

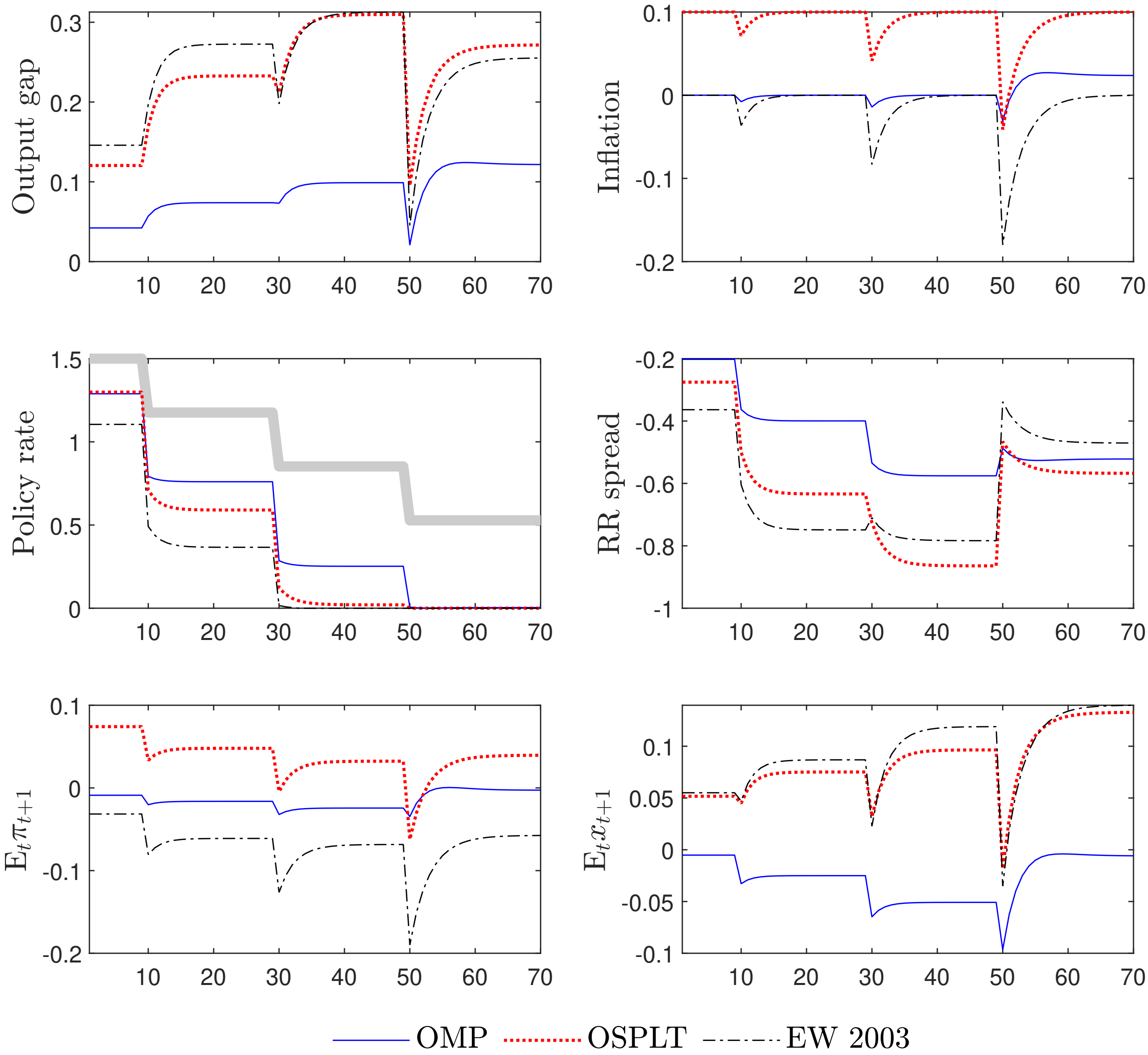

Note: the figure shows the results of the same simulation underlying Figure 2. The solid blue line corresponds to the solid blue line in Figure 2. The dotted red and black lines indicate outcomes under price level targeting with an optimal price level trend, and with a price level trend equal to zero, respectively. The solid grey line in the policy rate panel indicates the natural rate. The response of the output gap is in percentage points; all other responses are expressed in annualised percentage points.

The relative merit of the two rules can be analyzed in terms of their welfare performance. The left hand side panel of Figure 4 reports the results of this analysis for different values of π^* between zero and 0.25%.

The figure shows that the rule with positive trend attains a lower welfare loss. Welfare is maximized when the price level trend is equal to 10 basis points (annualized). Nevertheless, the difference in performance between the two rules is small. The welfare loss of the rule with constant price target is approximately 12% larger than for the rule with the optimal price level trend. By contrast, the loss under optimal policy is 43% smaller than under the rule with the optimal price level trend.[15]

The right panel of Figure 4 shows the equilibrium distribution of inflation under price level targeting for two values of the price level trend π^*: zero and 10 basis points. The panel demonstrates that a positive price level trend has the advantage of reducing the mass of the left tail of the inflation distribution.

Figure 5 shows the economy's responses to the same sequence of $\bar{r}_t$ shocks as in Figure 2 when monetary policy follows the two price level targeting rules (8) and (9). The responses under optimal policy already shown in Figure 2 are also displayed again in Figure 5 – solid blue line – for ease of comparison.

By design, the key characteristic of the Eggertsson and Woodford price level targeting rule (8) is to stabilize the gap-adjusted price level. This is consistent with optimal policy outcomes as long as the policy rate does not hit the ZLB constraint. Compared to optimal policy, however, the price level targeting rule (8) is less sophisticated in managing expectations. As a result, the output gap is larger than in the optimal policy case. The spread between the neutral rate and the long-run natural rate must be more pronounced, thus the neutral rate falls to zero faster than under optimal policy, as the long-run natural rate goes down. The neutral rate is in fact zero as soon as the long-run natural rate falls below 1%, in period 30 of the simulation.

Once $r^* = 0$, the central bank has no room to reduce the policy rate further. After a negative permanent shock (see period 50) the gap-adjusted price level falls and undershoots its target. The central bank continues promising to bring the gap-adjusted price level back to the original target under rule (8), while optimal policy prescribes to increase the gap-adjusted price level to a point higher than the original target, by an amount proportionate to the previous target shortfall. As a result, the real rate increases more on impact under rule (8), and the economic slowdown is much more pronounced than under optimal policy, even if comparably short-lived. In

[15] The unconditional welfare loss amounts to 0.0477, 0.0425, and 0.0241 under the constant gap adjusted price level (GAPL) targeting rule, the optimal growing GAPL targeting rule, and the optimal policy respectively.

addition, expectations incorporate a deflation bias under rule (8) – see the panels in the last row of Figure 5 – while they are consistent with zero inflation under optimal policy.

Economic dynamics under the modified price level targeting rule (9) are closer to those under optimal policy. This is the case because, before the ZLB is reached, rule (9) is consistent with a significantly higher level of the neutral rate. At the beginning of the simulation, when $\bar{r}_t = 1.5\%$, the neutral rate is 30 basis points higher under rule (9) than under rule (8). Consequently, the central bank has more space to reduce the policy rate in reaction to shocks under rule (9). The deterministic price trend also allows for superior stabilization outcomes once the ZLB becomes binding – again as of period 50. In contrast to what we observe under rule (8), inflation expectations remain positive in this case. As a result, the real rate can be lower and provide a higher degree of policy accommodation.

In sum, a price level targeting rule continues to provide a reasonably good approximation of optimal policy also if permanent natural rate shocks are present in the model. This is the case especially if the original Eggertsson and Woodford (2003) rule is complemented by a price level trend. Both price level targeting rules, however, imply that the neutral policy rate should reach zero earlier than under optimal policy. This would be the case as soon as the long-run natural rate falls below 1%.

7 Concluding remarks

Empirical research suggests that the long-run natural interest rate is not constant over time, but varies unpredictably. In this paper we have constructed a model which accounts for such evidence. We have shown that the risk of future reductions in the long-run natural rate tends to impart a downward bias on output and inflation expectations. To offset this bias in expectations, the neutral rate will be lower than the long-run natural rate. Obviously, this approach is no longer feasible once the policy rate hits the ZLB. At that point, the central can only stabilise inflation by promising to create inflation in the future.

The paper also shows that price level targeting rules can approximately implement optimal policy, especially if they incorporate an exogenous drift in the price level. Such rules have the advantage of not requiring knowledge of the long-run natural rate.

As in most of the related literature, we have conducted our analysis in the context of a variant of the baseline new Keynesian model. This is due to the complexity of solving the model nonlinearly under uncertainty. In future research, it would however be interesting to analyse the robustness of our results to the case of a more complex and realistic framework.

While allowing for uncertainty as to the future long-run natural rate, we have preserved the assumption that current and past values of the natural rate are common knowledge within the model. Exploring the implications of the model when the natural rate is not observable is an interesting avenue for future research.

References

Adam, K. and Billi, R. M. (2006). Optimal monetary policy under commitment with a zero bound on nominal interest rates. *Journal of Money, credit and Banking*, pages 1877–1905.

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

Aksoy, Y., Basso, H. S., Smith, R. P., and Grasl, T. (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics*, pages 193—-222.

Andrade, P., Le Bihan, H., Galí, J., and Matheron, J. (2019). The optimal inflation target and the natural rate of interest. *Brookings Papers on Economic Activity*, (Fall 2019 Edition).

Andrade, P., Le Bihan, H., Galí, J., and Matheron, J. (2021). Should the ecb adjust its strategy in the face of a lower r*. *Journal of Economic Dynamics and Control*, 14:104207.

Bianchi, F., Melosi, L., and Rottner, M. (2021). Hitting the elusive inflation target. *Journal of Monetary Economics*, 124:107–122.

Billi, R. M. (2011). Optimal inflation for the us economy. *American Economic Journal: Macroeconomics*, 3(3):29–52.

Billi, R. M., Galí, J., and Nakov, A. (2023). Optimal monetary policy with $r^* < 0$. *Journal of Monetary Economics*, forthcoming.

Carvalho, C., Ferrero, A., and Nechio, F. (2016). Demographics and real interest rates: Inspecting the mechanism. *European Economic Review*, 88:208–226.

Clarida, R., Galí, J., and Gertler, M. (1999). The science of monetary policy: a new keynesian perspective. *Journal of Economic Literature*, XXXVII:1661–1707.

Coeurdacier, N., Rey, H., and Winant, P. (2011). The risky steady state. *American Economic Review*, 101(3):398–401.

Davig, T. and Leeper, E. M. (2007). Generalizing the taylor principle. *American Economic Review*, 97(3):607–635.

Debortoli, D., Maih, J., and Nunes, R. (2014). Loose commitment in medium-scale macroeconomic models: Theory and applications. *Macroeconomic Dynamics*, 18:175–198.

Del Negro, M., Giannone, D., Giannoni, M. P., and Tambalotti, A. (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*, 2017(1):235–316.

Edge, R. M., Kiley, M. T., and Laforte, J.-P. (2008). Natural rate measures in an estimated dsge model of the us economy. *Journal of Economic Dynamics and control*, 32(8):2512–2535.

Eggertsson, G. B. and Woodford, M. (2003). Zero bound on interest rates and optimal monetary policy. *Brookings papers on economic activity*, 2003(1):139–233.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Citeseer.

Fernández-Villaverde, J., Marbet, J., Nuño, G., and Rachedi, O. (2021). Inequality and the zero lower bound. Technical report, mimeo.

Fisher, J. D. M. (2015). On the structural interpretation of the smets–wouters “risk premium” shock. *Journal of Money, Credit and Banking*, 47(2–3):511–516.

Gagnon, E., Johannsen, B. K., and López-Salido, D. (2016). Understanding the new normal: The role of demographics. *IMF Economic Review*, 69:357—390.

Giannoni, M., Patterson, C., Del Negro, M., et al. (2015). The forward guidance puzzle. In *2015 Meeting Papers*, number 1529. Society for Economic Dynamics.

Hamilton, J. D., Harris, E. S., Hatzius, J., and West, K. D. (2016). The equilibrium real funds rate: Past, present, and future. *IMF Economic Review*, 64(4):660–707.

IMF (2022). *World Economic Outlook: Countering the Cost-of-Living Crisis*. International Monetary Fund, Washington, DC. October.

Jung, T., Teranishi, Y., and Watanabe, T. (2005). Optimal monetary policy at the zero-interest-rate bound. *Journal of Money, credit, and Banking*, 37(5):813–835.

Justiniano, A. and Primiceri, G. E. (2010). Measuring the equilibrium real interest rate. *Economic Perspectives*, 34(1).

Kiley, M. T. and Roberts, J. M. (2017). Monetary policy in a low interest rate world. *Brookings Papers on Economic Activity*, 2017(1):317–396.

King, R. G., Plosser, C. I., and Rebelo, S. T. (1988). Production, growth and business cycles: I. the basic neoclassical model. *Journal of monetary Economics*, 21(2-3):195–232.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2012). The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267.

Krugman, P. (1998). It's baaack: Japan's slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, 1998(2):137–187.

Laubach, T. and Williams, J. C. (2003). Measuring the natural rate of interest. *Review of Economics and Statistics*, 85(4):1063–1070.

Levin, A., López-Salido, D., Nelson, E., and Yun, T. (2010). Limitations on the effectiveness of forward guidance at the zero lower bound. *International Journal of Central Banking*.

Lindé, J., Smets, F., and Wouters, R. (2016). Challenges for central banks' macro models. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 2183–2262.

Maliar, L. and Maliar, S. (2015). Merging simulation and projection approaches to solve high-dimensional problems with an application to a new keynesian model. *Quantitative Economics*, 6(1):1–47.

Masolo, R. M. and Winant, P. E. (2019). The stochastic lower bound. *Economics Letters*, 180:54–57.

Mertens, T. M. and Williams, J. C. (2019). Monetary policy frameworks and the effective lower bound on interest rates. In *AEA Papers and Proceedings*, volume 109, pages 427–32.

Michaillat, P. and Saez, E. (2021). Resolving new keynesian anomalies with wealth in the utility function. *Review of Economics and Statistics*, 103(2):197–215.

Nakata, T. and Schmidt, S. (2019). Conservatism and liquidity traps. *Journal of Monetary Economics*, 104:37–47.

Nakov, A. (2008). Optimal and simple monetary policy rules with zero floor on the nominal interest rate. *International Journal of Central Banking*, 4(2):73–127.

Obstfeld, M. (2023). Natural and neutral real interest rates: Past and future. Prepared for the 10th Asian Monetary Policy Forum, Singapore, May 25-26.

Platzer, J. and Peruffo, M. (2022). Secular drivers of the natural rate of interest in the united states: A quantitative evaluation. IMF Working Paper No. 2022/030.

Powell, J. H. (2018). *Monetary policy in a changing economy.* Speech at the Federal Reserve Bank of Kansas City's Symposium "Changing Market Structure and Implications for Monetary Policy", Jackson Hole.

Reifschneider, D. and Williams, J. C. (2000). Three lessons for monetary policy in a low-inflation era. *Journal of Money, Credit and Banking*, pages 936–966.

Richter, A. W. and Throckmorton, N. A. (2015). The zero lower bound: frequency, duration, and numerical convergence. *The BE Journal of Macroeconomics*, 15(1):157–182.

Roulleau-Pasdeloup, J. (2020). Optimal monetary policy and determinacy under active/passive regimes. *European Economic Review*, 130:103582.

Schmitt-Grohé, S. and Uribe, M. (2022). The macroeconomic consequences of natural rate shocks: An empirical investigation. NBER Working Paper 30337.

Sidrauski, M. (1967). Rational choice and patterns of growth in a monetary economy. *The American Economic Review*, 57(2):534–544.

Vestin, D. (2006). Price-level versus inflation targeting. *Journal of Monetary Economics*, 53(7):1361–1376.

Woodford, M. (2003). Optimal interest-rate smoothing. *Review of Economic Studies*, 70:861–886.

Woodford, M. (2010). Optimal monetary stabilization policy. In Friedman, B. M. and Woodford, M., editors, *Handbook of Monetary Economics*, volume 3, pages 723–828.

APPENDIX - NOT FOR PUBLICATION

A The model

A.1 Households' and firms' optimization problems

The problem of household j is to

$$\max_{C_{j,t},H_{j,k,t},M_{j,t},B_{j,t+1}} \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t U_t \left(C_{j,t}, H_{j,k,t}, \frac{M_{j,t}}{P_t} \right)$$

where $C_{j,t}$, $H_{j,t}$, and $M_{j,t}$ are consumption, hours worked in firm k and government bonds,

$$U_{j,t} = \bar{C}_t \left(\log C_{j,t} + S_t \upsilon \left(\frac{M_{j,t}}{P_t} - \frac{M_t}{P_t} \right) - \frac{\gamma}{1+v} \bar{H}_t^{-v} \int_0^1 H_{j,k,t}^{1+v} \mathrm{d}k \right)$$

where the function $\upsilon(\cdot)$ is increasing and concave and S_t and $\bar{H}$ are shocks.

Utility maximisation is subject to the budget constraint

$$\mathrm{E}_t Q_{t,t+1} B_{j,t+1} + M_{j,t} \leq I_{t-1}^m M_{j,t-1} + B_{j,t} + \int_0^1 W_{k,t} H_{j,k,t} \mathrm{d}k + \Pi_t^{firm} + T_t - P_t C_{j,t}$$

where we assume complete markets. $B_{j,t+1}$ is a portfolio of state contingent assets and Π_t^{firm} and T_t are firms' distributed profits and taxes/transfers.

We assume $\bar{C}_t = \Delta_t \bar{C}_{t-1}$, for $t > 1$ and $\bar{C}_0 = 1$. We will also assume that productivity includes a stochastic trend $\bar{A}_t$, such that $A_t = \bar{A}_t Z_t$. Before linearising, we detrended consumption as $\tilde{C}_t = C_t / \bar{A}_t$ and detrended real money as $\widetilde{m}_t = (M_t/P_t)/\bar{A}_t$. We also assume that $\upsilon'(\cdot)$ is homogeneous of degree -1, so that $\upsilon'\left([\widetilde{m}_{j,t} - \widetilde{m}_t] \bar{A}_t\right) = \upsilon'(\widetilde{m}_{j,t} - \widetilde{m}_t) \bar{A}_t^{-1}$. Since everyone has same preferences and same initial wealth, we can drop the $j's$. The first order conditions include:

$$\begin{aligned}
\frac{W_{k,t}}{P_t} &= \gamma \bar{H}_t^{-v} H_{k,t}^v \tilde{C}_t \bar{A}_t \\
1 - \frac{I_t^m}{I_t} &= S_t \tilde{C}_t \upsilon'(0) \\
1 &= \mathrm{E}_t \left[\beta \Delta_{t+1} \frac{\tilde{C}_t}{\tilde{C}_{t+1}} \frac{\bar{A}_t}{\bar{A}_{t+1}} \frac{I_t}{\Pi_{t+1}} \right]
\end{aligned}$$

where we assume that $\upsilon'(0) > 0$.

We will assume that $\delta_t = \ln \Delta_t$, $s_t = \ln S_t$ and $z_t = \ln Z_t$ follow AR(1) processes

$$\begin{aligned}\delta_{t+1} &= \rho_\delta \delta_t + \sigma_\delta \varepsilon_{\delta,t+1},\\ s_{t+1} &= \rho_s s_t + \sigma_s \varepsilon_{s,t+1},\\ z_{t+1} &= \rho_z z_t + \sigma_z \varepsilon_{z,t+1},\end{aligned}$$

and that productivity growth, $\Xi_{t+1} = \frac{\bar{A}_{t+1}}{\bar{A}_t}$ is itself integrated, i.e., the productivity growth rate $\xi_t = \log \Xi_t$ follows

$$\begin{aligned}\xi_t &= \xi_{t-1} + \psi_t\\ \psi_t &= (1-\rho_\psi)\,\psi + \rho_\psi \psi_{t-1} + \sigma_\psi \varepsilon_t^\psi\end{aligned}$$

where ψ_t is the (rate of) change in productivity growth.

Then, we detrended nominal interest rates as $\breve{I}_t = \frac{I_t}{\Xi_t}$ and $\breve{I}_t^m = \frac{I_t^m}{\Xi_t}$ and we used $\Delta_t^m \equiv \frac{\breve{I}_t - \breve{I}_t^m}{\breve{I}_t}$ to denote the convenience yield on treasury bonds. In a steady state with zero inflation, $\Delta^m = \tilde{C}v'(0)$ and $\breve{I}^m = \frac{\Psi}{\beta}$, where Ψ denotes the exponential value of the rate of change in productivity growth. Up to first order

$$\frac{1-\Delta^m}{\Delta^m}\left(\breve{\imath}_t - \breve{\imath}_t^m\right) = \tilde{c}_t + s_t$$

and

$$\tilde{c}_t = \mathrm{E}_t \tilde{c}_{t+1} - \left(\breve{\imath}_t^m - \mathrm{E}_t \pi_{t+1}\right) + \mathrm{E}_t \widehat{\psi}_{t+1} - \mathrm{E}_t \delta_{t+1}$$

where small case letters denote (log) deviations from the steady state. These two equations can be combined to obtain

$$\tilde{c}_t = (1-\Delta^m)\left[\mathrm{E}_t \tilde{c}_{t+1} - \left(\breve{\imath}_t^m - \mathrm{E}_t \pi_{t+1}\right) + \mathrm{E}_t \widehat{\psi}_{t+1} - \mathrm{E}_t \delta_{t+1}\right] - \Delta^m s_t$$

In a natural equilibrium where prices are flexible

$$\tilde{c}_t^n = (1-\Delta^m)\left(\mathrm{E}_t \tilde{c}_{t+1}^n - \breve{r}_t^n + \mathrm{E}_t \widehat{\psi}_{t+1} - \mathrm{E}_t \delta_{t+1}\right) - \Delta^m s_t$$

If we define the output gap as $x_t = \tilde{c}_t - \tilde{c}_t^n$, we can therefore obtain

$$x_t = (1-\Delta^m)\left(\mathrm{E}_t x_{t+1} - \left(\breve{\imath}_t^m - \mathrm{E}_t \pi_{t+1} - \breve{r}_t^n\right)\right)$$

where the natural rate in deviation from its long-run level is

$$\breve{r}_t^n = \mathrm{E}_t \tilde{c}_{t+1}^n - \frac{1}{1-\Delta^m}\tilde{c}_t^n + \mathrm{E}_t \widehat{\psi}_{t+1} - \mathrm{E}_t \delta_{t+1} - \frac{\Delta^m}{1-\Delta^m} s_t$$

There is a continuum of firms indexed by $k \in [0,1]$ producing intermediate goods under monopolistic competition and sticky prices. The production function is

$$Y_{k,t} = A_t \left(H_{k,t}\right)^{\frac{1}{\phi}}.$$

At any point in time, firms may reset their prices $p_{k,t}$ with probability α. Profit maximisation yields the first order condition

$$\left(\frac{1-\alpha\Pi_t^{\theta-1}}{1-\alpha}\right)^{\frac{1+\omega\theta}{1-\theta}} = \frac{\gamma\phi\frac{\theta}{\theta-1}\mathrm{E}_t\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}\frac{\bar{C}_T}{\bar{C}_t}\mu_T^W\bar{H}_T^{-v}\left(\frac{P_T}{P_t}\right)^{\theta(1+\omega)}\left(\frac{Y_T}{A_T}\right)^{1+\omega}}{\mathrm{E}_t\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}\frac{\bar{C}_T}{\bar{C}_t}(1-\tau)\left(\frac{P_T}{P_t}\right)^{\theta-1}}$$

Detrending real variables and following standard derivations we can linearize this condition to obtain $\pi_t = \kappa\left(\tilde{y}_t - z_t\right) + \beta\mathrm{E}_t\pi_{t+1}$, for $\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\frac{1+\omega}{1+\omega\theta}$. In the natural equilibrium we obtain $\tilde{y}_t^n = z_t$, so that the Phillips curve can be rewritten as

$$\pi_t = \kappa x_t + \beta\mathrm{E}_t\pi_{t+1}$$

for $x_t \equiv \tilde{y}_t - \tilde{y}_t^n$.

In the text, we assume that $z_t = s_t = 0$ at all times and that $\psi = \rho_\psi = 0$. Replacing this into the expression of the natural rate in deviation from its long-run level, we obtain

$$\breve{r}_t^n = -\rho_\delta\delta_t + \mathrm{E}_t\psi_{t+1}$$

We then define $\bar{\delta}_t = -\rho_\delta\delta_t$ and work directly with $\bar{\delta}_t \equiv \rho_{\breve{r}^n}\bar{\delta}_{t-1} + \sigma_{\breve{r}^n}\varepsilon_t^{\bar{\delta}}$.

A.2 Second order welfare approximation

Household period utility can be rewritten as

$$U_t = \bar{C}_t\left[\ln\tilde{C}_t + \bar{a}_t + s_t v(0) - \frac{\gamma}{1+v}\int_0^1 H_{k,t}^{1+v}\mathrm{d}k\right]$$

where $\bar{a}_t = \log\bar{A}_t$ and where we used $\tilde{m}_{j,t} = \tilde{m}_t$. Using $H_{k,t} = \left(\frac{Y_{k,t}}{Z_t\bar{A}_t}\right)^{\phi}$ and the demand schedule $Y_{k,t} = \left(\frac{p_{k,t}}{P_t}\right)^{-\theta}Y_t$, we obtain

$$U_t = \bar{C}_t\left[\ln\tilde{Y}_t + \bar{a}_t + s_t v(0) - \frac{\gamma}{1+v}\left(\frac{\tilde{Y}_t}{Z_t}\right)^{\phi(1+v)}d_t\right]$$

for $d_t = \int_0^1 \left(\frac{p_{k,t}}{P_t}\right)^{-\theta(1+\omega)} \mathrm{d}k$. Note that $\bar{C}_t$, $\bar{a}_t$, s_t and $\upsilon(0)$ are independent of policy to write

$$U_t = \bar{C}_t \left[\ln \tilde{Y}_t - \frac{\gamma}{1+\upsilon}\left(\frac{\tilde{Y}_t}{Z_t}\right)^{\phi(1+\upsilon)} d_t\right] + t.i.p.$$

where $t.i.p.$ are terms independent of policy.

Expand to second order noting that $\bar{C} = 1$ and using also $1+\omega = \phi(1+\upsilon)$:

$$\begin{aligned} U_t - \left(\ln \tilde{Y} - \frac{\gamma}{1+\upsilon}\tilde{Y}^{1+\omega} d\right) \simeq & -\frac{1}{2}\gamma\phi^2(1+\upsilon)\tilde{Y}^{1+\omega} d\tilde{y}_t^2 - \frac{\gamma}{1+\upsilon}\tilde{Y}^{1+\omega} d\hat{d}_t \\ & + \left(1-\gamma\phi\tilde{Y}^{1+\omega} d\right)\tilde{y}_t + \left(1-\gamma\phi\tilde{Y}^{1+\omega} d\right)\bar{c}_t\tilde{y}_t \\ & - \frac{1}{2}\frac{\gamma}{1+\upsilon}\tilde{Y}^{1+\omega} d\widehat{d}_t^2 - \frac{\gamma}{1+\upsilon}\tilde{Y}^{1+\omega} d\bar{c}_t\widehat{d}_t - \gamma\phi\tilde{Y}^{1+\omega} d\tilde{y}_t\widehat{d}_t \end{aligned}$$

where $\bar{c}_t = \delta_t + \bar{c}_{t-1}$.

The rest of the derivations are standard. Imposing the steady state subsidy $1-\tau = \frac{\theta}{\theta-1}$ to ensure that steady state output $\tilde{Y} = \left(\frac{\gamma\phi}{1-\tau}\frac{\theta}{\theta-1}\right)^{-\frac{1}{1+\omega}}$ becomes efficient, we can write intertemporal utility as of the beginning of time t_0 as

$$\frac{1-\alpha}{\alpha\theta}\frac{1-\alpha\beta}{1+\theta\omega}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left[U_t + \frac{1+\ln(\gamma\phi)}{1+\omega}\right] \simeq \frac{1}{2}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left(-\frac{(1-\alpha)(1-\alpha\beta)}{\alpha\theta}\frac{1+\omega}{1+\theta\omega}\tilde{y}_t^2 - \pi_t^2\right)$$

so that, in the absence of technology shocks z_t, the period utility to maximise for the CB can be written as

$$U_t^{CB} = -\pi_t^2 - \lambda x_t^2$$

for $\lambda = \frac{(1-\alpha)(1-\beta\alpha)}{\alpha\theta}\frac{1+\omega}{1+\theta\omega}$.

A.3 Optimal steady state policy

Consider the steady state of our economy and assume purely deterministic growth, i.e. $\bar{A}_{t+1}/\bar{A}_t = \Psi$.

Note that in a steady state with generic (gross) inflation rate Π, steady state utility is

$$U \propto \ln \tilde{Y} - \frac{\gamma}{1+\upsilon}\tilde{Y}^{(1+\omega)} d$$

If we consider the limiting case $\beta \to 1$, we can choose steady state inflation Π to maximise the resulting utility

$$U = \frac{\theta}{1-\theta}\ln\left(1-\alpha\Pi^{\theta-1}\right) + \frac{1}{1+\omega}\ln\left(1-\alpha\Pi^{\theta(1+\omega)}\right) + t.i.p.$$

subject to the ZLB constraint $I^m \geq 1$ or, equivalently,

$$s.t. \qquad \Pi \geq \frac{1}{R^n}$$

The first order condition require, using $\pi = \log \Pi$ and $r^n = \log R^n$, either

$$\pi = 0, \qquad \text{if} \qquad r^n \geq 0$$

or

$$\pi = -r^n, \qquad \text{if} \qquad r^n < 0$$

B The optimal policy commitment

The results in this section are based on a slightly different model calibration.

B.1 Numerical solution

This section describes the numerical procedure used for solving the model under the optimal policy commitment. We solve both the full model (section B.1.2) and a version without time-variation in the long-run natural rate (section B.1.1). The latter version allows us to compare our results to the existing literature.

B.1.1 Stationary natural rate of interest

Productivity growth

We assume that productivity growth is constant, i.e., $\Xi_{t+1} = \frac{\bar{A}_{t+1}}{\bar{A}_t} = \Psi$, and we denote by ψ the log of productivity growth.

System of equilibrium equations

$$\pi_t = \kappa x_t + \beta \mathrm{E}_t \pi_{t+1} \tag{11}$$
$$x_t = (1 - \Delta^m) \left[\mathrm{E}_t x_{t+1} - (\breve{\imath}_t^m - \mathrm{E}_t \pi_{t+1} - \breve{r}_t^n)\right] \tag{12}$$
$$\breve{r}_t^n = \bar{\delta}_t \tag{13}$$
$$2\lambda x_t = -\lambda_{x,t} + \beta^{-1} (1 - \Delta^m) \lambda_{x,t-1} + \kappa \lambda_{p,t} \tag{14}$$
$$2\pi_t = \beta^{-1} (1 - \Delta^m) \lambda_{x,t-1} - \lambda_{p,t} + \lambda_{p,t-1} \tag{15}$$
$$\lambda_{x,t}(\breve{\imath}_t^m + \ln (1 - \Delta^m) + \psi - \ln \beta) = 0 \tag{16}$$
$$\breve{\imath}_t^m \geq \log(\beta) - \psi - \log (1 - \Delta^m) \tag{17}$$
$$\lambda_{x,t} \geq 0 \tag{18}$$

Solution algorithm

We use a projection approach. To discretize the state space $\mathcal{S} \subset R^3$, we form a grid defined by three N-vectors of evenly spaced points, namely λ_{p-1}, λ_{x-1}, and $\bar{\delta}$. The initial range of values considered for each state variable is $[-0.005, 0.005]$, $[0, 0.005]$, and +/- 5 unconditional standard deviations of $\bar{\delta}_t$ respectively. Then, we proceed iteratively: we solve the model, simulate it, update the boundaries for λ_{p-1} and the upper bound for λ_{x-1} so as to cover all possible values, and solve the model again until both the solution and the grid converge. In our application, we set N=50.

We use the piecewise linear interpolation for approximating $x(s)$ and $\pi(s)$ off the grid, where $s = (\lambda_{p,t-1}, \lambda_{x,t-1}, \bar{\delta}_t)$ denotes the vector of state variables at time t, and a fixed-point iteration for solving the system on the grid.

Define $s_{+1} = (\lambda_{p,t}, \lambda_{x,t}, \breve{r}^n_{t+1})$ the vector of state variables at time t+1, and $f^c(.)$ the local polynomial approximating the control variable $c \in \{x, \pi\}$. Expectation terms are of the form: $\mathrm{E}_t[f^c(s_{+1})] = \int_{-\infty}^{\infty} g^c(\varepsilon^{\bar{\delta}}_{t+1}) \exp(-\varepsilon^{\bar{\delta}2}_{t+1}) d\varepsilon^{\bar{\delta}}_{t+1}$, which we approximate using a 9 node Gauss-Hermite (GH) quadrature.

The solution algorithm proceeds in four steps.

Step 1: Choose an initial $x_0(s)$ and $\pi_0(s)$, and a tolerance level τ

Step 2: Iteration j. For each possible state s, given $x_{j-1}(s)$ and $\pi_{j-1}(s)$, compute $\lambda_p(s)$ using (15), guess that $\lambda_x(s) = 0$, compute $\mathrm{E}x_{j-1}(s_{+1})$ and $\mathrm{E}\pi_{j-1}(s_{+1})$, and retrieve

$$\begin{aligned} \pi_j(s) &= \kappa x_{j-1}(s) + \beta \mathrm{E}\pi_{j-1}(s_{+1}) \\ x_j(s) &= \frac{1}{2\lambda}\Big[-\lambda_x(s) + \beta^{-1}(1-\Delta^m)\lambda_{x,-1} + \kappa\lambda_p(s)\Big] \end{aligned}$$

Step 3: Adjust if this allocation does not satisfy the ELB constraint. Let $\bar{\imath}^m$ denote the ELB on the policy rate. If $x_j(s) > (1-\Delta^m)\Big[\mathrm{E}_t x_{j-1}(s') - (\bar{\imath}^m - \mathrm{E}_t\pi_{j-1}(s') - \breve{r}^n_t)\Big]$, compute $\lambda_x(s)$ using (14), adjust $\mathrm{E}x_{j-1}(s_{+1})$ and $\mathrm{E}\pi_{j-1}(s_{+1})$ accordingly, and retrieve

$$\begin{aligned} \pi_j(s) &= \kappa x_{j-1}(s) + \beta \mathrm{E}\pi_{j-1}(s_{+1}) \\ x_j(s) &= (1-\Delta^m)\left[\mathrm{E}x_{j-1}(s_{+1}) - \left(\bar{\imath}^m - \mathrm{E}\pi_{j-1}(s_{+1}) - \bar{\delta}\right)\right] \end{aligned}$$

Step 4: Let $e^\pi_j(s) = |\pi_j(s) - \pi_{j-1}(s)|$, $e^x_j(s) = |x_j(s) - x_{j-1}(s)|$ and $e_j(s) = e^\pi_j(s) + e^x_j(s)$ denote different measures of approximation error. Stop if $\sum_s e_j(s) < \tau$. Otherwise, update the guess, and repeat step 2.

Accuracy

The accuracy of the solution is evaluated based on the mean and the maximum residuals obtained over a simulation of 10000 economies, each 1000 periods long. For each possible state of the economy, we use our solution to compute a piecewise linear interpolation of the output gap and of the rate of inflation $\Big(x^*(s_t), \pi^*(s_t)\Big)$. Then, as described in the algorithm, we plug these values in the system to retrieve the output gap and the rate of inflation implied by the equations $\Big(x^{\mathrm{IMP}}(s_t), \pi^{\mathrm{IMP}}(s_t)\Big)$, and we compute the residuals as follows:

$$e_t^x \equiv \left| x_t^{\mathrm{IMP}} - x_t^* \right| \cdot 100 \tag{19}$$

$$e_t^\pi \equiv \left| \pi_t^{\mathrm{IMP}} - \pi_t^* \right| \cdot 400 \tag{20}$$

Table 4 reports the mean and the maximum residuals, for each calibration of the steady state value of the natural rate. For the rate of inflation, we observe that the mean approximation error is negligeable. For the output gap, we observe that the mean approximation error is one order of magnitude higher than the mean of this variable, but in most cases it remains below or around one basis point.

Table 4: Stationary r_t^n and optimal commitment: simulation moments and accuracy indicators

r^n	$\breve{x}$	π	max$[e^x]$	E$[e^x]$	max$[e^\pi]$	E$[e^\pi]$
3,499	0	0,001	0,069	0,001	0,005	0
2,999	0	0,002	0,089	0,001	0,006	0
2,499	0	0,004	0,116	0,002	0,008	0
1,999	0	0,011	0,149	0,004	0,011	0
1,499	0	0,025	0,196	0,007	0,014	0
0,999	0,001	0,058	0,253	0,011	0,018	0,001
0,499	0,003	0,133	0,319	0,012	0,023	0,001
-0,001	0,007	0,293	0,395	0,01	0,028	0,001
0,997	0,006	0,267	0,637	0,025	0,045	0,002

This table reports simulation moments (annualized and in percent) along with accuracy indicators for each calibration of the steady state value of the natural rate. The last row contains the results for an economy with a steady state value of the natural rate at 1%, and a standard deviation of shocks two times larger than the baseline calibration.

Robustness

We test the robustness of the results in two ways. On the one hand, in the spirit of Maliar and Maliar (2015), we use an adaptative grid instead of an evenly spaced grid. On the other hand, we change the solution algorithm.

Adaptative grid

For each state variable except $\lambda_{x,t-1}$, we place relatively more points in the middle 95% of the distribution of this variable. For $\lambda_{x,t-1}$, given that the distribution is truncated in and concentrated near zero, we place the points using a multiplicative sequence of the form: $\lambda_{x,k} = \frac{\lambda_{x,k-1}}{1-\delta}$ with $0 < \delta < 1$. Figure 6 provides an illustrative example of the grid based on the solution obtained when the steady state natural rate of interest is 3.5%.

Table 5 and Table 6 compare accuracy measures and simulation moments when using the adaptative grid instead of the evenly spaced grid. We find that approximation errors diminish significantly when using the adaptative grid. However, the simulation moments are essentially unchanged. Only when the steady state value of the natural rate is equal to zero is the average duration of an ELB episode about one year longer. From this, we conclude that using a denser grid would certainty improve the accuracy of the solution, but it would not substantially affect the results that are reported in the main text.

Figure 6: Stationary r_t^n and optimal commitment: An adaptative grid

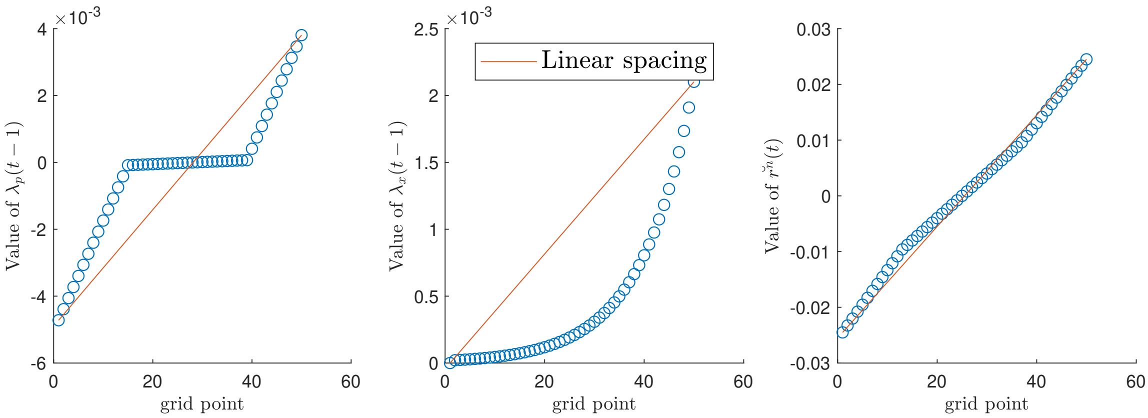

Table 5: Stationary r_t^n and optimal commitment: Robustness with respect to the grid, accuracy indicators

	Evenly spaced grid				Adaptative grid			
r^n	max[e^x]	E[e^x]	max[e^π]	E[e^π]	max[e^x]	E[e^x]	max[e^π]	E[e^π]
3,499	0,069	0,001	0,005	0	0,035	0	0,005	0
2,999	0,089	0,001	0,006	0	0,044	0	0,004	0
2,499	0,116	0,002	0,008	0	0,057	0,001	0,004	0
1,999	0,149	0,004	0,011	0	0,073	0,001	0,005	0
1,499	0,196	0,007	0,014	0	0,096	0,002	0,007	0
0,999	0,253	0,011	0,018	0,001	0,123	0,003	0,009	0
0,499	0,319	0,012	0,023	0,001	0,156	0,003	0,011	0
-0,001	0,395	0,01	0,028	0,001	0,194	0,003	0,013	0
0,997	0,637	0,025	0,045	0,002	0,312	0,006	0,022	0,001

See table 4 for details.

Table 6: Stationary r_t^n and optimal commitment: Robustness with respect to the grid, simulation moments

	Evenly spaced grid						Adaptative grid					
r^n	$\breve{x}$	π	i	$r - r^n$	ELB freq	ELB dur	$\breve{x}$	π	i	$r - r^n$	ELB freq	ELB dur
3,499	0	0,001	3,499	0	4,962	2,408	0	0,001	3,499	-0,001	4,93	2,745
2,999	0	0,002	3	-0,001	8,101	2,852	0	0,002	2,999	-0,001	9,22	3,05
2,499	0	0,004	2,502	-0,001	14,954	3,289	0	0,004	2,502	-0,001	16,075	3,522
1,999	0	0,011	2,009	0	22,9	4,003	0,001	0,01	2,007	-0,002	24,928	4,533
1,499	0	0,025	1,523	0	34,812	5,424	0,001	0,024	1,521	-0,001	36,595	6,055
0,999	0,001	0,058	1,057	0	49,867	7,185	0,002	0,056	1,053	-0,002	52,208	8,404
0,499	0,003	0,133	0,631	-0,001	67,987	11,291	0,004	0,13	0,627	-0,001	69,861	13,439
-0,001	0,007	0,293	0,29	-0,002	84,619	22,222	0,007	0,289	0,286	-0,001	86,012	25,776
0,997	0,006	0,267	1,263	-0,001	67,806	11,182	0,007	0,26	1,255	-0,003	69,465	13,486

See table 4 for details.

Alternative solution algorithm

We use essentially the same algorithm except that, instead of iterating on $x(s)$ and $\pi(s)$, we iterate on the lagrange multipliers $\lambda_x(s)$ and $\lambda_p(s)$.

Table 7 compares the simulation moments obtained under our baseline solution method (iteration on $\pi(s)$ and $x(s)$ and evenly spaced grid) with those obtained under the alternative solution algorithm when using the adaptative grid. For the output gap and the rate of inflation, the results do not change. The policy rate and the real rate gap tend to be slightly lower on average, but the difference never exceeds 3 basis points. This in turn affects the ELB frequency and the average duration of an ELB episode. They tend to be slightly higher : the ELB frequency is at most 6 percentage points higher; for most cases, the difference in the average duration of an ELB episode does not exceed one year. This suggests that, if anything, the baseline method underestimates slightly these moments. Overall, though, we conclude that the baseline method provides a good approximation to the solution since the alternative algorithm generates very similar results.

Table 7: Stationary r_t^n and optimal commitment: Robustness with respect to the solution algorithm, simulation moments

	Baseline algorithm						Alternative algorithm					
r^n	$\breve{x}$	π	i	$r-r^n$	ELB freq	ELB dur	$\breve{x}$	π	i	$r-r^n$	ELB freq	ELB dur
3,499	0	0,001	3,499	0	4,962	2,408	0	0,001	3,498	-0,001	5,705	2,69
2,999	0	0,002	3	-0,001	8,101	2,852	0,001	0,002	2,998	-0,002	10,109	3,142
2,499	0	0,004	2,502	-0,001	14,954	3,289	0,001	0,005	2,5	-0,003	17,022	3,755
1,999	0	0,011	2,009	0	22,9	4,003	0,002	0,011	2,005	-0,004	26,924	4,706
1,499	0	0,025	1,523	0	34,812	5,424	0,003	0,025	1,515	-0,008	40,013	6,319
0,999	0,001	0,058	1,057	0	49,867	7,185	0,005	0,057	1,043	-0,011	55,991	9,173
0,499	0,003	0,133	0,631	-0,001	67,987	11,291	0,007	0,131	0,616	-0,013	73,194	15,057
-0,001	0,007	0,293	0,29	-0,002	84,619	22,222	0,011	0,29	0,276	-0,011	87,978	31,064
0,997	0,006	0,267	1,263	-0,001	67,806	11,182	0,015	0,262	1,232	-0,025	73,159	15,048

See table 4 for details.

B.1.2 Drifting natural rate of interest

Productivity growth

We assume that the rate of change in productivity follows a bounded unit root process, i.e., $\log(\Xi_{t+1}) = \xi_{t+1} = \xi_t + \psi_{t+1}$ and

$$\xi_{t+1} \in [\xi_L, \xi_H]$$
$$\psi_{t+1} = \sigma_\psi \varepsilon_{\psi,t+1}$$

where ψ_{t+1} denotes the rate of change in productivity growth, and $\varepsilon_{\psi,t+1}$ denotes a realization of the truncated standard normal distribution between $\varepsilon_{\psi,L}(t) = \frac{\xi_L - \xi_t}{\sigma_\psi}$ and $\varepsilon_{\psi,H}(t) = \frac{\xi_H - \xi_t}{\sigma_\psi}$.

System of equilibrium equations

The set of equilibrium conditions includes equations 11 to 12, 14 to 15, and 21 to 24.

$$\breve{r}_t^n = \bar{\delta}_t + \mathrm{E}_t(\psi_{t+1}) \tag{21}$$

$$\mathrm{E}_t(\psi_{t+1}) = \sigma_\psi \frac{\phi(\varepsilon_{\psi,L}(t) - \phi(\varepsilon_{\psi,H}(t)}{\Phi(\varepsilon_{\psi,H}(t) - \Phi(\varepsilon_{\psi,L}(t)} \tag{22}$$

$$\lambda_{x,t}(\breve{\imath}_t^m + \ln(1-\Delta^m) + \xi_t - \ln\beta) = 0 \tag{23}$$

$$\breve{\imath}_t^m \geq \log(\beta) - \xi_t - \log(1-\Delta^m) \tag{24}$$

where $\phi(.)$ and $\Phi(.)$ denote the pdf and the cdf of the standard normal distribution respectively.

Solution algorithm

The main changes with respect to the procedure described in section B.1.1 are twofold. First, there is an additional (exogenous) state variable. We use a grid defined by four N-vectors of evenly spaced points including the rate of productivity growth ξ. The range of values considered for ξ is $[\xi_L, \xi_H]$. Moreover, we set N=40.

Second, we use a combination of Gaussian quadratures to approximate expectation terms. Define $s_{+1} = (\lambda_{p,t}, \lambda_{x,t}, \bar{\delta}_{t+1}, \xi_{t+1})$ the vector of state variables at time t+1. We use the equivalence $\varepsilon_{\psi,t+1} = \frac{\varepsilon_{\psi,H}(t)-\varepsilon_{\psi,L}(t)}{2} y_{t+1} + \frac{\varepsilon_{\psi,H}(t)+\varepsilon_{\psi,L}(t)}{2}$, where y_{t+1} denotes a realization of the truncated standard normal between -1 and 1, to express expectation terms as follows:

$$\mathrm{E}_t[f^c(s_{+1})] = \int_{-\infty}^{\infty} \Big[\int_{-1}^{1} g^c(\varepsilon_{t+1}^{\bar{\delta}}, y_{t+1}) dy_{t+1} \Big] \exp(-\varepsilon_{t+1}^{\bar{\delta}2}) d\varepsilon_{\delta,t+1} \tag{25}$$

Then, we use a 20 node Gauss-Legendre (GL) quadrature to approximate the integral in square brackets, and a 9 node Gauss-Hermite quadrature to approximate the first integral.

Accuracy

Table 8 reports simulation moments for various calibrations along with the different measures of accuracy. We observe that, in the case where the natural rate is equal to 1% on average, the approximation error for the output gap amounts to 2 basis points on average, which is one order of magnitude higher than the mean of this variable.

Table 8: Drifting r_t^n and optimal commitment: simulation moments and accuracy indicators

r^n	$\breve{x}$	π	max$[e^x]$	E$[e^x]$	max$[e^\pi]$	E$[e^\pi]$
3,503	0	0,001	0,112	0,001	0,008	0
1,003	0,001	0,082	0,405	0,019	0,029	0,001

See table 4 for details.

Robustness

Adaptative grid

We use the adaptative grid described in section B.1.1 extended to include possible realizations of the rate of productivity growth ξ_t. Given that the unconditional distribution of ξ_t looks like a uniform distribution, we keep a vector of N evenly spaced points for this dimension.

Table 9 and 10 report, respectively, the accuracy indicators and simulation moments obtained when using the adaptative grid. The adaptative grid is effective in reducing measurement errors. For the output gap: when the long run natural rate is 1% on average, the mean approximation error is one order of magnitude lower; the maximum approximation error is more than twice lower when using the adaptative grid. Regarding simulation moments, they are essentially unchanged. From this, we draw the same conclusion than in section B.1.1: a denser grid would certainly improve the accuracy of the solution, but it would not substantially affect our results.

Table 9: Drifting r_t^n and optimal commitment: Robustness with respect to the grid, accuracy indicators

	Evenly spaced grid				Adaptative grid			
r^n	max$[e^x]$	E$[e^x]$	max$[e^\pi]$	E$[e^\pi]$	max$[e^x]$	E$[e^x]$	max$[e^\pi]$	E$[e^\pi]$
3,503	0,112	0,001	0,008	0	0,053	0	0,005	0
1,003	0,405	0,019	0,029	0,001	0,159	0,004	0,011	0

See table 4 for details.

Table 10: Drifting r_t^n and optimal commitment: Robustness with respect to the grid, simulation moments

	Evenly spaced grid						Adaptative grid					
r^n	$\breve{x}$	π	i	$r-r^n$	ELB freq	ELB dur	$\breve{x}$	π	i	$r-r^n$	ELB freq	ELB dur
3,503	0	0,001	3,504	-0,001	5,228	2,599	0	0,001	3,503	-0,001	5,844	2,932
1,003	0,001	0,082	1,086	0,001	49,346	7,531	0,002	0,077	1,078	-0,002	52,887	8,808

See table 4 for details.

Alternative solution algorithm

Table 11 reports the simulation moments obtained under the alternative solution algorithm when using the adaptative grid. Overall, we observe that the results do not change substantially. However, the results suggest that, if anything, the baseline method may overestimate slightly the policy rate and the real rate gap, and underestimate slightly the ELB frequency and the average duration of an ELB episode.

Table 11: Drifting r_t^n and optimal commitment: Robustness with respect to the solution algorithm, simulation moments

	Baseline algorithm						Alternative algorithm					
r^n	$\breve{x}$	π	i	$r-r^n$	ELB freq	ELB dur	$\breve{x}$	π	i	$r-r^n$	ELB freq	ELB dur
3,503	0	0,001	3,504	-0,001	5,228	2,599	0,001	0,001	3,502	-0,002	6,691	2,953
1,003	0,001	0,082	1,086	0,001	49,346	7,531	0,008	0,079	1,062	-0,019	56,976	10,119

See table 4 for details.

B.2 Comparison of unconditional moments

Table 12 shows the unconditional moments of all the variables in the full model, and in two version of the model without time-variation in the long-run natural rate. Consistently with the small variance of shocks to the rate of growth of productivity, unconditional outcomes are relatively similar in all these model variants.

Table 12: Sample moments when $\mathrm{E}[r^n]$=1%

	Constant $(r^n)^L$	Integrated $(r^n)^L$	Persistent $(r^n)^L$
$\breve{\Delta}^m$	$\frac{3.2}{400}$	$\frac{3.2}{400}$	$\frac{3.2}{400}$
r^n	0,999	1,003	1,009
x	0,001	0,001	-0,003
π	0,058	0,082	0,115
i	1,057	1,086	1,092
RR spread	0	0,001	-0,033
ELB frequency (x100)	49,867	49,346	55,354
ELB duration (quarters)	7,185	7,531	10,799

This table reports the sample mean of key macro aggregates (in annualized terms), together with the ELB frequency and the average duration of an ELB episode. The columns indicate the stochastic process followed by the natural rate of interest: stationary with high frequency component only (column 2); integrated with both a high and a low frequency component (column 3); stationary with both a high and a low frequency component (column 4). In the latter case, the low frequency component of the natural rate follows a very persistent AR(1) process instead of the integrated process described in the main text: $\Xi_t = \Psi \cdot \Psi_t$ where $\psi = \log(\Psi) = 2\%$ annually, $\mathrm{E}_t[\log \Psi_{t+1}] = \bar{\psi}_t = \rho_\psi \bar{\psi}_{t-1} + \sigma_\psi \varepsilon_t$, and $\rho_\psi = 0.99$.

B.3 Additional impulse responses

Figure 7: Simulations of a sequence of permanent shocks under optimal policy

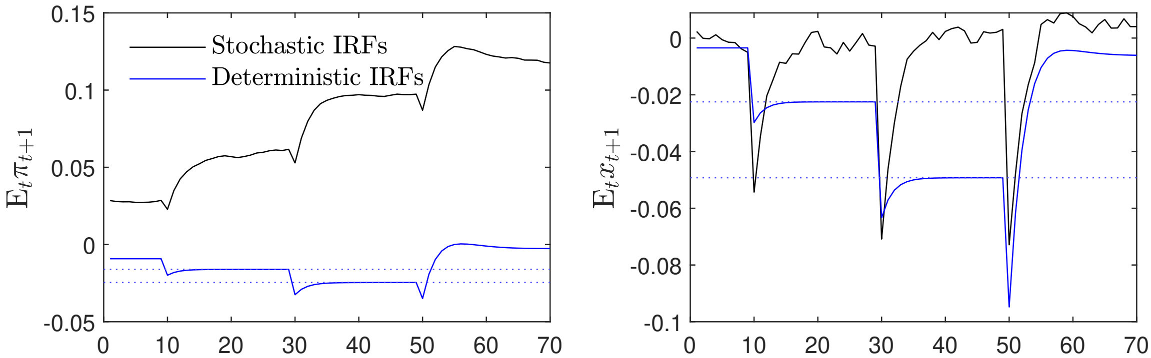

C Price level targeting

The results in this section are based on a slightly different model calibration.

C.1 The optimal price level targeting rule

Define the gap adjusted price level (GAPL) $\tilde{P}_t \equiv p_t + \frac{\lambda}{\kappa} x_t$. Following Eggertsson and Woodford 2003, the optimal gap adjusted price level target can be written as

$$P_t^* = P_{t-1}^* + \beta^{-1}(1 - \Delta^m)(1 + \kappa)\Delta_{t-1}^{\tilde{P}} - \beta^{-1}(1 - \Delta^m)\Delta_{t-2}^{\tilde{P}}$$

C.2 Numerical solution of the model under a *simple* price level targeting rule

This appendix presents the numerical solution of the model under the simple price level targeting rule when the rate of change in productivity growth follows a bounded unit root process.

A simple price level targeting rule

We assume that the central bank has a gap adjusted (log) price level (GAPL) target defined by

$$p_t + \frac{\lambda}{\kappa} x_t = P_t^*$$

where the time-varying target P_t^* follows an exogenous, deterministic trend π^* such that

$$P_t^* = P_{t-1}^* + \pi^*$$

It can be shown that, in absence of the ELB constraint, the central bank would stimulate output as much as necessary to reach the target at any point in time. As a result, the (log) price level would grow at the same pace as the GAPL target. Based on this, we can rewrite the rule in terms of the detrended (log) price level $\bar{p}_t \equiv p_t - P_t^*$

$$\bar{p}_t + \frac{\lambda}{\kappa} x_t = 0$$

System of equilibrium equations

Substituting the rate of inflation by $\bar{p}_t - \bar{p}_{t-1} + \pi^*$ in both the Euler equation and the Phillips curve, we obtain the following system of equilibrium equations.

$$\bar{p}_t = \frac{1}{1+\beta}\Big[\bar{p}_{t-1} + \kappa x_t + \beta \mathrm{E}_t \bar{p}_{t+1} + (\beta - 1)\pi^*\Big] \tag{26}$$

$$x_t = \Big(1 - \Delta^m\Big)\Big[\mathrm{E}_t x_{t+1} - \Big(\breve{i}_t^m - (\pi^* + \mathrm{E}_t \bar{p}_{t+1} - \bar{p}_t) - \breve{r}_t^n\Big)\Big] \tag{27}$$

$$0 = \Big(\breve{i}_t^m - \log(\beta) + \xi_t + \log\left(1 - \Delta^m\right)\Big)\Big(\bar{p}_t + \frac{\lambda}{\kappa} x_t\Big) \tag{28}$$

$$\breve{i}_t^m \geq \log(\beta) - \xi_t - \log\left(1 - \Delta^m\right) \tag{29}$$

$$\bar{p}_t + \frac{\lambda}{\kappa} x_t \leq 0 \tag{30}$$

Solution algorithm

The main changes with respect to the procedure described in section B.1 are twofold. First, there are only three predetermined variables $s_t = (\bar{p}_{t-1}, \bar{\delta}_t, \xi_t)$. We form a grid defined by three N-vectors of evenly spaced points, with N=50. Second, we use a fixed point iteration on $\bar{p}(s)$ and on $x(s)$ for solving the system on the grid. The solution algorithm proceeds in four steps.

Step 1: Choose an initial $\bar{p}_0(s)$ and $x_0(s)$, and a tolerance level τ

Step 2: Iteration j. For each possible state s, given $\bar{p}_{j-1}(s)$ and $x_{j-1}(s)$, compute $\mathrm{E}_s \bar{p}_{j-1}(s_{+1})$ and $\mathrm{E}_s x_{j-1}(s_{+1})$, and retrieve the value of the detrended log price level and of the output gap implied by both the Phillips curve and the policy rule:

$$x_j(s) = -\frac{\kappa}{\lambda}\bar{p}_{j-1}(s) \tag{31}$$

$$\bar{p}_j(s) = \frac{1}{1+\beta}[\bar{p}_{-1} + \kappa x_{j-1}(s) + \beta \mathrm{E}_s \bar{p}_{j-1}(s') + (\beta - 1)\pi^*] \tag{32}$$

Step 3: Adjust if this allocation does not satisfy the ELB constraint. Let $\bar{\imath}^m(\xi)$ denote the ELB on the policy rate. If $x_j(s) > \Big(1 - \Delta^m\Big)\Big[\mathrm{E}_s x_{j-1} s' - \Big(\bar{\imath}^m(\xi) - (\pi^* + \mathrm{E}_s \bar{p}_{j-1}(s') - \bar{p}_{j-1}(s)\Big) - \bar{\delta} - \mathrm{E}_s(\psi')\Big)\Big]$, retrieve

$$x_j(s) = \Big(1 - \Delta^m\Big)\Big[\mathrm{E}_s x_{j-1} s' - \Big(\bar{\imath}^m(\xi) - (\pi^* + \mathrm{E}_s \bar{p}_{j-1}(s') - \bar{p}_{j-1}(s)\Big) - \bar{\delta} - \mathrm{E}_s(\psi')\Big)\Big]$$

Step 4: Let $e_j^\pi(s) = |\bar{p}_j(s) - \bar{p}_{j-1}(s)|$, $e_j^x(s) = |x_j(s) - x_{j-1}(s)|$ and $e_j(s) = e_j^\pi(s) + e_j^x(s)$ denote different measures of approximation error. Stop if $\sum_s e_j(s) < \tau$. Otherwise, update the guess, and repeat step 2.

Accuracy

The main change with respect to the procedure described in section B.1 is the approximation error for the rate of inflation. We use the difference between the interpolated value of the detrended log price level $\bar{p}_t^*$ and the value implied by the equations $\bar{p}_t^{\text{IMP}}$

$$e_t^\pi \equiv \left| \bar{p}_t^{\text{IMP}} - \bar{p}_t^* \right| \cdot 400 \tag{33}$$

Table 13 reports approximation errors for various calibrations of the average value of the long run natural rate of interest (first column) and of the GAPL target growth (second column). For the rate of inflation, both the mean and the maximum approximation error are negligeable. For the output gap, the mean approximation error is negligeable, while the maximum approximation error amounts up to around 8.5 basis points, which is one order of magnitude higher than the average value of this variable.

Table 13: Drifting r_t^n and price level targeting: simulation moments and accuracy indicators

r^n	π^*	$\breve{x}$	π	max$[e^x]$	E$[e^x]$	max$[e^\pi]$	E$[e^\pi]$
3,503	0	0	0	0,038	0	0,006	0
1,003	0	0,003	0	0,085	0,004	0,003	0
1,003	0,15	0,006	0,15	0,078	0,003	0,003	0

See table 4 for details.

Robustness check based on an adaptative grid

Table 14 and 15 report, respectively, the approximation errors and the simulation moments when using an adaptative grid instead of an evenly-spaced grid. For $\bar{p}_{t-1}$, given that the distribution is asymmetric and concentrated near its maximum value, we place the points using a multiplicative sequence of the form: $\bar{p}_k = (1-\delta)\bar{p}_{k-1}$ with $0 < \delta < 1$. Figure 8 provides an illustrative example of the grid based on the solution obtained when the natural rate of interest is equal to 3.5% on average and trend inflation is equal to zero.

Overall, we observe that the approximation errors do not diminish substantially when using the adaptative grid. This suggests that using a denser grid would not improve the accuracy of the solution. Moreover, the simulation moments are essentially unchanged.

Figure 8: Drifting r_t^n and price level targeting: An adaptative grid

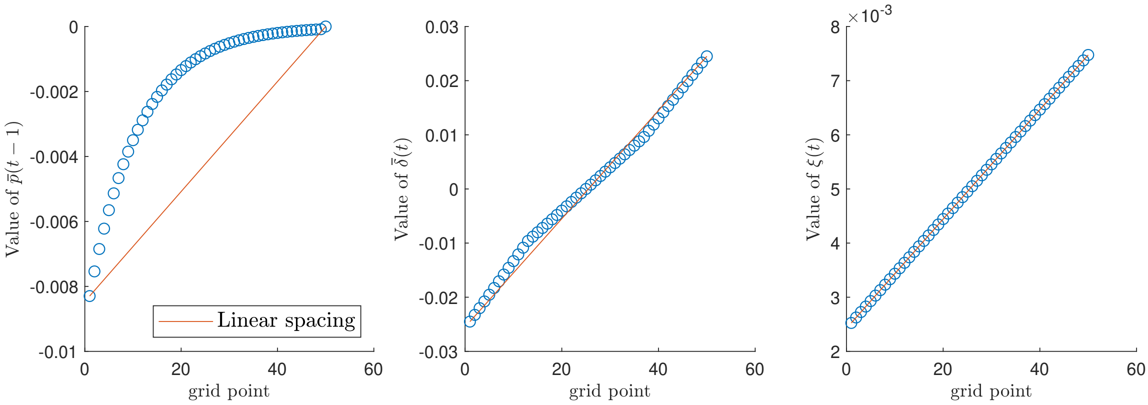

Table 14: Drifting r_t^n and price level targeting: Robustness with respect to the grid, accuracy indicators

		Evenly spaced grid				Adaptative grid			
r^n	π^*	max[e^x]	E[e^x]	max[e^π]	E[e^π]	max[e^x]	E[e^x]	max[e^π]	E[e^π]
3,503	0	0,038	0	0,006	0	0,043	0	0,006	0
1,003	0	0,085	0,004	0,003	0	0,062	0,002	0,003	0
1,003	0,15	0,078	0,003	0,003	0	0,06	0,001	0,003	0

See table 4 for details.

Table 15: Drifting r_t^n and price level targeting: Robustness with respect to the grid, simulation moments

		Evenly spaced grid						Adaptative grid					
r^n	π^*	$\breve{x}$	π	i	$r - r^n$	ELB freq	ELB dur	$\breve{x}$	π	i	$r - r^n$	ELB freq	ELB du
3,503	0	0	0	3,502	-0,001	5,56	2,28	0	0	3,502	-0,001	5,78	2,297
1,003	0	0,003	0	0,996	-0,007	54,707	7,852	0,002	0	0,999	-0,004	56,584	8,189
1,003	0,15	0,006	0,15	1,147	-0,007	49,35	6,79	0,005	0,15	1,15	-0,004	50,979	7,077

See table 4 for details.

C.3 The optimal *simple* price level targeting rule

This appendix describes the method used to identify the optimal rule within the class of simple price level targeting rules we are focusing on.

Section A.2 shows that, up to second order, aggregate welfare can be approximated as

$$\int_j \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t U_t(C_{j,t}, H_{j,k,t}, \frac{M_{j,t}}{P_t}) dj \approx -\frac{1}{2}\frac{\alpha\theta(1+\theta\omega)}{(1-\alpha)(1-\alpha\beta)} \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t(\pi_t^2 + \lambda x_t^2) + tip$$

We solved the model for various calibrations of π^* going from 0 to 25% in annualised terms. Then, as a measure of efficiency, we considered the unconditional welfare loss function [16]. We computed it based on the two different methods described below. Figure 9 reports the results and shows that the optimal GAPL target growth, i.e., the value that minimizes the unconditional welfare loss function, is the same regardless of the method used.

Method 1: Recursion

Define $\mathcal{E}_t = \mathrm{E}_t \sum_{j=t}^{\infty} \beta^{j-t}(\pi_j^2 + \lambda x_j^2)$, the expected (unweighted) welfare loss conditional on information available at time t, and $\mathcal{W}_t = \frac{1}{2}\frac{\alpha\theta(1+\theta\omega)}{(1-\alpha)(1-\alpha\beta)}\mathcal{E}_t$, the expected

[16] See for example Adam and Billi (2006) and Andrade et al. (2019) among others

welfare loss conditional on information available at time t. Then:

$$\begin{aligned}
\mathcal{E}_t = \quad & \pi_t^2 + \lambda x_t^2 + \mathrm{E}_t \sum_{j=t+1}^{\infty} \beta^{j-t}(\pi_j^2 + \lambda x_j^2) \\
= \quad & \pi_t^2 + \lambda x_t^2 + \beta \mathrm{E}_t \sum_{j=t+1}^{\infty} \beta^{j-(t+1)}(\pi_j^2 + \lambda x_j^2) \\
= \quad & \pi_t^2 + \lambda x_t^2 + \beta \mathrm{E}_t \mathrm{E}_{t+1} \sum_{j=t+1}^{\infty} \beta^{j-(t+1)}(\pi_j^2 + \lambda x_j^2) \\
= \quad & \pi_t^2 + \lambda x_t^2 + \beta \mathrm{E}_t \mathcal{E}_{t+1}
\end{aligned}$$

We compute this object numerically using a fixed point method. Then, we compute the sample average:

$$\bar{\mathcal{W}} = \frac{1}{N} \sum_n \mathcal{W}_n \tag{34}$$

over $N = 10000$ possible states of the economy.

Method 2: Simulation

Alternatively, we approximate the unconditional welfare loss by using the sample average of $L_t = \frac{1}{2}\frac{\alpha\theta(1+\theta\omega)}{(1-\alpha)(1-\alpha\beta)} \sum_{j=t}^{T} \beta^{j-t}(\pi_j^2 + \lambda x_j^2)$:

$$\bar{L} = \frac{1}{N} \sum_n L_n \tag{35}$$

We compute this object across N=10000 simulations, each T=1000 periods long.

Figure 9: Unconditional welfare loss under price level targeting

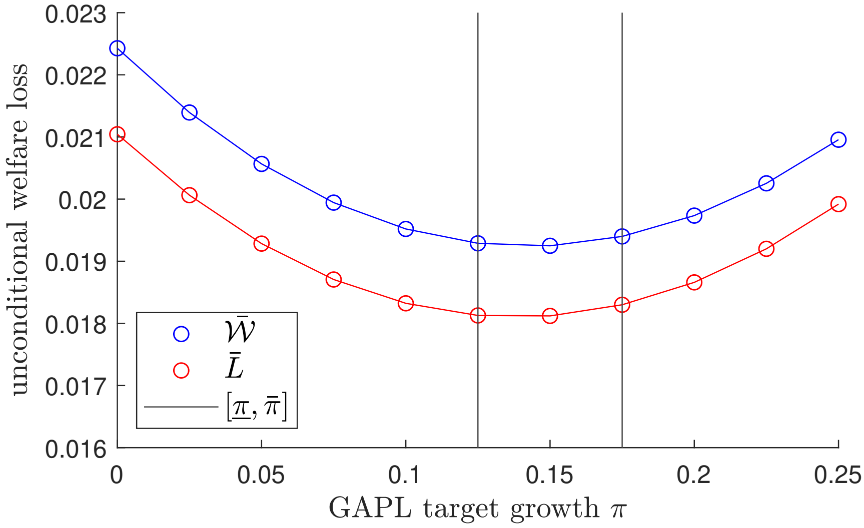

C.4 Impulse response functions to shocks under various policy rules

Figure 10 compares outcomes between optimal policy and the two price level targeting rules in a stochastic simulation.

Figure 10: Simulations of a sequence of permanent shocks under various policy rules

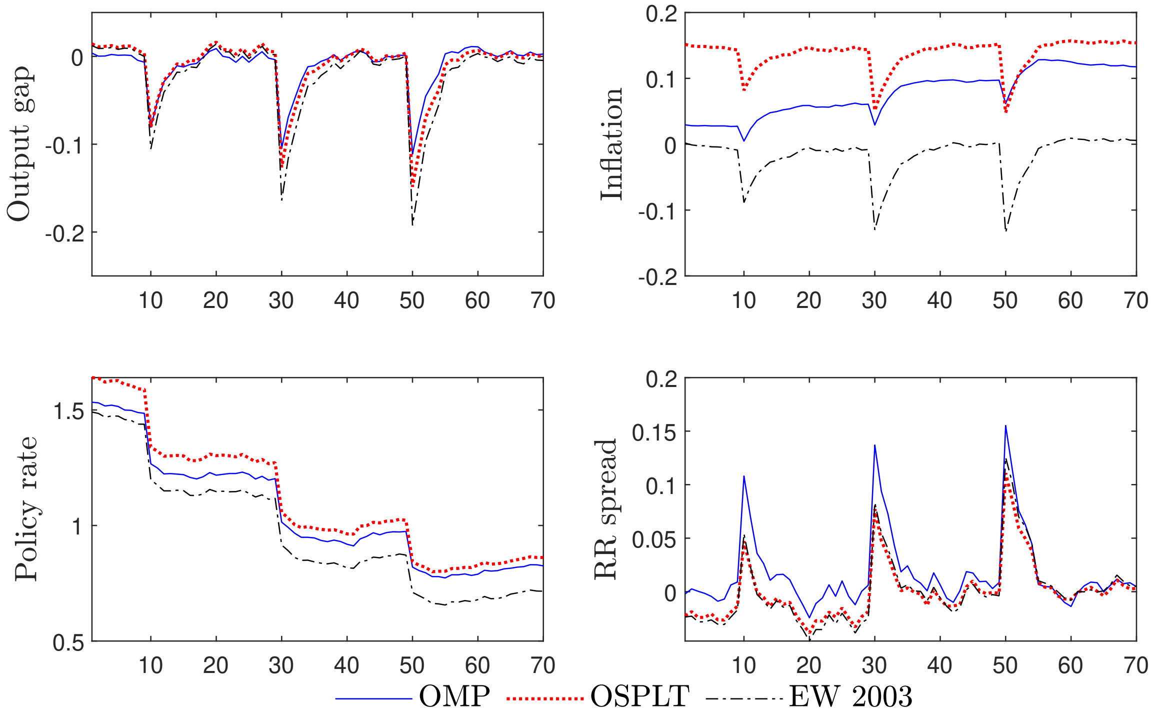

Figure 11: Impulse responses to a positive transitory shock under various policy rules starting from $(r_t^n)^L = 0.5\%$

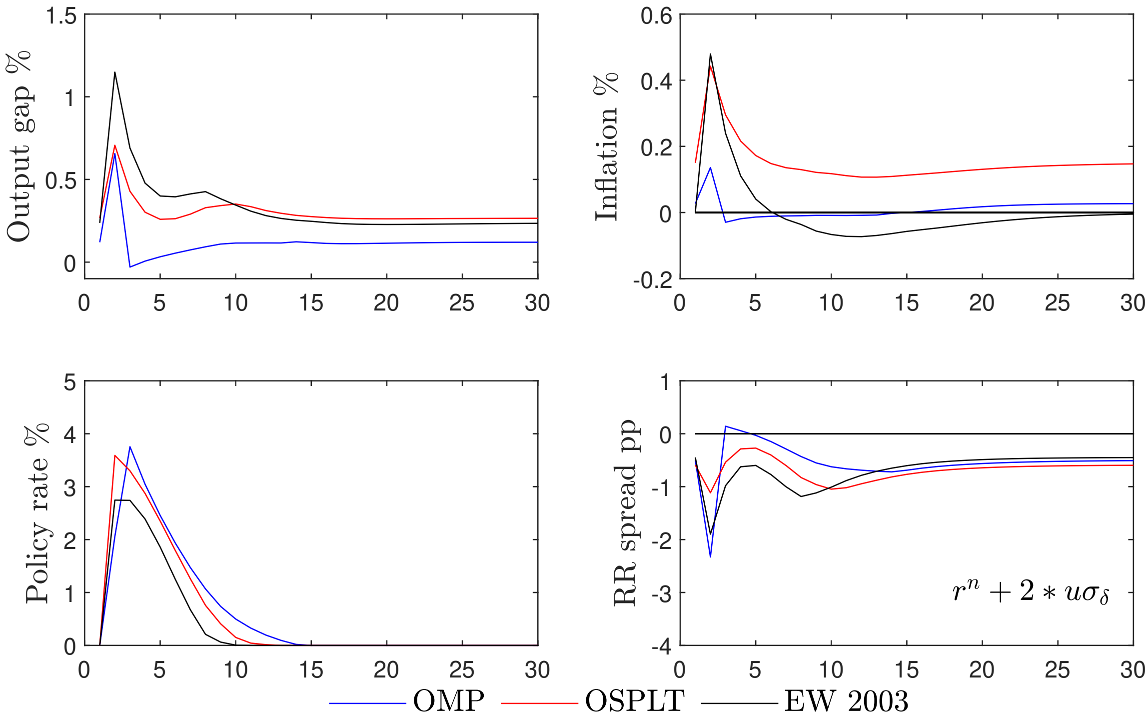

Acknowledgements

We wish to thank Mathieu Boullot, Nuno Coimbra, Wouter Den Haan, Stéphane Dupraz, Hubert Kempf, Julien Matheron, Jean-Stéphane Mésonnier, Xavier Ragot, Stephanie Schmitt-Grohé, Lars Svensson, Karl Walentin and Pablo Winant for useful comments and suggestions. We also thank seminar participants at the European Central Bank, Ghent University and the Riksbank, as well as participants in the CEF2022 conference, the EEA2022 Congress, the 3rd ILMA workshop 2021 and the 4th Warsaw Money-Macro-Finance conference.
The views expressed here are personal and do not necessarily reflect those of the European Central Bank or the Eurosystem.

Sandra Daudignon
Ghent University, Ghent, Belgium; email: sandra.daudignon@ugent.be

Oreste Tristani
European Central Bank, Frankfurt am Main, Germany; Centre for Economic Policy Research, London, United Kingdom; email: oreste.tristani@ecb.europa.eu

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

PDF ISBN 978-92-899-5984-1 ISSN 1725-2806 doi:10.2866/35939 QB-AR-23-025-EN-N