WORKING PAPER

MONETARY POLICY AND THE DRIFTING NATURAL RATE OF INTEREST

Sandra Daudignon
Oreste Tristani

Department of Economics

GHENT
UNIVERSITY

Monetary policy and the drifting natural rate of interest[*]

Sandra Daudignon[†]　　Oreste Tristani[‡]
Ghent University　　ECB and CEPR

November 27, 2022

Abstract

Empirical analyses starting from Laubach and Williams (2003) find that the natural rate of interest is not constant in the long-run. This paper studies the optimal response to stochastic changes of the long-run natural rate in a suitably modified version of the new Keynesian model. We show that, because of the zero lower bound (ZLB) on nominal interest rates, movements towards zero of the long-run natural rate cause an increasingly large downward bias in expectations. To offset this bias, the central bank should aim to keep the real interest rate systematically below the long-run natural rate, as long as policy is not constrained by the ZLB. The neutral rate – the level of the policy rate consistent with stable inflation and the natural rate at its long-run level – will be lower than the long-run natural rate. This is the case both under optimal policy, and under a price level targeting rule. In the latter case, the neutral rate is equal to zero as soon as the long-run natural rate falls below 1%.

Keywords: nonlinear optimal policy, zero lower bound, commitment, liquidity trap, New Keynesian, natural rate of interest.

JEL Codes: C63, E31, E52.

1 Introduction

Recent empirical research, starting from Laubach and Williams (2003), has found that the long-run *natural* interest rate – the real interest rate consistent with output at its long-run equilibrium and stable inflation – has a time-varying low-frequency component. Econometric estimates suggest that the long-run natural rate has declined over the past few decades (e.g. Laubach and Williams (2016)) and may have hovered around zero both in the U.S. and in the euro area over the 2010s (Holston, Laubach and Williams, 2017, Fiorentini et al. (2018)). The future level of the long-run natural rate is uncertain.

These findings raise questions for the conduct of monetary policy, due to the zero lower bound (ZLB) constraint on nominal interest rates.[1] Many studies have analysed the optimal monetary policy response to *temporary* reductions of the natural rate in models where the natural rate is constant in the long-run. But how should policy respond to a *permanent* reduction (or increase) of the natural interest rate – that is, to changes in its expected long-run level? Can optimal policy be implemented, at least approximately, through the same simple rules that have been shown to work well in models with a constant, long-run natural rate?

We provide answers to these questions by analysing optimal policy under commitment in the context of a modified version of the benchmark new Keynesian model. We modify the model along two dimensions, to make it consistent with the aforementioned empirical findings.

First, to account for variation in the long-run natural rate, we allow for permanent shocks to the rate of *growth* of productivity – not only to the *level* of productivity, as is common, for example, in Altig et al. (2011), Christiano et al. (2014), Christoffel et al. (2008). Declines (increases) in the long-run rate of growth of productivity will induce declines (increases) in the long-run natural rate of interest. A simple way of allowing for permanent effects of shocks would be to assume that productivity growth follows a random walk. This assumption would, however, imply that long-run productivity growth can reach arbitrarily large, positive or negative, values. Historically, a moving average of productivity growth has instead fluctuated within a relatively narrow, positive range in the period following World War II. To rule out the possibility that long-run productivity growth reaches implausible levels, we impose finite upper and lower boundaries on the support of permanent shocks. As a result, the long-run natural rate will also vary between predefined upper and lower

---

[1] 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.

bounds.

Our second modification of the new Keynesian model aims to capture the evidence that the natural rate may have reached near-zero levels in recent years. Accounting for this empirical features purely through technological factors would require productivity growth to be negative in the long-run – an implausible assumption. Motivated by the findings in Del Negro et al. (2017),[2] we allow for a liquidity premium on nominal bonds. More specifically, we follow Michaillat and Saez (2021) in postulating that households derive utility from their holdings of government bonds relative to everyone else. As a result, government bonds will incorporate a (negative) convenience yield, which will allow us to rationalise very low values for the long run natural rate of interest. This specification also has the advantage of reducing the power of forward guidance in our model – see Michaillat and Saez (2021).

As in the rest of the new Keynesian literature, we will also allow for transitory shocks to the natural rate of interest. We calibrate the standard deviations of the transitory and permanent components to match the empirical results in Fiorentini et al. (2018) and Adam and Billi (2006). We solve our model using projection methods.

Our main, novel result is that, as long as the ZLB does not bind, monetary policy ought to be *over-expansionary* at low levels of the long-run natural rate, compared to optimal policy prescriptions that abstract from the ZLB. More specifically: 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. Only when the ZLB binds, so that this policy option becomes infeasible, should the central bank promise future inflation in response to further reductions in the long-run natural rate.

To understand this result, it is useful to start from the observation that very low levels of the long-run natural rate, coupled with the ZLB constraint, reduce the scope for monetary policy easing in the face of adverse, transitory shocks. The lack of easing space tends to impart a downward bias on output and inflation expectations – a bias that will be larger, the closer to zero the long-run natural rate. It is to offset this bias in expectations that, in the absence of shocks, the central bank should aim to maintain a lower policy rate than the long-run natural rate, as long as this is feasible (i.e. as long as the ZLB does not bind). In other words, the "neutral rate" – i.e. the level of the policy rate consistent with stable inflation and the natural rate at its long-run level – will be lower than the long-run natural rate. The spread between the two interest rates, and the ensuing degree of monetary accommodation, should be tailored to produce an output gap that is positive and large enough, to offset the deflationary bias in inflation expectations. This approach would ensure that price

---

[2]See also Caballero and Farhi (2018) and Krishnamurthy and Vissing-Jorgensen (2012).

3

stability is eventually restored – i.e. that inflation returns to zero – in the absence of additional shocks.

The gap between the neutral rate and the long-run natural rate will increase – hence the degree of policy stimulus should be larger – after any reduction of the long-run natural rate. A lower long-run natural rate further reduces the easing space available for monetary policy. So, the downward bias in expectations becomes larger, and the degree of policy accommodation should increase in a commensurate fashion. We also show that, under optimal policy, the policy rate should adjust slowly to the new, neutral rate, even if the long-run natural rate moves with a discrete jump. This is in contrast to optimal policy prescriptions abstracting from the ZLB constraint, according to which the policy rate should "track", or follow one-to-one, movements in the natural rate.

The approach described above causes the neutral rate to reach the zero level earlier than the long-run natural rate, when the latter falls. The difference between these two interest rates is quantitatively significant under optimal policy. It is even larger when monetary policy is implemented through a simple rule. The reason is that simple rules are less sophisticated than optimal policy in their ability to manage expectations in reaction to shocks. They must therefore rely more heavily on a conventional policy stimulus.

Amongst simple rules, we specifically analyse the performance of variants of the price level targeting rule put forward in Eggertsson and Woodford (2003), because they have the advantage of not requiring knowledge of the long-run natural rate. Conditional on this type of rule, the neutral rate would be equal to zero as soon as the long-run natural rate falls below 1% – the realistic case in the recent past, according to empirical estimates.

Once the policy rate hits the ZLB, the central bank loses any ability to cut policy rates further in reaction to any new, negative shocks. As also shown in previous studies, its only option is to be over-expansionary in a different way, i.e. promise to create more positive inflation (and a positive output gap) as soon as possible in reaction to positive shocks. The promise will be reflected in expectations and it will stimulate the economy through the ensuing reduction of the real rate. We show that, at very low levels of the long-run natural rate, this approach cannot be tailored to ensure that actual inflation eventually returns to zero in the absence of additional shocks. Due to the forward looking nature of both the Euler equation and the Phillips curve, the promise to create inflation in reaction to future shocks will also pull up *current* inflation and output. Actual inflation will thus be positive even after shocks have disappeared.

In a stochastic simulation of all the shocks in the model, the over-expansionary

attitude of the central bank is consequential for economic outcomes. Average inflation will become more and more positive, the more the long-run natural rate falls towards zero. However, the central bank will optimally tolerate only a small increase in inflation, because this is very costly for welfare.

Our paper contributes to the literature on the consequences of the ZLB for optimal monetary policy – see Dominquez et al. (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 natural rate in the long-run (the steady-state). They also rely on calibrations consistent with a relatively high level of the nominal interest rate (typically 3.5%). The promise to maintain policy rate 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. (2022) 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 argues 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), 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. Once again, in contrast to all these papers, we allow for time-variation in the long-run natural rate. 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. The final section 7 offers some concluding remarks.

2 The empirical evidence on the natural rate of interest

This section briefly summarises the empirical evidence that motivates our theoretical model.

It is useful to note that different literatures rely on different notions of natural rate of interest. The theoretical literature defines the natural rate as the real interest level that would be observed at any point in time *in the absence of nominal rigidities* (Woodford, 2003).[3] It is a magnitude that can vary at high frequency as a result of *transitory* preference and technology shocks (see, for example, Edge et al. (2008), Justiniano and Primiceri (2010)). Most of the empirical literature in reduced form has instead defined the natural rate as the value of the real short-term interest rate *expected to prevail in the distant future* (Laubach and Williams, 2003; Hamilton et al. (2016); and many others). By construction, this magnitude will not be affected by transitory shocks, but only by shocks with permanent effects. In a standard theoretical model, the empirical notion would be constant over time and it would coincide with the steady state real rate.

In the rest of our paper we refer to the empirical notion as *long-run natural rate*. We simply write natural rate to refer to the theoretical notion. Thus, in a model where the long-run natural rate is time-varying (due to permanent shocks), the natural rate will be affected by both transitory and permanent shocks.

The very low inflation rates observed over the 2010s, while policy interest rates remained close to zero, was considered as suggestive evidence of a very low level of the natural rate. This generated renewed interest in measuring the long-run natural rate empirically – see, for example, Hamilton et al. (2016), Holston et al. (2017), Fiorentini et al. (2018). In these econometric studies the long-run natural rate is typically modelled as a random walk process, since it does not appear to converge to a constant value over time. This piece of evidence is inconsistent with standard assumptions in theoretical models. In Section 3, we will build a variant of the new Keynesian model which allows for plausible variation in the long-run natural rate.

---

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

6

All the aforementioned empirical estimates find that the long-run natural rate has fallen in recent decades. For example, Holston, Laubach and Williams (2017) finds that, in 2016, the long-run natural rate 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) 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) forecasted the real rate in the U.S. to asymptote to a value slightly lower than 0.5% by 2021.

Nevertheless, there is considerable uncertainty surrounding the evolution of the natural rate in the long-run. Some recent estimates for the U.S. suggest that the long-run natural rate may be again in safely positive territory.[4] Other papers forecast the natural rate to reach a trough of 0.38% by 2030 and then rise again to 1% in the very long run.[5] Uncertainty over the long-run natural rate will also be a feature of our model.

Empirical models disagree on the exact determinants of the time-variation in the long-run natural rate. The approach pioneered in Laubach and Williams (2003) emphasises time-variation in trend productivity growth. Hamilton et al. (2016), however, 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)), 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)).

In our model we do not aim to capture all possible determinants of variation in the long-run natural rate. In line with Laubach and Williams (2006), we will attribute fluctuations in the long-run natural rate to permanent shocks to productivity growth, which can be more easily calibrated based on the TFP-data from Fernald (2014). As long as the natural rate is independent of monetary policy, the precise structural determinants of the natural rate are, however, not crucial to our optimal policy results. In our policy analysis, we will simply focus on temporary vs. permanent natural rate shocks, abstracting from their structural determinants.

---

[4]See IMF (2022), p.21.

[5]See Platzer and Peruffo (2022).

3 The model

The model we employ in our analysis is standard, so we only highlight in this section 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. We conclude the section with a description of the optimal policy problem.

3.1 Distinguishing features of the model

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, intermediate goods. 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 Fisher (2015) and 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

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

where the function $\upsilon(\cdot)$ is increasing and concave and $\bar{C}_t$ is a preference 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.

As demonstrated in Michaillat and Saez (2021), this specification leads to a form of discounting in the linearised Euler equation of the model – see equation (1) below. This feature can help mitigate the so-called forward guidance puzzle (Giannoni et al. (2015)).

8

We also assume that economy-wide productivity evolves around a stochastic trend $\overline{A}_t$. More specifically, within large but finite boundaries ξ^H and ξ^L, the rate of growth of productivity is such that $\bar{A}_{t+1}/\bar{A}_t = \Xi_{t+1}$ and

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

where $\xi_t = \log \Xi_t$ and $\mathcal{TN}(.)$ denotes the truncated standard normal distribution. This assumption on the distribution of ε_t^ψ ensures that the productivity growth rate is bounded within the $[\xi_L; \xi_H]$ interval. Under this assumption, the expected rate of change of productivity growth will be zero over most of the state space. Only when the productivity growth rate is close to either ξ^H or ξ^L, will $\sigma_\psi \mathrm{E}_t \varepsilon_{t+1}^\psi$ depart from zero.

Following a standard approach, the Euler equation of the model can be linearized around the non-stochastic steady state with zero-inflation.[6] As in all models with productivity growth, real quantities first need to be detrended by the level of productivity $\bar{A}_t$. In our case, however, the rate of growth of productivity does not have a unique non-stochastic steady state. We therefore detrend interest rates by the productivity growth rate Ξ_t before linearization. We obtain

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

where $x_t \equiv \tilde{c}_t - \tilde{c}_t^n$ is the detrended logarithm of output in deviation from its natural level and $\breve{r}_t^n = \mathrm{E}_t \tilde{c}_{t+1}^n - \frac{1}{1-\Delta^m}\tilde{c}_t^n - \mathrm{E}_t \delta_{t+1} + \sigma_\psi \mathrm{E}_t \varepsilon_{t+1}^\psi$ is the natural rate of interest in deviation from its long-run level.

Note that equation (1) is very similar to the standard linearised Euler equation of the new Keynesian model apart from the 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] Such coefficient acts as a discount factor in the Euler equation and, as such, it will tend to mitigate the forward guidance puzzle for reasons discussed in McKay et al. (2016). As discussed above, output and the natural rate have different stochastic trends. Output will inherit the stochastic trend of the level of productivity, $\bar{A}_t$, while the nominal interest rate and the natural rate will both inherit the trend of the growth rate of productivity, Ξ_t. This is consistent with the empirical specification in Laubach and Williams (2003).

---

[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.

[7]$\Delta^m \equiv 1 - \breve{I}^m/\breve{I}$, where $\breve{I}$ and $\breve{I}^m$ are the gross, detrended return on a complete portfolio of state contingent bonds and on government bonds, respectively.

The supply side of the model can be summarised as usual by a Phillips curve

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} \tag{2}$$

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

The natural rate in deviation from its long-run level can be solved out explicitly as

$$\breve{r}_t^n = \bar{\delta}_t + \sigma_\psi E_t \varepsilon_{t+1}^\psi \tag{3}$$

where $\bar{\delta}_t = -\rho_\delta \delta_t$. In the rest of the paper, we do not use δ_t, but work directly with the derived process $\bar{\delta}_t = \rho_{\breve{r}^n} \bar{\delta}_{t-1} + \sigma_{\breve{r}^n} \varepsilon_t^{\bar{\delta}}$.

The long-run natural rate, which we denote as $(r_t^n)^L$, can be written as

$$(r_t^n)^L = -\ln \beta + \ln (1 - \Delta^m) + \xi_t, \tag{4}$$

which is not constant, but time-varying 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 the long-run natural rate. For given long-run natural rate, persistent, but stationary fluctuations will be induced by shocks $\bar{\delta}_t$.

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

$$\breve{i}_t^m \geq - (r_t^n)^L. \tag{5}$$

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 long-run natural rate. 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 the long run natural rate.

3.2   Optimal policy

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 $\lambda = \kappa/\theta$.

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 $\breve{\imath}_t^m > - \left(r_t^n\right)^L$ and $\lambda_{x,t} > 0$ when the nominal rate is at the ZLB. In the above equations $\lambda_{x,t}$ is the lagrange multiplier associated to equation (1) and it is also proportional to the multiplier on the ZLB constraint; $\lambda_{p,t}$ is the multiplier associated to the Phillips curve.

The equilibrium will solve the system given by equations (1), (2), (5), (6) and (7). Note that if the ZLB were not binding at any point in time, $\lambda_{x,t} = 0$ for any t and optimal policy would boil down to the conditions that hold in the standard new Keynesian model with constant, long-run natural rate and without convenience yield on government bonds

$$\pi_t = -\frac{\lambda}{\kappa} \left(x_t - x_{t-1}\right) \tag{8}$$

Hence permanent shocks to the rate of growth of productivity would *per se* not prevent the maintenance of price stability at all times. They only require that the long-run nominal interest rate tracks the long-run natural rate of interest.

4 Solution and calibration

This section provides an overview of the numerical method used to solve the model in Section 3. More details are provided in appendix B.1 and C.

4.1 Solution algorithm

We solve our model using projection techniques. The model has four state variables: two exogenous shocks and two lagged, Lagrange multipliers. Our baseline algorithm proceeds in four steps: select a grid of points in the state space $\mathcal{S} \in R^4$; solve the model on the grid; simulate the solution; update the grid. We perform a few iterations on these steps until the grid includes all possible states realized during the simulation.

The numerical solution essentially involves two difficulties. The first one is the kink(s) in the policy functions due to the occasionally binding ELB constraint on the policy interest rate. We choose a solution method which allows to take this specific type of non-linearity into account. More specifically, we use a *piecewise* linear interpolation of the policy functions to approximate them off the grid and a

fixed-point iteration algorithm to solve the model on the grid. We also use a dense grid to obtain a sufficiently accurate solution.

The second difficulty has to do with the dimension of the state space. For the solution to be sufficiently accurate, the grid needs to be dense. However, the running time for solving the model when using a dense state space increases exponentially with the number of state variables. We use a grid defined by four N-vectors of evenly-spaced points, where $N = 40$. Then, instead of increasing the number of grid points to obtain a more accurate solution, we use a grid of the same size but modified in the spirit of Maliar and Maliar (2015). The alternative grid delivers a more accurate solution and shows that our results are robust to a change of grid.[8]

The quadrature method used to approximate expectation terms also changes compared to the standard case in which exogenous variables follow an autoregressive process. The productivity growth rate in our model follows a unit root process with upper and lower boundaries. These boundaries will induce some reflecting behaviour in the natural rate process. Once trend productivity growth reaches its lower boundary, the long-run natural rate can only increase. Conversely, it can only fall once trend productivity growth reaches its upper boundary. Under this assumption, expectation terms take the form of a combination of two integrals, including one with a bounded support. We take this into account by combining both Gauss-Hermite quadrature and Gauss Legendre quadrature.

4.2   Calibration

The values of key parameters are reported in Table 1.

We set the boundaries, ξ_L and ξ_H at 1% and 3%, respectively. This is broadly consistent both with the long time-series for utilization-adjusted TFP growth 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 somewhat below 1% in the 1990s, before increasing again in the early 2000s. Over the whole 2010s, average TFP growth remains stable around the 1% mark. 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.

---

[8]We also check the robustness of the results to a change of solution algorithm.

Table 1: Calibration of structural parameters

Parameter	Definition	Value	Source
Specific			
γ	Coefficient of relative risk aversion	1	King et al. (1988)
σ_ψ	Standard deviation of shocks to ξ	0.108 (%, annualised)	Fiorentini et al. (2018)
$[\xi^L, \xi^H]$	Boundaries of ξ	[1,3] (%, annualised)	Fernald (2014)
Δ^m	Convenience yield	3.2 (%, annualised)	Del Negro et al. (2017)
β	Subjective discount factor	0.9945	Platzer & Peruffo (2022)
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

As shown in expression (4), the long-run natural rate will also be affected by the parameters Δ^m and β. For Δ^m, we follow Del Negro et al. (2017), which finds that the convenience yield increased markedly and persistently in recent decades, especially since the early 2000s. Taking into account estimation uncertainty, the estimates in that paper suggest that the convenience yield may be as large as 3.0% (in annualised terms).[9] In terms of our model specification, this implies $\Delta^m = 0.008$. We finally set $\beta = 0.9945$ to produce an unconditional mean of the long run natural rate equal to 1%. This is predicted to be the value which the natural rate may reach in the future in recent estimates (see for example Platzer and Peruffo (2022)).

The unconditional distribution of the long-run natural rate of interest $(r_t^n)^L$ consistent with our calibration is shown in the left panel of figure 1. The distribution is uniform over most of its support, and it has somewhat lower mass close to its boundaries, where the boundaries for the rate of growth of productivity induce some reflecting behavior in the long-run natural rate.

---

[9]This value corresponds to the top of a 90% confidence interval for the estimated convenience yield.

Figure 1: Unconditional distribution of the natural rate

Note: both the natural rate and the long-run natural rate are expressed in annualised terms.

We finally need to calibrate the standard deviations of permanent productivity shocks, σ_ψ – which drive fluctuations in the long-run natural rate – and of transitory shocks $\sigma_{\breve{r}^n}$. 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 $\sigma_{\breve{r}^n}$ found in Adam and Billi (2006), which is equal to 1.18% in annualised terms. This suggests that temporary shocks play a dominant role on the distribution of the natural rate.

The right panel in figure 1 confirms this conjecture. The panel shows the ergodic distribution of the natural rate r_t^n (where by definition $r_t^n = (r_t^n)^L + \breve{r}_t^n$). 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 distribution of the natural rate.

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

5   Optimal policy

This section highlights the implications of the model under optimal monetary policy. We start describing the impact of fluctuations in the long-run natural rate on the unconditional distribution of all endogenous variables. To provide an intuition for these results, we next study the optimal policy response to individual shocks, both of transitory and of permanent nature. We also provide an in-depth discussion of the short-run economic adjustment following a permanent shock, which has not yet been studied in the theoretical literature. We conclude the section by showing the results of a stochastic simulation, which is illustrative of the average dynamics of economic variables across different initial conditions.

As a benchmark for comparison, it is useful to note that, given the system of equations (1), (2), (3), (6) and (7), all (permanent and transitory) shocks to natural rate are efficient in our model. Absent the ZLB constraint, it would be optimal for monetary to fully stabilize inflation in reaction to such shocks. Absent the ZLB constraint, inflation and the output gap would always be equal to zero under optimal policy.

5.1   Sample moments

Table 2 shows sample moments for selected model variables when the model is solved taking into account the ZLB constraint. Consistently with the distribution displayed in Figure 1, the unconditional mean of the natural rate of interest is 1%. However, time-variation in the long-run natural rate implies that there are persistent periods in which the natural rate will be close to zero.

Table 2: Simulation moments

Unconditional means					Other moments	
r^n	x	π	i	RR spread	ELB frequency (x100)	ELB duration (quarters)
1,003	0,001	0,082	1,086	0,001	49,346	7,531

When this happens, monetary policy has very little, or no space to reduce policy rates in reaction to negative temporary shocks. Nevertheless, as already demonstrated in new Keynesian models with constant long-run natural rate, the central bank is not powerless in this situation. It can still control inflation, if it can "credibly promise to be irresponsible" (Dominquez et al. (1998)). More specifically, it can promise to be over-expansionary, i.e. create excess inflation, in reaction to any *future*

shocks. The promise will be reflected into private sector expectations and feed back into higher, actual inflation.

Table 2 shows the effects of this type of promise on average, across the various possible realizations of the natural rate. Both inflation and the nominal interest rate are higher than in the non-stochastic steady state, on average. From the quantitative perspective, this effect is small: the two variables increase by about 10 basis points (annualised). This is the case in spite of the fact that the ZLB is binding about 50% of the time in the model, and the average ZLB spell is of almost 2 years. The reason is that optimal policy is highly sophisticated in managing expectations and can therefore minimise inflation deviations from zero, that are costly due to price stickiness.

Nevertheless, even if the unconditional mean of inflation is only mildly positive, inflation must increase more conditional on a sequence of negative shocks that drive the long-run natural rate to particularly low levels. This will become apparent in our impulse response analysis.

5.2   Impulse responses to individual shocks

In this section, we study impulse responses to individual shocks, of either permanent or transitory nature.

It is well known that impulse responses are state and size dependent in a nonlinear model. One possible approach to study economic dynamics is to rely on generalised impulse responses, i.e. average responses across different initial conditions, weighted by their likelihood. Generalised impulse responses are representative of the average adjustment of economic variables after a selected shock of a given size. In our case, however, they are hard to interpret, because they reflect a convolution of different factors, depending on the initial level of the long-run natural rate, and of the policy rate vis-a-vis the ZLB.

To develop an intuition for the mechanisms at work in our model, we therefore compute standard impulse responses at selected points in the state space. By construction, after any shock economic variables will converge to the risky steady state. This is the point visited when the effects of all current shocks have been produced, but, in contrast to the non-stochastic steady state, agents are aware that other shocks could also occur in the future – see Coeurdacier et al. (2011).

The characteristics of the risky steady state are of independent interest, since they highlight the effects on endogenous variables of the risk of hitting the ZLB in the future – see also Hills et al. (2019). Since that risk changes, depending on the level of the long-run natural rate, so will the risky steady.

16

In the rest of this section, we first show impulse responses to transitory shocks starting from a particular value of the long-run natural rate. The results are useful to illustrate the impact of the ZLB constraint on the risky steady state. We subsequently analyse impulse responses to permanent shocks, which is the main contribution of our paper. We will specifically highlight the dependence of such responses on the initial value of the long-run natural rate and the characteristics of the subsequent adjustment process.

5.2.1 Transitory shocks

Due to the high incidence of ZLB risk in our model, output and inflation expectations will tend to incorporate a negative skew. This can be appreciated through an analysis of the impulse responses to negative transitory shocks. The ZLB constrains the central bank's ability to respond to negative shocks. Once the policy rate hits the ZLB, the real rate will necessarily be higher, hence more contractionary than the central bank would implement in the absence of the ZLB constraint. As a result, output and inflation expectations will tend to incorporate a negative bias. Such negative bias will not only appear at the ZLB, but also when the policy rate is positive, because of the *risk* that future adverse shocks will drive the economy to the ZLB.

To counter the downward bias in expectations, the central bank should implement an overly expansionary monetary policy even in the absence of shocks. Away from the ZLB, this requires keeping the risky steady state of the policy rate, or equivalently the neutral rate (i.e., the level of the policy rate consistent with stable inflation and the natural rate at its long-run level) below the long-run natural rate. Once the ZLB is reached, the central bank should be over-expansionary in a different way. It should promise to create more positive inflation and a positive output gap in reaction to positive shocks. The promise will be reflected in expectations, thus exerting an expansionary effect both on output through the Euler equation and on inflation through the Phillips curve.

The (risky) steady state in Figure 2 illustrates the mechanisms. Impulse responses to positive and negative transitory shocks are shown in the figure. For the policy rate, the initial (and final) point of the impulse responses is the neutral rate. The figure shows that the negative bias in expectations is sufficiently important to drive the neutral rate down to zero already when the long run natural rate is equal to 0.5%. As a result, the risky steady state of inflation and the output gap are somewhat positive.

17

Figure 2: Impulse responses to transitory shocks under optimal policy starting from $(r_0^n)^L = 0.5\%$
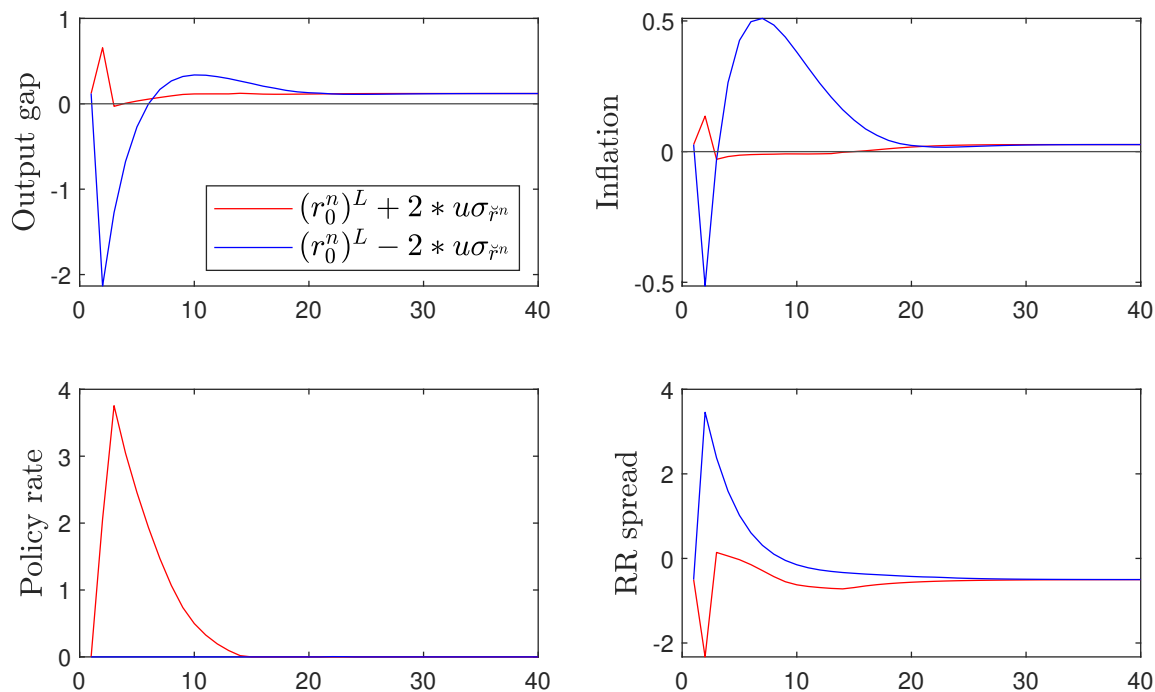
Note: the response of the output gap is in percentage points; all other responses are expressed in annualised percentage points.

The optimal policy responses to positive and negative shocks are strongly asymmetric.

In response to positive shocks, it would be feasible for the central bank to stabilise inflation perfectly and immediately, by increasing the policy rate one-to-one with the natural rate. However, the central bank does not follow this policy course, but it allows for some positive inflation. The policy rates increases by less than the natural rate.

In response to negative shocks, there is no scope for lowering the policy rate, since it is already at the ZLB. However, monetary policy is not powerless. The shock leads to an immediate increase in the spread between real interest rates and the natural rate but the deflationay and recessionary episode is short-lived. To stabilize the economy, the central bank can promise to keep the policy rate low (i.e. at zero) even after the natural rate returns to its long-run level of 0.5%. In addition, the central bank promises to create inflation as soon as possible after positive shocks.

18

The promise is reflected in expectations and generates an overshooting in the output gap and inflation. This result is qualitatively consistent with those in Eggertsson and Woodford (2003), but it is quantitatively amplified when the long run natural rate falls closer to zero.

In sum, the risk of hitting the ZLB implies that the central bank should implement an overly expansionary monetary policy. When the long run natural rate is equal to 0.5%, the neutral policy rate should already be at the ZLB. Around this point, the central bank will react asymmetrically to positive and negative transitory shocks.

5.2.2 Permanent shocks

Figure 3 displays impulse responses to permanent shocks starting from three different initial values of the long-run natural rate of interest: 1.5%, 1.18% and 0.85%.

Consider first the properties of the risky steady state at these different, initial values of the long-run natural rate. The downward bias in expectations is already present when $(r_t^n)^L = 1.5\%$ – see the thick, solid line in Figure 3. At this initial level of the long-run natural rate, the risky steady state of the policy rate, or equivalently the neutral rate (i.e. the level of the policy rate consistent with stable inflation and the natural rate at its long-run level), is only equal to 1.3%. At the risky steady state, the spread between the real rate and the natural rate is therefore negative (and equal to -0.2%). The output gap is sufficiently positive that inflation can be kept at zero.

Similar, but stronger, effects are visible at the risky steady states corresponding to $(r_t^n)^L = 1.18\%$ and $(r_t^n)^L = 0.85\%$. In the former case, the neutral rate falls to 0.8% and in the latter to 0.2% – see the solid and dotted lines in Figure 3, respectively. At the risky steady state, the spread between the real rate and the natural rate increases (in absolute value), as the long-run natural rate draws closer to zero. It is equal to -0.4% when $(r_t^n)^L = 1.18\%$ and to -0.6% when $(r_t^n)^L = 0.85\%$. The output gap also increases at the risky steady state, when the long-run natural rate falls towards zero. At all three levels of the long-run natural rate, however, inflation can be stabilised at zero in the risky steady state.

Consider now the adjustment process to the permanent shocks, which is especially interesting, since it could not be analysed in models where the natural rate is constant in the long-run. A notable feature of Figure 3 is that the adjustment of all endogenous variables to the new risky steady state is slow. This may be surprising, given the purely forward-looking nature of the standard new-Keynesian model. Indeed, equations (1) and (2) may suggest that, further to a shock to the long-run

natural rate shock, the economy will immediately jump to its new risky steady state.

Figure 3: Impulse responses to permanent shocks under optimal policy starting from different values of $(r_0^n)^L$
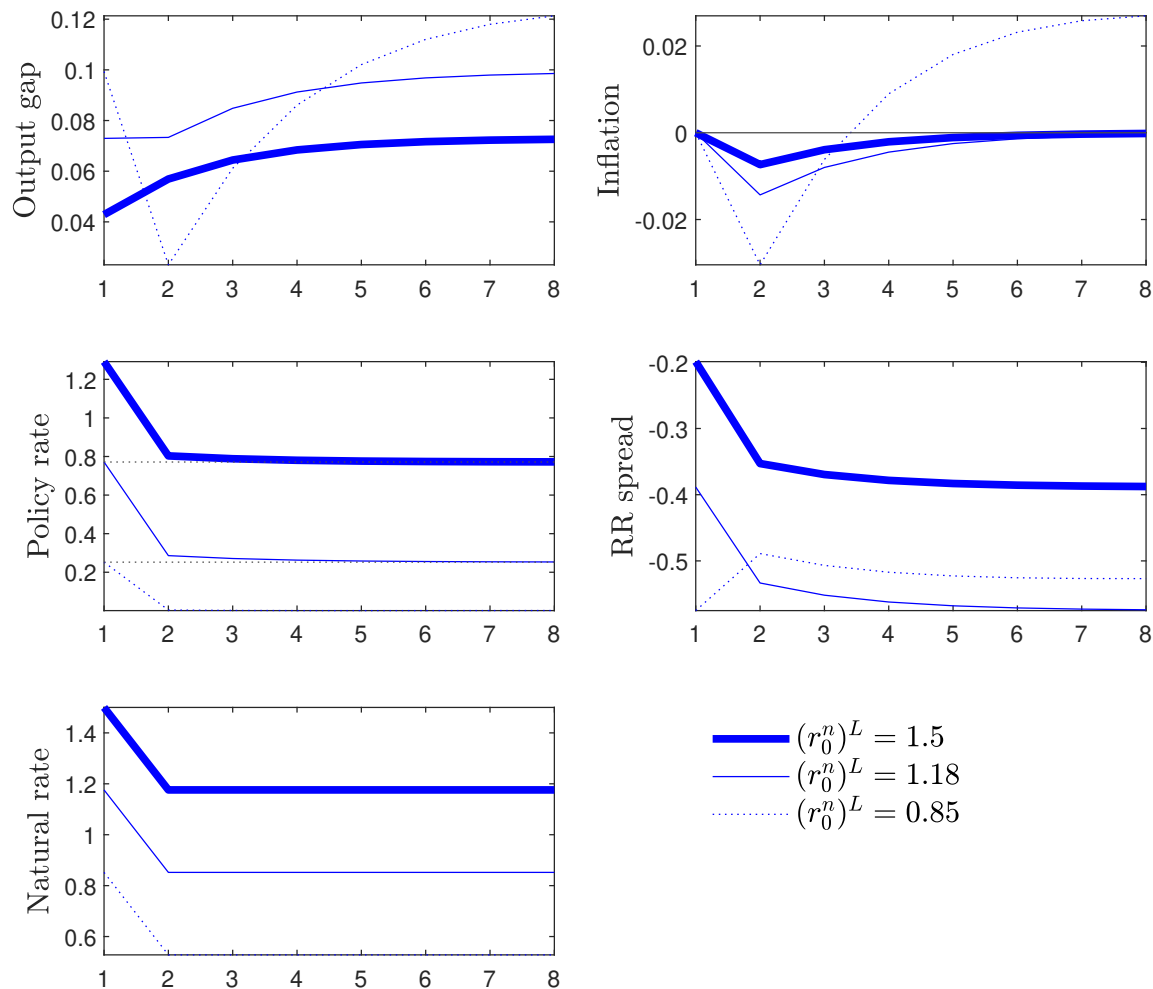
Note: the figure shows impulse responses to permanent shocks of a given size (0.324 percentage points, i.e., three times the standard deviation of the shock σ_ψ) starting from different initial values of the long-run natural rate. The response of the output gap is in percentage points; all other responses are expressed in annualised percentage points.

The slow adjustment process is indeed not due to the structure of the economy, but it is a policy choice. It is a result of the history-dependence of optimal policy

under commitment.

Take first the case when the shock occurs starting from $(r_t^n)^L = 1.5\%$. Since the shock is deflationary, the Lagrange multiplier associated with the Phillips curve, $\lambda_{p,t}$, turns positive – see equation (7) for the case in which the policy rate is positive and therefore $\lambda_{x,t-1} = \lambda_{x,t} = 0$. The central bank wants to push up inflation and, to do so, it expands the output gap and promises to keep expanding it in the future as long as inflation is negative – see equation (6) when $\lambda_{p,t} > 0$ and $\lambda_{x,t-1} = \lambda_{x,t} = 0$. The output gap keeps increasing until it reaches a higher risky steady state, when inflation is back to zero.

Interestingly, the slow adjustment 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) 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 and in output. These results are qualitatively consistent with ours in Figure 3.[10]. 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.

Going back to Figure 3, impulse responses are qualitatively similar when the shock occurs starting from either $(r_t^n)^L = 1.5\%$ or $(r_t^n)^L = 1.18\%$. However, economic dynamics change when the initial value of the long-run natural rate is below 1%. In that case, monetary policy can no longer respond with a sufficient reduction in the policy rate, because the neutral rate reaches zero and the ZLB constraint binds. As shown by the dotted line in Figure 3, the spread between the real rate and the policy rate increases. The initial fall in the output gap and inflation is therefore larger than in the previous two cases. The only option for the central bank to provide monetary accommodation is to promise an overly expansionary reaction to future shocks. Through its impact on expectations, this promise will also pull up output, and inflation expectations. However, due to the forward looking nature of households' and firms' decisions, expectations will also have an impact on actual outcomes. Not only will the output gap be more positive in the new risky steady state, but also inflation will become positive.

Billi et al. (2022) also characterises a situation in the new Keynesian model in

---

[10]Even if the output gap increases in the figure, output falls due to the permanent reduction in productivity caused by the shock.

which optimal inflation is positive and the policy rate is at zero in the absence of shocks. In Billi et al. (2022), however, this situation occurs because the natural rate is negative in the non-stochastic steady state. The central bank sets the policy rate at zero and promises positive inflation to ensure that the real rate equals the natural rate. By contrast, in our model the neutral rate can optimally reach zero in the risky steady state even if the long-run natural rate is positive. Our results are determined by the nonlinear expectation effects induced by the ZLB constraint.

In sum, the optimal monetary policy reaction to a reduction in the long-run natural rate is to implement an over-expansionary monetary policy stance. If there is enough room for reduction in the policy rate, policy accommodation can be produced by ensuring that the long-run real interest rate is lower than the long-run natural rate. The central bank can then stabilise inflation at zero. Once the policy rate reaches the ZLB, and further reductions of the policy rate become infeasible, policy accommodation requires the promise to allow for excessive inflation in the future. In this case, actual inflation cannot be stabilised at zero, but it will become positive.

5.3   The impact of shocks in a stochastic simulation

In the previous section, we have highlighted that, once the long-run natural rate falls towards zero, the optimal monetary policy stance is optimally over-expansionary. In a rational expectations model, such policy approach will be reflected into private sector expectations and, since inflation is forward looking, it will also influence actual inflation. In this section we quantify this effect by looking at average outcomes of an illustrative stochastic simulation, in which the natural rate is hit by a sequence of three permanent, negative shocks which bring it from 1.5% to 0.5% and, in addition, it is continuously buffeted by transitory and permanent shocks drawn from their unconditional distribution. Note that the three permanent shocks correspond to those already shown individually in Figure 3.

As already discussed above, the sequence of permanent shocks in isolation requires monetary policy to be increasingly over-expansionary (blue lines in Figure 4). Outcomes are significantly different if the three permanent shocks are accompanied by other random shocks drawn from their distribution (black lines in Figure 4). In the latter case, there is some positive probability that each permanent shock will happen in combination with a negative, transitory shock that drives the policy rate to the ZLB. In those contingencies, the central bank will be forced to respond through the promise of future inflation, and of a positive output gap – a promise which will be reflected in higher, average output and inflation expectations, compared to the case in which the three shocks happen in isolation. In those same contingencies, the

policy rate will instead not fall much – due to the ZLB – while it will increase after positive, transitory shocks. On average, the policy rate will therefore be higher in the stochastic simulation, compared to the case in which the three permanent shocks happen in isolation. The average spread between the real rate and the natural rate will thus *not* fall with the long-run natural rate.

Figure 4: 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. 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. Following this sequence of shocks, the long run natural rate is equal to 0.528%. In each panel: the blue line reports outcomes following the sequence of permanent shocks in isolation; the black lines shows outcomes when transitory and permanent shocks are drawn from their distribution at the same time. The output gap is expressed in percentage points; all other variables in annualised percentage points.

All in all, in spite of the expectations of a higher output gap, the actual output gap will on average be close to zero, after recording a short-lived recession on impact

23

after each permanent shock. The higher inflation expectations will instead push up actual inflation (after an impact reduction following each permanent shock – see the top-right panel in the figure). Already at the beginning of the simulation, when the long-run natural rate is 1.5%, the conditional mean of inflation is not zero, but slightly positive (around 3 basis points in annualised terms). Following to the first, negative permanent shock in period 10, average inflation falls on impact, but it eventually increases to about 0.05% on average. Similar, but comparatively larger adjustments are observed after the second and third permanent shocks. At the end of the simulation, when the long-run natural rate is just above 0.5%, inflation reaches almost 0.15%.

Our results on inflation are related to those in Andrade et al. (2018, 2021), which study the optimal inflation target to be assigned to a central bank following a Taylor-type rule for different (constant) values of the steady state natural rate. Similarly to those papers, we find that average inflation should optimally increase, as the long-run natural rate falls. However, the optimal inflation adjustment is small in our case, and it should only take place once the long-run natural rate falls below 2%. By contrast, the adjustment in the inflation target is sizable in Andrade et al. (2018, 2021). The reason is that, when following a simple policy rule, the central bank has a limited ability to make promises about its future policy course. It will therefore tolerate higher inflation, so as to increase the scope for reductions of the policy rate in reaction to adverse shocks.

6 Implementation: price level targeting

In this section we investigate how the optimal equilibrium described in Section 5 can be implemented in practice. Implementation through a simple policy rule would be especially desirable in our case, if the rule did not require any estimate or knowledge of the long-run natural rate of interest. This magnitude is in fact difficult to filter from the data, especially in real time.

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

$$p_t + \frac{\lambda}{\kappa} x_t = P^*$$ (9)

where P^* is the target gap-adjusted (log-)price level (a given constant). Equation (9) commits the CB to counteract shocks destabilizing the price level. When perfect stabilization is not feasible because of the ELB, it commits the CB to undo actual de-

flation with future inflation so as to bring the gap adjusted (log-) price level (GAPL) back to target. Over long time periods, therefore, inflation will be zero on average.

As emphasized in Eggertsson and Woodford (2003), this rule does have the advantage of not requiring knowledge of the statistical process for natural rate. However, it limits the ability of the CB to manage private sector expectations and thereby its ability to mitigate the effect of the ELB. Eggertsson and Woodford (2003) study its properties within a stationary model, in which optimal inflation is zero on average. By contrast, optimal policy in our model calls for a positive rate of inflation on average, which indicates that the central bank actively manages expectations. 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 (9), 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

$$p_t + \frac{\lambda}{\kappa}x_t = P_t^*$$ (10)

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

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

Clearly, the adjusted rule (10) boils down to rule (9) when $\pi^* = 0$. It does not improve the ability of the CB to manage private sector expectations, but it increases the room for reductions of the policy rate. Absent the lower bound on the policy rate, equation (10) commits the central bank to create a positive output gap as high as necessary to ensure that the rate of inflation is equal to π^* at any point in time. This outcome requires to keep the policy rate above the natural rate.

The relative merit of the two rules can be analyzed in terms of their welfare performance. The left panel of Figure 5 reports the results for different values of π^*, including zero.

Figure 5: Welfare loss under price level targeting with drift
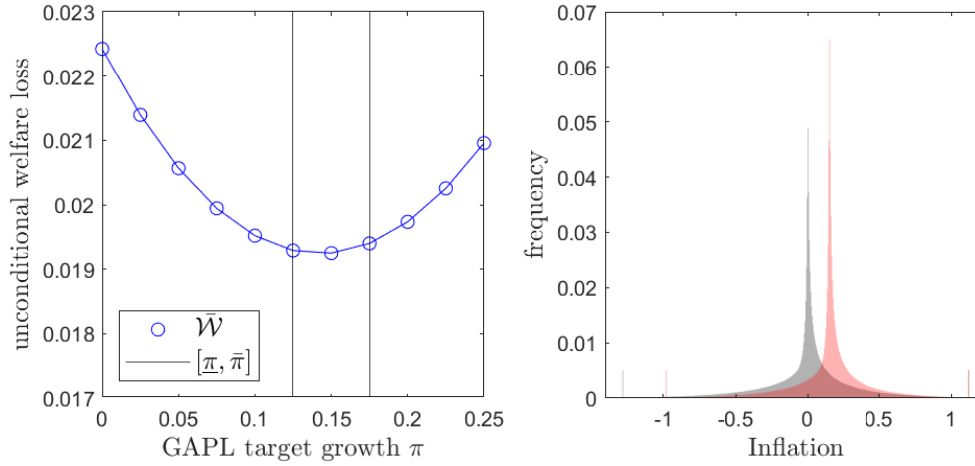
Note: the left panel shows the welfare-theoretic central bank loss function for different values of the GAPL target growth – see appendix C for further details. The right panel shows the unconditional distribution of inflation winsorized at 1% under price level targeting when π^* is equal to zero (grey) and 15 basis points (red). Both the price level target π^* and inflation are expressed in annualised percentage terms.

The figure shows that the rule with positive trend attains a lower welfare loss. Any values of the trend growth of up to 25 basis point per year would be strictly preferable to the original rule with a constant price level target. Welfare is maximized when the price level trend is equal to 15 basis points (annualized). This is approximately twice as large as the unconditional inflation rate under optimal policy. Nevertheless, the differences in the performance between the two rules are small. The welfare loss of the rule with constant price target is approximately 17% larger than the rule with the optimal price level trend. By contrast, the loss under optimal policy is 39% smaller than under the rule with the optimal price level trend.[11]

The right panel of Figure 5 shows the equilibrium distribution of inflation under price level targeting for two values of the price level trend π^*: zero and 15 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.

---

[11]The unconditional welfare loss, as measured by $\bar{\mathcal{W}}$, amounts to 0.0224, 0.0192, and 0.0117 under the constant gap adjusted price level (GAPL) targeting rule, the optimal growing GAPL targeting rule, and the optimal policy respectively.

26

In the rest of this section we describe the performance of the price level targeting rules via impulse response analysis. Focusing again on the case of permanent shocks, we compare the performance of price level targeting rules (9) and (10) with outcomes under optimal policy. Figure 6 shows the results of a sequence of negative, permanent shocks in isolation, where the solid blue line corresponds to the solid blue line already shown in Figure 4.

By design, the key characteristic of the Eggertsson and Woodford price level targeting rule (9) is to stabilize the 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 (9) is less sophisticated in managing expectations. As a result, *ceteris paribus*, the spread between the neutral rate and the long-run natural rate must be more pronounced than in the optimal policy case. The output gap must be larger and 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.

At that point, the central bank has no room to reduce the policy rate after a negative permanent shock (see period 50). As a result, the gap-adjusted price level falls and undershoots its target. To bring it back to target, the central bank promises to create an output boom and inflation as soon as possible in the future. This promise is reflected in expectations and allows the central bank to mitigate the impact of the shock. At the new risky steady state, the real rate spread is substantially higher than at the previous risky steady state, by an amount almost equal to the fall in the long-run natural rate. The rate of inflation is stabilized at zero even though the output gap is lower. This happens because private agents expect more inflation than in the initial risky steady state due to the central bank's promise to bring the GAPL back to target.

The monetary policy stance implemented as of period 50 under rule (9) displays two main differences with respect to optimal policy. First, already on impact, the central bank has no room to reduce the policy rate after the permanent shock. Second, the central bank is less aggressive in the management of expectations. It only promises to bring the gap-adjusted price level back to the original target in the future, 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, and the economic slowdown is much more pronounced than under optimal policy, even if equally short-lived. In addition, expectations incorporate a deflation bias under rule (9), while they are consistent with zero inflation under optimal policy.

27

Figure 6: Impulse responses to a sequence of permanent shocks under various rules
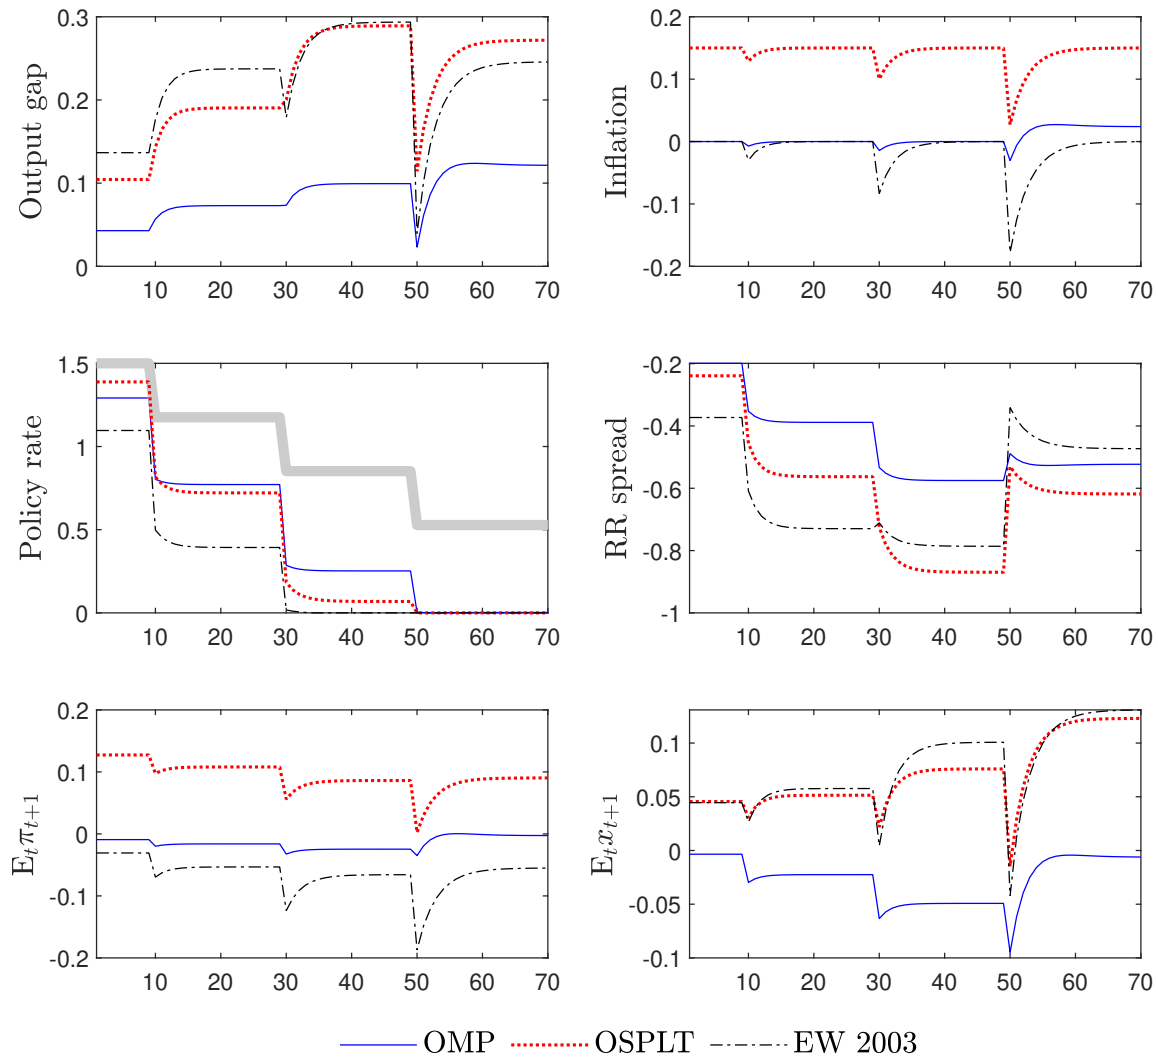
Note: *the figure shows the results of the same simulation underlying Figure 4. The focus is on the case in which shocks occur in isolation. The solid blue line corresponds to the solid blue line in Figure 4. The dotted red and black lines indicate outcomes under price level targeting with an optimal 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.*

Economic dynamics under the modified price level targeting rule (10) are closer

to those under optimal policy. This is the case because, before the ZLB is reached, rule (10) is consistent with a significantly higher level of the neutral rate, even though the target rate of inflation amounts to only 15 basis points. The initial risky steady state illustrates the mechanism. We observe that, conditional on a long run natural rate equal to 1.5%, the risky steady state output gap under rule (10) is lower than under rule (9), and the policy rate is 30 basis points higher. Everything else equal, the central bank has more space to reduce the policy rate under rule (10). Once the ZLB binds, in period 50, the additional space allows for superior stabilization outcomes after negative shocks. In contrast to what we observe under rule (9), inflation expectations remain positive when the price level target incorporate a positive trend. 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 when 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 central bank should aim to maintain the policy rate below the natural rate. The neutral rate will thus 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, consistently with the existing literature, 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. 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.

Altig, D., Christiano, L. J., Eichenbaum, M., and Linde, J. (2011). Firm-specific capital, nominal rigidities and the business cycle. *Review of Economic dynamics*, 14(2):225–247.

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.

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

Caballero, R. J. and Farhi, E. (2018). The safety trap. *The Review of Economic Studies*, 85(1):223–274.

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

Christiano, L. J., Motto, R., and Rostagno, M. (2014). Risk shocks. *American Economic Review*, 104(1):27–65.

Christoffel, K. P., Coenen, G., and Warne, A. (2008). The new area-wide model of the euro area: a micro-founded open-economy model for forecasting and policy analysis.

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

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.

Dominquez, K. M., Krugman, P. R., and Rogoff, K. (1998). It's baaack: Japan's slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, 1998(2):137–205.

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.

Fiorentini, G., Galesi, A., Pérez-Quirós, G., and Sentana, E. (2018). The rise and fall of the natural interest rate. Banco de Espana Working Paper 1822.

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.

Hills, T. S., Nakata, T., and Schmidt, S. (2019). Effective lower bound risk. *European Economic Review*, 120:103321.

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.

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

Laubach, T. and Williams, J. C. (2016). Measuring the natural rate of interest redux. *Business Economics*, 51(2):57–67.

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*.

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.

McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–58.

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.

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.

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.

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.

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.

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,

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

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 $\upsilon'([\widetilde{m}_{j,t} - \widetilde{m}_t]\bar{A}_t) = \upsilon'(\widetilde{m}_{j,t} - \widetilde{m}_t)\bar{A}_t^{-1}$. Since everyone has same preferences and same initial wealth, we can drop the $j's$. The first order conditions include:

$$\frac{W_{k,t}}{P_t} = \gamma \bar{H}_t^{-\upsilon} H_{k,t}^{\upsilon} \tilde{C}_t \bar{A}_t$$

$$1 - \frac{I_t^m}{I_t} = S_t \tilde{C}_t \upsilon'(0)$$

$$1 = \mathrm{E}_t \left[\beta \Delta_{t+1} \frac{\tilde{C}_t}{\tilde{C}_{t+1}} \frac{\bar{A}_t}{\bar{A}_{t+1}} \frac{I_t}{\Pi_{t+1}} \right]$$

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

$$\delta_{t+1} = \rho_\delta \delta_t + \sigma_\delta \varepsilon_{\delta,t+1},$$

$$s_{t+1} = \rho_s s_t + \sigma_s \varepsilon_{s,t+1},$$

$$z_{t+1} = \rho_z z_t + \sigma_z \varepsilon_{z,t+1},$$

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

$$\xi_t = \xi_{t-1} + \psi_t$$

$$\psi_t = (1 - \rho_\psi)\psi + \rho_\psi \psi_{t-1} + \sigma_\psi \varepsilon_t^\psi$$

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

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

and

$$\tilde{c}_t = E_t \tilde{c}_{t+1} - (\breve{i}_t^m - E_t \pi_{t+1}) + E_t \widehat{\psi}_{t+1} - E_t \delta_{t+1}$$

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

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

In a natural equilibrium where prices are flexible

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

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

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

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

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

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

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

Detrending real variables and following standard derivations we can linearize this condition to obtain $\pi_t = \kappa \left(\tilde{y}_t - z_t \right) + \beta E_t \pi_{t+1}$, 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

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1}$$

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

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

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

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

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

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

36

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

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

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

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

$$U_t - \left(\ln \tilde{Y} - \frac{\gamma}{1+v}\tilde{Y}^{1+\omega}d\right) \simeq -\frac{1}{2}\gamma\phi^2(1+v)\tilde{Y}^{1+\omega}d\tilde{y}_t^2 - \frac{\gamma}{1+v}\tilde{Y}^{1+\omega}d\hat{d}_t$$

$$+ \left(1 - \gamma\phi\tilde{Y}^{1+\omega}d\right)\tilde{y}_t + \left(1 - \gamma\phi\tilde{Y}^{1+\omega}d\right)\bar{c}_t\tilde{y}_t$$

$$- \frac{1}{2}\frac{\gamma}{1+v}\tilde{Y}^{1+\omega}d\hat{d}_t^2 - \frac{\gamma}{1+v}\tilde{Y}^{1+\omega}d\bar{c}_t\hat{d}_t - \gamma\phi\tilde{Y}^{1+\omega}d\tilde{y}_t\hat{d}_t$$

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

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

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

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: $E_t[f^c(s_{+1})] = \int_{-\infty}^{\infty} g^c(\varepsilon^{\bar{\delta}}_{t+1}) \exp(-\varepsilon^{\bar{\delta}^2}_{t+1}) d\varepsilon^{\bar{\delta}}_{t+1}$, 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 $Ex_{j-1}(s_{+1})$ and $E\pi_{j-1}(s_{+1})$, and retrieve

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

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

$$\pi_j(s) = \kappa x_{j-1}(s) + \beta E\pi_{j-1}(s_{+1})$$
$$x_j(s) = (1 - \Delta^m)\left[Ex_{j-1}(s_{+1}) - \left(\bar{\iota}^m - E\pi_{j-1}(s_{+1}) - \bar{\delta}\right)\right]$$

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 $\left(x^*(s_t), \pi^*(s_t)\right)$. 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 $\left(x^{\text{IMP}}(s_t), \pi^{\text{IMP}}(s_t)\right)$, and we compute the residuals as follows:

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

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

Table 3 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 3: Stationary r_t^n and optimal commitment: simulation moments and accuracy indicators

r^n	$\breve{x}$	π	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$
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 7 provides an illustrative example of the grid based on the solution obtained when the steady state natural rate of interest is 3.5%.

Table 4 and Table 5 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 certainly improve the accuracy of the solution, but it would not substantially affect the results that are reported in the main text.

Figure 7: Stationary r_t^n and optimal commitment: An adaptative grid

Table 4: Stationary r_t^n and optimal commitment: Robustness with respect to the grid, accuracy indicators

r^n	Evenly spaced grid				Adaptative grid			
	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$
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 3 for details.

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

r^n	Evenly spaced grid						Adaptative grid					
	$\breve{x}$	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 3 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 6 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 6: Stationary r_t^n and optimal commitment: Robustness with respect to the solution algorithm, simulation moments

r^n	Baseline algorithm						Alternative algorithm					
	$\breve{x}$	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 3 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:

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

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 7 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 7: Drifting r_t^n and optimal commitment: simulation moments and accuracy indicators

r^n	$\breve{x}$	π	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$
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 3 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 8 and 9 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 8: Drifting r_t^n and optimal commitment: Robustness with respect to the grid, accuracy indicators

r^n	Evenly spaced grid				Adaptative grid			
	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$
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 3 for details.

46

Table 9: Drifting r^n_t and optimal commitment: Robustness with respect to the grid, simulation moments

	Evenly spaced grid						Adaptative grid					
r^n	$\check{x}$	π	i	$r - r^n$	ELB freq	ELB dur	$\check{x}$	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 3 for details.

Alternative solution algorithm

Table 10 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 10: Drifting r^n_t and optimal commitment: Robustness with respect to the solution algorithm, simulation moments

	Baseline algorithm						Alternative algorithm					
r^n	$\check{x}$	π	i	$r - r^n$	ELB freq	ELB dur	$\check{x}$	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 3 for details.

B.2   Comparison of unconditional moments

Table 11 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 11: Sample moments when 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, $E_t[\log \Psi_{t+1}] = \bar{\psi}_t = \rho_\psi \bar{\psi}_{t-1} + \sigma_\psi \varepsilon_t$, and $\rho_\psi = 0.99$.

B.3 Additional impulse responses

Figure 8: Simulations of a sequence of permanent shocks under optimal policy

C Price level targeting

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

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

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.

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

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

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

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

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

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 $\bar{\iota}^m(\xi)$ denote the ELB on the policy rate. If $x_j(s) > \left(1 - \Delta^m\right)\Big[\mathrm{E}_s x_{j-1}s' - \left(\bar{\iota}^m(\xi) - (\pi^* + \mathrm{E}_s\bar{p}_{j-1}(s') - \bar{p}_{j-1}(s)\right) - \bar{\delta} - \mathrm{E}_s(\psi')\Big]$, retrieve

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

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 12 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 12: Drifting r_t^n and price level targeting: simulation moments and accuracy indicators

r^n	π^*	$\breve{x}$	π	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$
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 3 for details.

Robustness check based on an adaptative grid

Table 13 and 14 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 9 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 adaptive grid. This suggests that using a denser grid would not improve the accuracy of the solution. Moreover, the simulation moments are essentially unchanged.

Figure 9: Drifting r_t^n and price level targeting: An adaptative grid

Table 13: Drifting r_t^n and price level targeting: Robustness with respect to the grid, accuracy indicators

r^n	π^*	Evenly spaced grid				Adaptative grid			
		$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$	$\max[e^x]$	$\mathrm{E}[e^x]$	$\max[e^\pi]$	$\mathrm{E}[e^\pi]$
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 3 for details.

52

		Evenly spaced grid						Adaptative grid					
r^n	π^*	$\check{x}$	π	i	$r-r^n$	ELB freq	ELB dur	$\check{x}$	π	i	$r-r^n$	ELB freq	ELB dur
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 3 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 [12]. We computed it based on the two different methods described below. Figure 10 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

---

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

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

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

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 10: Unconditional welfare loss under price level targeting

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

Figure 11 compares outcomes between optimal policy and the two price level targeting rules in a stochastic simulation.

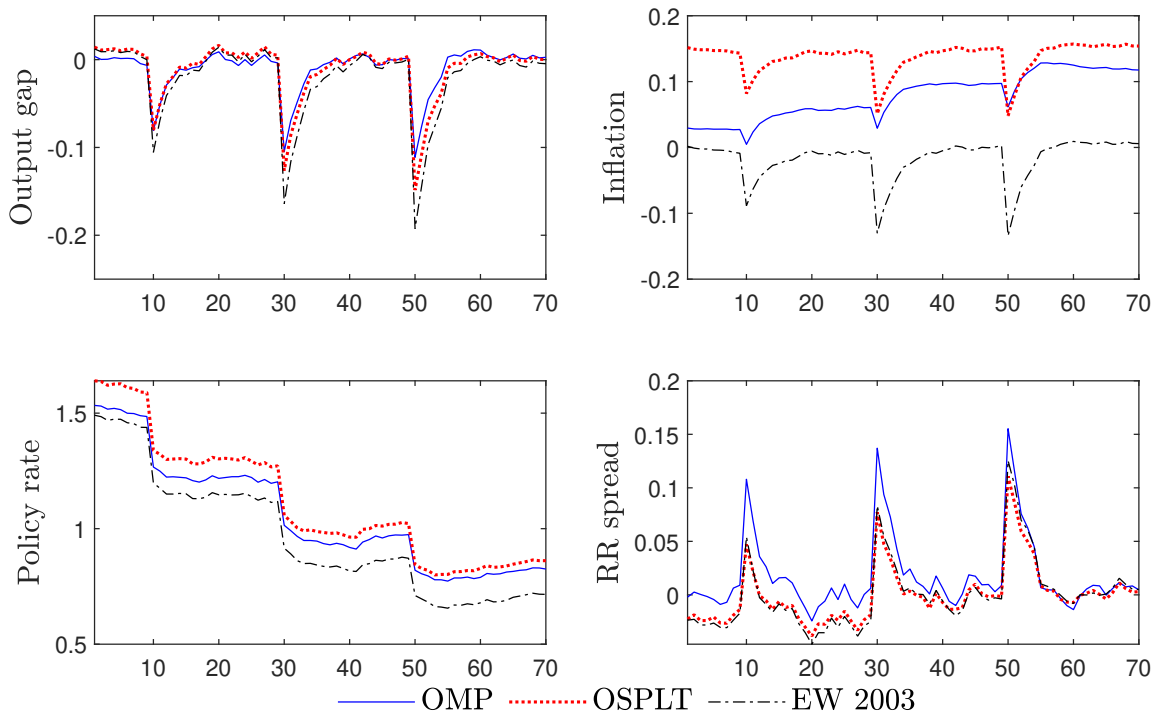Figure 11: Simulations of a sequence of permanent shocks under various policy rules

Figure 12: Impulse responses to a positive transitory shock under various policy rules starting from $(r_t^n)^L = 0.5\%$

$r^n + 2 * u\sigma_\delta$

OMP ——— OSPLT ——— EW 2003