

Optimal Monetary Policy with $r^* < 0$

Roberto Billi
Sveriges Riksbank

Jordi Galí
CREI, UPF and BSE

Anton Nakov
European Central Bank

July 2023

Abstract

We study the optimal monetary policy problem in a New Keynesian economy with a zero lower bound (ZLB) on the nominal interest rate, when the steady state natural rate (r^*) becomes permanently negative. We show that the optimal policy aims to approach *gradually* a new steady state with positive average inflation. Around that steady state, the optimal policy implies well defined (second-best) paths for inflation and output in response to shocks to the natural rate. Under plausible calibrations, the optimal policy implies that the nominal rate remains at its ZLB most of the time. Despite the latter feature, the central bank can implement the optimal outcome as a unique equilibrium by means of an appropriate nonlinear interest rate rule. In order to establish that result, we derive sufficient conditions for local determinacy in a general model with endogenous regime switches.

JEL Classification Numbers: E32, E52

Keywords: zero lower bound, New Keynesian model, decline in r^* , equilibrium determinacy, regime switching models, secular stagnation

We thank Danila Smirnov for excellent research assistance. We have benefited from comments by Salvatore Nistico, Albert Satorra, Gabor Lugosi, Albert Marcet, Sebastian Schmidt, Jean Barthélemy, Magali Marx, Philippe Andrade, Hervé le Bihan, Fernando Álvarez, Boragan Aruoba and Johannes Wieland. Galí acknowledges financial support from the European Research Council under the European Union's Horizon 2020 research and innovation program (Grant 882332-HEMPEF), from the Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S) and from the Generalitat de Catalunya, through the CERCA Programme. The views expressed herein are solely the responsibility of the authors and should not be interpreted as reflecting the views of the European Central Bank or the Sveriges Riksbank.

1 Introduction

Over the past decade, a growing consensus has emerged among academic economists and policymakers pointing to a substantial decline in the *average natural rate of interest*, a variable often referred to as r^* . Some of the likely sources of that decline –including lower productivity growth, demographic factors, higher inequality or enhanced precautionary savings induced by higher uncertainty– suggest that such a downward trend is unlikely to be reversed in the near future.¹

A low r^* has important implications for monetary policy, due to the presence of a zero lower bound (ZLB) on the nominal interest rate. Thus, and given the inflation target, a low r^* will generally hamper the ability of monetary policy to stabilize the economy, bringing about more frequent episodes in which the ZLB becomes binding and the economy plunges into a protracted recession with below-target inflation. Not surprisingly, the evidence of a decline in r^* has been a key motivation behind the monetary policy strategy reviews undertaken by many central banks in recent years.

On the research front, and as discussed in the literature review below, several authors have studied the problem of optimal monetary policy in the face of shocks that drive the natural rate of interest *temporarily* into negative territory. A common finding of those analyses is that an optimizing central bank will keep the short-term nominal rate at zero during those episodes, and even for some time after the natural rate has returned to positive values –with the latter feature often referred to as "lower for longer" policy. In all of those analyses, however, the natural rate tends to gravitate towards a positive mean, i.e. $r^* > 0$. By contrast, in the present paper we study the problem of optimal monetary policy under the ZLB constraint when the mean of the natural rate becomes permanently *negative*, i.e. $r^* < 0$.

As discussed below, that environment is of particular interest since the coexistence of a

¹See, e.g. Eggertsson et al. (2019) for a model-based analysis of some of the forces underlying the decline in r^* . Rachel and Smith (2017) and Rachel and Summers (2019) argue for the likely permanent nature of recent trends in those forces, especially those that manifest themselves as an outward shift in global savings. Despite the strong global inflationary pressures at the time of writing this paper, we believe that the factors behind the decline in r^* not only have not disappeared, but they may have been enhanced by the impact on uncertainty of the COVID pandemic or the Ukrainian war. If that is the case, the consequences of a low r^* and its interaction with the zero lower bound constraint are likely to take again center stage in the policy debate once inflation returns to levels close to target.

negative r^* with the ZLB constraint makes it impossible to support the (first-best) zero inflation outcome even in the deterministic case, i.e. in the absence of fluctuations in the natural rate. In the latter case, the optimal policy implies positive inflation and a binding ZLB constraint in the deterministic steady state, a feature that is absent from conventional analyses that assume a positive r^* , in which the deterministic steady state is characterized by zero inflation and a strictly positive nominal rate. The focus of our analysis lies, however, on the stochastic case, i.e. in the optimal policy in the presence of fluctuations in the natural rate around $r^* < 0$, and on the implications of that policy for the nominal rate, inflation and the output gap.

While the assumption of a negative r^* is at odds with the predictions of the standard macro framework with an infinite-lived representative consumer, it can be microfounded once the latter assumption is relaxed. Thus, for instance, models with overlapping generations, or heterogeneous agents and idiosyncratic shocks, can generate a negative r^* under certain parameterizations. Furthermore, we believe the assumption of a negative r^* is more than a theoretical curiosity: recent estimates of the evolution of the natural rate in advanced economies display a downward trend that has already attained negative territory in some cases.² In any event, the relevance of a negative r^* can hardly be dismissed as a real possibility in a not too distant future, if the trends in some of the fundamental forces behind the recent decline in the natural rate were to persist or even strengthen further.

As much of the related literature, we cast our analysis of the optimal monetary policy problem in the context of an otherwise standard New Keynesian model subject to a ZLB constraint and a central bank loss function characterized by a conventional dual mandate.³ A number of interesting results emerge from our analysis.

Focusing first on the deterministic case, we show that in response to an unanticipated decline in r^* which brings the latter *permanently* into negative territory, the optimal policy aims at

²A recent paper by Davis et al. (2023) uses a market-based approach to estimate the Postwar evolution of r^* in ten industrialized economies. Their estimates of current values of r^* are negative in 8 out of the 10 countries. Evidence on global r^* in Del Negro et al. (2019) cannot reject the null of negative values in recent years. Brand and Mazelis (2019) use a semi-structural model incorporating a Taylor rule, and also uncover negative estimates of r^* in the US and the euro area in recent years.

³We use the textbook New Keynesian model as a framework in which we revisit the optimal policy problem in the presence of a negative r^* . This is meant to highlight in a most transparent way the key *qualitative* implications of a negative r^* for monetary policy. We believe that adding additional "realistic" features to the model (e.g. imperfect credibility, parameter uncertainty, investment, etc.) would complicate the analysis without qualitatively altering or shedding additional light on those key implications.

steering the economy *gradually* towards a new steady state characterized by positive inflation. The choice of a gradual transition (rather than an immediate jump to the new steady state) makes it possible for inflation to remain closer to zero –its efficient value– for a longer period, which is welfare improving.

Secondly, we solve for the paths of inflation and the output gap implied by the optimal (second-best) policy in the presence of fluctuations in the natural rate of interest around $r^* < 0$. Not surprisingly, the presence of the ZLB constraint prevents the central bank from fully stabilizing inflation and the output gap, so the first-best outcome cannot be attained. Most interestingly, we show that if either the volatility of the natural rate is not too large (for any given r^*) or if r^* is low enough (for any assumed volatility of the natural rate), then the optimal policy implies a *persistently binding ZLB constraint*, with the nominal rate remaining at zero most of the time (*all* the time, in some of our simulations). Behind the appearance of extreme passivity suggested by a near-constant policy rate, however, there is still a meaningful optimal policy problem facing the central bank, which yields unique optimal paths for inflation and the output gap.⁴

Thirdly, we show that average inflation under the optimal policy is decreasing and convex in r^* . The resulting relation balances three requirements: (i) the intrinsic desirability of price stability, which calls for inflation being as close to zero as possible, (ii) the equilibrium requirement that, on average, inflation must be no lower than $-r^*$ due to the ZLB constraint, and (iii) a precautionary motive linked to the desire to limit the incidence of binding ZLB episodes. Thus, when r^* is positive and large the precautionary motive is negligible and optimal average inflation is zero. As r^* approaches zero from above, optimal average inflation becomes positive due to a more significant precautionary motive, but it remains very low and responds less than one-for-one to changes in r^* . The more r^* moves into negative territory, the more optimal average inflation approaches $-r^*$, its minimum average value consistent with the ZLB constraint, due to the increasing weight of the price stability motive resulting from the convexity of the loss function. The convergence of optimal average inflation to $-r^*$ mirrors the convergence of the average nominal rate to zero, and is thus associated with a near-permanently binding ZLB

⁴This is because the constant interest rate policy is consistent with a continuum of paths for output and inflation, which can be welfare-ranked.

constraint.

In order to characterize that finding more precisely, we introduce the concept of *precautionary inflation*, which we define as the difference between optimal average inflation in the presence of natural rate shocks and optimal inflation in the deterministic case. That measure can be interpreted as capturing the central bank's willingness to accept a higher average inflation in order to limit the incidence of binding ZLB episodes. We show that precautionary inflation displays a non-monotonic relation with r^* . Thus, when r^* is very high, the risk of a binding ZLB is low, and there is no need to deviate from the first-best outcome of zero inflation at all times. At the other extreme, when r^* is sufficiently negative and, hence, the lower bound on average inflation is already high, the central bank has little incentive to raise average inflation further above that lower bound, and thus chooses to keep average inflation at the same level as in the deterministic case. By contrast, precautionary inflation is strictly positive for a range of r^* values closer to zero, for which optimal inflation in the deterministic case is either zero (if $r^* \gtrsim 0$) or positive but low (if $r^* \lesssim 0$), since in that case the costs of deviations from full price stability are relatively low, and are outweighed by the gains from a lower incidence of a binding ZLB made possible by the additional policy space created by a higher average inflation and nominal rate.

Fourthly, we describe one particular way in which the central bank can implement the optimal (second best) policy. More specifically we propose a nonlinear policy rule which calls for one-sided adjustments in the nominal rate in response to (off-equilibrium) deviations from the desired inflation and output gap paths. In order to establish the implementability of those paths as a *unique* equilibrium under the proposed rule, we derive and exploit a sufficient condition for local determinacy for a relatively general class of models with *endogenous* regime switches. We believe the latter finding has some independent interest, beyond the application at hand, and complements existing results in the literature for exogenous regime switching models.

The rest of the paper is organized as follows. The remaining of the present section provides a brief review of the related literature. Section 2 formulates the optimal policy problem and derives the associated optimality conditions. Section 3 analyzes the economy's (deterministic) transitional dynamics under the optimal policy. Section 4 characterizes the fluctuations of inflation and output around the steady state, in response to natural rate shocks. Section 5 discusses

the implementation of the optimal plan, deriving sufficient conditions on the coefficients of a proposed interest rate rule to support the optimal plan as a unique equilibrium. Section 6 concludes.

1.1 Related Literature

Our paper is related to a branch of the literature that studies the optimal design of monetary policy in the presence of a ZLB constraint on the nominal rate. Since Krugman (1998), a number of articles have studied optimal monetary policy with an occasionally binding zero lower bound (ZLB) on the nominal interest rate. Closest to us is the work by Eggertsson and Woodford (2003), Jung, Teranishi and Watanabe (2005), Adam and Billi (2006), and Nakov (2008), who analyze the problem of optimal policy under commitment in the basic New Keynesian model with a ZLB constraint.

A different line of work has focused on the implications of the ZLB for the optimal choice of an inflation target, *conditional on an assumed simple interest rate rule*. Relevant papers include Coibion et al. (2012), Bernanke et al. (2019), and Andrade et al. (2020, 2021).⁵ In all the papers above, however, the natural interest rate remains negative only temporarily, with the binding ZLB being a transitory phenomenon. In contrast, the analysis of the present paper assumes a negative r^* , and hence a permanent “secular stagnation” environment, with a ZLB that is binding most of the time, with the possible exception of brief periods in the wake of large increases in the natural rate.⁶

A branch of the literature has uncovered the possibility of multiple equilibria in the presence of the ZLB. A seminal contribution in that literature is Benhabib et al. (2001), which shows the necessary existence of two steady states when the central bank follows a Taylor rule with a ZLB constraint, with the low inflation, low interest rate steady state being globally stable.⁷

⁵In that literature, the inflation target is defined as a parameter of the assumed interest rate rule which has a natural interpretation as an inflation target. Thus, for example, π^* is interpreted as the inflation target in the simple interest rate rule

$$i_t = \max\{0, r^* + \pi^* + \phi_\pi(\pi_t - \pi^*)\}$$

With an occasionally binding ZLB, equilibrium average inflation is generally below the target π^* under interest rate rules of this type. By contrast, in the present paper we do not assume a simple rule and consider instead the fully optimal policy and report the average inflation associated with the implied equilibrium.

⁶The environment analyzed in the present paper is reminiscent of that described in Summers’s celebrated speech on secular stagnation at the 2013 IMF annual Research Conference (Summers (2015)).

⁷Aruoba et al. (2018) estimate a small-scale New Keynesian model with that multiplicity property, in which

Armenter (2018) and Nakata and Schmidt (2019, 2022) show that the multiplicity property may arise also when the central bank follows an optimal discretionary (time-consistent) policy, which leads to the accommodation of occasional declines in private sector’s confidence, triggering a persistent liquidity trap episode. Our analysis, by contrast, focuses on the optimal policy with commitment, which yields a unique (second-best) equilibrium allocation.

Our finding that the optimal policy requires that the nominal rate remains constant at the ZLB most of the time raises the possibility of equilibrium indeterminacy and the challenge of finding a way to implement the constrained-efficient outcome chosen by the central bank. This leads us to propose a nonlinear policy rule which generates a representation of the deviations from the optimal plan in the form of a system with switches between regimes, and for which we study the conditions for uniqueness of its solution. From that perspective, the present paper is related to a branch of the literature that studies the conditions for equilibrium determinacy in regime-switching models. Applications of this literature have typically focused on regime switches driven by *exogenous* stochastic variations in the coefficients of a Taylor-type interest rate rule, which are often assumed to follow a finite-state Markov process. Prominent examples include Davig and Leeper (2007), Farmer et al. (2009) and Barthélemy and Marx (2019). The main difference in our approach is that under our assumed interest rate rule the model’s implied regime switches are *endogenous*, i.e. the regime is a function of the state.⁸ That endogeneity arises as a consequence of the particular nonlinearity embedded in the interest rate rule that implements the optimal allocation, which makes the effective coefficients of the corresponding linear model depend on the (off-equilibrium) deviations of inflation and output from their optimal paths.⁹ We believe our analysis may be of interest beyond the present application, since its validity should carry over to a wide range of linear stochastic models with endogenous

the economy potentially fluctuates between the two steady states in response to sunspot shocks. For Japan (though not the U.S.) they find evidence of an expectations-driven transition to the liquidity trap steady state. Their evidence for the U.S. on the other hand suggests that the ZLB episode in that economy may instead have been the result of adverse fundamental shocks Bullard (2020) makes a case for the relevance of that analysis to the Japanese and U.S. economies. Mertens and Ravn (2014) examine the differential implications of fiscal policy interventions in a neighborhood of the two steady states.

⁸Barthélemy and Marx (2017) also allow for endogeneity of the regime switches but only of a sort with continuous transition probabilities, which rules out the threshold switches that arise naturally in models with a ZLB constraint like ours.

⁹One drawback of our approach, of limited consequence in our particular application, is that it only allows us to derive *sufficient* conditions for determinacy, i.e. we cannot establish necessity, in contrast with the papers mentioned above.

regime switches.

2 The Optimal Monetary Policy Problem

The equilibrium conditions describing the economy’s non-policy block are assumed to be given by

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

$$y_t = \mathbb{E}_t \{ y_{t+1} \} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \{ \pi_{t+1} \} - r_t^n) \quad (2)$$

for $t = 0, 1, 2, \dots$ where π_t denotes inflation, y_t is the output gap, i_t is the short-term nominal rate and r_t^n is the natural rate of interest.¹⁰ Equation (1) is the familiar New Keynesian Phillips curve, which can be derived from the aggregation of firms’ price setting decisions in an environment with price rigidities à la Calvo (1983). Equation (2) is the so-called dynamic IS equation, which results from combining an Euler equation for (log) aggregate consumption, a goods market clearing condition and an equation describing the evolution of output and the real interest rate under flexible prices.¹¹

Variations in the natural rate of interest r_t^n are assumed to be described by

$$r_t^n = r^* + z_t \quad (3)$$

where $\{z_t\}$ follows an exogenous $AR(1)$ process with zero mean, autoregressive coefficient ρ_z and innovation variance σ_z^2 . The unconditional mean of the natural rate is given by r^* , which coincides with the real interest rate, $r_t \equiv i_t - \mathbb{E}_t \{ \pi_{t+1} \}$, in the deterministic steady state. In much of the analysis below we assume

$$r^* < 0 \quad (4)$$

In a companion appendix, we formally describe an environment where (1) and (2) obtain as equilibrium conditions, and where the steady state real interest rate may be negative. The proposed environment is a version of a New Keynesian model with overlapping generations

¹⁰See, e.g., Woodford (2003) or Galí (2015) for a derivation of (1) and (2) in a standard New Keynesian model. In a companion appendix, we show that similar equilibrium conditions obtain in an OLG version of the New Keynesian model that allows for a negative steady state real rate, as considered below.

¹¹Note that we write the previous equations in levels –as opposed to deviations from steady state values– since the steady state is endogenous in our model, and the result of a policy choice. While (1) is derived as a first-order approximation around a zero inflation steady state, we assume the approximation remains valid for small deviations from that steady state, as considered in our analysis.

(NK-OLG) à la Blanchard-Yaari, as developed in Galí (2021).¹² In that environment the steady state real interest is not fully pinned down by the discount rate; instead it also depends on the extent to which income of any given cohort declines over time as a result of retirement or other shocks that make individuals leave employment permanently (e.g. skill obsolescence). That phenomenon tends to enhance savings, lowering the steady state real rate, which may take a negative value.¹³

The monetary authority is assumed to choose at $t = 0$ a state-contingent sequence $\{y_t, \pi_t\}_{t=0}^{\infty}$ that minimizes the welfare loss function

$$\frac{1}{2} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \vartheta y_t^2)$$

subject to the sequence of constraints (1) and (2), as well the ZLB constraint

$$i_t \geq 0 \tag{5}$$

all for $t = 0, 1, 2, \dots$ ¹⁴

Note that the ZLB constraint can be rewritten in terms of inflation and the output gap as:

$$r_t^n + \mathbb{E}_t\{\pi_{t+1}\} + \sigma(\mathbb{E}_t\{y_{t+1}\} - y_t) \geq 0 \tag{6}$$

for $t = 0, 1, 2, \dots$

The (discounted) Lagrangian is given by:

$$\mathcal{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{2} (\pi_t^2 + \vartheta y_t^2) - \xi_{1,t} (\pi_t - \kappa y_t - \beta \pi_{t+1}) - \xi_{2,t} [\pi_{t+1} + \sigma(y_{t+1} - y_t)] \right]$$

The associated optimality conditions are:

$$\pi_t = \xi_{1,t} - \xi_{1,t-1} + \beta^{-1} \xi_{2,t-1} \tag{7}$$

¹²The analysis in Galí (2021) focuses on the possibility of rational bubbles in that environment. Here we assume away that possibility and focus instead on a bubbleless version of the NK-OLG model.

¹³As is well known, other departures from the representative consumer assumption are also consistent with a negative steady state real rate, e.g., models with heterogenous households subject to idiosyncratic income shocks, as in Aiyagari (1994) or Huggett (1993). In contrast with the NK-OLG model, those models do not generally yield an aggregate Euler equation like (2), though the latter has been shown to constitute a good approximation under plausible calibrations (see, e.g., Debortoli and Galí (2022)).

¹⁴As discussed in the companion appendix, the previous loss function can be microfounded as the second order approximation to the expected welfare losses of individuals currently alive in a New Keynesian model with overlapping generations.

$$\vartheta y_t = -\kappa \xi_{1,t} - \sigma \xi_{2,t} + \sigma \beta^{-1} \xi_{2,t-1} \quad (8)$$

$$\xi_{2,t} \geq 0 \quad (9)$$

$$\xi_{2,t} [r_t^n + \mathbb{E}_t\{\pi_{t+1}\} + \sigma(\mathbb{E}_t\{y_{t+1}\} - y_t)] = 0 \quad (10)$$

which should be interpreted as holding for each date and state of nature. The previous conditions, combined with (1), (2), (3), (6) and initial values for $\xi_{1,-1}$ and $\xi_{2,-1}$ (which will depend on the particular problem analyzed) describe the economy's equilibrium under the optimal policy.

In the next two sections, we characterize that equilibrium and provide simulations for a calibrated version of the model. First we study the transitional dynamics in a deterministic environment after an unanticipated shock to r^* . Then we introduce shocks to the natural rate of interest and we look at the economy's response to those shocks in a neighborhood of the new (stochastic) steady state, as implied by the optimal policy.

3 Transitional Dynamics under the Optimal Policy

In the present section we focus on the equilibrium implied by the optimal policy with full commitment (Ramsey) in a deterministic environment. More specifically, we assume that the economy had been in a (deterministic) steady state for some time, with $r_t^n = r^* > 0$, $\pi_t = 0$ and $i_t = r^*$, for $t = -1, -2, \dots$. This is of course the (trivial) outcome of the optimal policy when $r^* > 0$ and in the absence of shocks.¹⁵

At $t = 0$ the economy is assumed to be hit by an unanticipated (MIT-type) shock that lowers r^* permanently, turning it negative, i.e. $r^* < 0$ for $t = 0, 1, 2, \dots$. We start by characterizing the new steady state under the optimal policy. In that steady state we must have $i = \pi + r^* \geq 0$ or, equivalently, $\pi \geq -r^* > 0$. In addition, it follows from (7)-(10) that under the optimal policy:

$$\pi = \beta^{-1} \xi_2 \geq 0$$

$$\vartheta y = -\kappa \xi_1 + \sigma(\beta^{-1} - 1) \xi_2$$

$$\xi_2 \geq 0 ; r^* + \pi = 0 ; \xi_2(r^* + \pi) = 0$$

¹⁵Formally, this can be determined by evaluating (1) and the optimality conditions (7) through (10) at a steady state with $r^* > 0$. The only solution to that system is given by $y = \pi = \xi_1 = \xi_2 = 0$.

It is easy to check that the optimal policy requires that $i = 0$ in the new steady state. To see this, note that if $i > 0$ then $\xi_2 = 0$ implying $\pi = 0$, which is inconsistent with a steady state. Thus the steady state under the optimal policy must satisfy:

$$\begin{aligned}\pi &= -r^* > 0 \\ y &= \frac{1-\beta}{\kappa}\pi = -\frac{1-\beta}{\kappa}r^* > 0 \\ \xi_2 &= \beta\pi = -\beta r^* > 0 \\ \xi_1 &= -\frac{\vartheta}{\kappa}y + \frac{\sigma(\beta^{-1}-1)}{\kappa}\xi_2 \\ &= -\frac{(1-\beta)}{\kappa}\left(\sigma - \frac{\vartheta}{\kappa}\right)r^*\end{aligned}$$

Note that this steady state is (globally) unique. This contrasts with the multiplicity of steady states that generally arise in the presence of the ZLB constraint when the central bank follows a Taylor-type interest rate rule as opposed to the optimal policy under commitment that characterizes our analysis.¹⁶

Next we study the transitional dynamics, i.e. we characterize the equilibrium paths that satisfy

$$\begin{aligned}\widehat{\pi}_t &= \beta\widehat{\pi}_{t+1} + \kappa\widehat{y}_t \\ \widehat{\pi}_t &= \widehat{\xi}_{1,t} - \widehat{\xi}_{1,t-1} + \beta^{-1}\widehat{\xi}_{2,t-1} \\ \vartheta\widehat{y}_t &= -\kappa\widehat{\xi}_{1,t} - \sigma\widehat{\xi}_{2,t} + \sigma\beta^{-1}\widehat{\xi}_{2,t-1} \\ \widehat{\xi}_{2,t} + \xi_2 &\geq 0 \\ \widehat{\pi}_{t+1} + \sigma(\widehat{y}_{t+1} - \widehat{y}_t) &\geq 0 \\ (\widehat{\xi}_{2,t} + \xi_2) [\widehat{\pi}_{t+1} + \sigma(\widehat{y}_{t+1} - \widehat{y}_t)] &= 0\end{aligned}$$

for $t = 0, 1, 2, \dots$ where a " $\widehat{}$ " symbol on a variable denotes deviations from its value in the new steady state. Note also that $\xi_{1,-1} = \xi_{2,-1} = 0$, implying initial conditions $\widehat{\xi}_{1,-1} = -\xi_1$ and $\widehat{\xi}_{2,-1} = -\xi_2$. We restrict ourselves to paths that converge to the new steady state, i.e. $\lim_{t \rightarrow \infty} \widehat{x}_t = 0$ for $\widehat{x}_t \in \{\widehat{\pi}_t, \widehat{y}_t, \widehat{\xi}_{1,t}, \widehat{\xi}_{2,t}\}$.

¹⁶See, e.g., Benhabib et al. (2001) for an analysis of the "perils of multiplicity" when the central bank follows a conventional Taylor rule under a ZLB constraint. Bullard (2020) makes a case for the relevance of their analysis to the Japanese and U.S. economies.

Figure 1 illustrates the transitional dynamics for a calibrated version of our economy.¹⁷ In particular, we assume $\sigma = 1$, $\beta = 0.99$, $\kappa = 0.1717$, $\vartheta = 0.0191$, which are values consistent with the baseline calibration in Galí (2015). In addition, we set $r = -0.0025$, implying an annualized steady state natural rate of minus 1 percent. Interest rates and the inflation rate are shown in annualized terms in all figures.

As shown in Figure 1, the transition to the steady state under the optimal policy is not immediate. Instead, the initial values of inflation and the output gap are significantly below their long run values of 1 and 0.058 percent, respectively, and adjust only gradually towards the new steady state. In fact, inflation is negative for a few periods under our baseline calibration.¹⁸ By choosing a path like the one depicted in Figure 1, the central bank succeeds in keeping inflation close to the first best temporarily, even though it is at the cost of a persistently negative output gap. Given the relative small weight of the latter in the central bank’s loss function under our baseline calibration ($\vartheta \simeq 0.02$), that choice turns out to be more desirable than jumping immediately to the new steady state (which would be perfectly feasible). The persistent low inflation and output gaps are consistent with the observed path for the real rate, which remains above its long run value r during the transition. Most interestingly, the path for the real rate is entirely driven by expected inflation, since the nominal rate remains at the ZLB throughout the transition. Thus, the central bank must implement its nontrivial optimal plan while keeping the setting for its policy instrument unchanged. In section 5 below, we discuss how the central bank may succeed in doing so, given the multiplicity of equilibrium paths consistent with a constant nominal rate.

Our previous analysis of the Ramsey policy made the simplifying assumption that low trend inflation (1%) would not alter much the linearized New Keynesian Phillips curve or the loss function of the central bank (both derived around zero inflation). We can relax this assumption by using the correct approximation of the New Keynesian Phillips curve following Ascari and Ropele (2007) and Ascari and Sbordone (2014), and by lowering the weight on the output gap in the loss function following Lago Alves (2014) by a factor of 0.8. In Figure A1 in the Appendix

¹⁷We use Dynare’s perfect foresight solver, based on Kanzow and Petra (2004), to compute the transition paths.

¹⁸The result of an optimal negative inflation in the short run is not general. In particular, it doesn’t obtain when the weight on the output gap is raised sufficiently (e.g. when $\vartheta = 1$).

we show the counterpart to our Figure 1, based on this modified analysis. One can appreciate small differences in the implied paths (e.g. for the output gap), but overall the results are quite robust to allowing for positive trend inflation.

4 Aggregate Fluctuations under the Optimal Policy

In this section, we characterize the behavior of inflation and the output gap under the optimal policy in a neighborhood of the (stochastic) steady state, in the presence of shocks to the natural rate (i.e. fluctuations in z_t). The (local) equilibrium dynamics are described by the system of stochastic difference equations given by:

$$\begin{aligned}\widehat{\pi}_t &= \beta \mathbb{E}_t \{\widehat{\pi}_{t+1}\} + \kappa \widehat{y}_t \\ \widehat{\pi}_t &= \widehat{\xi}_{1,t} - \widehat{\xi}_{1,t-1} + \beta^{-1} \widehat{\xi}_{2,t-1} \\ \vartheta \widehat{y}_t &= -\kappa \widehat{\xi}_{1,t} - \widehat{\xi}_{2,t} + \beta^{-1} \widehat{\xi}_{2,t-1} \\ \widehat{\xi}_{2,t} + \xi_2 &\geq 0 \\ \sigma(\mathbb{E}_t \{\widehat{y}_{t+1}\} - \widehat{y}_t) + \mathbb{E}_t \{\widehat{\pi}_{t+1}\} + z_t &\geq 0 \\ [\widehat{\xi}_{2,t} + \xi_2][\sigma(\mathbb{E}_t \{\widehat{y}_{t+1}\} - \widehat{y}_t) + \mathbb{E}_t \{\widehat{\pi}_{t+1}\} + z_t] &= 0\end{aligned}$$

for $t = 0, 1, 2, \dots$. We are interested in the equilibrium generated as an outcome of the optimal policy under the timeless perspective, i.e. once the transition to the new steady state has been completed. Accordingly, we assume the initial Lagrange multipliers are at their steady state value, thus implying initial conditions $\widehat{\xi}_{1,-1} = 0$ and $\widehat{\xi}_{2,-1} = 0$. Appendix A describes our approach to determining the solution to the system above.

Figure 2 displays the equilibrium path for inflation and the output gap under the optimal policy, given a sequence of realized values of the shock $\{z_t\}$, drawn from an $AR(1)$ process with autoregressive coefficient $\rho_z = 0.5$ and Gaussian innovations with standard deviations $\sigma_z = 0.0025$. This calibration implies an unconditional standard deviation for the (annualized) natural rate of 1.15 percent. Accordingly, r_t^n remains negative about 80% of the time. The remaining parameters are kept at their baseline settings. The top-left box of the Figure displays the simulated path of the natural rate (in black) and the actual real rate (in blue) over 100

periods. Note that the latter is much smoother than the former, which reflects the central bank’s inability to match fluctuations in the natural rate one-for-one, due to the ZLB constraint. As a result, monetary policy can’t prevent some fluctuations in inflation and the output gap, as illustrated in the two bottom plots. The resulting outcome of the optimal policy is thus clearly second-best.

Most interestingly, we see that the nominal rate remains at the ZLB throughout the simulation, as shown on the top-right plot of figure 2. Thus, the central bank must steer the economy along the optimal path without changing the settings for its policy instrument, and keeping it instead constant at zero.¹⁹ The reason why it does not lower the nominal rate in the face of a negative natural rate is clear: the ZLB prevents it from doing so. Perhaps less obvious is why it chooses to keep the nominal rate at zero even when the natural rate rises above zero. Intuitively, the anticipation that the central bank will keep the interest rate lower than the natural rate when the latter is high helps stabilize inflation and the output gap when the natural rate is low (and can thus not be matched due to the ZLB). More precisely, the stabilizing gains in periods with a low natural rate from the anticipation of a constant zero nominal rate in future periods when the natural rate is positive, more than offset the losses from not matching the natural rate in the latter periods. As a result, the nominal rate remains at the ZLB throughout the simulation. That strategy, which relies on the forward looking nature of aggregate demand and inflation, can thus be viewed as a form of forward guidance.

The property of a constant nominal rate at zero is not general, however. In particular, the central bank may find it desirable to deviate from the constant zero nominal rate policy in response to an increase in the natural rate of interest that is sufficiently large, and which would induce very high inflation if not counteracted at least partly by an increase in the nominal rate. This is illustrated in Figure 3, which shows a simulation of equilibrium fluctuations under the optimal policy, based on a calibration identical to that underlying the simulations of Figure 2 except for a higher shock volatility, with $\sigma_z = 0.0075$. Thus, in the simulation shown in Figure 3 there are three episodes in which the central bank optimally chooses to raise the nominal rate

¹⁹Given the assumed unbounded support of the natural rate, we cannot rule out that the nominal rate could rise above zero temporarily, given a sufficiently long simulation. But as discussed below in the context of Figure 4, this is not even the case under our baseline calibration when we simulate the economy over 10,000 periods. On the other hand, the nominal rate rises occasionally above zero when we increase the variance of the shocks, as discussed below.

above zero, even if only briefly. Roughly speaking, those episodes can be seen to take place when two conditions are met simultaneously: (i) the natural interest rate is unusually high, and (ii) this has not been preceded by a recent episode with an unusually low natural rate, for in the latter case it would have been desirable to keep the nominal rate "low for longer" for the reasons discussed above. Note, however, that the nominal rate remains at the ZLB for much of the simulation, despite the high incidence of a positive natural rate.

In Figure 4 we display the fraction of time that the economy remains at the ZLB under the optimal policy, as a function of r^* , based on a simulation with 10,000 observations for each value of r^* , and under the baseline calibration for the shock volatility ($\sigma_z = 0.0025$) and the remaining parameters. As Figure 4 makes clear, when r^* is sufficiently high (above 3%, roughly), the incidence of the ZLB falls to zero. As r^* decreases, the ZLB incidence starts rising significantly above zero, with the mapping between the two variables becoming quite steep as r^* approaches zero, and reaching unity (i.e. a permanently binding ZLB, effectively) when r^* is about -0.5% or below.

In Figures 5a and 5b, we display, respectively, the mean and standard deviation of inflation under the optimal policy as a function of r^* , under the same baseline calibration as Figure 4. We note that the range of r^* values for which the first best is attained (i.e., for which the mean and standard deviation of inflation are both zero) corresponds to that for which the ZLB is never binding (r^* above 3%, roughly). On the other hand, for a range of r^* roughly between 1% and 3% the optimal policy is associated with an average inflation very close to zero, without being able to fully stabilize that variable (and hence the output gap), due to the positive incidence of binding ZLB episodes.²⁰ For values of r^* below 1%, average inflation becomes positive, and keeps increasing as we lower r^* further. The more r^* moves into negative territory, the more optimal average inflation approaches $-r^*$. That convergence of optimal average inflation to $-r^*$ mirrors the convergence of the average nominal rate to zero, and is thus associated with a

²⁰Our findings for positive values of r^* imply average inflation rates somewhat below those typically found in the literature on the optimal inflation target in the presence of the ZLB (see, e.g., Coibion et al. (2012) and Andrade et al. (2020)). The reason for this is that the previous literature assumes a simple interest rate rule, while we analyze the fully optimal policy. The latter makes it possible to stabilize the economy with a smaller "inflation cushion." Note also that in the above mentioned literature the inflation target is defined as the parameter of the assumed interest rate rule that provides a reference value for inflation. In that context the equilibrium *average* inflation is generally slightly lower than the inflation *target*, due to the occasionally binding ZLB constraint that prevents the central bank from counteracting deflationary episodes.

permanently binding ZLB constraint.²¹

In order to understand that result, note that it follows from equation (2) that average inflation must be equal to the average nominal rate minus r^* . Thus, the ZLB constraint implies a lower bound for average inflation given by $-r^*$. In the *deterministic* economy, zero steady state inflation is feasible and optimal when $r^* \geq 0$. On the other hand, when $r^* < 0$, zero steady state inflation is no longer feasible, and the optimal policy implies average inflation equal to $-r^*$, the lowest feasible value. In the stochastic economy, on the other hand, there may be an incentive to deviate from the previous prescription by allowing for a higher average inflation and nominal rates, in order to build some "policy space" that allows the central bank to better counteract adverse demand shocks. The previous motive can be characterized more precisely by introducing the notion of *precautionary inflation*, denoted by $\bar{\pi}^p$, which we define as the component of average inflation that results from a precautionary motive, i.e. from the desire to limit the incidence of the ZLB. More specifically, we define precautionary inflation for any given r^* as the difference between average inflation under the optimal policy, $\bar{\pi}(r^*)$, and the optimal steady state inflation in the corresponding deterministic economy, which is given by $\max\{0, -r^*\}$. Formally,

$$\bar{\pi}^p(r^*) = \bar{\pi}(r^*) - \max\{0, -r^*\}$$

Figure 6 displays precautionary inflation as a function of r^* under our baseline calibration. Note that the implied mapping is clearly non-monotonic. Thus, for r^* sufficiently high, the risk of a binding ZLB is low, and there is no need to deviate from the first-best outcome of zero inflation at all times. At the other extreme, when r^* is sufficiently negative and, hence, the lower bound on average inflation (given by $-r^*$) is already high, the central bank has little incentive to raise average inflation further above that lower bound, so it chooses to keep average inflation at the same level as in the deterministic case. By contrast, precautionary inflation is strictly positive for a range of r^* values closer to zero, for which optimal inflation in the

²¹Our finding of a one-to-one (inverse) mapping between optimal average inflation and r^* is reminiscent of a similar finding in Andrade et al. (2020). In the latter paper, however, that finding emerges for positive but relatively low values of r^* , while in the present paper it does so only for negative values of r^* . The reason for the difference lies in the different assumptions made on the nature of policy (simple rule vs. fully optimal with commitment). In our analysis, the optimal policy with commitment makes it possible to limit the incidence of costly ZLB episodes while maintaining an average inflation close to zero; this is not feasible under the simple rule assumed in Andrade et al. (2020), so a higher inflation target is desirable in that case.

deterministic case is either zero (if $r^* \gtrsim 0$) or positive but low (if $r^* \lesssim 0$), since in that case the costs of deviations from full price stability are relatively low, and are outweighed by the gains from a lower incidence of a binding ZLB made possible by the choice of a higher average inflation.

As illustrated previously by Figures 2 and 3, the extent of ZLB incidence does not only depend on r^* but also on the volatility of the natural rate. This is confirmed and shown more clearly in Figure 7, which displays (in the shaded area) the set of values for r^* and σ_z for which the ZLB is (near) permanently binding.²² Three observations are worth making. First, we see that an equilibrium with a (near) permanently binding ZLB emerges under the optimal policy only if $r^* < 0$. Secondly, for any given negative r^* , the ZLB constraint becomes (near) permanently binding under the optimal policy as long as σ_z is sufficiently low. Finally, we see that the lower is r^* the larger is the volatility of the natural rate required in order to observe, even if only occasionally, a positive nominal rate under the optimal policy.

Similar qualitative findings to those discussed in the present section emerge when we replace shocks to the natural rate with cost-push shocks, i.e. exogenous disturbances to the New Keynesian Phillips curve (1). As is well known, in that case a trade-off between inflation stabilization and output gap stabilization emerges independently of the presence of a ZLB (see, e.g., Clarida et al. (1999)), with the optimal policy calling for output gap variations in order to dampen fluctuations in inflation. As in the environment analyzed above, with a negative r^* and relatively small shocks, the (second-best) management of output and inflation fluctuations is consistent with a nominal rate that remains at zero throughout our simulation. In response to sufficiently large positive (i.e. inflationary) cost push shocks, on the other hand, the policy rate under the optimal policy temporarily rises above zero. This is illustrated in Figures A2a and A2b in the Appendix, which display simulations of the equilibrium outcomes under the optimal policy in the presence of cost-push shocks. See Appendix B for a presentation of the modified model and calibration.

How the central bank manages to steer the economy as required by the solution to its

²²We use the *near* qualifier to stress the fact that we cannot rule out in theory the possibility of brief episodes with positive nominal rates in response to a sequence of large positive realizations of the natural rate, even though that event has not materialized even once over a 10,000 period simulation for the calibrations for which a unit incidence of a binding ZLB is reported.

optimal policy problem while keeping the nominal rate unchanged at zero is the subject of the next section.

5 Optimal Monetary Policy Implementation under a ZLB Constraint

Let (i_t^*, y_t^*, π_t^*) denote the central bank's optimal plan, i.e. the solution to the policy problem analyzed in the previous sections. Next we consider one particular way of implementing that plan as the unique equilibrium, through an interest rate rule that responds to eventual deviations from that plan. We then impose conditions on the rule coefficients that guarantee that such deviations do not occur in any bounded equilibrium (i.e. they become off-equilibrium paths).

Consider thus deviations from the optimal plan satisfying the equilibrium conditions (1), (2) and (5). Formally, and letting $\tilde{\pi}_t \equiv \pi_t - \pi_t^*$, $\tilde{y}_t \equiv y_t - y_t^*$ and $\tilde{i}_t \equiv i_t - i_t^*$, we have

$$\tilde{\pi}_t = \beta \mathbb{E}_t \{ \tilde{\pi}_{t+1} \} + \kappa \tilde{y}_t \quad (11)$$

$$\tilde{y}_t = \mathbb{E}_t \{ \tilde{y}_{t+1} \} - \frac{1}{\sigma} (\tilde{i}_t - \mathbb{E}_t \{ \tilde{\pi}_{t+1} \}) \quad (12)$$

as well as the ZLB constraint

$$\tilde{i}_t \geq -i_t^* \quad (13)$$

for all t .²³

We complement the previous equations with the following piece-wise linear interest rate rule

$$i_t = i_t^* + \phi_\pi^{(q)} \tilde{\pi}_t + \phi_y^{(q)} \tilde{y}_t \quad (14)$$

where $q \in \{1, 2, 3, 4\}$ denotes the "regime" prevailing at each point in time, which is determined by the sign configuration of the deviations of inflation and the output gap from their values under the optimal plan. Thus, the central bank conducts monetary policy by setting the nominal rate to i_t^* as prescribed by the optimal plan, unless inflation and/or the output gap deviate from their corresponding optimal paths, in which case the nominal rate responds to those deviations according to (14).

²³Note that the previous representation in terms of equilibrium deviations from the optimal plan holds independently of the underlying source of fluctuations (natural rate shocks or cost-push shocks). More generally, (i_t^*, y_t^*, π_t^*) can be interpreted as the central bank's desired equilibrium path, which may or may not coincide with the solution to the optimal policy problem analyzed above.

More specifically, we define the following regimes, with the corresponding sign restrictions on their associated rule coefficients:

$$\begin{aligned}
q = 1 & : \tilde{\pi}_t \geq 0, \tilde{y}_t \geq 0 \Rightarrow \phi_\pi^{(1)} \geq 0, \phi_y^{(1)} \geq 0 \\
q = 2 & : \tilde{\pi}_t < 0, \tilde{y}_t < 0 \Rightarrow \phi_\pi^{(2)} \leq 0, \phi_y^{(2)} \leq 0 \\
q = 3 & : \tilde{\pi}_t \geq 0, \tilde{y}_t < 0 \Rightarrow \phi_\pi^{(3)} \geq 0, \phi_y^{(3)} \leq 0 \\
q = 4 & : \tilde{\pi}_t < 0, \tilde{y}_t \geq 0 \Rightarrow \phi_\pi^{(4)} \leq 0, \phi_y^{(4)} \geq 0
\end{aligned} \tag{15}$$

Note that the specification of rule (14) together with the sign restrictions in (15) guarantee that $i_t \geq i_t^* \geq 0$ for all t , thus meeting the ZLB constraint (13) at all times, *even on any off-equilibrium path*.

Note that $\tilde{\pi}_t = \tilde{y}_t = \tilde{i}_t = 0$ for all t is always a solution to the system (11)-(13), and the one which corresponds to the desired outcome, i.e. the optimal plan. Our objective is to study the conditions on $\phi_\pi^{(q)}$ and $\phi_y^{(q)}$, for $q \in \{1, 2, 3, 4\}$ that guarantee that the previous solution is (locally) unique or, equivalently, that the optimal plan is effectively implemented.²⁴

We tackle this problem by treating (11)-(12) as a regime switching model, with *endogenous* regime switches. Then we apply a novel result that allows us to establish sufficient conditions for the (local) uniqueness of the solution of an endogenous regime switching model. The advantage of our approach is that we do not need to specify a law of motion describing the transition across regimes. Given the potential interest of the latter result beyond the problem at hand, we first state it for a more general setting before we apply it to the model above.

5.1 A Sufficient Condition for Equilibrium Determinacy of an Endogenous Regime Switching Model

Consider a regime switching model whose equilibrium is described by a system of difference equations of the form:

$$\mathbf{x}_t = \mathbf{A}_t \mathbb{E}_t \{ \mathbf{x}_{t+1} \} \tag{16}$$

where \mathbf{x}_t is an $(n \times 1)$ vector of non-predetermined variables and A_t is an $(n \times n)$ matrix. We assume $\mathbf{A}_t \in \mathcal{A}$ where $\mathcal{A} \equiv \{ \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \dots, \mathbf{A}^{(Q)} \}$ is a finite set of $(n \times n)$ nonsingular matrices. The evolution of \mathbf{A}_t over time is left unspecified. It may evolve exogenously, e.g. according to a

²⁴The fact that the proposed interest rate rule includes only contemporaneous values should not be interpreted as suggesting the the optimal policy is not history-dependent, since that history-dependence is already embedded in the "targets" π_t^* , y_t^* and i_t^* .

Markov process. Alternatively, \mathbf{A}_t may vary endogenously, i.e. it may be a function of current and lagged values of \mathbf{x}_t .

It is clear that $\mathbf{x}_t = 0$ for all t is a solution to (16). Our goal is to establish *sufficient* conditions on \mathcal{A} that guarantee that $\mathbf{x}_t = 0$ for all t is the only bounded solution to (16). We take this to be the case if $\lim_{T \rightarrow +\infty} \mathbb{E}_t\{\|\mathbf{x}_{t+T}\|\} > M\|\mathbf{x}_t\|$ for any scalar $M > 0$ and $\mathbf{x}_t \neq 0$, and where $\|\cdot\|$ is the usual L^2 norm.

Let us define the induced matrix norm $\|\mathbf{A}\| \equiv \max_{\mathbf{x}} \|\mathbf{A}\mathbf{x}\|$ subject to $\|\mathbf{x}\| = 1$. In addition, define $\alpha \equiv \max\{\|\mathbf{A}^{(1)}\|, \|\mathbf{A}^{(2)}\|, \dots, \|\mathbf{A}^{(Q)}\|\}$. Note that nonsingularity of $\mathbf{A}^{(q)}$ for $q = 1, 2, \dots, Q$ implies $\alpha > 0$.

Theorem [*sufficient condition for determinacy*]: If $\alpha < 1$, then $\mathbf{x}_t = 0$ for all t is the only bounded solution to (16)

Proof: See Appendix C

Remark: the previous condition is sufficient but not necessary. As a counterexample consider a switching regime model given by (16) with $\mathbf{A}_t = \mathbf{A}^{(1)}$ for odd t and $\mathbf{A}_t = \mathbf{A}^{(2)}$ for even t , where

$$\mathbf{A}^{(1)} = \begin{bmatrix} 1.1 & 0 \\ 0 & 0.5 \end{bmatrix} \quad ; \quad \mathbf{A}^{(2)} = \begin{bmatrix} 0.5 & 0 \\ 0 & 1.1 \end{bmatrix}$$

Note that the previous model does not satisfy the sufficiency condition since $\alpha = 1.1 > 1$. Yet, $\mathbf{x}_t = 0$ can be shown to be the only bounded solution. See Appendix D for a proof.²⁵

Remark: note that $\|\mathbf{A}\| < 1$ implies that all the eigenvalues of \mathbf{A} lie within the unit circle, though the converse is not true. See Appendix E for a proof. Hence our sufficient condition $\alpha < 1$ also implies that $\mathbf{x}_t = 0$ is the unique bounded solution for each of the "single regime" models $\mathbf{x}_t = \mathbf{A}^{(q)}\mathbb{E}_t\{\mathbf{x}_{t+1}\}$, for $q = 1, 2, \dots, Q$. By contrast, *under the usual eigenvalue criterion*, the equilibrium may be unique for each of the "single regime" models but indeterminate for the regime-switching model. The latter possibility is discussed in Barthélemy and Marx (2019) in the context of a New Keynesian model with exogenous switches in the interest rate rule coefficients: they show how indeterminacy may emerge even if each of the regimes adheres to the Taylor principle when considered in isolation (i.e. it satisfies the eigenvalue condition for uniqueness in the corresponding single regime economy). Our strengthened condition, in terms

²⁵We thank Danila Smirnov for suggesting this counterexample.

of the norm of the $\mathbf{A}^{(q)}$ matrices as opposed to their eigenvalues, rules out such a possibility: if the norm condition is satisfied for each of the regimes in isolation, then it is also satisfied "globally" for the regime-switching model.

Remark: an alternative sufficient condition for determinacy is given by $\rho(\mathcal{A}) < 1$, where $\rho(\mathcal{A}) \equiv \lim_{T \rightarrow +\infty} \max\{\|A_{i_1} A_{i_2} \cdots A_{i_T}\|^{\frac{1}{T}} : A_i \in \mathcal{A}\}$ is the joint spectral radius of \mathcal{A} . The proof is almost identical to that in Appendix C. Note that this alternative condition is weaker than $\alpha < 1$ but is not necessary either. In particular, the counterexample above also applies, since $\rho(\mathcal{A}) > 1.1$. We prefer to work with the norm condition since it is easier to check computationally.

5.2 Application to the Problem of Optimal Monetary Policy Implementation

Next, we apply the result of the previous subsection to the problem of implementation of the optimal monetary policy analyzed above. Recall that feasible deviations from the optimal outcome are described by (11), (12) and (14), with the latter effectively defining four regimes. Plugging (14) into (12) to eliminate \tilde{i}_t , and after some straightforward substitutions, we can represent the dynamics for $\mathbf{x}_t \equiv [\tilde{y}_t, \tilde{\pi}_t]'$ as in (16), with

$$\mathbf{A}^{(q)} \equiv \frac{1}{\sigma + \phi_y^{(q)} + \kappa \phi_\pi^{(q)}} \begin{bmatrix} \sigma & 1 - \beta \phi_\pi^{(q)} \\ \sigma \kappa & \kappa + \beta(\sigma + \phi_y^{(q)}) \end{bmatrix}$$

for $q \in \{1, 2, 3, 4\}$, corresponding to the four regimes defined above.

The coloured areas in Figure 8 display the configurations of $(\phi_\pi^{(q)}, \phi_y^{(q)})$ values for which $\alpha < 1$, i.e. for which $\|\mathbf{A}^{(q)}\| < 1$, for $q \in \{1, 2, 3, 4\}$. Note that each regime corresponds to a different quadrant of the Figure, with the corresponding configurations of determinacy-inducing coefficients depicted in a different color for each regime. Thus, to the extent that the central bank adopts rule (14) with coefficients that fall within the depicted regions under each regime, no deviations from the desired allocation will be consistent with a (bounded) equilibrium, and hence the adopted rule will indeed implement the desired allocation (y_t^*, π_t^*) , while satisfying the ZLB constraint. For completeness, Figure 8 also displays in light grey the set of $(\phi_\pi^{(q)}, \phi_y^{(q)})$ values for which the two eigenvalues of $\mathbf{A}^{(q)}$ fall within the unit circle, which correspond to the necessary and sufficient condition for (local) uniqueness in a single regime economy. Note that,

for each regime the light grey area subsumes the coloured regions, consistent with the fact that the former represent necessary and sufficient conditions, while the latter only sufficient, for each single regime model.

Finally, a word about some of the rule’s implications. As discussed above, the rule instructs the central bank to deviate from the interest rate i_t^* implied by the optimal policy if and only if inflation and/or output deviate from their optimal values, π_t^* and y_t^* . If the rule coefficients satisfy the sufficient condition for a unique equilibrium (as assumed in our simulations), those deviations remain off-equilibrium, i.e. they never materialize ex-post. While the previous feature is often found in interest rate rules that implement a desired feasible allocation,²⁶ a specific characteristic of our nonlinear rule is that, by construction, all its implied off-equilibrium deviations for the nominal rate are *positive*, i.e. they involve raising the nominal interest rate above i_t^* . That property guarantees that the ZLB constraint is never violated, not even on off-equilibrium paths, given that $i_t^* \geq 0$ for all t .

Needless to say, some of the off-equilibrium interest rate movements called for by the rule may be perceived *ex-post* as being suboptimal (e.g. raising the interest rate if inflation falls below its level under the optimal plan), but this sort of time inconsistency is inherent to optimal policies under commitment even in the absence of the ZLB constraint, their benefits arising from the (desirable) effects of their anticipation (as it is the case here).²⁷

6 Concluding Remarks

The analysis in the present paper has shown that in response to a permanent decline in the natural rate of interest, so that the latter’s mean, r^* , becomes negative, a central bank may optimally choose to keep the policy rate at zero permanently . We have also shown that in such an environment, and despite the possible constancy of the policy rate, there is still a meaningful optimal policy problem: a fully credible central bank operating under commitment can keep influencing macro outcomes and implement the constrained-efficient allocation in the face of

²⁶See, e.g., the discussion in Galí (2015, chapters 4 and 5) regarding the implementation of optimal policies through interest rate rules, in the context of a baseline New Keynesian model without a ZLB constraint.

²⁷Departures from the assumption of full credibility adopted here will generally have implications on the optimal policy outcomes. Given the absence of a widely accepted model of imperfect credibility we do not pursue this avenue here.

continuous shocks that may impinge on the economy.

More specifically, we have studied the optimal monetary policy problem in a New Keynesian economy with a zero lower bound (ZLB) on the nominal interest rate, and in which r^* becomes permanently negative. In the deterministic case the optimal policy aims to approach *gradually* the new steady state with positive average inflation, while keeping the policy rate at zero. A gradualist approach minimizes welfare losses by keeping inflation close to zero for longer.

In the presence of shocks to the natural rate of interest, and once the new (stochastic) steady state has been attained, the optimal policy problem yields unique optimal paths for inflation and the output gap. If r^* is sufficiently negative and the shocks to the natural rate are not too large, the optimal policy requires that the nominal rate remains at its ZLB permanently.

Finally we have shown that the central bank can implement the optimal policy as a (locally) unique equilibrium by means of an appropriate nonlinear state-contingent rule consistent with the ZLB. In order to establish that result, we derive a sufficient condition for local determinacy in a more general model of endogenous regime switches. That result may be of interest beyond the problem studied in the present paper.

In order to keep the analysis as close as possible to that of the standard monetary policy problem in the New Keynesian model, we have abstracted from both quantitative easing (QE) and fiscal policy, among other possible instruments. Those additional policy instruments may help improve the outcome in the face of a permanently negative r^* . In the case of QE, the analysis of its role would require modifying the standard New Keynesian environment in order to overcome the well-known irrelevance result (Eggertsson and Woodford (2003)) and render it effective independently of interest rate policy.²⁸ We plan to pursue that analysis in future work.

²⁸See, e.g., Nisticò and Seccareccia (2022).

REFERENCES

- Aiyagari, Rao (1994): "Uninsured Idiosyncratic Risk and Aggregate Savings," *Quarterly Journal of Economics* 109 (3), 659–684,
- Adam, Klaus and Roberto Billi (2006): "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates," *Journal of Money, Credit and Banking* 38, 1877-1905.
- Andrade, Philippe, Jordi Galí, Hervé Le Bihan, and Julien Matheron (2020): "The Optimal Inflation Target and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, Fall issue, 173-230
- Andrade, Philippe, Jordi Galí, Hervé Le Bihan, and Julien Matheron (2021): "Should the ECB Adjust its Strategy in the Face of a Lower r^* ?", *Journal of Economic Dynamics and Control*, Fall issue, 173-230
- Armenter, Roc (2019): "The Perils of Nominal Targets," *Review of Economic Studies* 85, 50-86.
- Aruoba, Boragan, Pablo Cuba-Borda, Frank Schorfheide (2018): "Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries," *Review of Economic Studies* 85, 87-118.
- Ascari, Guido and Ropele, Tiziano (2007): "Optimal Monetary Policy under Low Trend Inflation," *Journal of Monetary Economics* 54(8), 2568-2583.
- Ascari, Guido, and Argia M. Sbordone (2014): "The Macroeconomics of Trend Inflation," *Journal of Economic Literature* 52(3), 679-739.
- Benhabib, Jess, Stephanie Schmitt-Grohe, and Martin Uribe (2001): "The Perils of Taylor Rules," *Journal of Economic Theory* 96, 40-69.
- Bernanke, Ben S., Michael T. Kiley, and John M. Roberts (2019): "Monetary Policy Strategies for a Low Rate Environment," *AEA Papers and Proceedings* 109, 421-426.
- Barthélemy, Jean and Magali Marx (2017): "Solving endogenous regime switching models," *Journal of Economic Dynamics and Control* 77, 1-25.
- Barthélemy, Jean and Magali Marx (2019): "Monetary Policy Switching and Indeterminacy," *Quantitative Economics* 10, 353-385.
- Billi, Roberto, Jordi Galí and Anton Nakov (2022): "Appendix to Optimal Monetary Policy with $r^* < 0$," mimeo.

Brand, Claus, and Falk Mazelis (2019): "Taylor-rule consistent estimates of the natural rate of interest." ECB Working Paper no. 2257.

Bullard, James (2020): "Seven Faces of the 'Peril'," Federal Reserve Bank of St. Louis Review, 339-352.

Calvo, Guillermo (1983): "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics*, 12, 383-398.

Clarida, Richard, Jordi Galí, and Mark Gertler (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, vol. 37, 1661-1707

Coibion, Olivier Yuriy Gorodnichenko, and Johannes Wieland (2012): "The Optimal Rate of Inflation in New Keynesian Models: Should Central Banks Raise their Inflation Target in Light of the Zero Lower Bound," *Review of Economic Studies* 20, 1-36.

Davig, Troy and Eric M. Leeper (2007): "Generalizing the Taylor Principle," *American Economic Review* 97, 607-635.

Davis, Josh, Cristian Fuenzalida, Leon Huetsch, Benjamin Mills and Alan M. Taylor (2023): "Global Natural Rates in the Long Run: Postwar Macro Trends and the Market Implied r^* in 10 Advanced Economies," paper presented at the NBER ISOM 2023 conference.

Debortoli, Davide and Jordi Galí (2022): "Idiosyncratic Income Risk and Aggregate Fluctuations," mimeo.

Del Negro, Marco, Domenico Giannone, Marc Giannoni and Andrea Tambalotti (2019): "Global trends in interest rates," *Journal of International Economics* 118, 248-262.

Eggertsson, Gauti, and Michael Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, vol. 1, 139-211.

Eggertsson, Gauti, Neil R. Mehrotra and Jacob A. Robbins (2019): "A Model of Secular Stagnation: Theory and Quantitative Evaluation," *American Economic Journal: Macroeconomics* 11(1), 1-48.

Farmer, Roger E.A., Daniel F. Waggoner, and Tao Zha (2009): "Understanding Markov-Switching Rational Expectations Models," *Journal of Economic Theory* 144, 1849-1867.

Galí, Jordi (2015): *Monetary Policy, Inflation and the Business Cycle. An Introduction to the New Keynesian Framework*, Second edition, Princeton University Press (Princeton, NJ).

Galí, Jordi (2021): "Monetary Policy and Bubbles in a New Keynesian Model with Over-

lapping Generations," *American Economic Journal: Macroeconomics* 13(2), 121-167

Holston, Kathryn, Thomas Laubach and John C. Williams (2017): "Measuring the natural rate of interest: International trends and determinants," *Journal of International Economics* 108, S59-S75.

Kanzow, Christian and Stefania Petra (2004): "On a semismooth least squares formulation of complementarity problems with gap reduction," *Optimization Methods and Software*, 19, 507-525.

Huggett, Mark (1993): "The Risk Free Rate in Heterogeneous Agent Incomplete Insurance Economies," *Journal of Economic Dynamics and Control* 17 (5-6), 953-969.

Jung, Taehun, Yuki Teranishi, and Tsutomu Watanabe, (2005): "Optimal Monetary Policy at the Zero Interest Rate Bound," *Journal of Money, Credit and Banking* 37 (5), 813-835.

Krugman, Paul (1998): "It's Baaaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, vol. 2, 137-187.

Lago Alves, Sergio (2014): "Lack of Divine Coincidence in New Keynesian Models," *Journal of Monetary Economics* 67(C), 33-46

Mertens, Karel and Morten O. Ravn (2014): "Fiscal Policy in an Expectations Driven Liquidity Trap," *Review of Economic Studies* 81, 1637-1667.

Miranda, Mario J. and Paul L. Fackler (2002): *Applied Computational Economics and Finance*, MIT Press (Cambridge, MA)

Nakata, Taisuke and Sebastian Schmidt (2019): "Conservatism and Liquidity Traps," *Journal of Monetary Economics* 104, 37-47.

Nakata, Taisuke and Sebastian Schmidt (2022): "Expectations-Driven Liquidity Traps: Implications for Monetary and Fiscal Policy," *American Economic Journal: Macroeconomics* 14(4), 68-103.

Nakov, Anton (2008): "Optimal and Simple Monetary Policy Rules with a Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking* vol 4(2), 73-127.

Nisticò, Salvatore and Marialaura Seccareccia (2022): "Unconventional Monetary Policy and Inequality," mimeo.

Rachel, Lucasz and Thomas D. Smith (2017): "Are Low Interest Rates Here to Stay?" *International Journal of Central Banking*, 1-42.

Rachel, Lucasz and Lawrence H. Summers (2019): "On Secular Stagnation in the Industrialized World," *Brookings Papers on Economic Activity*, Spring issue, 1-54.

Summers, Lawrence H. (2015): "Have we Entered an Age of Secular Stagnation?" *IMF Economic Review* 63(1), 277-280.

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

APPENDIX A: Solving for the local equilibrium dynamics under the optimal policy

We use the numerical algorithm for solving rational expectations models as implemented in the CompEcon toolkit of Miranda and Fackler (2002). In particular, we solve for the optimal policy x as a function of the state s , when equilibrium is governed by a system of the form

$$f[s_t, x_t, E_t h(s_{t+1}, x_{t+1})] = \xi_t$$

where s follows the state transition function

$$s_{t+1} = g(s_t, x_t, \varepsilon_{t+1})$$

and x_t and ξ_t in our case satisfy the following Kuhn-Tucker condition

$$i_t \geq 0, \quad \xi_{2t} \geq 0, \quad i_t > 0 \Rightarrow \xi_{2t} = 0.$$

The solution is obtained with the collocation method, which consists of approximating the expectation functions by linear combinations of known basis functions, θ_j . The corresponding coefficients, c_j , are determined by requiring the approximating function to satisfy the equilibrium equations *exactly* at n collocation nodes:

$$h[s, x(s)] \approx \sum_{j=1}^n c_j \theta_j(s)$$

For a given value of the coefficient vector c , the equilibrium policies x_i are computed at the n collocation nodes s_i by solving a standard root-finding problem. The coefficient vector c is updated solving the n -dimensional linear system

$$\sum_{j=1}^n c_j \theta_j(s_i) = h(s_i, x_i)$$

The previous iterative procedure is repeated until the distance between successive values of c becomes sufficiently small. To approximate the expectation functions, we discretize the innovation to r_t^n using a K -node Gaussian quadrature scheme:

$$Eh[s, x(s)] \approx \sum_{k=1}^K \sum_{j=1}^n \omega_k c_j \theta_j [g(s_i, x, \varepsilon_k)]$$

where ε_k and ω_k are Gaussian quadrature nodes and weights chosen so that the discrete distribution approximates the continuous univariate normal distribution $N(0, \sigma^2)$. We use linear splines on a uniform grid of 200 points for values of the natural rate of interest between -10 percent and $+10$ percent, so that each point on the grid corresponds to 10 basis points.

APPENDIX B: The Case of Cost-Push Shocks

As a robustness exercise we have also analyzed the optimal policy in response to inefficient or "cost-push" shocks. In that case the equilibrium conditions describing the economy's non-policy block are given by:

$$\pi_t = \beta \mathbb{E}_t \{\pi_{t+1}\} + \kappa y_t + u_t \quad (\text{A.1})$$

$$y_t = \mathbb{E}_t \{y_{t+1}\} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \{\pi_{t+1}\} - r^*) \quad (\text{A.2})$$

where $\{u_t\}$ is a cost push shock that is assumed to follow an $AR(1)$ process with autoregressive coefficient ρ_u and white noise Gaussian innovations with variance σ_u^2 .

The analysis of the optimal monetary policy proceeds unchanged, except for the two equations above. Figure A.2a displays simulations associated with the optimal policy under our baseline calibration, with $\rho_u = 0.5$ and $\sigma_u^2 = (0.00125)^2$ and with the remaining parameters unchanged. Figure A.2b displays analogous results for the case of "large" shocks, with $\sigma_u^2 = (0.0075)^2$. As discussed in the main text, the qualitative findings are analogous to those obtained under the assumption of shocks to the natural rate of interest.

APPENDIX C: Proof of Theorem [sufficiency conditions for determinacy]

By the law of iterated expectations

$$\begin{aligned} \mathbf{x}_t &= \mathbf{A}_t \mathbb{E}_{t+T-1} \{\mathbf{x}_{t+1}\} \\ &= \mathbb{E}_t \{\mathbf{A}_t \mathbf{A}_{t+1} \cdots \mathbf{A}_{t+T-1} \mathbf{x}_{t+T}\} \end{aligned}$$

Thus,

$$\begin{aligned} \|\mathbf{x}_t\| &= \|\mathbb{E}_t \{\mathbf{A}_t \mathbf{A}_{t+1} \cdots \mathbf{A}_{t+T-1} \mathbf{x}_{t+T}\}\| \\ &\leq \mathbb{E}_t \{\|\mathbf{A}_t \mathbf{A}_{t+1} \cdots \mathbf{A}_{t+T-1} \mathbf{x}_{t+T}\|\} \\ &\leq \mathbb{E}_t \{\|\mathbf{A}_t \mathbf{A}_{t+1} \cdots \mathbf{A}_{t+T-1}\| \|\mathbf{x}_{t+T}\|\} \\ &\leq \alpha^T \mathbb{E}_t \{\|\mathbf{x}_{t+T}\|\} \end{aligned}$$

where the last inequality uses the fact that

$$\|A_{i_1}A_{i_2}\cdots A_{i_T}\| \leq \|A_{i_1}\| \|A_{i_2}\| \cdots \|A_{i_T}\| \leq \alpha^T$$

where $A_i \in \mathcal{A}$.

Accordingly, $\alpha < 1$ implies that $\lim_{T \rightarrow +\infty} \mathbb{E}_t\{\|\mathbf{x}_{t+T}\|\} > M \|\mathbf{x}_t\|$ for any arbitrarily large $M > 0$ and $\mathbf{x}_t \neq 0$. *QED*.

APPENDIX D [A Counterexample]

Letting $\mathbf{A} \equiv \mathbf{A}^{(1)}\mathbf{A}^{(2)} = \mathbf{A}^{(2)}\mathbf{A}^{(1)}$ we can write

$$\mathbf{x}_t = \mathbf{A}^T \mathbb{E}_t\{\mathbf{x}_{t+2T}\}$$

Thus,

$$\begin{aligned} \|\mathbf{x}_t\| &\leq \|\mathbf{A}^T\| \mathbb{E}_t\{\|\mathbf{x}_{t+2T}\|\} \\ &= \|\mathbf{A}\|^T \mathbb{E}_t\{\|\mathbf{x}_{t+2T}\|\} \end{aligned}$$

In our numerical example $\|\mathbf{A}\| = 0.55 < 1$. Accordingly,

$$\mathbb{E}_t\{\|\mathbf{x}_{t+2T}\|\} = 0.55^{-T} \|\mathbf{x}_t\|$$

which implies $\lim_{T \rightarrow +\infty} \mathbb{E}_t\{\|\mathbf{x}_{t+T}\|\} > M \|\mathbf{x}_t\|$ for any arbitrarily large $M > 0$ and $\mathbf{x}_t \neq 0$. *QED*.

APPENDIX E [Eigenvalue vs. Norm Criteria]

Let \mathbf{A} be a nonsingular matrix with $\|\mathbf{A}\| < 1$. Thus, $0 < \mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x} < 1$ for all \mathbf{x} such that $\|\mathbf{x}\| = 1$. Let \mathbf{Q} be the matrix of (orthonormal) eigenvectors of $\mathbf{A}'\mathbf{A}$ and let Υ be the corresponding (diagonal) matrix with (real) eigenvalues on its diagonal. Thus, $\mathbf{A}'\mathbf{A}\mathbf{Q} = \mathbf{Q}\Upsilon$ with $\mathbf{Q}'\mathbf{Q} = \mathbf{I}$. Hence $\mathbf{Q}'\mathbf{A}'\mathbf{A}\mathbf{Q} = \Upsilon$, with all diagonal elements of Υ between zero and one. Thus we can write $\mathbf{A}'\mathbf{A} = \mathbf{Q}\Upsilon\mathbf{Q}'$ or, equivalently, $\mathbf{A}'\mathbf{Q}\mathbf{Q}'\mathbf{A} = (\mathbf{Q}\Upsilon^{\frac{1}{2}})(\Upsilon^{\frac{1}{2}}\mathbf{Q}')$ implying $\mathbf{A}'\mathbf{Q} = \mathbf{Q}\Upsilon^{\frac{1}{2}}$. Thus the eigenvalues of \mathbf{A}' (and, hence, of \mathbf{A} , since both share the same characteristic polynomial) are given by the diagonal elements of $\Upsilon^{\frac{1}{2}}$ and are thus real and between zero and one. This is precisely the condition for determinacy in a single regime model.

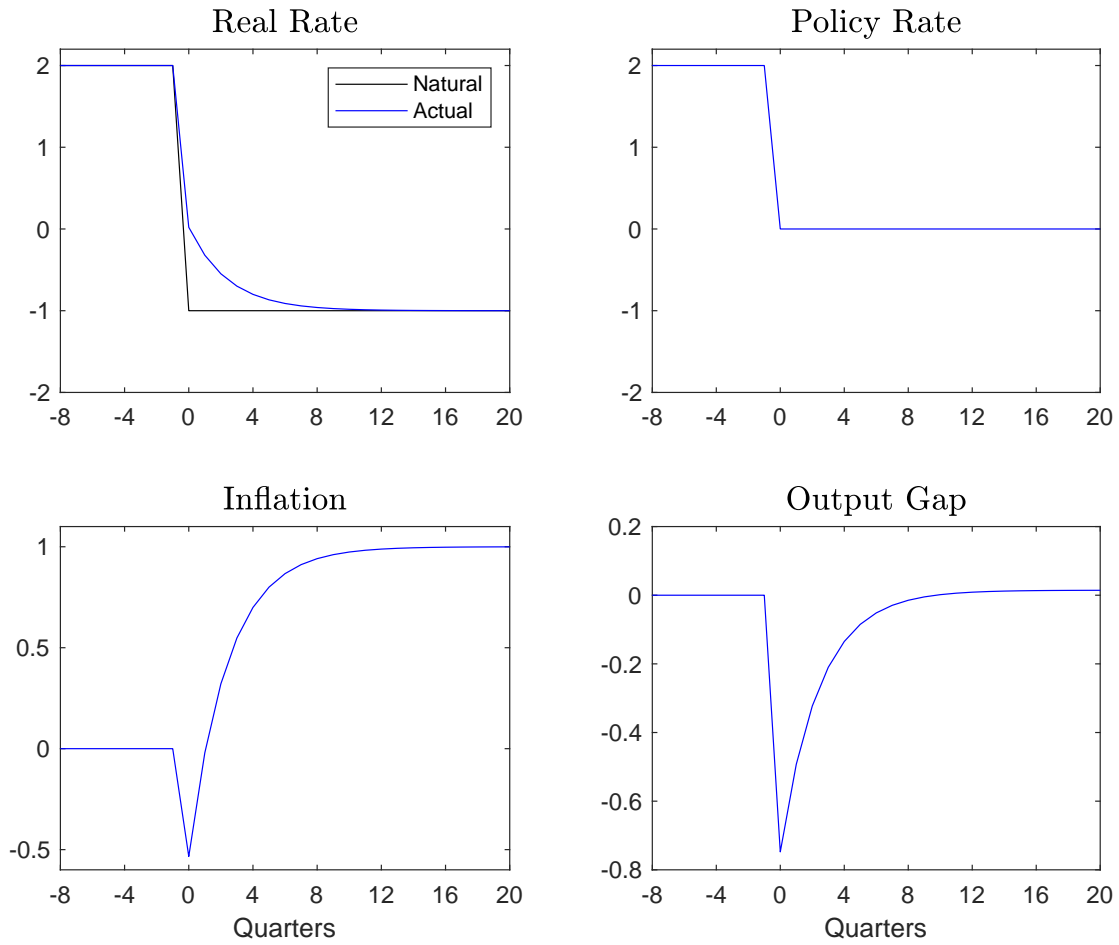


Figure 1: Transitional dynamics under the optimal monetary policy. Inflation and interest rates in annualized terms.

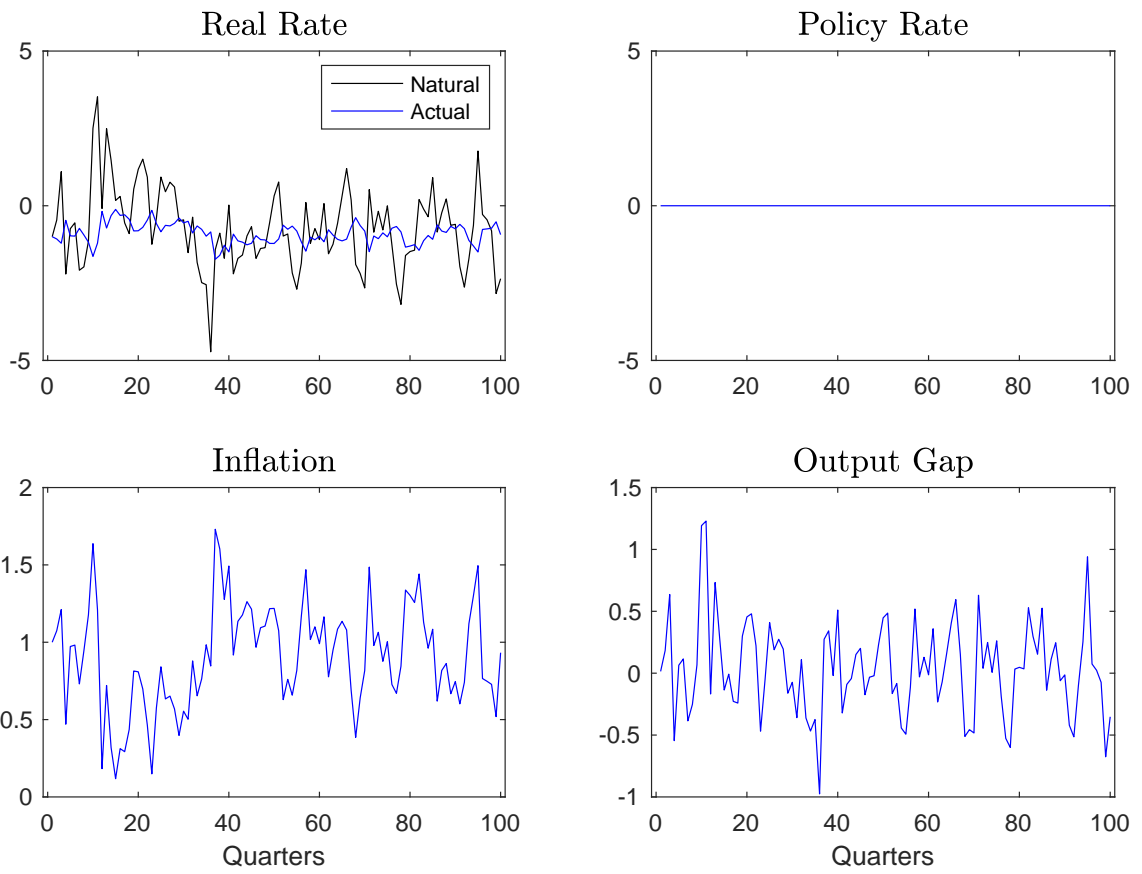


Figure 2: Aggregate fluctuations under the optimal monetary policy and baseline calibration. Inflation and interest rates in annualized terms.



Figure 3: Aggregate fluctuations under the optimal monetary policy with higher shock volatility. Inflation and interest rates in annualized terms.

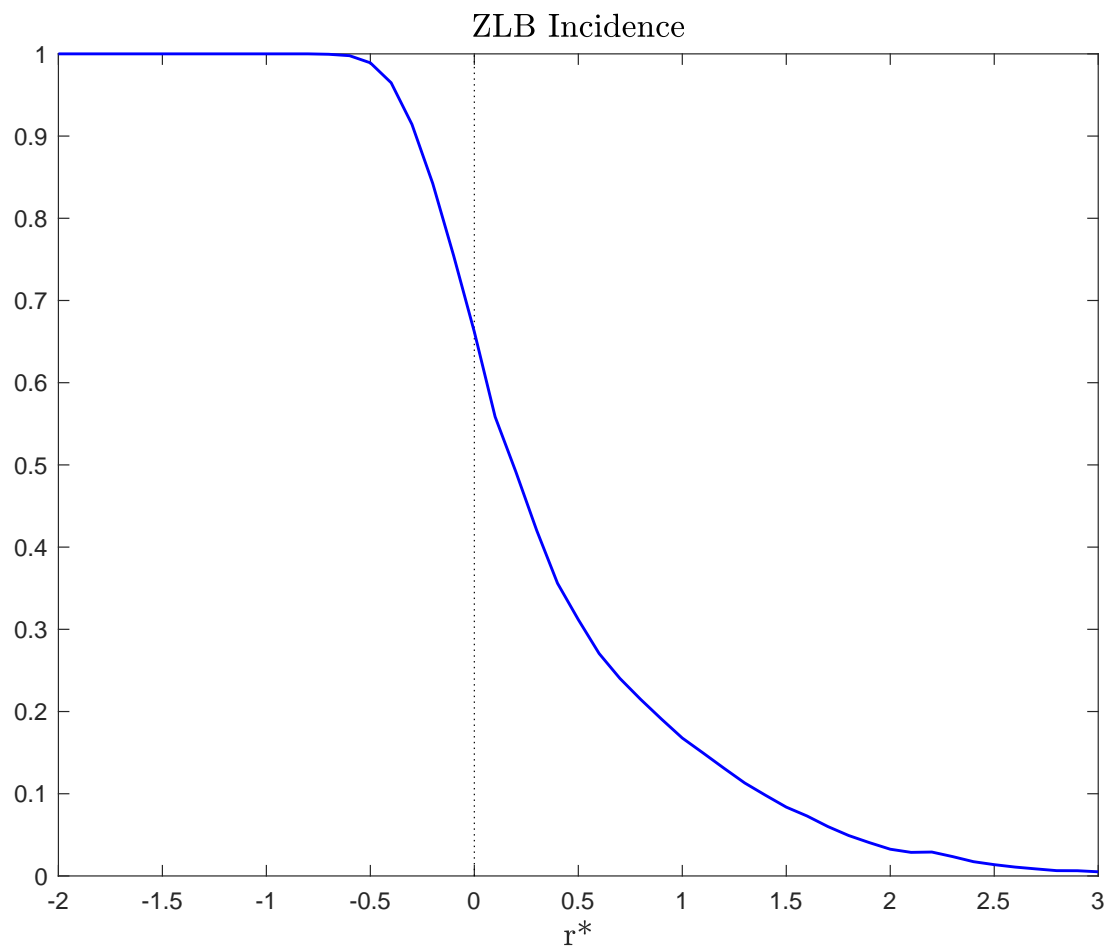


Figure 4: ZLB incidence under the optimal monetary policy.

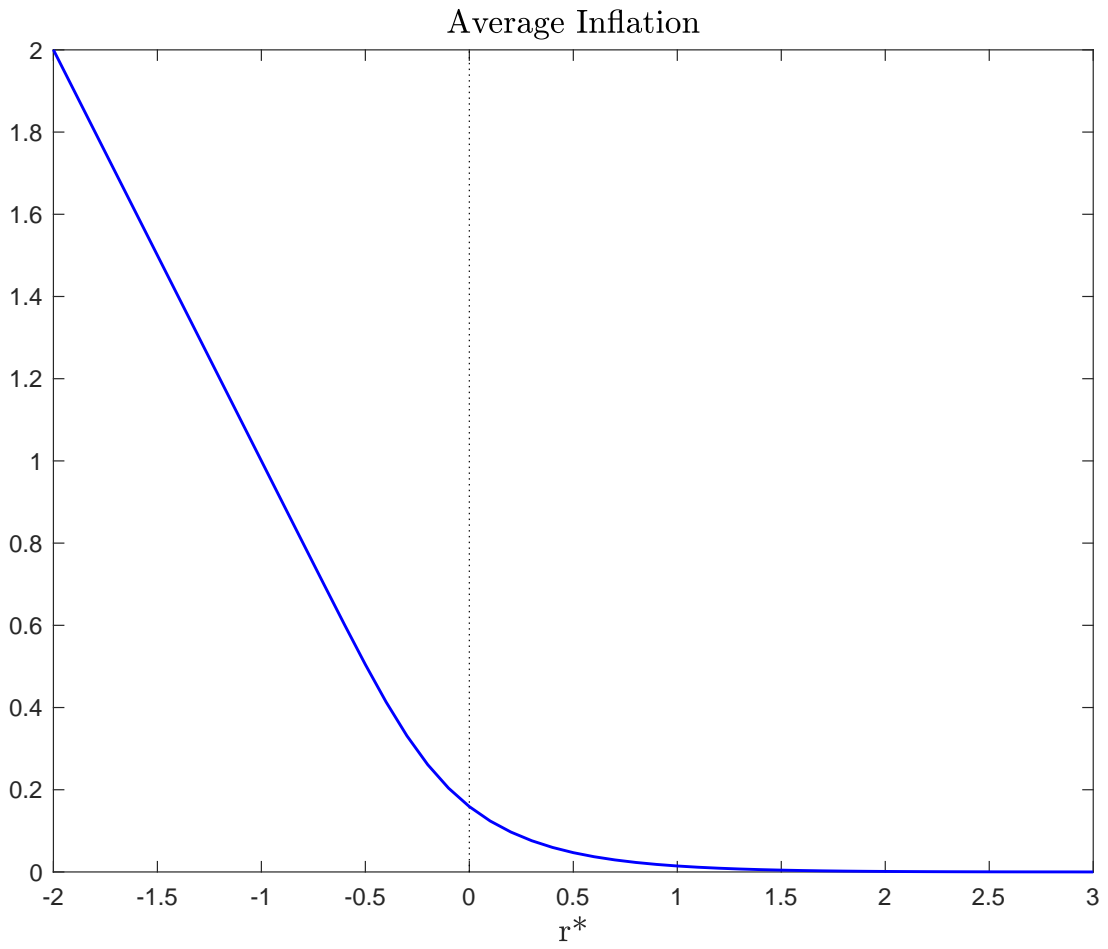


Figure 5a: Average inflation under the optimal monetary policy in annualized terms.

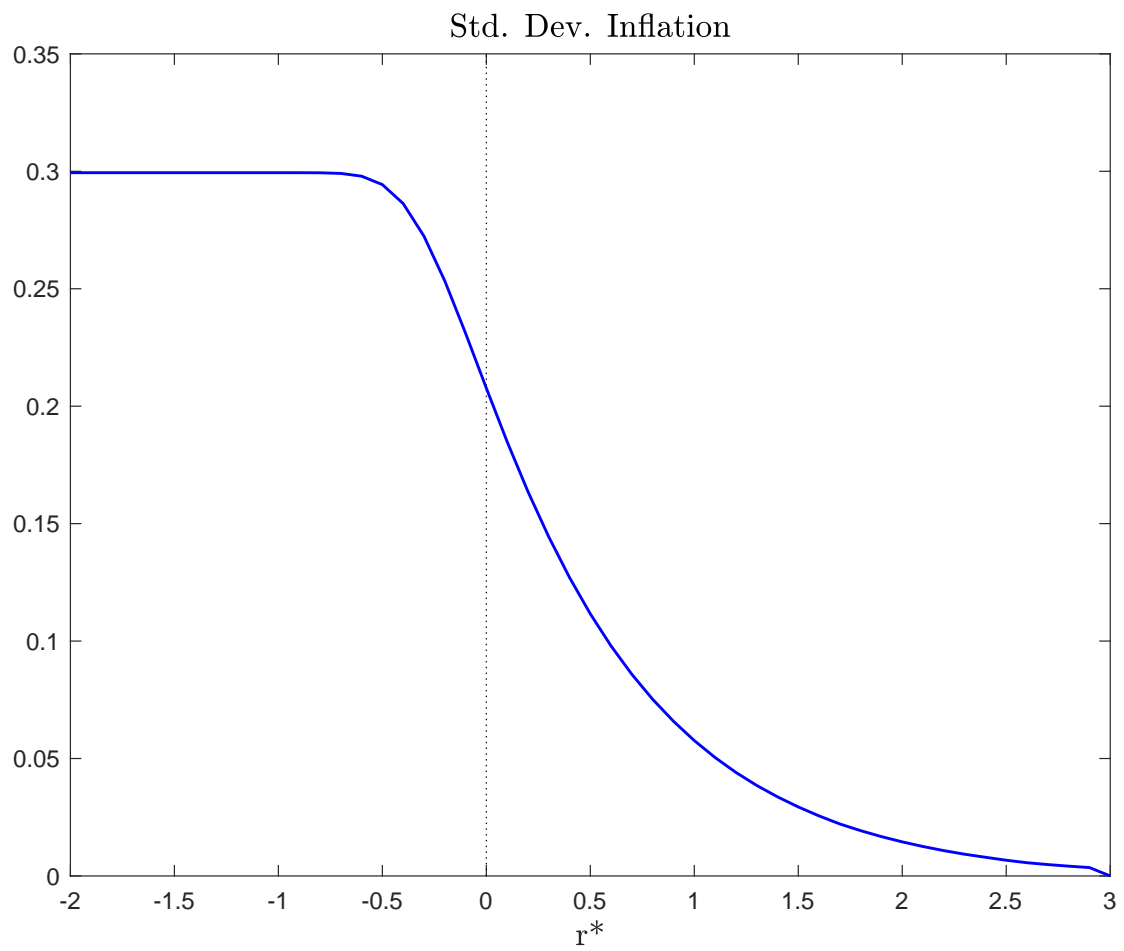


Figure 5b: Volatility of inflation under the optimal monetary policy in annualized terms.

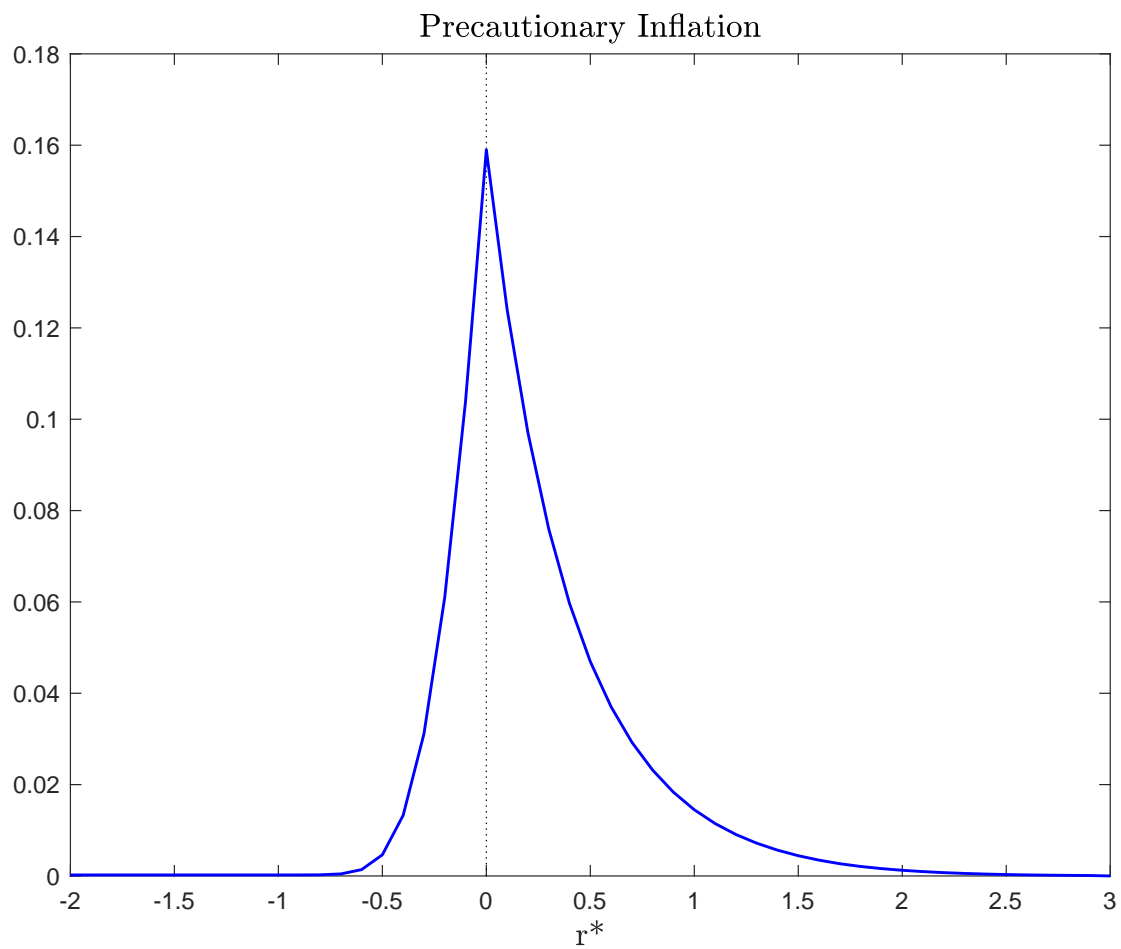


Figure 6: Precautionary inflation under the optimal monetary policy in annualized terms.

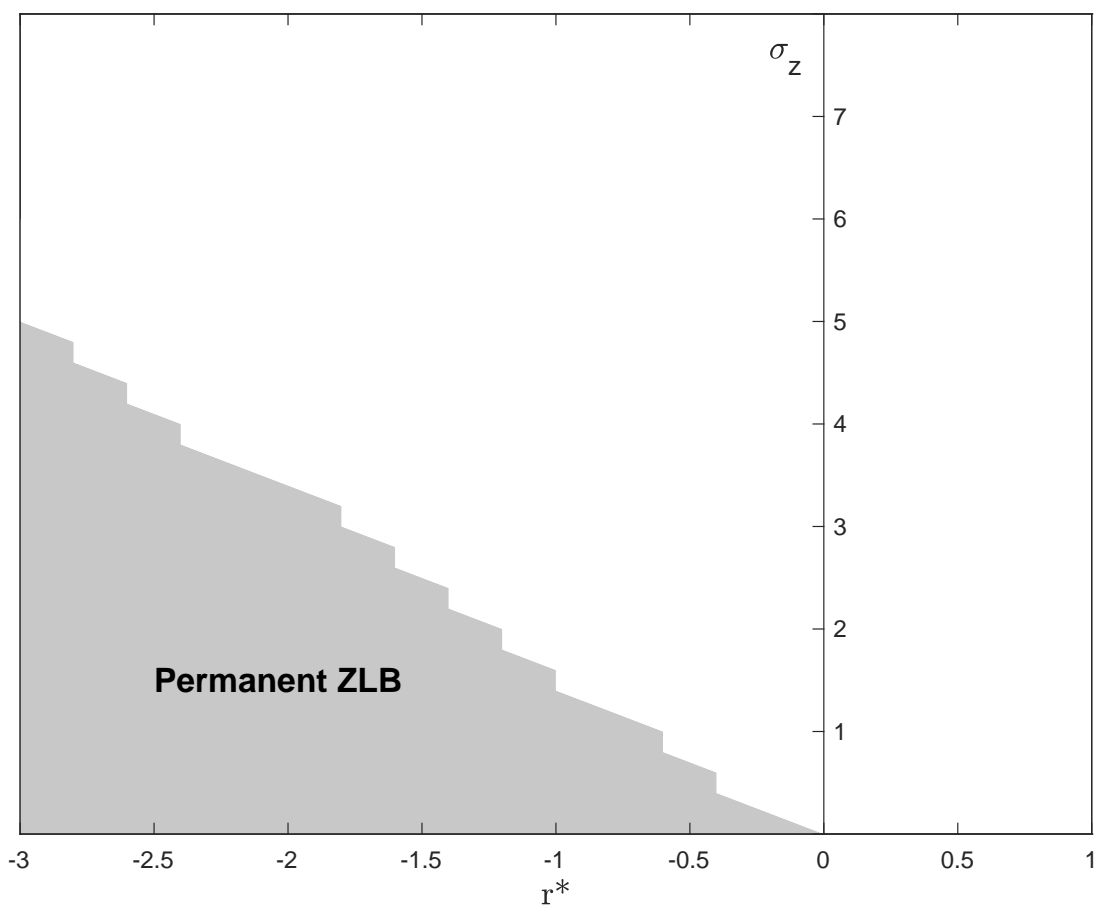


Figure 7: ZLB near-permanently binding under the optimal monetary policy.

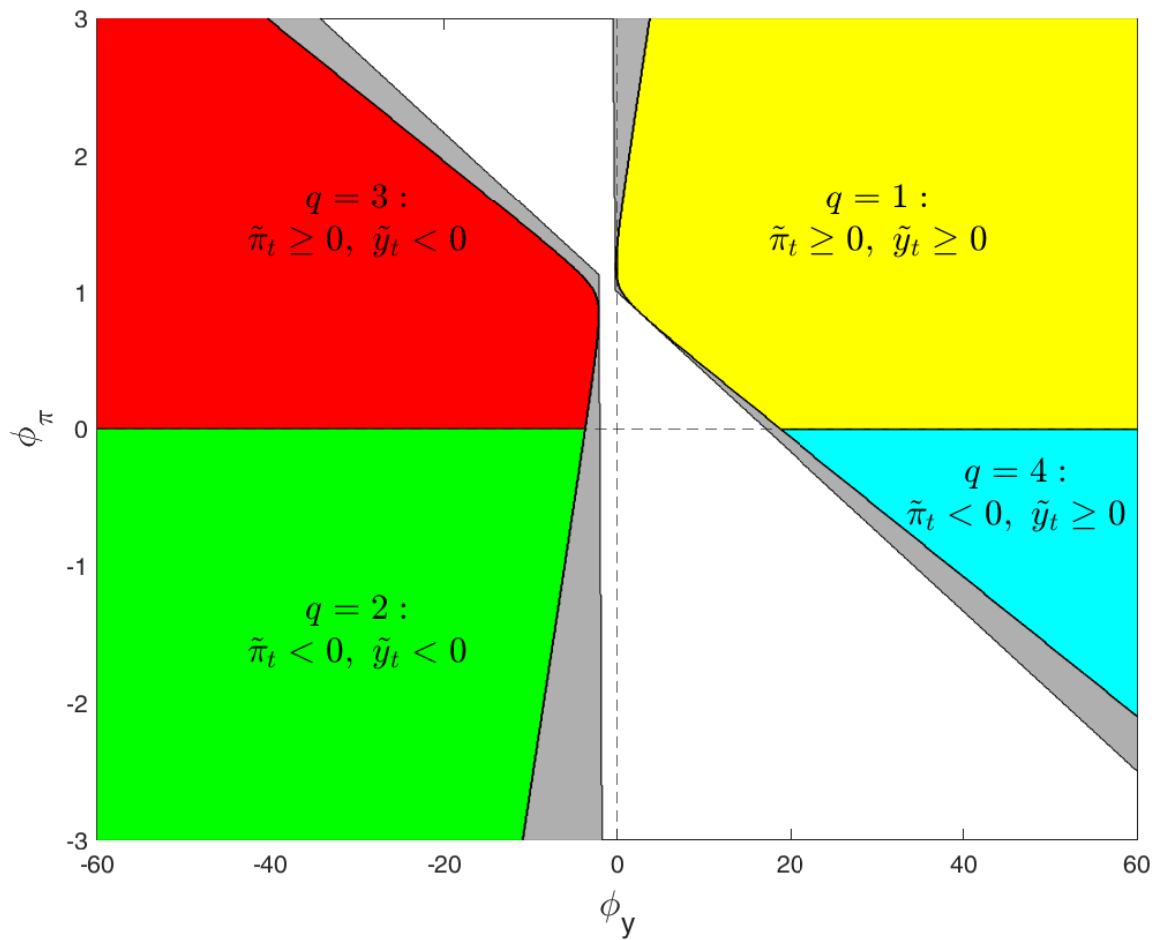


Figure 8: Implementation of the optimal monetary policy with state-contingent interest rate rule. Colored areas show values of the rule coefficients consistent with the sufficient condition for determinacy.

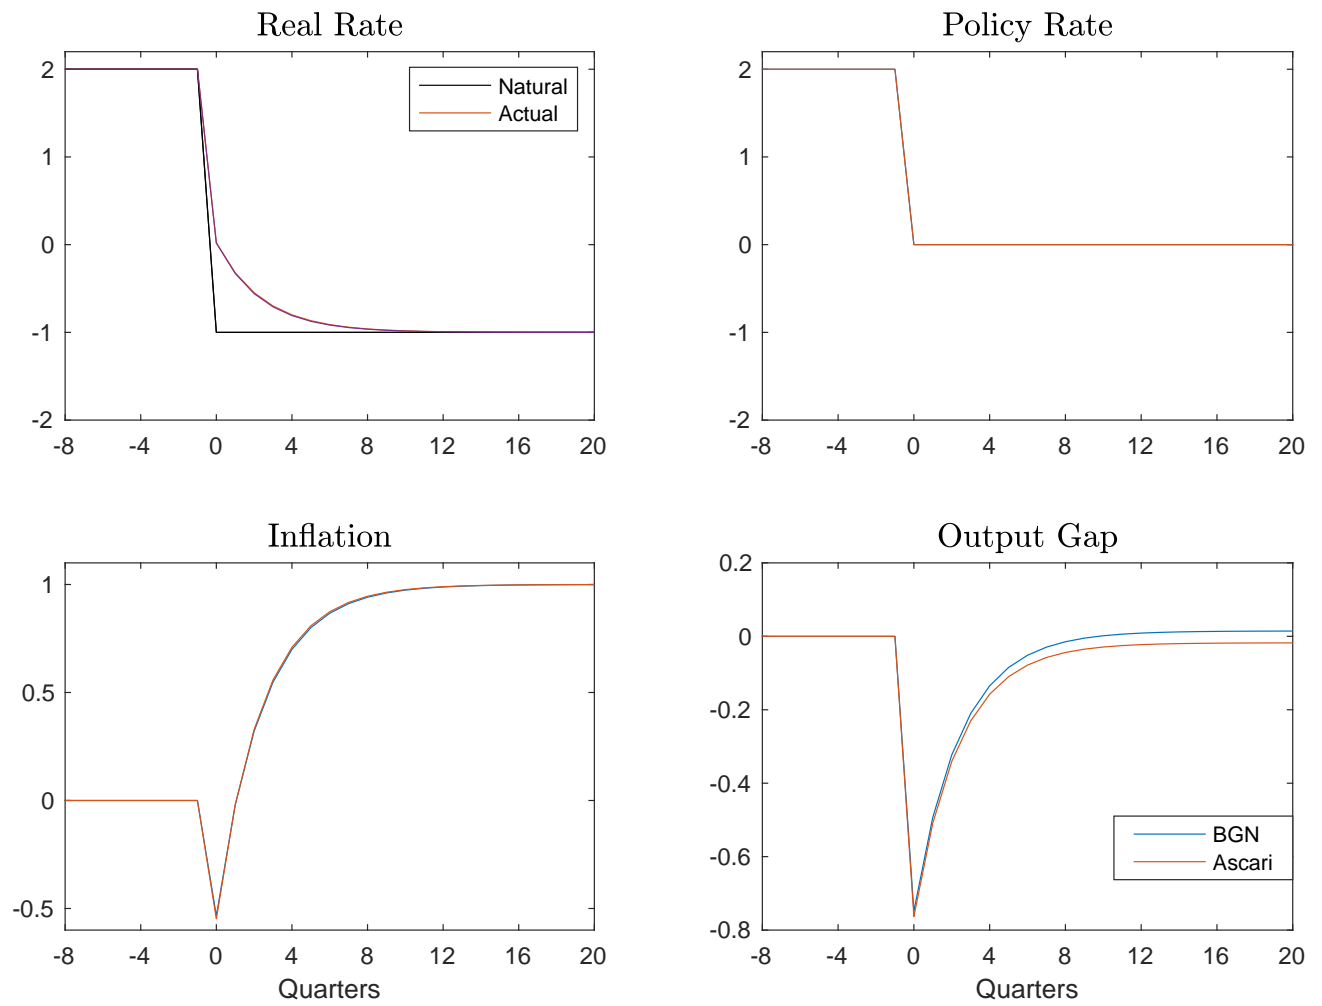


Figure A1: Optimal transition paths with and without Ascari (2014) and Alves (2014) corrections for positive trend inflation.

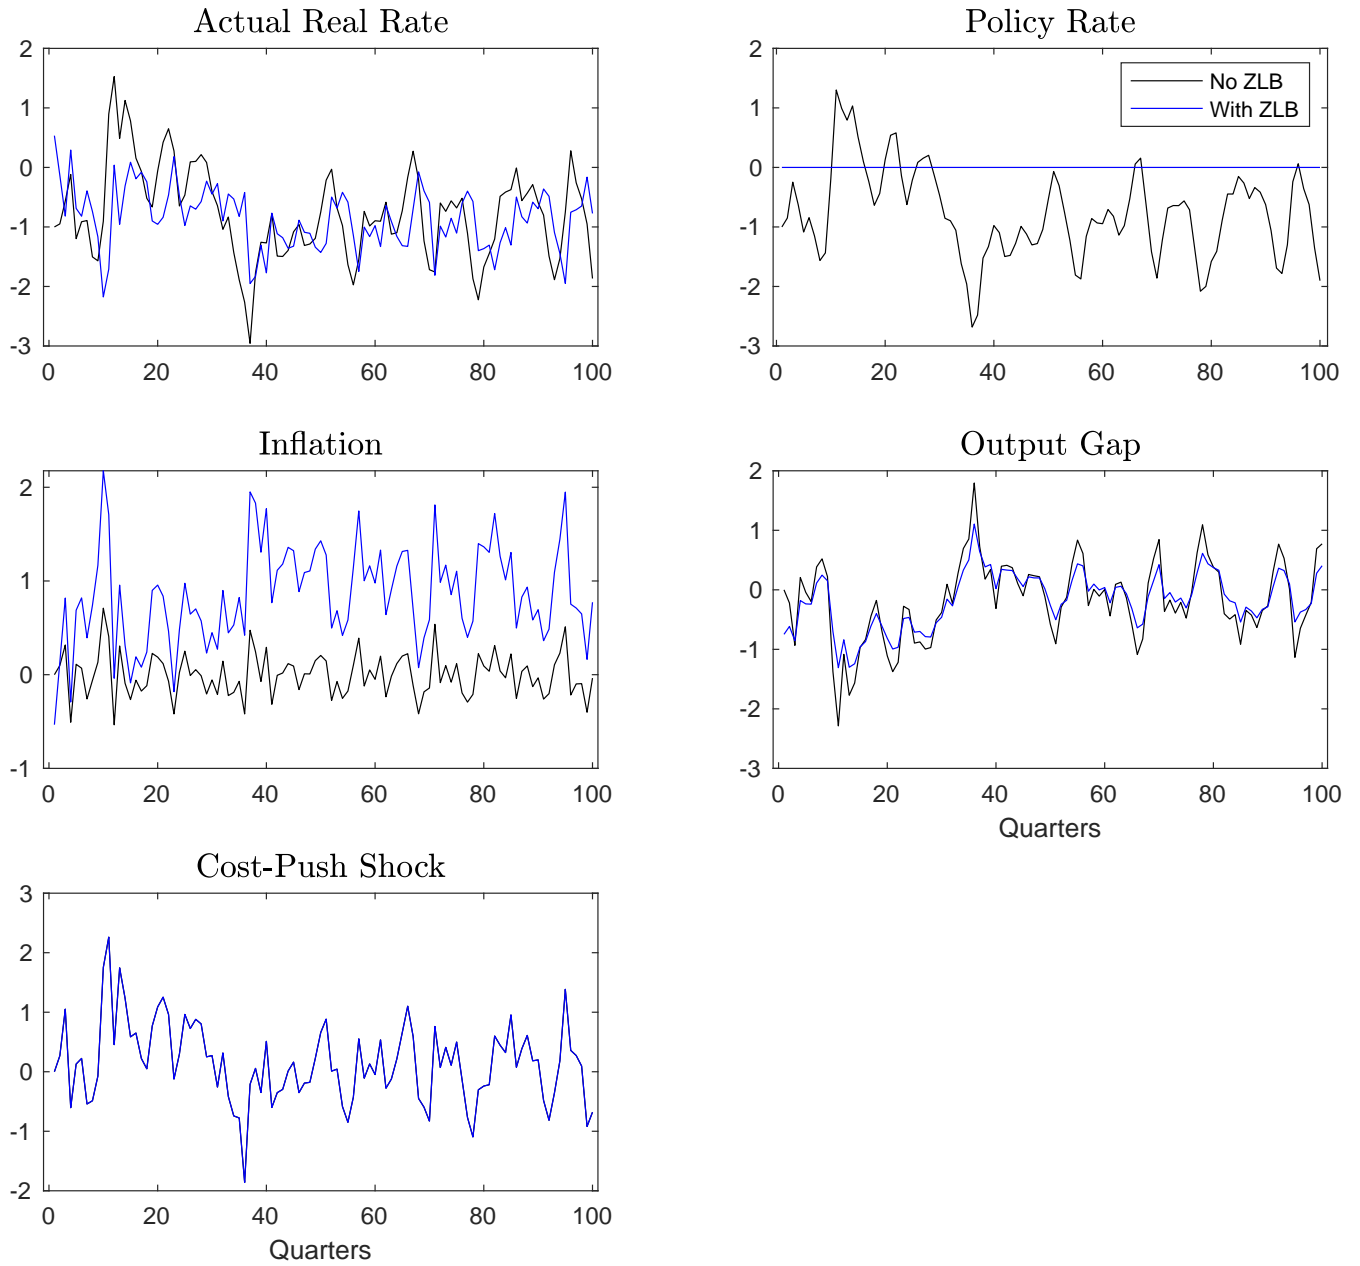


Figure A2a: Aggregate fluctuations under the optimal monetary policy with cost-push shocks. Inflation and interest rates in annualized terms.

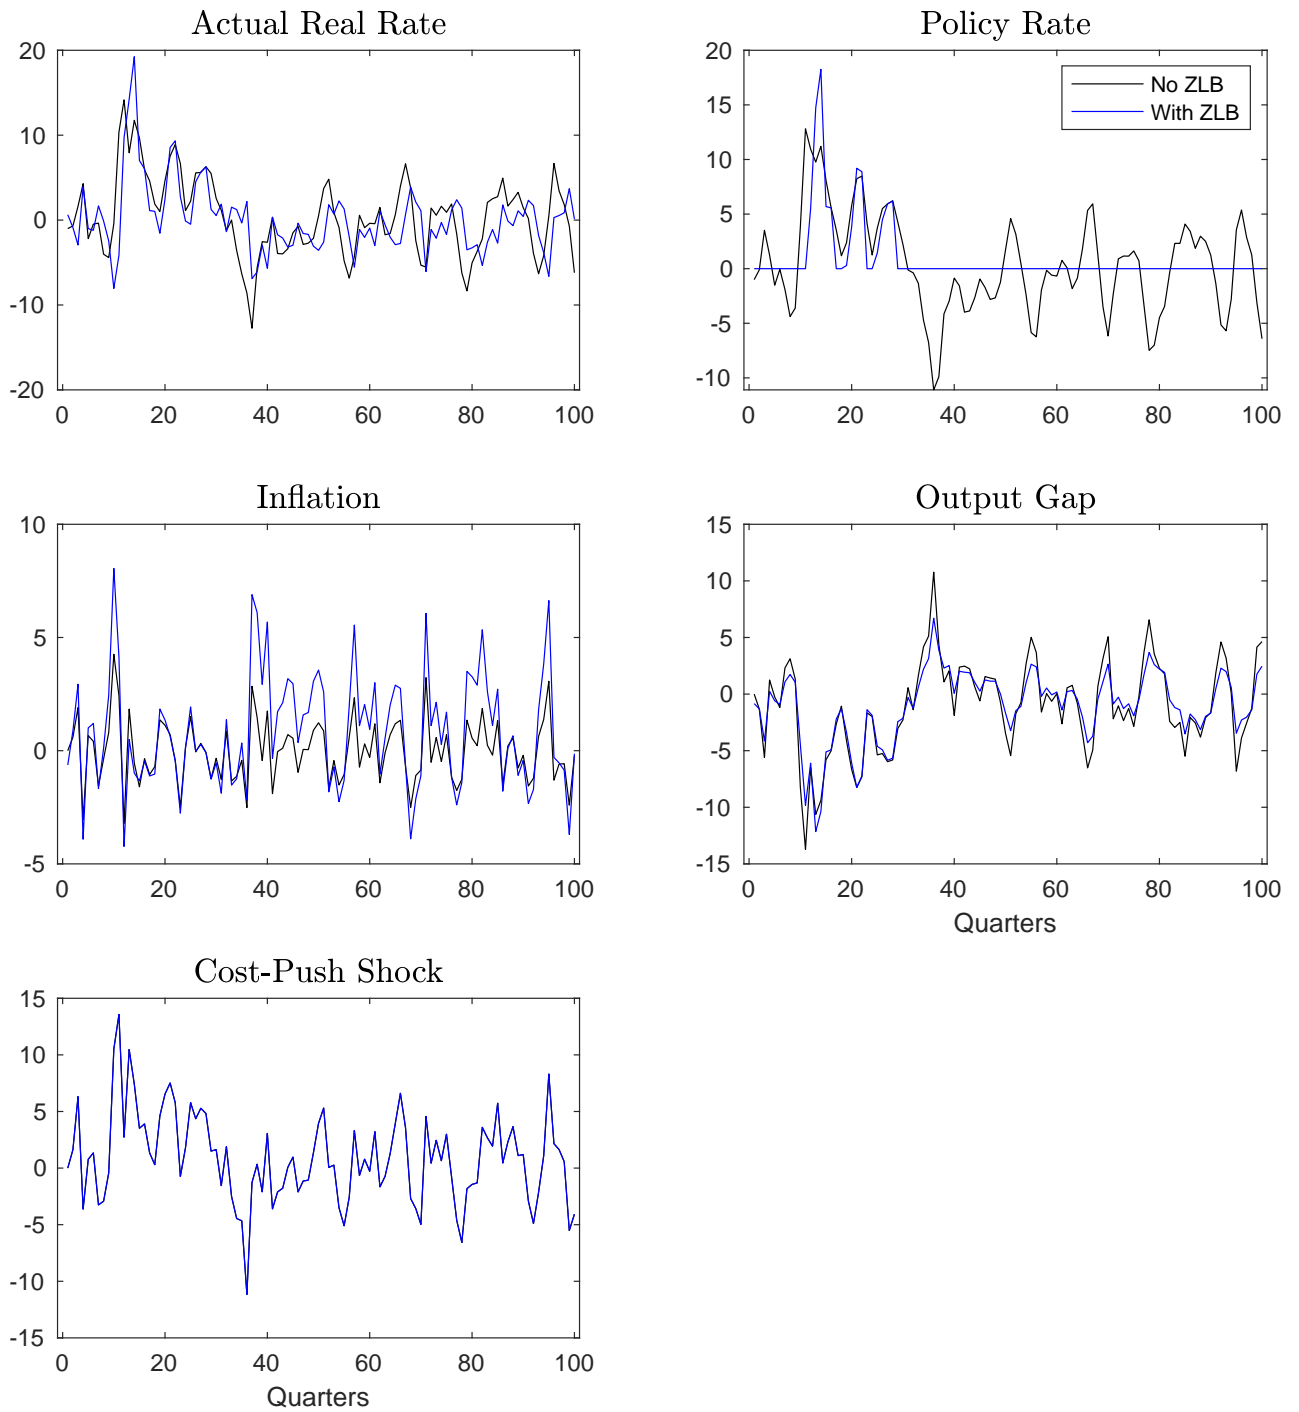


Figure A2b: Aggregate fluctuations under the optimal monetary policy with cost-push shocks with higher volatility. Inflation and interest rates in annualized terms.