



Should the ECB Adjust its Strategy in the Face of a Lower r^* ?

**Philippe Andrade
Jordi Galí
Hervé Le Bihan**

**This version: April 2021
(February 2021)**

Barcelona GSE Working Paper Series

Working Paper n° 1236

Should the ECB Adjust its Strategy in the Face of a Lower r^* ? ¹

Philippe Andrade

Federal Reserve Bank of Boston

Jordi Galí

CREI, UPF and Barcelona GSE

Hervé Le Bihan

Banque de France

Julien Matheron

Banque de France

April 13, 2021

Abstract

We address this question using an estimated New Keynesian DSGE model of the Euro Area with trend inflation, imperfect indexation, and a lower bound on the nominal interest rate. In this setup, a decrease in the steady-state real interest rate, r^* , increases the probability of hitting the lower bound constraint, which entails significant welfare costs and warrants an adjustment of the monetary policy strategy. Under an unchanged monetary policy rule, an increase in the inflation target of eight tenth the size of the drop in the real natural rate of interest is warranted. Absent an increase in the inflation target, and assuming the effective lower bound prevents the ECB from implementing more aggressive negative interest rate policies, adjusting the monetary strategy requires considering alternative instruments or policy rules, such as committing to make-up for recent, below-target inflation realizations.

Keywords: inflation target, effective lower bound, monetary policy strategy, euro-area.

JEL Codes: E31, E52, E58

¹Andrade: philippe.andrade@bos.frb.org. Galí: jgali@crei.cat. Le Bihan: herve.lebihan@banque-france.fr. Matheron: julien.matheron@banque-france.fr. The views expressed herein are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Boston, the Federal Reserve System, the Banque de France or the Eurosystem. We thank Thomas Lubik, John Roberts, and participants at the 2020 Sintra conference, the European Commission-JEDC conference on “Secular Stagnation, Low Interest Rates and Low Inflation: Causes and Implications for Policy” (4-6 November 2020) and at seminars at Banque de France and Toulouse School of Economics, for comments and suggestions. Galí acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S) and from the Generalitat de Catalunya, through CERCA and SGR Programme (2017-SGR-1393).

1 Introduction

The ECB launched a review of its monetary policy strategy in January 2020. As a motivation for the strategy review, the ECB stated that “declining trend growth, on the back of slowing productivity and an aging population, as well as the legacy of the financial crisis, have driven interest rates down, reducing the scope for the ECB and other central banks to ease monetary policy by conventional instruments in the face of adverse cyclical developments” (ECB, 2020).

The view that interest rates will remain structurally lower than what they used to be is consistent with a recent but sizable literature that has documented a permanent—or, at least very persistent—decline in the “natural” rate of interest in advanced economies, including the euro area (Brand and Mazelis, 2019, Del Negro et al., 2018, Holston et al., 2017). The on-going COVID-19 crisis could reinforce these downward pressures on the natural interest rate as agents revise upward their views on the fundamental economic risks they face, inducing larger precautionary savings (Kozłowski et al., 2020). A structurally lower real interest rate matters for monetary policy as, everything else being constant, it will cause the nominal interest rate to hit its effective lower bound (ELB) more frequently, hampering the ability of monetary policy to stabilize the economy, and bringing about more frequent (and potentially protracted) episodes of recessions and below-target inflation.

Central banks can contemplate three main reactions to this new environment. The first option is to keep the same strategy, while at the same time including permanently in their toolbox the unconventional policies adopted since the Great Recession, on the grounds that they can alleviate the ELB constraint and succeed in stabilizing the economy (Bernanke, 2020, Coenen et al., 2020). The second option is to raise the inflation rate that the central bank deems to be consistent with the objective of price stability. A higher inflation target would compensate for the decline in the steady-state nominal interest rate induced by the for the lower real natural rate of interest, and hence contain the increase in the probability of hitting the ELB and maintain the “policy space” (Ball, 2014, Blanchard et al., 2010). The third option is to change their reaction function. In particular the monetary authority can attempt to go (further) into negative interest rate policies and to lower its effective lower bound (Rogoff, 2017). Another modification of the reaction function is to adopt policies that promise to make up for the inflation lost during ELB episodes so that, thanks to the induced higher inflation expectations, the real interest rate declines despite the constraint on the nominal rate (Bernanke et al., 2019). A particular form of such make-up strategy is the Average Inflation Targeting (AIT) strategy that the Federal Reserve announced in late August 2020 (Mertens and Williams, 2019a).

The present paper contributes to this debate by asking two questions. First, to what extent does a lower steady-state real interest rate (r^*) call for a change in the optimal inflation target (π^*) if the central bank keeps its policy rule unchanged? Second, to what extent can a change in the policy rule be an alternative to increasing the inflation target? To address these questions, we conduct a quantitative welfare-based analysis which relies on an estimated structural macroeconomic model of the euro area. Our main findings can be

summarized as follows: (i) Not changing the monetary policy strategy is suboptimal; (ii) a 1 percentage point decrease in r^* from its pre-2008 level calls for an increase in the the inflation target of roughly a 0.8 percentage point when the policy rule is unchanged; and (iii) a change in the policy rule can be an alternative to increasing the target if the commitment to making up for inflation lost during ELB episodes is strong and credible enough.

Our results are obtained from extensive simulations of a New Keynesian DSGE model. The model is estimated for the euro area over the 1985Q1-2008Q3 sample, a period preceding the Great Recession, the euro area sovereign crisis, and the Covid-19 crisis that triggered a protracted period of zero and negative interest rates in the euro area. This is intended to capture the “Great Moderation” period which we use as a benchmark for comparison with the “new normal” characterized by, inter alia, a lower natural rate of interest. The framework is similar to that in [Andrade et al. \(2019b\)](#), which contains a similar analysis based on U.S. data. Among other features, it assumes: (i) price stickiness and imperfect indexation of prices to non-zero trend inflation, (ii) wage stickiness and imperfect indexation of wages to both inflation and technical progress, and (iii) a lower bound constraint on the nominal interest rate. The first two features imply the presence of potentially substantial costs associated with non-zero wage or price inflation. The third feature warrants a strictly positive average inflation rate, in order to mitigate the incidence and adverse effects of the ELB.

According to our baseline simulations, the optimal inflation target, conditional on the “pre-crisis” estimate of $r^* = 2.8$ percent, was around 1.4 percent for the euro area (in annual terms). An important feature of this analysis is to allow for an ELB equal to -0.5 percent, which reflects the recent euro-area experience. Alternatively, if one postulates a strict zero lower bound on the nominal interest rate, as was arguably internalized by the ECB when it conducted its strategy review in 2003, the optimal inflation target is at 1.8%, a value very much consistent with the “below but close to 2 percent” inflation aim adopted by the ECB at that time. Targeting an optimal inflation level helps keeping the frequency of hitting the ELB at levels lower than 5% in the pre-crisis regime. This relatively low frequency of ELB episodes also results from the relatively small volatility of the shocks in the pre-crisis sample, as it does not include the large shocks associated with the Great Recession and the euro-area sovereign crisis episodes.

Our simulations show that, starting from levels around this baseline, a one percentage point drop of r^* dramatically increases the probability of hitting the ELB, if the monetary authority keeps its inflation target unchanged.¹ This makes an unchanged strategy suboptimal, even if our set-up already takes on board some degree of “low for long” monetary policy reaction when the ELB is binding.

When considering the policy option of increasing the inflation objective while keeping the reaction func-

¹Note our exercise is different from assessing what would be the optimal response to a time-varying steady state—a specification consistent with econometric work like that of e.g. [Holston et al. \(2017\)](#). Our exercise is arguably consistent with “secular stagnation” understood as a permanently lower real rate of interest – without having to assume a unit root process in the real rate of interest.

tion unchanged, we obtain that, in response to a 1 percentage point drop in r^* , it is optimal for the central bank to increase the inflation target by about 0.8 percentage point, that is to roughly 2.2 percent. This optimal reaction does not fully offset the increase in the probability to be at the ELB that, ceteris paribus, the drop in r^* implies. Indeed, the benefit of providing a better hedge against hitting the ELB, comes at the cost of a higher steady-state inflation which, because of imperfect indexation, entails permanent price distortions and output loss. The above assessment of the “post-crisis” optimal inflation rate is arguably conservative as (i) we consider a 1 percentage point decrease in r^* for illustrative purposes, which stands on the conservative side in view of the existing empirical estimates of the decline in r^* and (ii) it does not take into account the possibility of a structurally larger volatility of shocks. When considering a 2 percentage point decline in r^* , an amount well within the range of estimates in the literature (see for instance [Constâncio, 2016](#)), the optimal inflation rate then becomes close to 3.2 percent.

We also compute changes in the optimal inflation target under alternative scenarios. First, we assume that fiscal authorities in the euro-area can automatically launch a coordinated and aggressive fiscal emergency package when the economy is hit by a sequence of shocks leading to substantially negative cumulated output gaps. Everything else being constant, under that scenario, the baseline optimal inflation target drops to as low as .4 percent. As emphasized in, among others, [Christiano et al. \(2011\)](#), fiscal policy is a very effective way to escape a trap. As a result, the possibility of launching such an emergency fiscal package mitigates the frequency with which the ELB binds. This reduces the benefits of having a positive inflation target for mildly positive values of r^* . However, under that set-up, it is still the case that a 1 percentage point drop in r^* from its pre-crisis value calls for a 0.8 percentage point increase in the optimal inflation target.

Second, we consider allowing for the possibility for policy rates to go even further in negative territory and set the effective lower bound to -0.8 percent instead of -0.5 percent. Under this scenario, the optimal inflation target stands at 1.1 percent for the baseline r^* , and should be increased to about 2 percent in reaction to a 1 percentage point drop in r^* .

Third, we consider changes in the monetary policy rule. We illustrate that an alternative to increasing the inflation target is to adopt a commitment to keep interest rate “lower for longer” at the end of the liquidity trap so as to make-up for past inflation lost at the ELB. This can be sufficient to maintain the optimal inflation target unchanged in reaction to a 1 percentage point decrease in r^* . Switching from an inflation targeting strategy to an average inflation targeting (AIT) strategy would also allow to maintain the optimal inflation target unchanged in reaction to a 1 percentage point drop in r^* if the window considered is as long as 8 years. Finally, a limit case is that of Price Level Targeting (PLT), for which the optimal inflation target turns out to be lower than in the baseline, even after allowing for a lower r^* .

Related Literature Our paper is related to the contributions quantifying the optimal inflation rate in NK models with an ELB constraint (Billi, 2011, Coibion et al., 2012, Schmitt-Grohé and Uribe, 2010).² In contrast to these works, and as in Andrade et al. (2019b), our focus is on the impact of changes in r^* on the optimal inflation target. In addition, while these papers study the US, we focus on the euro area, and in particular, we allow for specific features such as the availability of negative interest rate policies in the ECB monetary policy toolbox.

Our work also contributes to the literature assessing how monetary policy should adjust to a lower r^* environment. Eggertsson and Woodford (2003) and Adam and Billi (2006) showed that monetary authorities should adopt rules by which they commit to temporarily exceed their inflation target when the economy reaches the ELB. Reifschneider (2016), Kiley (2019), and Chung et al. (2019) for the US economy, and Coenen et al. (2020) for the euro area, advocate more frequent recourse to the non-conventional monetary policies that were implemented by the Fed and the ECB. Reifschneider and Williams (2000), Bernanke (2017), Kiley and Roberts (2017) and Bernanke et al. (2019) examine the benefits of temporary, or permanent, switches to a price-level targeting framework. Mertens and Williams (2019a,b), Arias et al. (2020), Budianto et al. (2020), and Hebden et al. (2020) study average inflation targeting.

Our baseline set-up does not include explicitly forward guidance or quantitative easing policies that can be used at the ELB. However, we consider that the central bank is committed to a simple interest rate rule which features some history dependence implying that interest rates are being kept lower for longer after ELB episodes. As we document, this rule delivers outcomes that are comparable to what Coenen et al. (2020) obtain for the euro area under the assumption that the ECB can conduct various form of forward guidance possibly combined with asset purchase programmes. We also study how the (r^*, π^*) relation is modified when the central bank commits more strongly to a strategy that will “make-up” for the inflation shortfall at the ELB, by either increasing the inertia of the baseline interest rule or by opting for either an AIT or a PLT strategy. In addition, we consider the possibilities of either lowering the policy rate even further below zero or launching a coordinated emergency fiscal package in case of large recessions, while maintaining the policy rule unchanged.

Overall, with respect to our earlier contribution Andrade et al. (2019b) this work differs not only by considering the euro area and some of its specificities, but also by exploring some additional policy options (AIT, negative interest ELB) that the central bank can contemplate to address the low r^* environment, as well as the possibility of an aggressive countercyclical fiscal policy in response to persistent recessions.

The remaining of the paper is organized as follows. Section 2 presents the New Keynesian model we use. Section 3 describes how the model is estimated and simulated, as well as how welfare is approximated, and presents estimation results. Section 4 presents some properties of the model, based on the baseline

²More recent references underly that the tolerance for inflation increases when the frequency of prices changes adjust to the level of inflation (Blanco, 2016, Budianto, 2021, L’Huillier and Schoenle, 2020). A higher inflation target might also be warranted when considering heterogeneous firms (Adam and Weber, 2019) or households (Cotton, 2020).

estimates and some counterfactuals. Section 5 is devoted to analyzing by how much π^* should change in reaction to a decline in r^* . Section 6 analyses changes in strategy that could be substitutes to increasing the inflation target. Section 7 provides some concluding remarks.

2 The Model

We use a relatively standard medium-scale New Keynesian model as a framework of reference, closely following [Andrade et al. \(2019b\)](#). Crucially, the model features elements that generate a cost to non-zero inflation: (1) nominal rigidities, in the form of staggered price and wage setting; (2) less than full price (and wage) indexation to past or trend inflation; and (3) trend productivity growth along, to which wages are imperfectly indexed.

As is well known, staggered price setting generates a positive relation between deviations from zero inflation and price dispersion (with the resulting inefficient allocation of resources). Moreover, the absence of systematic indexation to trend inflation magnifies these costs, as emphasized by [Ascari and Sbordone \(2014\)](#). Also, and ceteris paribus, price inflation induces (nominal) wage inflation, which in turn triggers inefficient wage dispersion in the presence of staggered wage setting. Imperfect indexation also magnifies the costs of non-zero price (or wage) inflation as compared to a set-up where price and wages mechanically catch up with trend inflation. Finally the lack of a systematic indexation of wages to productivity also induces an inefficient wage dispersion.

At the same time, there are benefits associated to a positive inflation rate, as interest rates are subject to a ELB constraint. This constraint hampers the stabilization potential of monetary policy as it limits the available “policy space”. By contrast, given the steady-state real interest rate, a larger average rate of inflation will decrease the incidence of binding ELB episodes and restore some policy space.

Overall, the model we use, and the trade-off between costs and benefits of steady-state inflation, are close to those considered by [Coibion et al. \(2012\)](#). However, as in [Andrade et al. \(2019b\)](#) we assume Calvo-style sticky wages, in addition to sticky prices.³

2.1 Households

The economy is inhabited by a (measure one) continuum of infinitely-lived, identical households. The representative household is composed of a continuum of workers, each specialized in a particular labor type indexed by $h \in [0, 1]$. The representative household’s objective is to maximize an intertemporal welfare function

$$\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \left\{ e^{\zeta_{c,t+s}} \log(C_{t+s} - \eta C_{t+s-1}) - \frac{\chi}{1+\nu} \int_0^1 N_{t+s}(h)^{1+\nu} dh \right\}, \quad (1)$$

³In their robustness analysis, [Coibion et al. \(2012\)](#) consider downward nominal wage rigidity, which entails different mechanisms than with Calvo-style rigidities.

where $\beta \equiv e^{-\rho}$ is the discount factor (ρ being the discount rate), $\mathbb{E}_t\{\cdot\}$ is the expectation operator conditional on information available at time t , C_t is consumption and $N_t(h)$ is the supply of labor of type h . The utility function features habit formation, with degree of habits η . The inverse Frisch elasticity of labor supply is ν and χ is a scale parameter in the labor disutility. The utility derived from consumption is subject to a preference shock $\zeta_{c,t}$.

The representative household maximizes (1) subject to the sequence of constraints

$$P_t C_t + e^{\zeta_{q,t}} Q_t B_t \leq \int_0^1 W_t(h) N_t(h) dh + B_{t-1} - T_t + D_t \quad (2)$$

where P_t is the aggregate price level, $W_t(h)$ is the nominal wage rate associated with labor of type h , $e^{\zeta_{q,t}} Q_t$ is the price at t of a one-period nominal bond paying one unit of currency in the next period, where $\zeta_{q,t}$ is a ‘‘risk-premium’’ shock, B_t is the quantity of such bonds acquired at t , T_t denotes lump-sum taxes, and D_t stands for the dividends rebated to the households by monopolistic firms.

2.2 Firms and Price Setting

The final good is produced by perfectly competitive firms according to the Dixit-Stiglitz production function

$$Y_t = \left(\int_0^1 Y_t(f)^{(\theta_p-1)/\theta_p} df \right)^{\theta_p/(\theta_p-1)},$$

where Y_t is the quantity of final good produced at t , $Y_t(f)$ is the input of intermediate good f , and θ_p the elasticity of substitution between any two intermediate goods. The zero-profit condition for labor-aggregating firms yields the relation

$$P_t = \left(\int_0^1 P_t(f)^{1-\theta_p} df \right)^{1/(1-\theta_p)}.$$

Intermediate goods are produced by monopolistic firms, each specialized in a particular good $f \in [0, 1]$. Firm f has technology

$$Y_t(f) = Z_t L_t(f)^{1/\phi}$$

where $L_t(f)$ is the input of aggregate labor, $1/\phi$ is the elasticity of production with respect to aggregate labor, and Z_t is an index of aggregate productivity. The latter evolves according to

$$Z_t = Z_{t-1} e^{\mu_z + \zeta_{z,t}}$$

where μ_z is the average growth rate of productivity. Thus, technology is characterized by a unit root in the model.

Intermediate goods producers are subject to nominal rigidities à la Calvo. Formally, firms face a constant probability α_p of not being able to re-optimize prices. In the event that firm f is not drawn to re-optimize at

t , it re-scales its price according to the indexation rule

$$P_t(f) = (\Pi_{t-1})^{\gamma_p} P_{t-1}(f)$$

where $\Pi_t \equiv P_t/P_{t-1}$, Π is the associated steady-state value and $0 \leq \gamma_p < 1$. Thus, in case firm f is not drawn to re-optimize, it mechanically re-scales its price by past inflation. Importantly, however, we assume that the degree of indexation is less than full since $\gamma_p < 1$. One obvious drawback of the Calvo set-up is that the probability of price change is assumed to be invariant, inter alia to the long run inflation rate. Drawing from the logic of menu cost models, the Calvo parameter of price fixity could be expected to decrease when trend inflation rises. However, in the range of values for trend inflation that we will consider, available micro economic evidence, such as that summarized in [Alvarez et al. \(2019\)](#), suggests there is no significant correlation between the frequency of price change and trend inflation.

If drawn to re-optimize in period t , a firm chooses P_t^* in order to maximize

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta \alpha_p)^s \Lambda_{t+s} \left\{ (1 + \tau_{p,t+s}) \frac{V_{t,t+s}^p P_t^*}{P_{t+s}} Y_{t,t+s} - \frac{W_{t+s}}{P_{t+s}} \left(\frac{Y_{t,t+s}}{Z_{t+s}} \right)^\phi \right\},$$

where Λ_t denotes the marginal utility of wealth, $\tau_{p,t}$ is a sales subsidy paid to firms and financed via a lump-sum tax on households, and $Y_{t,t+s}$ is the demand function that a monopolist who last revised its price at t faces at $t + s$; it obeys

$$Y_{t,t+s} = \left(\frac{V_{t,t+s}^p P_t^*}{P_{t+s}} \right)^{-\theta_p} Y_{t+s}$$

where $V_{t,t+s}^p$ reflects the compounded effects of price indexation to past inflation

$$V_{t,t+s}^p = \prod_{j=t}^{t+s-1} (\Pi_j)^{\gamma_p}.$$

We further assume that

$$1 + \tau_{p,t} = (1 + \tau_p) e^{-\zeta_{u,t}},$$

with $\zeta_{u,t}$ appearing in the system as a cost-push shock. Furthermore, we set τ_p so as to neutralize the steady-state distortion induced by price markups.

2.3 Aggregate Labor and Wage Setting

There is a continuum of perfectly competitive labor aggregating firms that mix the specialized labor types according to the CES technology

$$N_t = \left(\int_0^1 N_t(h)^{(\theta_w-1)/\theta_w} dh \right)^{\theta_w/(\theta_w-1)},$$

where N_t is the quantity of aggregate labor and $N_t(h)$ is the input of labor of type h , and where θ_w denotes the elasticity of substitution between any two labor types. Aggregate labor N_t is then used as an input in the production of intermediate goods. Equilibrium in the labor market thus requires

$$N_t = \int_0^1 L_t(f) df.$$

Here, it is important to notice the difference between $L_t(f)$, the demand for aggregate labor emanating from firm f , and $N_t(h)$, the supply of labor of type h by the representative household.

The zero-profit condition yields the relation

$$W_t = \left(\int_0^1 W_t(h)^{1-\theta_w} dh \right)^{1/(1-\theta_w)},$$

where W_t is the nominal wage paid to aggregate labor while $W_t(h)$ is the nominal wage paid to labor of type h .

Mirroring prices, we assume that wages are subject to nominal rigidities, à la Calvo, in the manner of [Erceg et al. \(2000\)](#). Formally, unions face a constant probability α_w of not being able to re-optimize wages. In the event that union h is not drawn to re-optimize at t , it re-scales its wage according to the indexation rule

$$W_t(h) = e^{\gamma_z \mu_z} (\Pi_{t-1})^{\gamma_w} W_{t-1}(h)$$

where, as before, wages are indexed to past inflation. However, we assume that the degree of indexation is here too only partial, by imposing $0 \leq \gamma_w < 1$. In addition, nominal wages are also indexed to average productivity growth with indexation degree $0 \leq \gamma_z < 1$.

If drawn to re-optimize in period t , a union chooses W_t^* in order to maximize

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta \alpha_w)^s \left\{ (1 + \tau_w) \Lambda_{t+s} \frac{V_{t,t+s}^w W_t^*}{P_{t+s}} N_{t,t+s} - \frac{\chi}{1 + \nu} N_{t,t+s}^{1+\nu} \right\}$$

where the demand function at $t + s$ facing a union who last revised its wage at t obeys

$$N_{t,t+s} = \left(\frac{V_{t,t+s}^w W_t^*}{W_{t+s}} \right)^{-\theta_w} N_{t+s}$$

and where $V_{t,t+s}^w$ reflects the compounded effects of wage indexation to past inflation and average productivity growth

$$V_{t,t+s}^w = e^{\gamma_z \mu_z (t+s)} \prod_{j=t}^{t+s-1} (\Pi_j)^{\gamma_w}.$$

Furthermore, we set τ_w so as to neutralize the steady-state distortion induced by wage markups.

2.4 Monetary Policy and the ELB

Monetary policy in "normal times" is assumed to be given by an inertial Taylor-like interest rate rule

$$\hat{i}_t = \rho_i \hat{i}_{t-1} + (1 - \rho_i) (a_\pi \hat{\pi}_t + a_x \hat{x}_t) + \zeta_{m,t}, \quad (3)$$

where $i_t \equiv -\log(Q_t)$, with \hat{i}_t denoting the associated deviation from steady state i.e, $\hat{i}_t \equiv i_t - i$. Also, $\pi_t \equiv \log \Pi_t$, $\hat{\pi}_t \equiv \pi_t - \pi$ is the gap between inflation and its target π , and $\hat{x}_t \equiv \log(Y_t/Y_t^n)$ where Y_t^n is the natural level of output, defined as the level of output that would prevail in an economy with flexible prices and wages and no cost-push shocks.⁴ Finally, $\zeta_{m,t}$ is a monetary policy shock.

Here π is the growth rate in the price index between $t - 1$ and t that is consistent with the inflation aim of the central bank. An annual inflation target of say 2 percent would thus imply $\pi = 2/400 = 0.005$ as the model will be parameterized and estimated on quarterly data. In the remaining of the paper, we refer to π when we use the word inflation target without other qualifications.

Crucially for our purpose, the nominal interest rate i_t is subject to an ELB constraint:

$$i_t \geq e.$$

In the baseline specification, we assume that $e = -0.50/400$, consistent with the rate on the Deposit Facility in the euro area at the time of writing this paper (-0.5 percent). The Deposit Facility rate is the relevant policy rate driving short-run interest rates in that environment.

The steady-state level of the real interest rate is defined by $r^* \equiv i - \pi$. Given logarithmic utility, it is related to technology and preference parameters according to $r^* = \rho + \mu_z$. Combining these elements, it is convenient to write the ELB constraint in terms of the deviation of the nominal interest rate

$$\hat{i}_t \geq e - (\mu_z + \rho + \pi) \tag{4}$$

The rule effectively implemented is given by:

$$\hat{i}_t^n = \rho_i \hat{i}_{t-1}^n + (1 - \rho_i)(a_\pi \hat{\pi}_t + a_y \hat{x}_t) + \zeta_{m,t} \tag{5}$$

$$\hat{i}_t = \max\{\hat{i}_t^n, e - (\mu_z + \rho + \pi)\} \tag{6}$$

An important feature is that – as in [Coibion et al. \(2012\)](#) and in a large share of the recent literature – the lagged rate that matters is a lagged “notional” or “shadow” interest rate \hat{i}_{t-1}^n , rather than the lagged actual rate \hat{i}_{t-1} .

Before proceeding, several remarks are in order. First, the inflation target, π , is not assumed to be optimal. Note also that the average realized inflation, $\bar{\pi}$, might be below the target as a consequence of ELB episodes, i.e. $\bar{\pi} < \pi$. In such instances of ELB, monetary policy fails to deliver the appropriate degree of accommodation, resulting in lower output gap and inflation than in an economy in which there would not be a ELB constraint.

⁴Note that even if there are optimal subsidies that undo price and wage markups, y_n should not be interpreted as the efficient level of output, because there are still the steady-state distortions induced by trend inflation combined with imperfect indexation.

Second, we assume the central bank policy is characterized by a simple interest rate rule rather than a Ramsey-type fully optimal policy of the type studied by e.g. [Khan et al. \(2003\)](#), [Adam and Billi \(2006\)](#), or [Schmitt-Grohé and Uribe \(2010\)](#). Such simple rules have been shown to be a good empirical characterization of the behavior of central banks in the last decades. From an institutional point of view, the setting the inflation objective (or a definition of price stability) appears to be done at a lower frequency than the definition of the rule operating the interest rate on a day-to-day basis. Note also that two features in our set-up, the inertia in the monetary policy rule, as well as the use of a lagged notional rate rather than the lagged actual rate, render the policy more persistent and thus closer to a Ramsey-like fully optimal interest rate rule. In particular the dependence of the lagged notional rate \hat{i}_t^n results in the nominal interest rate \hat{i}_t remaining “lower for longer” in the aftermath of ELB episodes (as \hat{i}_t^n will stay negative for a protracted period).

As equation (4) makes clear, μ_z , ρ , π have symmetric roles in the ELB constraint. Put another way, for given structural parameters and a given process for \hat{i}_t , the probability of hitting the ELB would remain unchanged if productivity growth or the discount rate decline by one percent and the inflation target is increased by a commensurate amount at the same time. Based on these observations, one may be tempted to argue that in response to a permanent decline in μ_z or ρ , the optimal inflation target π^* will necessarily change by the same amount (with a negative sign).

The previous conjecture is, however, incorrect. The reasons for this are twofold. First, any change in μ_z (or ρ) also translates into a change in the coefficients of the equilibrium dynamic system. Second, because there are welfare costs associated with increasing the inflation target, the policy maker would also have to balance the benefits of keeping the incidence of ELB episodes constant with the additional costs in terms of extra price dispersion and inefficient resource allocation. These costs can be substantial and more than compensate for the benefits of holding the probability of ELB constant. Assessing these forces in the context of the euro area economy is precisely this paper’s endeavor.

3 Estimation and Simulations

3.1 Estimation Without a Lower Bound on Nominal Interest Rates

Because the model has a stochastic trend, we first induce stationarity by dividing trending variables by Z_t . The resulting system is then log-linearized in the neighborhood of its deterministic steady state.⁵ We append to the system a set of equations describing the dynamics of the structural shocks, modeled as AR(1) processes, namely

$$\zeta_{k,t} = \rho_k \zeta_{k,t-1} + \sigma_k \epsilon_{k,t}, \quad \epsilon_{k,t} \sim N(0, 1)$$

for $k \in \{c, q, z, u, m\}$.

⁵See the Technical Appendix of [Andrade et al. \(2019b\)](#) for further details.

Absent the ELB constraint, the model can be solved and cast into the usual linear transition and observation equations:

$$s_t = \mathcal{T}(\theta)s_{t-1} + \mathcal{R}(\theta)\epsilon_t, \quad x_t = \mathcal{M}(\theta) + \mathcal{H}(\theta)s_t,$$

with s_t a vector collecting the model's state variables, x_t a vector of observable variables and ϵ_t a vector of innovations to the shock processes $\epsilon_t = (\epsilon_{c,t}, \epsilon_{q,t}, \epsilon_{z,t}, \epsilon_{u,t}, \epsilon_{m,t})'$. The solution coefficients are regrouped in the conformable matrices $\mathcal{T}(\theta)$, $\mathcal{R}(\theta)$, $\mathcal{M}(\theta)$, and $\mathcal{H}(\theta)$ which depend on the vector of structural parameters θ .

We estimate the model using data for a “pre-crisis” period over which the ELB constraint is not binding. This enables us to use the linear version of the model. The sample of observable variables is $X_T \equiv \{x_t\}_{t=1}^T$ with

$$x_t = [\Delta \log(\text{GDP}_t), \Delta \log(\text{GDP Deflator}_t), \Delta \log(\text{Wages}_t), \text{Short Term Interest Rate}_t]'$$

where the short term nominal interest rate is the Euribor 3 months rate.

We use a sample of quarterly data covering the 1985Q1-2008Q3 period.⁶ The sample starts approximately when the disinflation policies were simultaneously conducted in the main euro area countries (see [Fève et al. 2010](#)). It ends on the eve of the Great Recession, and before the implementation of non-conventional policies by the ECB. It can thus be viewed as corresponding to a period of relatively homogeneous monetary policy regime, and a reference “pre-crisis” economy.

The parameters ϕ , θ_p , and θ_w are calibrated prior to estimation. The parameter ϕ is set to 1/.7. Given the assumed subsidy correcting the steady-state price markup distortion, this results in a steady-state labor share of 70 percent. We follow [Christoffel et al. \(2008\)](#) who calibrate the parameters related to price and wage mark-up using euro area micro evidence. The parameter θ_p is set to 4.3, resulting in a steady-state price markup of 30 percent. The parameter θ_w is set to 3.9, resulting in a wage markup of 35 percent. These numbers fall into the arguably large ballpark of available values used in the literature.⁷

We rely on a full-system Bayesian estimation approach to estimate the other model parameters. After having cast the dynamic system in the state-space representation for the set of observable variables, we use the Kalman filter to evaluate the likelihood of the observed variables. We then form the joint posterior distribution of the structural parameters by combining the likelihood function $p(X_T|\theta)$ with a joint density characterizing some prior beliefs $p(\theta)$. The joint posterior distribution thus obeys

$$p(\theta|X_T) \propto p(X_T|\theta)p(\theta),$$

Given the specification of the model, the joint posterior distribution cannot be recovered analytically but may be computed numerically, using a Monte-Carlo Markov Chain (MCMC) sampling approach. More

⁶The data are obtained from the “Area Wide Model” database of [Fagan et al. \(2001\)](#) and Eurostat national accounts for the Euro Area. GDP is expressed in per capita terms.

⁷In [Andrade et al. \(2019b\)](#) we provide additional simulations showing that our baseline results are robust to alternative calibrations of these parameters.

specifically, we rely on the Metropolis-Hastings algorithm to obtain a random draw of size 1,000,000 from the joint posterior distribution of the parameters.

3.2 Computing the Optimal Inflation Target

Simulations with an ELB Constraint The model becomes non-linear when one allows the ELB constraint to bind. The solution method we implement follows the approach developed by [Bodenstein et al. \(2009\)](#) and [Guerrieri and Iacoviello \(2015\)](#). The approach can be described as follows. There are two regimes: the no-ELB regime $k = n$ and the ELB regime $k = e$ and the canonical representation of the system in each regime is

$$\mathbb{E}_t\{\mathcal{A}^{(k)}s_{t+1} + \mathcal{B}^{(k)}s_t + \mathcal{C}^{(k)}s_{t-1} + \mathcal{D}^{(k)}\epsilon_t\} + f^{(k)} = 0$$

where $\mathcal{A}^{(k)}$, $\mathcal{B}^{(k)}$, $\mathcal{C}^{(k)}$, and $\mathcal{D}^{(k)}$ are conformable matrices and $f^{(k)}$ is a vector of constants. In the no-ELB regime, the vector $f^{(n)}$ is filled with zeros. In the ELB regime, the row of $f^{(e)}$ associated with i_t is equal to $\mu_z + \rho + \pi$. Similarly, the rows of the system matrices associated with i_t in the no-ELB regime correspond to the coefficients of the Taylor rule while in the ELB regime, the coefficient associated with i_t is equal to 1 and all the other coefficients are set to zero.

In each period t , given an initial state vector s_{t-1} and vector stochastic innovations ϵ_t , we simulate the model under perfect foresight (i.e., assuming that no further shocks hit the economy) over the next N periods, for N sufficiently large. In case this particular draw is not conducive to an ELB episode, we find s_t using the linear solution stated above. In contrast, if this draw leads to an ELB episode, we postulate integers $N_e < N$ and $N_x < N$ such that the ELB is reached at time $t + N_e$ and left at time $t + N_x$. In this case, we solve the model by backward induction. We obtain the time varying solution

$$s_{t+q} = d_{t+q} + \mathcal{T}_{t+q}s_{t+q-1} + \mathcal{R}_{t+q}\epsilon_{t+q}$$

where, for $q \in \{N_e, \dots, N_x - 1\}$

$$\mathcal{T}_{t+q} = -\left(\mathcal{A}^{(e)}\mathcal{T}_{t+q+1} + \mathcal{B}^{(e)}\right)^{-1}\mathcal{C}^{(e)}, \quad \mathcal{R}_{t+q} = -\left(\mathcal{A}^{(e)}\mathcal{T}_{t+q+1} + \mathcal{B}^{(e)}\right)^{-1}\mathcal{D}^{(e)},$$

$$d_{t+q} = -\left(\mathcal{A}^{(e)}\mathcal{T}_{t+q+1} + \mathcal{B}^{(e)}\right)^{-1}\left(\mathcal{A}^{(e)}d_{t+q+1} + f^{(e)}\right)$$

and, for $q \in \{0, \dots, N_e - 1\}$

$$\mathcal{T}_{t+q} = -\left(\mathcal{A}^{(n)}\mathcal{T}_{t+q+1} + \mathcal{B}^{(n)}\right)^{-1}\mathcal{C}^{(n)}, \quad \mathcal{R}_{t+q} = -\left(\mathcal{A}^{(n)}\mathcal{T}_{t+q+1} + \mathcal{B}^{(n)}\right)^{-1}\mathcal{D}^{(n)},$$

$$d_{t+q} = -\left(\mathcal{A}^{(n)}\mathcal{T}_{t+q+1} + \mathcal{B}^{(n)}\right)^{-1}\left(\mathcal{A}^{(n)}d_{t+q+1} + f^{(n)}\right),$$

using $\mathcal{T}_{t+N_x} = \mathcal{T}$, $\mathcal{R}_{t+N_x} = \mathcal{R}$, and d_{t+N_x} set to a column filled with zeros as initial conditions of the backward recursion.

We then check that given the obtained solution, the system hits the ELB at $t + N_e$ and leaves the ELB at $t + N_x$. Otherwise, we shift N_e and/or N_x forward or backward by one period and start all over again until convergence. Once convergence has been reached, we use the resulting matrices to compute s_t and repeat the process for all the simulation periods.

Our approach is thus similar to the one used by [Coibion et al. \(2012\)](#) in their study of the optimal inflation target in a New Keynesian setup.⁸ A shortcoming of this approach is that the agents in the model are not assumed to expect that the ELB has some probability to bind again in the future, after the exit from (possible) current ELB episode.⁹ As Section 4 below illustrates, the parameter estimates and the associated empirical properties of the model that we obtain are comparable to what is found in related studies.

A welfare-based optimal inflation target A second-order approximation of the household expected utility derived from the structural model is used to quantify welfare, in a similar manner as in [Woodford \(2003\)](#).¹⁰ Let $\mathcal{W}(\pi; \theta)$ denote this welfare criterion. This notation emphasizes that welfare depends on the inflation target π together with the rest of the structural parameters θ . In our approach the structural parameters θ are fixed at reference values (the mean of the posterior likelihood) and taken to be known with certainty by the policy maker.¹¹

The optimal inflation target associated with a given vector of parameters θ , $\pi^*(\theta)$ is approximated via numerical simulations of the model allowing for an occasionally binding ELB constraint, using the algorithm outlined above.¹² The optimal inflation rate associated to a particular vector of parameters θ is then obtained as the one maximizing the welfare function, that is:

$$\pi^*(\theta) \equiv \arg \max_{\pi} \mathcal{W}(\pi; \theta).$$

⁸In practice we combine the implementation of the [Bodenstein et al. \(2009\)](#) algorithm developed by [Coibion et al. \(2012\)](#) with the solution algorithm and the parser from Dynare. Our implementation is in the spirit of [Guerrieri and Iacoviello \(2015\)](#), resulting in a less user-friendly yet faster suite of programs.

⁹Global solution methods, such as advocated and implemented by [Gust et al. \(2017\)](#) are in principle more accurate. However, given the size of our model, and the large set of inflation targets and real interest rates that we need to consider (and given that these have to be considered for each and every parameter configuration in our simulations), a global solution would be computationally prohibitive.

¹⁰See the Technical Appendix in [Andrade et al. \(2019b\)](#) for details.

¹¹See [Andrade et al. \(2019b\)](#) for an alternative exercise in which the policy maker maximizes welfare while recognizing the uncertainty associated with the model's parameters.

¹²More precisely, a sample of size $T = 100000$ of innovations $\{\epsilon_t\}_{t=1}^T$ is drawn from a Gaussian distribution (we also allow for a burn in sample of 200 points that we later discard). We use these shocks to simulate the model for given parameter vector θ . The welfare function $\mathcal{W}(\pi; \theta)$ is approximated by replacing expectations with sample averages. The procedure is repeated for each of $K = 51$ inflation targets on the grid $\{\pi^{(k)}\}_{k=1}^K$ ranging from $\pi = 0.5/4\%$ to $\pi = 5/4\%$ (expressed in quarterly rates). Importantly, we use the exact same sequence of shocks $\{\epsilon_t\}_{t=1}^T$ in each and every simulation over the inflation grid.

Table 1: Estimation Results

Parameter	Prior Shape	Prior Mean	Priod std	Post. Mean	Post. std	Low	High
ρ	Normal	0.30	0.10	0.31	0.09	0.16	0.45
μ_z	Normal	0.47	0.10	0.39	0.07	0.28	0.50
π^*	Normal	0.74	0.10	0.73	0.08	0.59	0.86
α_p	Beta	0.66	0.05	0.63	0.04	0.56	0.69
α_w	Beta	0.66	0.05	0.56	0.04	0.50	0.62
γ_p	Beta	0.50	0.15	0.13	0.05	0.05	0.22
γ_w	Beta	0.50	0.15	0.32	0.11	0.13	0.49
γ_z	Beta	0.50	0.15	0.51	0.15	0.27	0.77
η	Beta	0.70	0.15	0.77	0.03	0.71	0.82
v	Gamma	1.00	0.20	0.84	0.17	0.56	1.11
a_π	Gamma	2.00	0.15	2.06	0.13	1.83	2.27
a_y	Gamma	0.50	0.05	0.50	0.05	0.42	0.58
ρ_{TR}	Beta	0.85	0.10	0.87	0.02	0.84	0.90
σ_z	Inverse Gamma	0.25	1.00	0.83	0.14	0.60	1.05
σ_R	Inverse Gamma	0.25	1.00	0.11	0.01	0.10	0.13
σ_q	Inverse Gamma	0.25	1.00	0.27	0.07	0.17	0.37
σ_c	Inverse Gamma	0.25	1.00	0.19	0.03	0.14	0.24
σ_u	Inverse Gamma	0.25	1.00	0.21	0.14	0.06	0.39
ρ_R	Beta	0.25	0.10	0.40	0.07	0.29	0.53
ρ_z	Beta	0.25	0.10	0.22	0.09	0.08	0.35
ρ_c	Beta	0.85	0.05	0.98	0.01	0.97	0.99
ρ_q	Beta	0.85	0.05	0.93	0.02	0.90	0.96
ρ_u	Beta	0.80	0.05	0.80	0.05	0.72	0.88

Note: 'std' stands for Standard Deviation, 'Post.' stands for Posterior, and 'Low' and 'High' denote the bounds of the 90 percent probability interval for the posterior distribution. Values for parameters ρ , μ_z , π are expressed in percent, in quarterly (not annualized) terms.

4 Estimation Results and Some Pre-crisis Properties of the Model

4.1 Estimation Results

Table 1 presents the parameter's postulated priors (type of distribution, mean, and standard error) and estimation results, i.e., the posterior mean and standard deviation, together with the bounds of the 90 percent probability interval for each parameter.

For the parameters π , μ_z and ρ , we impose Gaussian prior distributions. The parameters governing the latter are chosen so that the model steady-state values match the mean values of inflation, real per capita GDP growth, and the real interest rate in our euro area samples. We use standard prior distributions for the other structural parameters. In particular, we use Beta distributions for parameters in $[0, 1]$, Gamma distributions for positive parameters, and Inverse Gamma distributions for the standard error of the structural shocks.

The estimation results are in the ballpark of parameter values reported in related studies. In particular, while our model does not lead to a closed-form Phillips curve, it can be shown that the coefficient in front of the real marginal cost in the supply curve - a parameter akin to the slope of the New Keynesian Phillips curve (NKPC)-, would be - given our estimated parameters and under the counterfactual assumption of perfect indexation - below 0.08. This value is close to the estimates obtained by [Christoffel et al. \(2008\)](#) over a similar sample. As shown in [Appendix A](#), with less than perfect indexation, the slope of the Phillips

Table 2: Properties of the Model

Parameters	Moments				
	$E(\pi_t)$	$std(\pi_t)$	$E(x_t)$	$std(x_t)$	P(ELB)
a. $\pi = 2\%$, $e = -0.5\%$					
$r^* = 1\%$	1.67	3.80	-0.19	1.32	16.30
$r^* = 2\%$	1.91	2.65	-0.07	0.87	9.80
$r^* = 3\%$	1.99	2.28	-0.02	0.71	5.51
b. $r^* = 2\%$, $e = -0.5\%$					
$\pi = 1\%$	0.67	3.75	-0.18	1.33	16.17
$\pi = 2\%$	1.91	2.65	-0.07	0.87	9.80
$\pi = 3\%$	2.99	2.31	-0.03	0.70	5.64
c. Pre-crisis regime					
$r^* = 2.8\%$, $\pi^* = 1.8\%$, $e = 0$	1.77	2.56	-0.06	0.84	9.06

Note: Results from simulations of the model under various values of r^* , π , and e , and the remaining model parameters at their estimated posterior mean. π_t denotes realized inflation, x_t the output gap, $E(\cdot)$ stands for mean, $std(\cdot)$ stands for Standard Deviation, and P(ELB) denotes the incidence of ELB episodes.


curve is even lower. Finally, an important feature is that the steady-state real rate of interest over the sample period is estimated to be (at the posterior mean, and in annualized terms) 2.8 percent, a value no longer deemed relevant in the current environment.

4.2 Properties of the Estimated Model

Table 2 reports key moments drawn from model simulations with an ELB of -0.5 percent, for counterfactual values for the real natural rate r^* and the inflation target π . When adjusting r^* , we freeze the discount rate ρ at its posterior mean and let μ_z vary accordingly. All the other model parameters are set to their posterior mean.

The moments that we obtain are reasonably close to the ones obtained by Coenen et al. (2020) using the NAWM model of the euro area under various policy packages such as a fully credible forward guidance, or an imperfectly credible forward guidance combined with large scale asset purchases or with fiscal accommodation. For instance with π at 2 percent and r^* at 2 percent, and an assumption of a fully credible forward guidance, they obtain a frequency of hitting the ELB of 8 percent, an average (annualized) inflation rate of 1.86 percent, an average output gap of -0.27 percent. The corresponding moments in our simulations are respectively 9.8 percent, 1.91 percent, and -0.07 percent. When simulating the model for a r^* equal to 1 percent, they obtain a frequency of hitting the ELB of 15 percent, an average (annualized) inflation rate of 1.80 percent, and an average output gap of -0.68 percent. We get respectively 16.3 percent, 1.67 percent, and -0.19 percent. As in Coenen et al. (2020), the constraint on the stabilization instrument induced by the ELB puts a wedge between the average realized inflation and the inflation target, as well as to an output gap that is on average negative. In sum, more frequent ELB episodes induced by a structurally lower r^* generate stabilization costs that are broadly consistent when comparing our model with the one of Coenen et al. (2020).

Figure 1: Frequency of the ELB




Note: The black dots indicate the unconditional probability of hitting the ELB constraint as a function of alternative r^* . This figure is plotted for a value of $\pi = 2$ percent and an ELB at $e = -0.5$.

That the stabilization performances of the policy rules in our model are comparable (and somehow even better) than the ones obtained in [Coenen et al. \(2020\)](#) might look surprising as, unlike what they do, we do not incorporate explicitly forward guidance or large scale asset purchases in our modeling exercise. However, the monetary policy rule that we consider, where the smoothing applies to the notional rate i_t^n that would prevail absent the lower bound constraint (see equation 5) introduces some history-dependence in the policy decision. As a consequence, in the instance of an ELB episode, the central bank is endogenously committed to keep rates lower for longer. As with forward guidance (possibly combined in [Coenen et al. \(2020\)](#)'s setup with a commitment to a prolonged period of asset purchases), this commitment to future accommodation generates higher expected inflation and output, which ex-post helps exiting the trap (or even sometimes not entering it). Another reason is that the effective lower bound constraint that we consider is lower than the one in [Coenen et al. \(2020\)](#) who set it at -0.35 percent instead of -0.5 percent in our case. Finally, their underlying economy is more volatile as, unlike ours, their estimation sample includes the most recent years.

To complement on these illustrative results, Figure 2 displays the probability of reaching the ELB as a function of the steady-state real interest rate, again, with an ELB at -0.5 percent, and with the parameter vector θ evaluated at the posterior mean, except for r^* , which takes values in $[0\%; 5\%]$, and π , which is set at exactly 2 percent (annualized) in this exercise. The probability of hitting the ELB increases quite rapidly when the steady-state real rate decreases. For instance, it increases from nearly 10 percent to 16 percent

Figure 2: Welfare and the Pre-Crisis Optimal Inflation Target



Note: The black line corresponds to the second-order approximation of welfare. The black dot indicates the maximal welfare and the dashed line is set at the corresponding optimal inflation target. This figure is computed using the pre-crisis natural rate of interest, assuming $e = 0$.

when r^* decreases from 2 percent to 1 percent. As will be illustrated below, changes in the ELB incidence are much smaller when the inflation target is allowed to adjust optimally following a decrease in r^* . In the Appendix, Figure B.1 displays the probability of reaching the ELB as a function of the annualized inflation target (for a given r^*), and convey similar insights: the probability increases steeply when the inflation target gets close to zero.

4.3 Pre-Crisis Optimal Inflation

In 2003, the ECB conducted a first review of its strategy. One of the outcomes of the review was to clarify its quantitative definition of price stability which was to be interpreted as an inflation rate of “below, but close to 2 percent” in the medium run. Figure 2 illustrates that, according to our model, this choice was consistent with what the optimal inflation target was given the pre-crisis estimated value of r^* . Indeed, given pre-crisis parameter estimates—in particular for a natural rate of $r^* = 2.8$ percent—and an ELB constraint at $e = 0$ —which was ECB’s perceived lower bound at that time—welfare is maximized at $\pi = 1.8$ percent, strikingly consistent with the choice that was made in 2003. The other way around, the quantitative trade-off between the incidence of the ELB and a positive steady-state inflation that our model delivers seems to align with the policy preferences at the time of the 2003 review. In particular, it does not overweight the costs of the ELB compared to those of a positive steady-state inflation.

The bottom panel of Table 2 illustrates that, according to our model, under such a policy, and considering that the model parameters were fixed at their pre-crisis value, the monetary policy authority was implicitly ready to tolerate an ELB incidence of about 10 percent. As Table 2 also illustrates, under such a low ELB incidence, the central bank could almost perfectly reach on average its targeted rates of inflation and output gap.

A pre-crisis estimates of $r^* = 2.8\%$ is broadly consistent with the views on that parameter at the time (see for example Coenen, 2003). However, as it became clear in the 2010s, r^* had been trending downward since the early 80s so that, in 2003, it was already below its value estimated over the pre-crisis and Great Moderation sample (see Brand and Mazelis, 2019). Thus, a potential issue with the policy decision that was made at the time is that, under a more conservative (that is lower) value of r^* , the ELB incidence would be significantly increased, which would call for a higher π^* . The link between r^* and π^* that we emphasize call for acknowledging that π^* is not a constant and depends on some key parameters which evolve over time and which value can be uncertain.¹³ In other words, it may make sense to reassess the inflation target of a central bank when those parameters are deemed to have changed sufficiently.

5 Adjusting the Inflation Target

The focus of this section is to investigate how the monetary authority should adjust its inflation target π^* in response to changes in the steady-state real interest rate, r^* . Intuitively, with a lower r^* the ELB is bound to bind more often, so one would expect a higher inflation target to be desirable in that case. But the answer to the practical question of by how much should the target be increased is not obvious. Indeed, the benefit of providing a better hedge against hitting the ELB, which is an infrequent event, comes at a cost of higher steady-state inflation which induces permanent costs, as argued by, e.g., Bernanke (2016).

5.1 Constructing the (r^*, π^*) Relation


To start with, we compute the relation linking the optimal inflation target to the steady-state real interest rate, based on simulations of the model. We show that the link between π^* and r^* depends to some extent on the reason underlying a variation in r^* , i.e. a change in growth rate of technology μ_z or a change in the discount rate ρ . In our set-up the first scenario roughly captures the “productivity slowdown” rationale for secular stagnation, while the second is meant to capture an increase in the demand for savings due to factors unrelated to changes in productivity growth. Again, we assume that the effective lower bound is at $e = -0.5$ percent in annualized terms, throughout the present section.

¹³In Andrade et al. (2019b) we analyse how a central bank should choose its inflation target when facing uncertainty about the parameters of the model describing the economy—in particular the ones determining r^* . It turns out that, because the welfare function is asymmetric and the costs of being below the optimal inflation target dominate the costs of being above, the central bank should choose a higher inflation target than in a situation where parameters are certain.

To characterize the link between r^* and π^* , the following simulation exercise is conducted. The structural parameter vector θ is fixed at its posterior mean, $\bar{\theta}$, with the exception of μ_z and ρ . These two parameters are varied – each in turn, keeping the other parameter, μ_z or ρ , fixed at its baseline posterior mean value. For both μ_z and ρ , we consider values on a grid ranging from 0.4 percent to 10 percent in annualized percentage terms. The model is then simulated for each possible values of μ_z or ρ and various values of inflation targets π using the same procedure as before.¹⁴ The optimal value π^* associated to each value of r^* is obtained as the one maximizing the welfare criterion $\mathcal{W}(\pi; \theta)$.

We finally obtain two curves. The first one links the optimal inflation target π^* to the steady-state real interest rate r^* for various growth rates of technology μ_z : $\pi^*(r^*(\mu_z))$, where the notation $r^*(\mu_z)$ highlights that the steady-state real interest rate varies as μ_z varies. The second one links the optimal inflation target π^* to the steady-state real interest rate r^* for various discount rates ρ : $\pi^*(r^*(\rho))$. Here, the notation $r^*(\rho)$ highlights that the steady-state real interest rate varies as ρ varies.

Figure 3: (r^*, π^*) locus (at the posterior mean)



Note: the black dots correspond to the $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus when r^* varies with μ_z ; the gray dots correspond to the $(r^*(\rho), \pi^*(r^*(\rho)))$ locus when r^* varies with ρ .

¹⁴In particular, we use the same sequence of shocks $\{\epsilon_t\}_{t=1}^T$ as used in the computation implemented in the baseline exercises of Section 3.2. Here again, we start from the same grid of inflation targets for all the possible values of μ_z or ρ . Then, for each value of μ_z or ρ , we refine the inflation grid over successive passes until the optimal inflation target associated with a particular value of μ_z or ρ proves insensitive to the grid.

5.2 The (r^*, π^*) Relation: Results

Figure 3 depicts the (r^*, π^*) relations thus obtained. The black dots correspond to the case when the real steady-state interest rate r^* varies with μ_z . The gray dots correspond to the case when the real steady-state interest rate r^* varies with ρ . For convenience, both the real interest rate and the associated optimal inflation target are expressed in annualized percentage rates. The dashed lines indicate the benchmark result corresponding to the optimal inflation target at the posterior mean of the structural parameter distribution. For this baseline scenario, we have $r^* = 2.8$ percent which, according to our quantitative exercise implies an optimal inflation target of $\pi^* = 1.4$ percent with and ELB set at $e = -0.5$.¹⁵

These results are complemented with Figure 4 which shows the relation between r^* and the probability of hitting the ELB, *evaluated at the optimal inflation target*. As with Figure 3, black dots correspond to the case when r^* varies with μ_z , while gray dots correspond to the case when it varies with ρ .

As expected the relation depicted in Figure 3 is decreasing. However, the slope varies with the value of r^* . It is relatively large in absolute value – although smaller than one, with a numerical value close to -0.8 —for moderate values of r^* (say below 4 percent). The slope declines in absolute value as r^* increases: lowering the inflation target to compensate for an increase in r^* becomes less and less desirable. This reflects the fact that, as r^* increases, the probability of hitting the ELB becomes smaller and smaller. For very large r^* values, the probability becomes almost zero, as Figure 4 shows.


At some point, the optimal inflation target becomes insensitive to changes in r^* when the latter originate from changes in the discount rate ρ . In this case, the inflation target stabilizes at a slightly negative value, in order to lower the nominal wage inflation rate required to support positive productivity growth, given the imperfect indexation of nominal wages to productivity. At the steady state, the real wage must grow at a rate of μ_z . It is optimal to obtain this steady-state growth as the result of a moderate nominal wage increase and a moderate price *decrease*, rather than as the result of a zero price inflation and a consequently higher nominal wage inflation.¹⁶

The previous tension is even more apparent when r^* varies with μ_z since, in this case, the effects of imperfect indexation of wages to productivity are magnified given that a higher μ_z calls for a higher growth in the real wage, which is optimally attained through greater price deflation, as well as a higher wage inflation. Notice however that even in this case, the optimal inflation target becomes little sensitive to

¹⁵As discussed above, the model delivers a pre-crisis benchmark that is very much in line with the ECB inflation aim of “below but close to 2 percent over the medium term” that was adopted after the 2003 strategy review when assuming that the ELB is at $e = 0$ instead, consistent with the fact that negative interest rates policies were considered not to be an option until the recent years.

¹⁶For very large r^* , as a rough approximation, we can ignore the effects of shocks and assume that the ELB is a zero-mass event. Assuming also a negligible difference between steady-state and efficient outputs and letting λ_p and λ_w denote the weights attached to price dispersion and wage dispersion, respectively, in the approximated welfare function, the optimal inflation obeys $\pi^* \approx -\lambda_w(1 - \gamma_z)(1 - \gamma_w)/[\lambda_p(1 - \gamma_p)^2 + \lambda_w(1 - \gamma_w)^2]\mu_z$. Given the low values of λ_w resulting from our estimation, it is not surprising that π^* is negative but close to zero. See [Amano et al. \(2009\)](#) for a similar point in the context of a model abstracting from ELB issues.

Figure 4: Relation between probability of ELB at optimal inflation and r^* (at the posterior mean)



Note: the black dots correspond to the $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus when r^* varies with μ_z ; the gray dots correspond to the $(r^*(\rho), \pi^*(r^*(\rho)))$ locus when r^* varies with ρ .

changes in r^* for very large values of r^* , typically above 6 percent.

For low values of r^* , on the other hand, the slope of the curve is steeper. In particular, in the empirically relevant region, the relation is not far from one-to-one. More precisely, starting from the posterior mean estimate of θ , a 1 percentage point decline in r^* should lead to a 0.8 percentage point increase, corresponding to an inflation target of about 2.2 percent instead of 1.4 percent. Importantly, this increase in the optimal inflation target is similar independently of the underlying factor causing the change in r^* : a drop in potential growth, μ_z , or a decrease in the discount factor, ρ . At the same time, the ELB incidence evaluated at the optimal inflation rate also increases when the real interest rate decreases. As Figure 4 shows, adjusting the inflation target limits the increase in the probability of an ELB regime to around 1 percentage point. That a 1 percentage point decline in r^* leads to an optimal adjustment of π^* of about 0.8 percentage point and limits the increase in the probability of hitting the ELB to about 1 percentage point should not be interpreted as a general rule. As our simulations illustrate, this optimal adjustment slope depends on the initial level of r^* one considers. It also depends on the underlying structural parameters. For instance in [Andrade et al. \(2019b\)](#) we obtain for the US economy that a 1 percentage point decline in r^* from its pre-crisis level should lead to an optimal adjustment of π^* of about 1 percentage point, and at the same time a subdued increase in the probability of hitting the ELB.

To conclude this section, we note that while our exercise has focused on the case a 1 percentage point decrease in r^* for illustrative purposes, this stands on the conservative side in view of the existing empirical

estimates of the decline in r^* , with several studies reporting an estimate of that parameter near zero or even negative in the recent period (for example, [Brand and Mazelis, 2019](#), [Holston et al., 2017](#)). Our baseline exercises which vary r^* by altering either μ_z or ρ and fixing the other parameters at their posterior mean cannot deliver very low levels of r^* . Figure C.1 in the Appendix, displays the (r^*, π^*) locus obtained when fixing μ_z (resp. ρ) at a value close to zero, and then varying r^* by varying ρ (resp. μ_z). This exercise allows to consider values of r^* that can get close to zero. The results show that, for values of r^* that are below 2 percent, the slope of the (r^*, π^*) curve stays relatively constant and close to the 0.8 value one gets with our baseline simulations. The optimal inflation target when $r^* = 0$ is about 3.7 percent. We also note that our analysis does not take into account the possibility of a structurally larger volatility of shocks in the “new normal”, that would also lead to an even higher π^* . In [Andrade et al. \(2019b\)](#), we however illustrate that the local *slope* of the (r^*, π^*) remains unchanged with more volatile shocks and that it can lead to much higher optimal inflation target for a given value of r^* .

6 Alternative Policy Setups

In this section, we alter the set-up considering in turn: a model in which fiscal policy has a substantial countercyclical role in large recessions (in general also implying ELB episodes); a lower level of the ELB; alternative monetary strategies embodied by different reaction functions. It has to be underlined that the first two types of scenarios, to a large extent, do not fall under the decision of the central bank.

6.1 Emergency Fiscal Package

So far, we considered economies in which the only tool available to the policy maker to mitigate the incidence of the ELB is adjusting the inflation target. One may however argue that in practice, emergency fiscal packages (EFP henceforth) are triggered in the event of a particularly dreadful crisis.¹⁷ It might this be the case that allowing for such a fiscal package will impinge on the slope of the (r^*, π^*) relation. To investigate this possibility, we consider a variation on our baseline setup in which we allow for such an EFP.

The assumed emergency fiscal package is triggered when the cumulative output gap over four consecutive quarters is below -6 percent. In that case, the fiscal package is modeled as an expansion of government purchases, amounting to 4 percent of the steady-state level of output. We also assume that after the crisis, when the cumulative output gap is above -6 percent, the fiscal package is withdrawn only gradually, with a coefficient of serial correlation equal to 0.85. This means that each quarter after the crisis (as defined here), only 15 percent of the package is withdrawn. This package is calibrated so as to prevent extremely adverse outcomes resulting from the deflation spiral embedded in the New Keynesian model considered here. It turns out that such events are generally associated with large output losses.¹⁸ As shown

¹⁷The fiscal reaction to the current covid-19 crisis in the US and the euro area is an obvious example.

¹⁸This package shares resemblance with the fiscal devices considered in [Reifschneider and Williams \(2000\)](#), [Williams \(2009\)](#), or

Table 3: Properties of the Model with an Emergency Fiscal Package

Parameters	Moments						
	$E(\pi_t)$	$std(\pi_t)$	$E(x_t)$	$std(x_t)$	$P(ELB)$	$E(EFP)$	$p_{95}(EFP)$
a. $\pi = 2\%$, $e = -.5\%$							
$r^* = 1\%$	2.03	2.26	0.02	0.72	12.96	1.02	7.83
$r^* = 2\%$	2.04	2.17	0.02	0.67	7.99	0.57	4.09
$r^* = 3\%$	2.03	2.13	0.02	0.64	4.59	0.34	1.52
b. $r^* = 2\%$, $e = -.5\%$							
$\pi = 1\%$	1.05	2.22	0.02	0.72	12.52	1.04	7.83
$\pi = 2\%$	2.04	2.17	0.02	0.67	7.99	0.57	4.09
$\pi = 3\%$	3.03	2.16	0.02	0.63	4.82	0.32	1.31

Note: Results from simulations of the model under various values of r^* , π , and e , and the remaining model parameters at their estimated posterior mean. π_t denotes the year-on-year inflation rate, x_t is the output gap, $E(\cdot)$ stands for mean, $std(\cdot)$ stands for Standard Deviation, $P(ELB)$ denotes the unconditional probability of hitting the ELB, EFP stands for emergency fiscal package (as a percentage of steady-state output), and $p_{95}(EFP)$ is the 95 percentile of the distribution of emergency fiscal packages.

by [Christiano et al. \(2011\)](#), in such circumstances, the type of fiscal packages considered here can prove very useful.

The effectiveness of fiscal policy as a stabilization instrument when the economy is at the ELB is illustrated in Table 3 which reports key moments from simulations of the model allowing for an EFP as outlined above. For instance, with $\pi = 2$ percent and $r^* = 1$ percent, the above mentioned inflation and output gap biases are almost non-existent—the average (annualized) inflation rate is 2.03 percent, and the average (annualized) output gap is 0.02 percent— even though the frequency of hitting the ELB approaches 13 percent. Note that, on average the package accounts for around 1 percent of the euro area GDP, which is sizable.

As Figure 5 illustrates, the possibility to trigger EFPs would lower the optimal inflation target from 1.4 percent in our baseline economy to .4 percent. However, the existence of the EFP does not completely offset the benefits of raising the inflation target in face of a lower r^* . A drop of 1 percentage point in r^* still calls for an increase by 0.8 percentage points in π^* .

6.2 Alternative Effective Lower Bound


The ECB’s deposit facility rate, which gears the overnight money market rate because of excess liquidity, was set at a negative value of -0.1 percent in June 2014 and has been further lowered down to -0.5 percent in September 2019. We have used this last value as a proxy for the ECB’s ELB.¹⁹ While the ELB is strictly speaking not a policy parameter (it should be determined by the point at which most economic agents would switch to cash to avoid negative interest rate), the choice to move to negative rate was arguably a significant policy decision.

To evaluate what the possibility of negative interest rate policies implies for the optimal inflation target,

in [Kiley and Roberts \(2017\)](#).

¹⁹Note the interest rate in our model is measured by the Euribor which typically stands several basis points above the Deposit Facility rate over the last years.

Figure 5: (r^*, π^*) locus (at the posterior mean)



Note: the black dots correspond to the baseline $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus obtained by letting r^* vary with μ_z ; the gray dots correspond to the $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus, allowing for an emergency fiscal package.


we simulate the estimated euro-area model with an effective lower bound set at zero instead of -0.5 percent. Results are presented in Figure 6. As expected, the (r^*, π^*) locus is shifted upwards, though by somewhat less than 50 basis points. For a baseline r^* , the optimal inflation target stands at 1.8 percent (arguably consistent with the ECB “below, but close to, 2%” definition of price stability). A 1 percentage point decline in r^* still calls for a 0.8 percentage point increase in π^* (thus to 2.6 percent, when moving down from the benchmark value of r^*).

We also quantify the (r^*, π^*) relation when allowing for an even lower effective lower bound which we set, for illustration, at -0.8 percent. We retain this value as it is line with the minimum value that has been historically observed in the trajectories of expected values for the Deposit Facility rate by professional forecasters and market participants in the euro area. As shown on Figure 6, the (r^*, π^*) locus is slightly shifted downward by less than 30 basis points. Moreover, under such a lower ELB, a drop in r^* of 1 percentage point below its baseline value calls for increasing the optimal inflation target from 1.2 percent to around 2 percent.

6.3 Lower for Longer

Our analysis is conditional on a specific reaction function of the central bank, described in our setup by the set of parameters α_π , α_x and ρ_i . Among these parameters, the smoothing parameter, ρ_i , has a key influence on the probability of experiencing a binding ELB.

Figure 6: Optimal inflation with negative ELB




Note: Each curve corresponds to a $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus. The black dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The gray dots correspond to the counterfactual simulation with an ELB set at 0. The gray squares correspond to the counterfactual simulation with an ELB set at -80 basis points.

A higher smoothing has, focusing on instances of exiting the effective lower bound, two effects in our model. The first effect is—through standard monetary policy rule inertia—to reduce the speed at which interest rates are raised when the economy exits the lower bound regime since the current rate inherits the past values of the effective nominal rate. The second effect comes from the fact that the smoothing applies to the notional rate i_t^n that would prevail absent the lower bound constraint (see equation 5) while the effective nominal interest rate is the maximum of zero and the notional rate (see equation 6). Thus the interest rate inherits the past negative values of the notional nominal rate. So, a higher smoothing results in maintaining the effective interest rate at zero for an extended period of time beyond that implied by the macroeconomic shocks that initially brought the economy at the zero lower bound constraint. As with FG, or with Price Level Targeting and Average Inflation Targeting strategies that we discuss below, our baseline monetary policy rule introduces history-dependence whereby, in the instance of an ELB episode, the central bank is committed to keeping rates lower for longer. The interest rate is set as a function of a weighted average of both past inflation gaps and output gaps, with the weights declining more rapidly for farther observations when ρ_i declines.²⁰ As this reaction function is known to the agents in the model, this commitment to future accommodation generates higher expected inflation and output, which ex-post helps exiting the trap

²⁰In contrast, in the case of AIT, the interest rate is determined as a function of a simple average of past inflation gaps over a bounded period.

Figure 7: (r^*, π^*) relation with alternative degrees of interest rate smoothing ρ_i



Note: Each curve corresponds to a $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus. The black dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\hat{\theta}$. The gray dots correspond to the counterfactual simulation with ρ_i set to 0.8. The gray squares correspond to the counterfactual simulation with ρ_i set to .95.

(or even sometimes not entering it).

A larger degree of smoothing reinforces the history-dependence of monetary policy, and, because of its impact on expectations, for given parameter values of the model, at the equilibrium, it tends to shorten the length of the ELB episodes and the probability of hitting the ELB constraint. Everything else being equal, one should therefore expect a lower optimal inflation rate for higher values of the smoothing parameters. This property of the model is illustrated in Figure 7 which depicts the (r^*, π^*) relation under three possible values of the smoothing parameter ρ_i . The values used under our baseline scenario, i.e. posterior mean estimates, is 0.87. We also consider two alternative setups: A higher value of $\rho_i = .95$ which is close to the persistence of the Fed reaction function over the 1960-2007 period calibrated by Coibion et al. (2012), and a lower value of $\rho_i = 0.8$. These two values are arguably an upper bound of the existing empirical uncertainty on the degree of smoothing, as they stand outside the 90 percent probability interval of our posterior parameter estimates.

The effect of a higher interest rate smoothing is to shift downward the (r^*, π^*) curve except for high values of r^* for which the probability of hitting the ELB is close to zero and the optimal inflation target is slightly negative. Under this strategy, the pre-crisis optimal inflation rate would be close to 0 percent in the euro area. Conversely, a lower interest rate smoothing shifts the (r^*, π^*) curve upward, even for relatively high values of r^* – because the probability of being in an ELB regime increases under this strategy. With a lower ρ_i , the pre-crisis optimal inflation rate would be close to 3 percent in the euro area.

As for the slope of the (r^*, π^*) curve, in the empirically relevant region, it is much less affected than the level of this locus. It is however more affected in this exercise than in other robustness experiments considered above. A very large smoothing parameter, due to its effect outlined above on the probability of ELB, somewhat alleviates the extent to which an increase in the inflation target is needed. The slope is indeed close to -0.7 in that case. For a strategy associated with a low smoothing parameter, the slope is close to -1 . For large values of r^* , the degree of smoothing is irrelevant.

Overall, increasing the length of the period during which the central bank is committed to keep its policy rate below the natural rate after a shock which calls for more accommodation can be an alternative to increasing the inflation target.

6.4 Adopting a Price Level Targeting rule

We also consider a PLT type of rule whereby the central bank reacts to deviations of the (log) price level to a target path, $\bar{p}_t = \bar{p}_{t-1} + \pi$ instead of the gap $\hat{\pi}_t$ between the inflation rate and its target. Formally, we assume that the central bank sets the policy rate according to the following rule:

$$\hat{i}_t^n = \rho_i \hat{i}_{t-1}^n + (1 - \rho_i) (a_p \hat{p}_t + a_x \hat{x}_t) + \zeta_{m,t} \quad (7)$$


with $\hat{i}_t = \max\{\hat{i}_t^n, e - (\mu_z + \rho + \pi)\}$ and where $\hat{p}_t = p_t - \bar{p}_t$.

We perform the same exercises as before, focusing on the case in which average productivity growth μ_z is the driver of changes in r^* . We consider two values for a_p : 0.1 and 1. We assume furthermore that $\rho_i = 0$. All the other parameters of the model are set to their posterior mean.

Figure 8 reports the (r^*, π^*) relation obtained under these two alternative scenarios. A striking feature of the new curve is that the pre-crisis optimal inflation target is slightly below 0 percent as opposed to 1.4 percent in the baseline. Price level targeting implies that the commitment to make-up for past inflation undershooting (or overshooting) is even stronger than when increasing the smoothing parameters in a rule which targets inflation instead. This commitment stabilizes inflation expectations lowering dramatically both the probability of hitting the zero lower bound and the average length of such episodes. As a consequence, there is no incentive to bear the costs of a positive steady-state inflation and the optimal inflation target is close to zero. This holds no matter whether the central bank reacts aggressively or not to the deviation of the price level to its targeted path.

Another striking result is that the (r^*, π^*) relation is much flatter in the vicinity of the pre-crisis level for r^* than under alternative inflation targeting monetary policy strategies. The slope is close to -0.4 , which is lower than the range of values obtained previously. Reacting to a one percentage point drop in r^* by adopting a full-blown price-level targeting strategy here only call for setting π^* around 0.2 percent. Thus even after a decline in r^* , the buffer provided by the current inflation aim of “below, but close, to 2 percent” would be, in our fully rational expectations set-up, unnecessarily large compared to the optimal one under an PLT strategy.

Figure 8: (r^*, π^*) relation with price level targeting strategy



Note: Each curve corresponds to the $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus obtained under the price level targeting policy rule given in equation 7 for a given a_p . The black (gray) dots correspond to the scenario wherein $a_p = .1$ (1). All the other structural parameters are set at their posterior mean $\hat{\theta}$.

6.5 Adopting an Average Inflation Targeting rule

We finally consider a rule whereby the central bank reacts to deviations of an average of past deviation from the inflation target. This alternative make-up strategy has been studied in several recent works, notably [Mertens and Williams \(2019a\)](#), and can be described as intermediate strategy between inflation targeting and price level targeting. Formally, we assume that the central bank sets the policy rate according to


$$\hat{i}_t^n = \rho_i \hat{i}_{t-1}^n + (1 - \rho_i) (a_\pi \bar{\pi}_t + a_x \hat{x}_t) + \zeta_{m,t} \quad (8)$$

with $\hat{i}_t = \max\{\hat{i}_t^n, e - (\mu_z + \rho + \pi)\}$ and $\bar{\pi}_t = (1/H) \sum_{h=0}^{H-1} \hat{\pi}_{t-h}$

We perform the same exercises as before, focusing on the case in which average productivity growth μ_z is the driver of changes in the natural rate. We consider two values for H : 16 and 32 quarters. All the other parameters of the model are set to their posterior mean.

Figure 9 illustrates how adopting such a strategy affects the (r^*, π^*) relation. As under a PLT strategy, the initial optimal inflation target is much lower in that case with the baseline inflation targeting rule. However, the slope remains of the same order as with the previous inflation targeting strategy. When inflation used in the rule is averaged over the window of the last 4 years, a drop in r^* of 1 percentage point still requires to increase the optimal inflation target by around 0.8 percentage point. Facing a decrease in r^* , increasing the length of the period can allow to maintain the optimal target unchanged, at its initial optimal

Figure 9: (r^*, π^*) relation with AIT strategy



Note: Each curve corresponds to the $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus obtained under the average inflation targeting policy rule given in equation 8 for a given H . The gray dots (squares) correspond to the scenario wherein the average is taken over the last $H = 16$ ($H = 32$) quarters. All the other structural parameters are set at their posterior mean $\bar{\theta}$.

level of, here, 1.4 percent. But this requires to contemplate very long windows of time, at least equal to 4 years, to compute the average inflation relevant for computing the deviation from target. Note that under the 8 year window AIT, the slope of the (r^*, π^*) curve is still around -0.8 in the region of interest. For instance, with a window of 8 years, π^* is initially close to 0 percent, and would increase to 0.7 percent in reaction to a 1 percentage point drop in r^* . Interestingly, if one were to extrapolate the previous slope to the area where r^* is close to zero, one would get a π^* that is close to 2 percent.

7 Conclusion

We have provided a quantitative assessment of how changes in the steady-state natural interest rate translate into changes in the optimal inflation target in a model of the euro area, subject to the ELB. Our main finding is that, starting from a baseline “pre-crisis” euro area economy, a 1 percentage point decline in the natural rate should be accommodated by an increase in the optimal inflation target of about 0.8 percentage point. While these results are obtained with a model estimated on euro-area data, they are overall close to the ones we have obtained for the U.S. economy in [Andrade et al. \(2019b\)](#). Interestingly, this main result holds true even if we allow for emergency fiscal packages as an additional response to large recessions (and a way to make the ELB less costly), or under an AIT strategy, two possibilities that we did not explore in our earlier study. Our results show that, in the low r^* environment that we focus on, the costs resulting from

the limits to stabilization induced by the ELB exceed the costs entailed by increasing the inflation target. As emphasized for example by [Nakamura et al. \(2018\)](#), the latter are relatively constant for low values of steady-state inflation.

Our analysis also highlights which policy options other than raising the inflation target are at the central bank's disposal in face of a lower r^* . They all involve a stronger commitment to make-up for foregone inflation during ELB episodes. Our simulations can be used to quantify the degree of "memory" in the policy rule that is required, as a substitute to an increase in the inflation target, in response to a 1% decline in r^* : It is at least as large as 4 years. This is obtained under the standard benchmark of full information and rational expectations. Make-up strategies are quite effective under that assumption as expectations fully internalize the policy and agents decisions fully react to their expectations. However, in practice, information and cognitive constraints attenuate the expectation channel (see for example [Pfajfar and Roberts 2018](#), [Andrade et al. 2020](#), [Coibion et al. 2020](#)) and that may call for an even larger degree of "memory" as illustrated by [Budianto et al. \(2020\)](#) or [Hebden et al. \(2020\)](#). This begs the question of how private agents would in practice react to such inertial policy rules and whether the implied long-lasting deviations of inflation realizations from target could lead to a counterproductive de-anchoring of inflation expectations. More generally, [Andrade et al. \(2019a\)](#), [Gabaix \(2020\)](#) or [Adam et al. \(2020\)](#) show that make-up strategies can be suboptimal when cognitive constraints are considered. In addition, under such limitations, inflation expectations may be less responsive to economic conditions, making inflation itself less volatile. This may as well change the costs and benefits of changing the inflation target, potentially making a higher inflation target even more appealing. That being said, a less effective expectation channel could be compensated by the additional option of implementing large asset purchases. As emphasized by [Iovino and Sergeyev \(2018\)](#) or [Woodford and Xie \(2020\)](#), deviations from the assumption of full information and rational expectation can make the economy non-Ricardian, which opens the way for quantitative easing to work in theory. Assessing whether adjusting the inflation target is preferable to moving to make-up strategies in a set-up departing from full information and fully rational expectations is left for further research.

Appendix

A Slope of the NKPC Under Less than Perfect Indexation

Following the derivation in [Ascari \(2004\)](#), the slope of the NKPC under less than perfect indexation in the context of our model is

$$\kappa_p(\Pi, \alpha_p; \theta) = \frac{(1 - \beta \alpha_p(\Pi)^{(1-\gamma_p)\phi\theta_p})(1 - \alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)})}{\alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)}[1 + \theta_p(\phi - 1)]}$$


where $\Pi = e^\pi$ and θ is a vector regrouping all the parameters appearing in the NKPC other than Π and α_p . The special case of perfect indexation to steady-state inflation boils down to setting $\Pi = 1$.

Since $\theta_p > 1$, $\kappa_p(\Pi, \alpha_p; \theta)$ is decreasing in both Π and α_p . It follows that for $\Pi > 1$, one has

$$\kappa_p(\Pi, \alpha_p; \theta) < \kappa_p(1; \theta).$$


B Probability of ELB as a function of π

Figure B.1: Probability of ELB



Note: The black dots indicate the unconditional probability of hitting the ELB constraint as a function of the inflation target π , for a value of $r^* = 1$ percent.

Figure C.1: (r^*, π^*) relation with near-zero r^*



Note: In the top panel, the black dots correspond to the baseline $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus obtained by letting r^* vary with μ_z with all the other parameters including ρ set at their posterior value. The gray dots correspond to the $(r^*(\mu_z), \pi^*(r^*(\mu_z)))$ locus when r^* varies with μ_z , with all the other parameters set at their posterior value, and ρ is set at a very low value of .1 per annum. In the bottom panel, the black dots correspond to the baseline $(r^*(\rho), \pi^*(r^*(\rho)))$ locus obtained by letting r^* vary with ρ with all the other parameters including μ_z set at their posterior value. The gray dots correspond to the $(r^*(\rho), \pi^*(r^*(\rho)))$ locus when r^* varies with ρ , with all the other parameters set at their posterior value, and μ_z is set at a very low value of 0 percent per annum.

C Very Low Real Natural Rates

In this section, we explore the robustness of the baseline (r^*, π^*) curve when we allow for even lower real interest rates. To do so, we either set ρ to .01 per annum and let vary μ_z over the same range as before or we set μ_z to zero and let vary ρ over the same range as before.