The Optimal Inflation Target and the Natural Rate of Interest [1]

Philippe Andrade

Banque de France

Jordi Galí

CREI, UPF and Barcelona GSE

Hervé Le Bihan

Banque de France

Julien Matheron

Banque de France

December 14, 2018

Abstract

We study how changes in the steady-state real interest rate affect the optimal inflation target in a New Keynesian DSGE model with trend inflation and a lower bound on the nominal interest rate. In this setup, a drop in the steady-state real interest rate increases the probability of hitting that constraint. This higher probability can be offset by an increase in the inflation target inducing a higher average nominal interest rate. However, a higher inflation target also entails greater distortion costs induced by steady-state inflation. We estimate the model on both U.S. and euro area data to quantify this trade-off. We find that the relation between the steady-state real interest rate and the optimal inflation target is downward sloping, but its slope is not necessarily one-for-one: increases in the optimal inflation rate are generally lower than declines in the steady-state real interest rate. However, in the currently empirically relevant region for the US as well as the euro area, the slope of the relation is close to -0.9. That latter finding is robust to considering several alternative parameter values as well as parameter uncertainty.

Keywords: inflation target, Effective lower bound.
JEL Codes: E31, E52 , E58

[1] Andrade: philippe.andrade@banque-france.fr. Galí: jgali@crei.cat. Le Bihan: herve.lebihan@banque-france.fr. Matheron: julien.matheron@banque-france.fr. We thank our discussants Isabel Correia, Kevin Lansing, Argia Sbordone, Raf Wouters and Johannes Wieland as well as participants at several conferences and seminars, for comments. The views expressed herein are those of the authors and do not necessarily reflect those of the Banque de France or the Eurosystem.

1 Introduction

A recent but sizable literature has pointed to a permanent –or, at least very persistent– decline in the "natural" rate of interest in advanced economies (Del Negro et al., 2018, Holston et al., 2017, Laubach and Williams, 2016). Various likely sources of that decline have been discussed, including a lower trend growth rate of productivity (Gordon, 2015), demographic factors (Eggertsson et al., 2017) , or an enhanced preference for safe assets (Caballero and Farhi, 2015, Summers, 2014).

A lower steady-state real interest rate matters for monetary policy. Given average inflation, a lower steady-state real rate will cause the nominal interest rate to hit its zero lower bound (ZLB) more frequently, hampering the ability of monetary policy to stabilize the economy, bringing about more frequent (and potentially protracted) episodes of recessions and below-target inflation.

In the face of that risk, and in order to counteract it, several prominent economists have forcefully argued in favor of raising the inflation target (see, among others, Ball 2014, Blanchard et al. 2010, Williams 2016). Since a lower natural rate of interest is conducive to a higher ZLB incidence, 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 ZLB, which is an infrequent event, comes at a cost of higher steady-state inflation which induces permanent costs, as recently argued by Bernanke (2016) among others. The answer to this question thus requires to assess how the tradeoff between the incidence of the ZLB and the welfare cost induced by steady-state inflation is modified when the natural rate of interest decreases. While the decrease in the natural rate of interest has been emphasized in the recent literature, such assessment has received surprisingly little attention.

The present paper contributes to this debate by asking two questions. First, to what extent does a lower steady-state real interest rate ($r^\star$) call for a higher optimal inflation target ($\pi^\star$)? Second, does the source of decline in $r^\star$ matter? Our main contribution is to characterize how the optimal inflation target is related to the steady-state real interest rate, using a structural, empirically estimated, macroeconomic model. Our main findings can be summarized as follows: (i) The relation between $r^\star$ and $\pi^\star$ is downward sloping, *but not necessarily one-for-one*; (ii) in the vicinity of the pre-crisis values for $r^\star$, the slope of the $(r^\star, \pi^\star)$ locus is close to -0.9 ; and (iii) for a plausible range of $r^\star$ values the relation is largely robust to the underlying source of variation in $r^\star$.

Our results are obtained from extensive simulations of a New Keynesian DSGE model estimated for both the US and the euro area over a Great Moderation sample. The framework features: (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 ZLB 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 inflation rate, in order to mitigate the incidence and adverse effects of the ZLB. To our knowledge, these three features have not been *jointly* taken into account in previous analyses of optimal inflation. Our analysis focuses on the trade-off between the costs attached to the probability of the hitting the ZLB constraint and the costs induced by a positive steady-state inflation rate. In particular one restriction is that our analysis does not incorporate alternative instruments that could be considered by policymakers facing to a change in $r^\star$ – such as changing the central bank reaction function, implementing non-conventional policies, or relying on active fiscal policies.

According to our simulations, the pre-crisis optimal inflation target obtained when the policymaker is assumed to know the economy's parameters with certainty (and taken to correspond to the mode of the posterior distribution) is around 2% for the US and around 1.5% for the euro area (in annual terms). This result is obtained in an environment with a relatively low probability of hitting the ZLB (6% for the US and slightly less than 10% for the euro area), given the small shocks estimated on our Great Moderation sample. Our simulations also show that, a 100 basis point drop of $r^\star$ from its pre-crisis level (respectively 2.5% in the US and 2.7% in the euro area) will almost double the probability of hitting the ZLB if the monetary authority keeps its inflation target unchanged. The optimal reaction of the central bank is to increase the inflation target by 99 basis points in the US and by 81 basis points in the euro area. Overall the slope of the $(r^\star, \pi^\star)$ relation close to -.9 in the vicinity of the pre-crisis parameter region. This optimal reaction mitigates the increase in the probability of hitting the ZLB.

A further noticeable feature of our approach is that we perform a full-blown Bayesian estimation of the model, using both US and euro area data. This allows us not only to assess the uncertainty surrounding $\pi^\star$, but also to derive an optimal inflation target taking into account the parameter uncertainty facing the policy maker, including uncertainty with regard to the determinants of the steady-state real interest rate. When that parameter uncertainty is allowed for, those values increase significantly: 2.40% for the US and 2.20% for the euro area The reason why the optimal targets under parameter uncertainty are higher has to do with the fact that the loss function is asymmetric so that choosing an inflation target that is lower than the optimal one is more costly than choosing an inflation target that is above. That being said it remains true that a Bayesian-theoretic optimal inflation target rises by about 90 basis points in response to a downward shift of the distribution in $r^\star$ of 100 basis points.

We also study the $(r^\star, \pi^\star)$ relation obtained under a variety of alternative assumptions: a central bank targeting the average realized inflation target instead of a objective inflation goal, a central bank constrained by an effective lower bound on the policy rate that can be below zero, a central bank uncertain about the parameter values of the model but certain about its reaction function, a central bank with a lower or a higher smoothing parameter in the interest-rate policy rule, structural shocks with higher variance, and higher goods and labor market markups. Strikingly, in the empirically relevant region, the slope of the $(r^\star, \pi^\star)$ curve is much less affected than the level of this locus. It remains in nearly all cases in the same ballpark as in the baseline scenario.

The remaining of the paper is organized as follows. Section 2 presents the New Keynesian model with a ZLB constraint that is used. Section 3 describes how the model is estimated and simulated, as well as how the welfare-based optimal inflation target is computed. Section 4 is devoted to the analysis of the $(r^\star, \pi^\star)$ relation. Section 5 discusses this locus for a set of the alternative exercises. Finally Section 6 provides some concluding remarks.

1.1 Related Literature

To our knowledge no paper has systematically investigated the $(r^\star, \pi^\star)$ relation. Coibion et al. (2012) (and its follow-up Dordal-i-Carreras et al. (2016)) and Kiley and Roberts (2017) are the papers most closely related to ours, as they study optimal inflation in quantitative set-ups that account for the ZLB. However, their analyses hold for a constant steady-state natural rate of interest. Relative to both papers, a key difference is our focus on eliciting the relation between the steady-state real interest rate and optimal inflation. Other differences are (i) our interest in the euro area, in addition to the US; (ii) we estimate, rather than calibrate, the model, and (iii) we allow for wage rigidity in the form of infrequent, staggered, wage adjustments. A distinctive feature with respect to Kiley and Roberts (2017) is that we use a model-consistent, micro-founded loss function to compute optimal inflation.

A series of papers assessed the probability of hitting the ZLB for a given inflation target (see, among others, Chung et al. 2012, Coenen 2003, Coenen et al. 2004). Interestingly, our own assessment of this pre-crisis ZLB incidence falls broadly in the ballpark of available estimates. A related recent contribution by Gust et al. (2017) emphasizes that the ZLB was indeed a significant constraint on monetary policy that exacerbated the recession and inhibited the recovery.[1] Lindé et al. (2017) offer a discussion of alternative methods to implement the ZLB at the estimation stage.

Other relevant references, albeit ones that put little or no emphasis on the ZLB, are the following. An early literature focuses on sticky prices and monetary frictions. In such a context, as shown by Khan et al. (2003) and Schmitt-Grohé and Uribe (2010), the optimal rate of inflation should be slightly negative. Similarly, a negative optimal inflation would result from an environment with trend productivity growth and prices and wages both sticky, as shown by Amano et al. (2009). In this kind of environment, moving from a 2% to a 4% inflation target would be extremely costly, as suggested by Ascari et al. (2015). By contrast, adding search and matching frictions to the setup, Carlsson and Westermark (2016) show that optimal inflation can be positive. Bilbiie et al. (2014) find positive optimal inflation can be an outcome in a sticky-price model with endogenous entry and product variety. Somewhat related, Adam and Weber (2017) show that, even without any ZLB concern, optimal inflation might be positive in the context of a

[1] Gust et al. (2017) rely on global solution methods while we resort to the lighter algorithms developed by Bodenstein et al. (2009) and Guerrieri and Iacoviello (2015). Given 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.

model with heterogeneous firms and systematic firm-level productivity trends. Finally, using a perpetual youth model, Lepetit (2017) shows that optimal inflation can be positive in the presence of heterogeneous discount factor, especially when the social planner is more patient than agents.

Among the recent papers with ZLB, Blanco (2016) studies optimal inflation in a state-dependent pricing model, i.e. a "menu cost" model. In this setup, optimal inflation is typically positive. Two forces explain the result. First, as in our analysis, positive inflation edges the economy against detrimental effects of ZLB.[2] Second, as shown by Nakamura et al. (2016), the presence of state-dependent pricing limits considerably the positive relationship between inflation and price dispersion, thus limiting the costs of inflation.

2 The Model

We use a relatively standard medium-scale New Keynesian model as a framework of reference. Crucially, the model features elements that generate a cost to inflation: (1) nominal rigidities, in the form of staggered price and wage setting; (2) less than perfect 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 lack of indexation to trend 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 ZLB constraint. In particular, and given the steady-state real interest rate, the incidence of binding ZLB episodes should decline with the average rate of inflation. Such episodes hamper the stabilization potential of monetary policy.

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 we assume Calvo-style sticky wages, in addition to sticky prices.[3]

[2] By contrast, see Burstein and Hellwig 2008 for a similar exercise without ZLB, which leads to negative optimal inflation rate.

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

2.1 Households

The economy is inhabited by a continuum of measure one 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

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

where $\beta \equiv \mathrm{e}^{-\rho}$ 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 $\zeta_{g,t}$.

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

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

where P_t is the aggregate price level, $W_t(h)$ is the nominal wage rate associated with labor of type h, $\mathrm{e}^{\zeta_{q,t}} Q_t$ 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

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

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 yields the relation

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

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

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

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 perfect 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 Golosov and Lucas (2007), suggests there is no significant correlation between the frequency of price change and trend inflation.

If drawn to re-optimize in period t, a firms chooses $P_t^\star$ in order to maximize

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

where Λ_t denotes the marginal utility of wealth, $\tau_{p,t}$ is a sales tax paid by firms and rebated to households in a lump-sum fashion, 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

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

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

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

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

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

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

$$N_t = \int_0^1 L_t(f) \mathrm{d}f.$$

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

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

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

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

where, as before, wages are indexed to past inflation. However, we assume that the degree of indexation is here too less than perfect 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 $W_t^\star$ in order to maximize

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

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

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

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

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

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

2.4 Monetary Policy and the ZLB

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

$$\hat{\imath}_t = \rho_i \hat{\imath}_{t-1} + (1-\rho_i)\left(a_\pi \hat{\pi}_t + a_y \hat{x}_t\right) + \zeta_{R,t} \tag{3}$$

where $i_t \equiv -\log(Q_t)$, with $\hat{\imath}_t$ denoting the associated deviation from steady state i.e, $\hat{\imath}_t \equiv i_t - i$. Also, $\pi_t \equiv \log \Pi_t$, $\hat{\pi}_t \equiv \pi_t - \pi$ is the gap between inflation and its target, 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. Finally, $\zeta_{R,t}$ is a monetary policy shock.

Here, π should be interpreted as the central bank target for change in the price index. An annual inflation target of 2% would thus imply $\pi = 2/400 = 0.005$ as the model will be parameterized and estimated with quarterly data. Note that the inflation target thus defined may differ from average inflation.

Crucially for our purpose, the nominal interest rate i_t is subject to a ZLB constraint:

$$i_t \geq 0$$

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

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

The rule effectively implemented is given by:

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

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

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 the lagged notional interest rate $\hat{\imath}_t$, rather that the lagged actual rate.

Before proceeding, several remarks are in order. First, the inflation target, π, is not assumed to be optimal. Note also that realized inflation might be on average below the target as a consequence of ZLB episodes, i.e. $\mathbb{E}\{\pi_t\} < \pi$. In such instances of ZLB, monetary policy fails to deliver the appropriate degree of accommodation, resulting in a more severe recession and lower inflation than in an economy in which there would not be a ZLB constraint.[4] Second, we assume the central bank policy is characterized by a monetary policy rule rather than a Ramsey-type fully optimal policy of the type studied by e.g. Khan et al.

[4] For convenience, Table A.1 in the Appendix summarizes the various notions of optimal inflation and long-run or target inflation considered in this paper.

(2003) or Schmitt-Grohé and Uribe (2010). Such 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 defining 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 notionnal rate rather than 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 $\hat{\imath}_t^n$ results in the nominal interest rate $\hat{\imath}_t$ being "lower for longer" in the aftermath of ZLB episodes (as $\hat{\imath}_t^n$ will stay negative for a protracted period).

As equation (4) makes clear, μ_z, ρ, π have symmetric roles in the ZLB constraint. Put another way, for given structural parameters and a given process for $\hat{\imath}_t$, the probability of hitting the ZLB 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 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. It turns out that this effect is non-negligible since, as we show later, after a one percentage point decline in r^* the inflation target has to be raised by more than one percent in order to keep the probability of hitting the ZLB unchanged. 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 ZLB 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 ZLB constant. Assessing these forces is precisely this paper's endeavor.

3 Estimation and simulations

3.1 Estimation without a lower bound on nominal interest rates

Estimation procedure 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, namely

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

for $k \in \{R, g, u, q, z\}$.

[5] See the Technical Appendix for further details.

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

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

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 $\epsilon_t = (\epsilon_{R,t}, \epsilon_{g,t}, \epsilon_{u,t}, \epsilon_{q,t}, \epsilon_{z,t})'$. 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 ZLB 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 effective Fed Funds Rate for the US and the Euribor 3 months rate for the Euro-Area.We use a sample of quarterly data covering the period 1985Q2-2008Q3. This choice is guided by two objectives. First, this sample strikes a balance between size and the concern of having a homogeneous monetary policy regime over the period considered. In the US case, the sample covers the Volcker and post-Volcker period, arguably one of relative homogeneity of monetary policy. For the euro area, the sample starts approximately when the disinflation policies were simultaneously conducted in the main euro area countries (see Fève et al. 2010) and then covers the single currency period. Here too, this corresponds to a period of relative monetary policy homogeneity. Second, we use a sample that coincides more or less with the so-called Great Moderation. Over the latter, as has been argued in the literature, we expect smaller shocks to hit the economy. In principle, this will lead to a conservative assessment of the effects of the more stringent ZLB constraint due to lower real interest rates. [6]

The parameters ϕ, θ_p, and θ_w are calibrated prior to estimation. The parameter θ_p is set to 6, resulting in a steady-state price markup of 20%. Similarly, the parameter θ_w is set to 3, resulting in a wage markup of 50%. These numbers fall into the argualby large ballpark of available values used in the literature. In the robustness section, we investigate the sensitivity of our results to these parameters. The parameter ϕ is set to $1/0.7$. Given the assumed tax correcting the steady-state price markup distortion, this results in a steady-state labor share of 70%.

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 measure 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

[6]The data are obtained from the Fred database for the US and from the "Area Wide Model" database of Fagan et al. (2001) and Eurostat national accounts for the Euro Area. In both cases, the GDP is expressed in per capita terms.

Table 1: Estimation Results - US

Parameter	Prior Shape	Prior Mean	Priod std	Post. Mean	Post. std	Low	High
ρ	Normal	0.20	0.05	0.19	0.05	0.11	0.27
μ_z	Normal	0.44	0.05	0.43	0.04	0.36	0.50
$\pi^\star$	Normal	0.61	0.05	0.62	0.05	0.54	0.69
α_p	Beta	0.66	0.05	0.67	0.03	0.61	0.73
α_w	Beta	0.66	0.05	0.50	0.05	0.43	0.58
γ_p	Beta	0.50	0.15	0.20	0.07	0.08	0.32
γ_w	Beta	0.50	0.15	0.44	0.16	0.21	0.68
γ_z	Beta	0.50	0.15	0.50	0.18	0.26	0.75
η	Beta	0.70	0.15	0.80	0.03	0.75	0.85
ν	Gamma	1.00	0.20	0.73	0.15	0.47	0.97
a_π	Gamma	2.00	0.15	2.13	0.15	1.89	2.38
a_y	Gamma	0.50	0.05	0.50	0.05	0.42	0.58
ρ_{TR}	Beta	0.85	0.10	0.85	0.02	0.82	0.89
σ_z	Inverse Gamma	0.25	1.00	1.06	0.22	0.74	1.38
σ_R	Inverse Gamma	0.25	1.00	0.10	0.01	0.09	0.11
σ_q	Inverse Gamma	0.25	1.00	0.39	0.11	0.16	0.61
σ_g	Inverse Gamma	0.25	1.00	0.23	0.04	0.16	0.29
σ_u	Inverse Gamma	0.25	1.00	0.24	0.05	0.06	0.46
ρ_R	Beta	0.25	0.10	0.51	0.06	0.41	0.61
ρ_z	Beta	0.25	0.10	0.27	0.13	0.09	0.45
ρ_g	Beta	0.85	0.10	0.98	0.01	0.97	1.00
ρ_q	Beta	0.85	0.10	0.88	0.04	0.80	0.95
ρ_u	Beta	0.80	0.10	0.80	0.10	0.65	0.96

Note: 'std' stands for Standard Deviation, 'Post.' stands for Posterior, and 'Low' and 'High' denote the bounds of the 90% probability interval for the posterior distribution.

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 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.

Estimation results Tables 1 and 2 present 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% 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 as to match the mean values of inflation, GDP growth, and the real interest rate in our US and euro area samples. Other than for these three parameters, we use the same prior distributions for the structural parameters in both the US and the euro area. Our choice of priors are standard. 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 suggest several key differences between the US and the euro area.

Table 2: Estimation Results - EA

Parameter	Prior Shape	Prior Mean	Priod std	Post. Mean	Post. std	Low	High
ρ	Normal	0.20	0.05	0.21	0.05	0.13	0.29
μ_z	Normal	0.50	0.05	0.47	0.05	0.40	0.55
$\pi^\star$	Normal	0.80	0.05	0.79	0.05	0.71	0.86
α_p	Beta	0.66	0.05	0.62	0.05	0.55	0.68
α_w	Beta	0.66	0.05	0.59	0.04	0.52	0.65
γ_p	Beta	0.50	0.15	0.12	0.04	0.04	0.19
γ_w	Beta	0.50	0.15	0.34	0.12	0.15	0.53
γ_z	Beta	0.50	0.15	0.51	0.18	0.26	0.76
η	Beta	0.70	0.15	0.74	0.04	0.69	0.80
ν	Gamma	1.00	0.20	0.96	0.18	0.65	1.25
a_π	Gamma	2.00	0.15	2.02	0.14	1.80	2.25
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	0.86	0.16	0.63	1.10
σ_R	Inverse Gamma	0.25	1.00	0.11	0.01	0.10	0.12
σ_q	Inverse Gamma	0.25	1.00	0.23	0.05	0.13	0.32
σ_g	Inverse Gamma	0.25	1.00	0.21	0.04	0.15	0.27
σ_u	Inverse Gamma	0.25	1.00	0.23	0.05	0.06	0.43
ρ_R	Beta	0.25	0.10	0.39	0.07	0.27	0.50
ρ_z	Beta	0.25	0.10	0.24	0.10	0.09	0.39
ρ_g	Beta	0.85	0.10	1.00	0.01	0.99	1.00
ρ_q	Beta	0.85	0.10	0.94	0.03	0.90	0.98
ρ_u	Beta	0.80	0.10	0.79	0.10	0.64	0.96

Note: 'std' stands for Standard Deviation, 'Post.' stands for Posterior, and 'Low' and 'High' denote the bounds of the 90% probability interval for the posterior distribution.

First, consistent with the sample period, we find a higher growth rate μ_z in the euro area than in the US. Given that both economies have similar discount rates, this will result in a higher steady-state real interest rate in the euro area than in the US. This difference will play an important role later when we assess (i) the level of the optimal inflation target and (ii) the effects of a lower steady-state real interest rate. Second, we find generically higher degrees of indexation to past inflation in the US than in the euro area. This will translate into a higher tolerance for inflation in the US in our subsequent analysis of the optimal inflation target. This is because a higher indexation helps to mitigate the distortions induced by a higher inflation target. Everything else equal, we thus expect a higher optimal inflation target in the US than in the euro area. Third, we obtain broadly similar parameters for the shocks processes. One exception, though, is the so-called risk-premium shock. The unconditional variance of the US shock is somewhat higher than its euro area counterpart.

For the US, most of our estimated parameters are in line with the calibration adopted by Coibion et al. (2012), with important qualifications. First, we obtain a slightly higher degree of price rigidity than theirs (0.67 versus 0.55). Second, our specification of monetary policy is different from theirs. In particular, they allow for two lags of the nominal interest rate in the monetary policy rule while we only have one lag. However, we can compare the overall degree of interest rate smoothing in the two setups. To this end, abstracting from the other elements of the rule, we simply focus on the sum of autoregressive coefficients. It amounts to 0.92 in their calibration while the degree of smoothing in our setup has a mean posterior value of 0.85. While this might not seem to be a striking difference, it is useful to cast these figures in terms of half-life of convergence in the context of autoregressive model of order 1. Our value implies twice as small a half-life than theirs. Third, our monetary policy shock and our shocks to demand have approximately twice as small an unconditional standard deviation as theirs.

3.2 Computing the Optimal Inflation Target

Simulations with a ZLB constraint The model becomes non-linear when one allows the ZLB 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-ZLB regime $k = n$ and the ZLB 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-ZLB regime, the vector $f^{(n)}$ is filled with zeros. In the ZLB 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-ZLB regime correspond to the coefficients of the Taylor rule while in the ZLB 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 a ZLB episode, we find s_t using the linear solution stated above. In contrast, if this draw leads to a ZLB episode, we postulate integers $N_e < N$ and $N_x < N$ such that the ZLB 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 $q \in \{N_e, ..., N_x - 1\}$

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

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

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

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

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

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 ZLB at $t + N_e$ and leaves the ZLB 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.[7] A shortcoming of this approach is that the agents in the model are not assumed to expect that the ZLB has some probability to bind again in the future, after the exit from (possible) current ZLB episode. This may bias estimates , as explained by in Gust et al. (2017), even when as in our case estimation is performed on a pre-ZLB concern. The scope of this concern is however dampened by the fact that in the pre crisis environment there is evidence that even experts severely underestimated the probability of the ZLB occuring, see Chung et al. (2012).

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).[8] Let $\mathscr{W}(\pi;\theta)$ denote this welfare criterion. This notation emphasizes that welfare depends on the inflation target π together with the rest of the structural parameters θ. Two cases are considered concerning the latter. In the baseline case the structural parameters θ are fixed at reference values and taken to be known with certainty by the policy maker. In an alternative exercise, the policy maker maximizes welfare while recognizing the uncertainty associated with the model's parameters.

The optimal inflation target associated with a given vector of parameters θ, $\pi^\star(\theta)$ is approximated via numerical simulations of the model allowing for an occasionally binding ZLB constraint, using the algorithm outlined above. [9] The optimal inflation rate associated to a particular vector of parameters θ is then obtained as the one maximizing the welfare function, that is:

$$\pi^\star(\theta) \equiv \arg\max_{\pi} \mathscr{W}(\pi;\theta).$$

Accounting for Parameter Uncertainty The location of the loss function $\mathscr{W}(\pi;\theta)$ evidently depends on the parameter of the economy. As a result of estimation uncertainty around of θ, $\pi^\star(\theta)$ will be subject to uncertainty. Further, a policy maker may wish take into account uncertainty surrounding θ when defining

[7]In 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.

[8]See the Technical Appendix for details.

[9]More 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 $\mathscr{W}(\pi;\theta)$ 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.

its optimal target. A relevant feature of the welfare functions in our set-up is that in general they will be markedly asymmetric: adopting an inflation target 1 percentage point below the optimal value generates welfare losses larger than setting it 1 percentage point above. As a result, the certainty-equivalence does not hold. A policy maker maximizing expected welfare while recognizing the uncertainty, will choose a an inflation target differing from that corresponding to the case where θ is set to its expected value, and taken to be known with certainty, as in our baseline analysis.

Formally, the estimated posterior distribution of parameters $p(\theta|X_T)$ can be exploited to quantify the impact of parameter uncertainty on the optimal inflation target and to compute a "Bayesian-theoretic optimal inflation target". We define the latter as the inflation target $\pi^{\star\star}$ which maximizes the expected welfare not only over the realizations of shocks but also over the realizations of parameters[10]

$$\pi^{\star\star} \equiv \arg\max_{\pi} \int_{\theta} \mathscr{W}(\pi;\theta)p(\theta|X_T)\mathrm{d}\theta.$$

We interpret the spread between the optimal inflation target at the posterior mean $\bar{\theta}$ and the optimal Bayesian inflation target as a measure of how uncertainty about the parameter value affects optimal inflation. Given the nature of asymmetry in the welfare function, the spread will turn out to be positive: a bayesian policy maker will tend to choose a higher inflation target than a policymaker taking θ to be known and equal to its expectation. A larger inflation target indeed acts as a buffer to hedge against particularly detrimental parameter values (either because they lead to more frequent ZLB episodes or because they lead to particularly acute inflation distortions). We define

$$\mathrm{Spr}(\theta) \equiv \pi^{\star\star} - \pi^{\star}(\theta)$$

and assess below $\mathrm{Spr}(\bar{\theta})$.

3.3 Pre-crisis Optimal inflation Targets and Probability of a ZLB Episode

Figures 1a and 1b display the welfare function for the US and the euro area – expressed as losses relative to the maximum social welfare – associated with three natural benchmarks for the parameter vector θ: the posterior mean (dark blue line), the median (light blue line), and the mode (lighter blue line). For convenience, the peak of each welfare function is identified with a dot of the same color. Also, to facilitate interpretations, the inflation targets are expressed in annualized percentage rates.

[10] This Bayesian inflation target is recovered from simulating the model under a ZLB constraint using the exact same sequence of shocks $\{\epsilon_t\}_{t=1}^{T}$ with $T = 100000$ as in the previous subsection (together with the same burn-in sample) and combining it with N draws of parameters $\{\theta_j\}_{j=1}^{N}$ from the estimated posterior distribution $p(\theta|X_T)$, with $N = 500$. As in the previous section, the social welfare function $\mathscr{W}(\pi;\theta)$ is evaluated for each draw of θ over a grid inflation targets $\{\pi^{(k)}\}_{k=1}^{K}$. The Bayesian welfare criterion is then computed as the average welfare across parameter draws. Here, we start with the same inflation grid as before and then run several passes. In the first pass, we identify the inflation target maximizing the Bayesian welfare criterion. We then set a finer grid of $K = 51$ inflation targets around this value. We repeat this process several times with successively finer grids of inflation targets until the identified optimal inflation target proves insensitive to the grid. In this particular exercise, some parameter draws for θ lead to convergence failure in the algorithm implementing the ZLB. These draws are discarded.

Figure 1: Examples of loss functions

(a) US

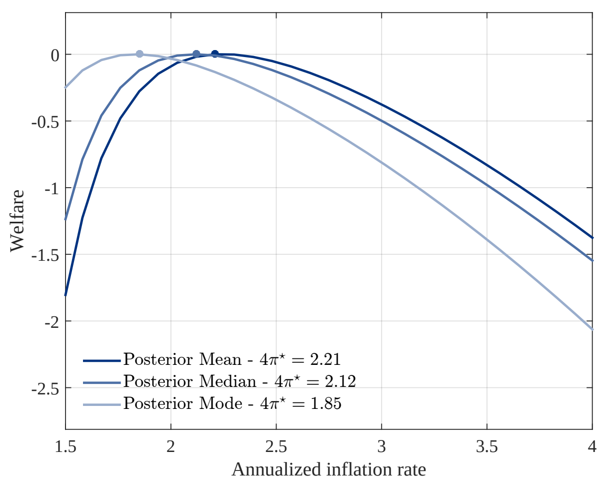

(b) EA

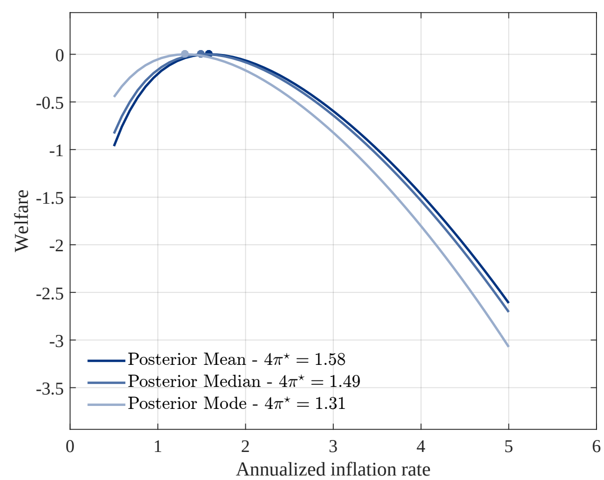

Note: Blue: parameters set at the posterior mean; light blue: parameters set at the posterior median; Lighter blue: parameters set at the posterior mode. $\pi^\star \equiv \log(\Pi^\star)$. In all cases, the welfare functions are normalized so as to peak at 0.

As Figure 1a illustrates, the US optimal inflation target is close to 2% and varies between 1.85% and 2.21% depending on which indicator of central tendency (mean, mode, or median) is selected. This range of values is consistent with the ones of Coibion et al. (2012) even though in the present paper it is derived from an estimated model over a much shorter sample.[11] Importantly, while the larger shocks in Coibion et al. (2012) ceteris paribus induce larger inflation targets, the high degree of interest rate smoothing in their analysis works in the other direction (as documented below in the last section) . In the euro area, as figure 1b reports, the optimal inflation target is close to 1.5% and varies between 1.31% and 1.58% across indicators of central tendency. Altogether, these numbers seem roughly consistent with the quantitative inflation targets adopted by the Fed and the ECB, respectively.

To complement on these illustrative results, figures 2a and 2b display the probability of reaching the ZLB as a function of the annualized inflation target (again, with the parameter vector θ evaluated at the posterior mean, median, and mode). For convenience, the circles in each curve mark the corresponding optimal inflation target.

According to the simulation exercise, $\pi^{\star\star} = 2.40\%$ for the US and $\pi^{\star\star} = 2.20\%$ for the euro area. In both cases, these robust optimal inflation targets are larger than the values obtained with θ sets at its central tendency. As expected, a Bayesian policy maker chooses a higher inflation target to hedge against particularly harmful states of the world (i.e., parameter draws) where the frequency of hitting the ZLB is high. Noticeably, this inflation buffer is substantially higher in the euro area where $\text{Spr}(\bar{\theta}) = 0.62\%$, while

[11] Coibion et al. (2012) calibrate their model on a post-WWII, pre-Great Recession US sample. By contrast, we use a Great Moderation sample.

Figure 2: Probability of ZLB

(a) US (b) EA

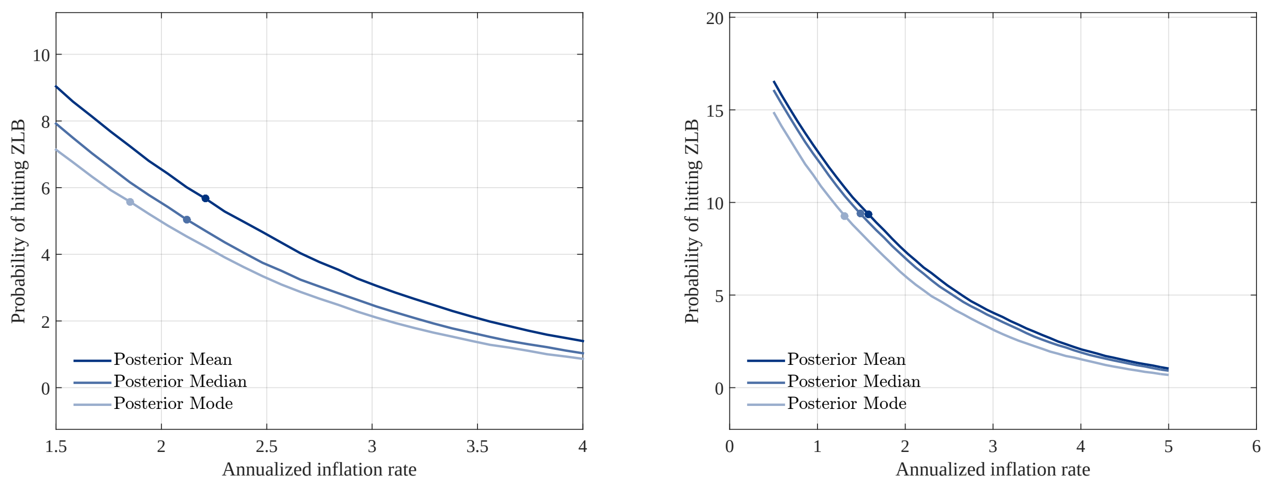

Note: Blue: parameters set at the posterior mean; light blue: parameters set at the posterior median; Lighter blue: parameters set at the posterior mode. $\pi^\star \equiv \log(\Pi^\star)$.

in the US, $\text{Spr}(\bar{\theta}) = 0.19\%$. This occurs in spite of the optimal inflation target being lower in the euro area than in the US when evaluated at the posterior mean of the parameter vector $\bar{\theta}$.[12]

The probability of hitting the ZLB associated to these positive optimal inflation targets is relatively low. The probabilities of reaching the ZLB is about 6% for the US and about 10% for the euro area. This result, as anticipated above, is the mere reflection of our choice of a Great Moderation sample. At the same time, our model is able to predict a fairly spread out distribution of ZLB episodes durations, with a significant fraction of ZLB episodes lasting more than say five years (see figures in the appendix). Given the existence of a single ZLB episode in the recent history, we do not attempt here to take a stand on what is a relevant distribution of ZLB episodes (see Dordal-i-Carreras et al. (2016) for some investigation in that direction).

Following a now standard approach when assessing the optimal policies, we now complement our characterization of optimal inflation by providing measures of consumption-equivalent welfare gains loss of choosing a non-the optimal inflation target (see Appendix *L* for a derivation). Applying such an approach here suggests that the welfare costs of moving from the above mentioned optimal inflation targets to a 4% inflation target are of the order of 0.30 percent of annual consumption for the euro area and 0.25 percent for the US (Results are reported in the Appendix Figures E.1a E.1b). By contrast the same respective welfare losses are obtained by reducing inflation target to 0.5% in the EA, and to around 1.6% in the US, illustrating the asymmetry of the loss function stemming from the ZLB. These computations overall suggest that the welfare costs associated to changes in the inflation target of magnitudes consistent with

[12]Figures D.1a and D.1b in Appendix show the posterior distribution of $\pi^\star(\theta)$. It is broadly symmetric in the US and shows substantial asymmetry in the euro area.

current policy debates are small, but still economically meaningful. The relatively low cost of a positive inflation target contrasts with the recent study of Ascari et al. (2015) who report a consumption-equivalent welfare loss of about 4% from moving trend inflation from 2 to 4% in a New Keynesian DSGE model with price and wage stickiness as well as capital but no ZLB. We obtain welfare costs that are of an order of magnitude smaller. The main determinant of this smaller cost of a positive inflation target is the fact that in our model, trend inflation brings some benefits as it reduces the probability of hitting the ZLB.

4 The Optimal Inflation Target and the Steady State Real Interest Rate

The focus of this section is to investigate how the monetary authority should adjust its optimal inflation target $\pi^\star$ in response to changes in the steady-state real interest rate, $r^\star$.[13] Intuitively, with a lower $r^\star$ the ZLB is bound to bind more often, so one would expect a higher inflation target should 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 ZLB, 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) .

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 and ignoring parameter uncertainty. We show that the link between $\pi^\star$ and $r^\star$ depends to some extent on the reason underlying a variation in $r^\star$, i.e. a change in the discount rate ρ or a change in growth rate of technology μ_z. In our set-up the first scenario roughly captures the "taste for safe asset" and "ageing population" rationale for secular stagnation, while the second one captures the "decline in technological progress" rationale . We also investigate the role of parameter uncertainty and, in particular, uncertainty about $r^\star$, in the determination of the Bayesian-theoretic optimal inflation target. Finally, we investigate how the relation between the optimal inflation target and the steady-state real interest rate depends on various features of the monetary policy framework, as well as on the size of shock or on the price and wage mark-ups.

4.1 The baseline $(r^\star, \pi^\star)$ relation

To characterize the link between $r^\star$ and $\pi^\star$, 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.

[13] Note our exercise here 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 Holston et al. (2017). Our exercise is arguably consistent with "secular stagnation" understood as a permanently lower real rate of interest – while doing without having to assume a unit root process in the real rate of interest.

For both μ_z and ρ, we consider values on a grid ranging from 0.4% to 10% 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 procedure as before.[14] The optimal value $\pi^\star$ associated to each value of $r^\star$ is obtained as the one maximizing the welfare criterion $\mathscr{W}(\pi;\theta)$.[15]

We finally obtain two curves. The first one links the optimal inflation target $\pi^\star$ to the steady-state real interest rate $r^\star$ for various growth rate of technology μ_z: $\pi^\star(r^\star(\mu_z))$, where the notation $r^\star(\mu_z)$ highlights that the steady-state real interest rate varies as μ_z varies. The second one links the optimal inflation target $\pi^\star$ to the steady-state real interest rate $r^\star$ for various discount rates ρ: $\pi^\star(r^\star(\rho))$. Here, the notation $r^\star(\rho)$ highlights that the steady-state real interest rate varies as ρ varies.[16]

Figure 3: $(r^\star,\pi^\star)$ locus (at the posterior mean)

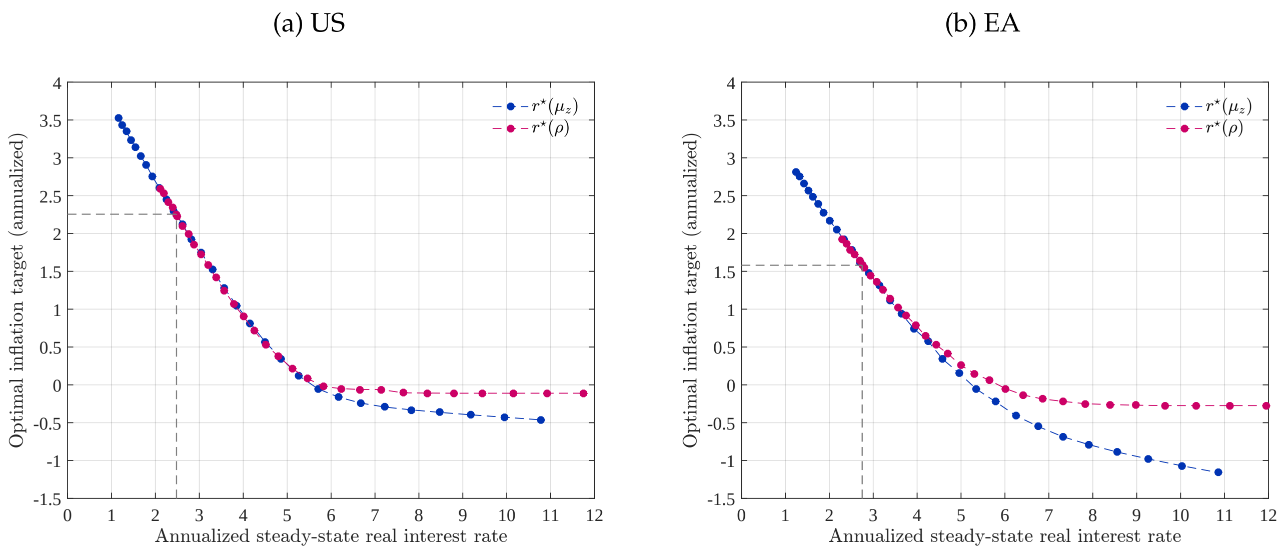

Note: the blue dots correspond to the $(r^\star,\pi^\star)$ locus when $r^\star$ varies with μ_z; the red dots correspond to the $(r^\star,\pi^\star)$ locus when $r^\star$ varies with ρ

Figures 3a and 3b depict the $(r^\star,\pi^\star)$ relations thus obtained for the US and the euro area, respectively. The blue dots correspond to the case when the real steady-state interest rate $r^\star$ varies with μ_z. The red dots correspond to the case when the real steady-state interest rate $r^\star$ varies with ρ. For convenience, both the real interest rate and the associated optimal inflation target are expressed in annualized percentage rates. The dashed grey lines indicate the benchmark result corresponding to the optimal inflation target at the

[14]In 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.

[15]To illustrate the construction of this figure, see Appendix E. There, we show how two particular points of this curve are derived from the welfare criteria.

[16]Figures F.1a, F.1b, F.2a, and F.2b report similar results at the posterior mode and at the posterior median.

Figure 4: Relation between probability of ZLB at optimal inflation and $r^\star$ (at the posterior mean)

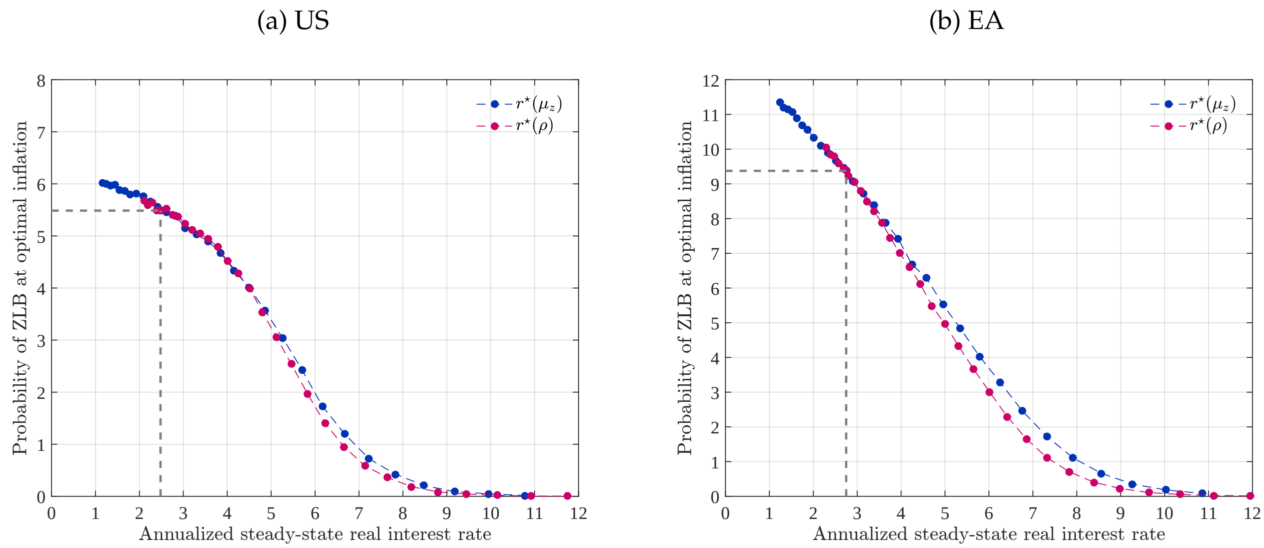

Note: the blue dots correspond to the $(r^\star, \pi^\star)$ locus when $r^\star$ varies with μ_z; the red dots correspond to the $(r^\star, \pi^\star)$ locus when $r^\star$ varies with ρ

posterior mean of the structural parameter distribution. These results are complemented with Figures 4a and 4b that show the relation between $r^\star$ and the probability of hitting the ZLB, *evaluated at the optimal inflation target*, for the US and the euro area, respectively. As with Figures 3a, blue dots correspond to the case when $r^\star$ varies with μ_z, while red dots correspond to the case when it varies with ρ.[17]

As expected the relations in Figures 3a and 3b are decreasing. However, the slope varies with the value of $r^\star$. For both the US and the euro area, the slope is relatively large in absolute value – although smaller than one – for moderate values of $r^\star$ (say below 4 percent). The slope declines in absolute value as $r^\star$ increases: Lowering the inflation target to compensate for an increase in $r^\star$ becomes less and less desirable. This reflects the fact that, as $r^\star$ increases, the probability of hitting the ZLB becomes smaller and smaller. For very large $r^\star$ values, the probability becomes almost zero, as Figures 4a and 4b show.

At some point, the optimal inflation target becomes insensitive to changes in $r^\star$ 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 the result of a zero price inflation and a consequently larger nominal wage inflation. [18]

[17]Figures G.1a and G.1b in Appendix show the relation between $r^\star$ and the nominal interest rate when the inflation target is set at its optimal value.

[18]For very large $r^\star$, as a rough approximation, we can ignore the effects of shocks and assume that the ZLB is a zero-mass

The previous tension is even more apparent when $r^\star$ 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 changes in $r^\star$ for very large values of $r^\star$, typically above 6%, both in the US and the euro area.

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, it shows that, starting from the posterior mean estimate of θ, a 100 basis points decline in $r^\star$ should lead to a +99 basis points increase in $\pi^\star$ in the US and to a +81 basis points increase in the euro area. Importantly, this increase in the optimal inflation target is the same no matter the underlying factor causing the change in $r^\star$: a drop in potential growth, μ_z, or a decrease in the discount factor, ρ. At the same time, the probability of ZLB evaluated at the optimal inflation rate also increases when the real rate decreases. In the US case, at some point, the speed at which this probability increases slows down, reflecting that the social planner would choose to increase the inflation target to almost compensate for the higher incidence of ZLB episodes. By contrast, in the euro area, the incidence of ZLB seems to increase substantially after a decline in the real interest rate, even at low values of the latter.

To gain insight into this striking difference, Figures 5a and 5b show how the probability of ZLB changes as a function of $r^\star$, holding the inflation target constant. We first set the inflation target at its optimal baseline value (i.e., the value computed at the posterior mean, 2.21 for the US and 1.58 for the euro area). This is reported below as the blue dots. Similarly, we also compute an analog relation assuming this time that the inflation target is held constant at the optimal value consistent with a steady-state real interest rate one percentage point lower (thus, inflation is set to 3.20 for the US and 2.39 for the euro area). Here again, the other parameters are set at their posterior mean. This corresponds to the red dots in the figure.

Consider first the blue line. At the level of the real interest rate prevailing before the permanent decline, assuming that the Central Bank sets its target to the associated optimal level, the probability of reaching the ZLB would be slightly below 6% in the US and close to 9% in the euro area. Imagine now that the real interest rates experiences a decline of 100 basis points. Keeping the inflation target at the same level as prior to the shock, the probability of reaching the ZLB would now climb up to approximately 11% in the US and 16% in the euro area. However, the change in the optimal inflation target brings the probability of reaching the ZLB back to approximately 6% in the US and 11% in the euro area. In the euro area, the social planner is willing to tolerate a smaller inflation target than the one that would fully neutralize the effects of

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 $\pi^\star \approx -\lambda_w(1-\gamma_z)(1-\gamma_w)/[\lambda_p(1-\gamma_p)^2 + \lambda_w(1-\gamma_w)^2]\mu_z$. Given the low values of λ_w resulting from our estimation, it is not surprising that $\pi^\star$ is negative but close to zero. See Amano et al. (2009) for a similar point in the context of a model abstracting from ZLB issues.

Figure 5: Relation between probability of ZLB and $r^{\star}$ (at the posterior mean)

(a) US (b) EA

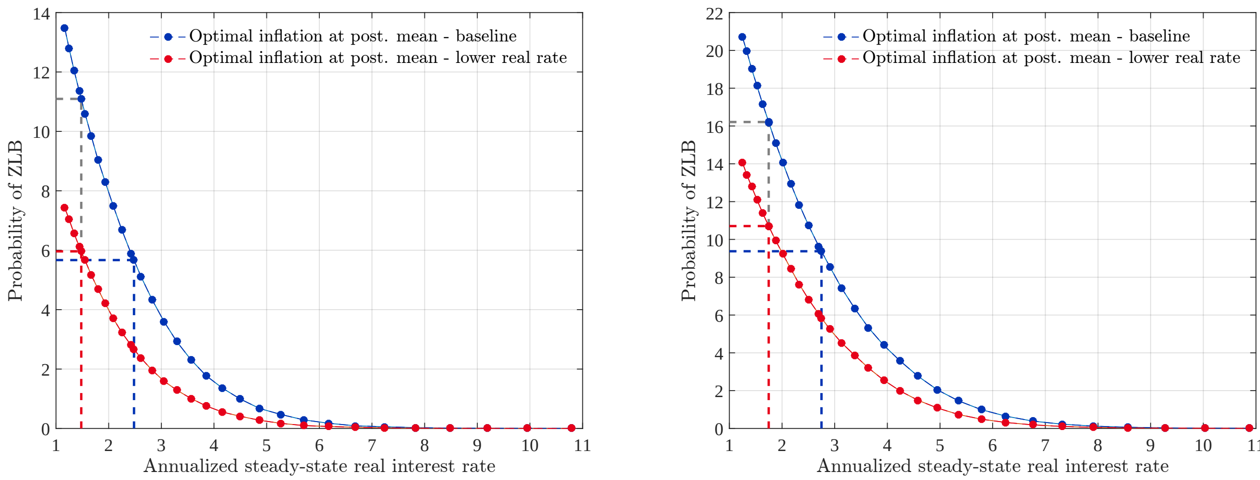

Note: The blue dots correspond to the relation linking $r^{\star}$ and the probability of ZLB, holding the optimal inflation target $\pi^{\star}$ at the baseline value. The red dots correspond the same relation when the optimal inflation target $\pi^{\star}$ is set at the value consistent with a steady-state real interest rate one percentage point lower.

the natural rate decline on the probability of hitting the ZLB. By way of contrast, the social planner in the US would almost neutralize this effect. In this sense, the US economy has a greater tolerance for steady-state inflation than the euro area. This is in part a consequence of the different estimates for the degree of indexation of prices to past inflation found at the estimation stage.

4.2 Accounting for Parameter Uncertainty

Next we investigate the impact of parameter uncertainty on the relation between the optimal inflation target and the steady-state real interest rate. Specifically, we want to determine how the Bayesian-theoretic optimal inflation target $\pi^{\star\star}$ reacts to a downward shift in the *distribution* of the steady-state real interest rate $r^{\star}$.

Assessing how such a change affects $\pi^{\star\star}$ for every value of $r^{\star}$ is not possible due to the computational cost involved. Such a reaction is thus investigated for a particular scenario: it is assumed that the economy starts from the posterior distribution of parameters $p(\theta|X_T)$ and that, everything else being constant, the mean of $r^{\star}$ decreases by 100 basis points. Such a 1 percentage point decline is chosen mainly for illustrative purposes. Yet, it is of a comparable order of magnitude, although relatively smaller in absolute value, as recent estimates of the drop of the natural rate after the crisis such as Laubach and Williams (2016) and Holston et al. (2017). The counterfactual exercise considered can therefore be seen as a relatively conservative characterization of the shift in steady-state real interest rate. Figures 6a and 6b depict the counterfactual shift in the distribution of $r^{\star}$ that is considered for, respectively, the US and the euro area.

Figure 6: Posterior distributions of $r^\star$ and counterfactual $r^\star$

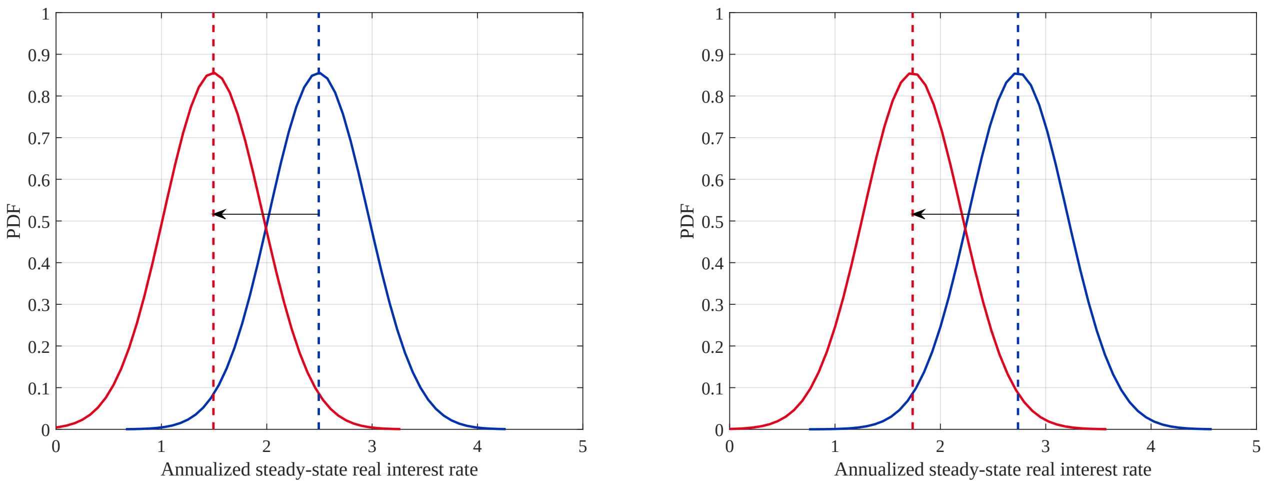

Note:: Plain curve: PDF of $r^\star$; dashed vertical line: mean value of $r^\star$.

The Bayesian-theoretic optimal inflation target corresponding to the counterfactual lower distribution of $r^\star$ is obtained from a simulation exercise that relies on the same procedure as before.[19] Given a draw in the posterior of parameter vector θ, the value of the steady-state real interest rate is computed using the expression implied by the postulated structural model $r^\star(\theta) = \rho(\theta) + \mu_z(\theta)$. From this particular draw, a counterfactual lower steady-state real interest rate, $r^\star(\theta_\Delta)$, is obtained by shifting the long-run growth component of the model μ_z downwards by 1 percentage point (in annualised terms). The welfare function $\mathscr{W}(\pi; \theta_\Delta)$ is then evaluated. Since there are no other changes than this shift in the mean value of μ_z in the distribution of the structural parameters, we can characterize the counterfactual distribution $p(\theta_\Delta|X_T)$ as a simple transformation of the estimated posterior $p(\theta|X_T)$. The counterfactual Bayesian-theoretic optimal inflation target is then obtained as

$$\pi_\Delta^{\star\star} \equiv \arg\max_{\pi} \int_{\theta_\Delta} \mathscr{W}(\pi; \theta_\Delta) p(\theta_\Delta|X_T) \mathrm{d}\theta_\Delta.$$

Figures 7a and 7b illustrate the counterfactual change in optimal inflation target obtained when the steady-state real interest rate declines by 100 basis points and its new value stays uncertain. For the US, the simulation exercise returns a value of $\pi_\Delta^{\star\star} = 3.30\%$ i.e. 90 basis points higher than the optimal value under uncertainty obtained with the posterior distribution of parameters obtained on a pre-crisis sample $\pi_\Delta^{\star\star} = 2.40\%$. For the euro area, $\pi_\Delta^{\star\star} = 3.10\%$, also 90 basis points higher than the optimal value $\pi_\Delta^{\star\star} = 2.20\%$ obtained with the baseline posterior distribution of parameters.[20]

[19] Again, we use the same sequence of shocks and the same parameter draws as in section 3.2.

[20] Figures H.1a and H.1b in Appendix show how the posterior distribution of $\pi^\star$ is shifted after the permanent decline in the mean of $r^\star$.

Figure 7: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$ in baseline and counterfactual

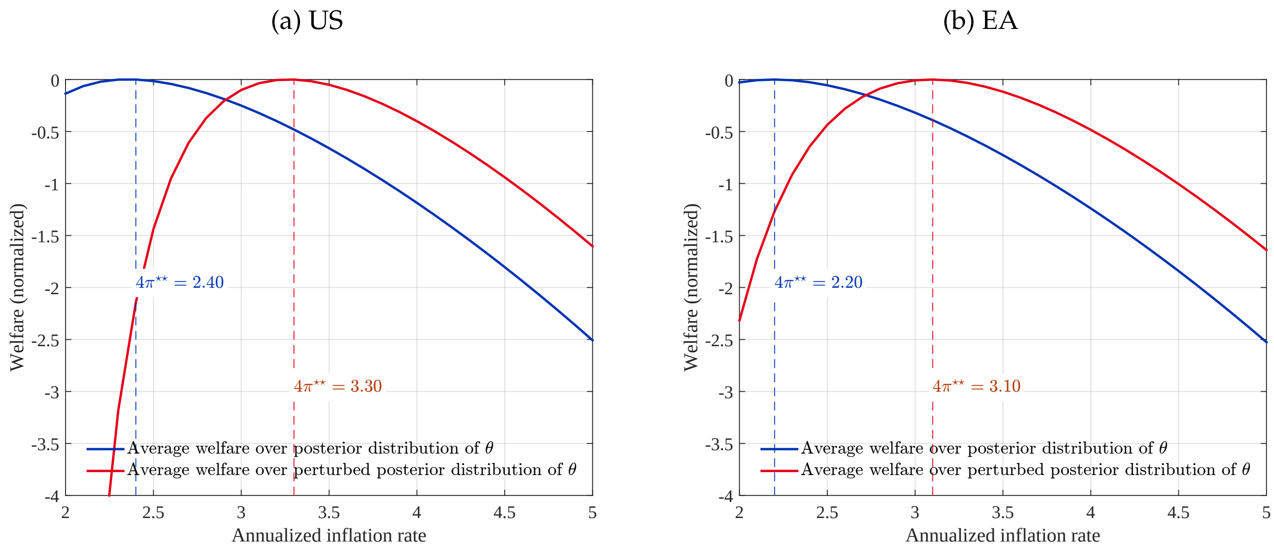

Note:: Blue curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$; Red curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$ with lower $r^\star$

Thus, we see that a monetary authority that is concerned about the uncertainty surrounding the parameters driving the costs and benefits of the inflation target raises the optimal inflation target but does not alter the reaction of this optimal inflation target following a drop in $r^\star$: in both cases, a 100 basis points decrease in the steady-state real interest rate calls for a roughly 90 basis point increase in the optimal inflation target in the vicinity of pre-crisis parameter estimates.

4.3 Further Experiments

In the present section we carry out a number of additional exercises related to the optimal adjustment of the inflation target in response to a change in the steady-state real interest rate. The first four exercises examine the implications of four alternative assumptions regarding monetary policy. The fifth exercise looks at the case of large shocks, while the sixth consider alternative calibrations for the price and wage mark-up.

Average vs Target Inflation As emphasized in recent works (see, notably, Hills et al. 2016, Kiley and Roberts 2017), when the probability of hitting the ZLB is non-negligible, realized inflation is on average significantly lower than the inflation rate that the central bank targets in the interest rate rule (and which would correspond to steady-state inflation in the absence of shocks). This results from the fact that anytime the ZLB is binding (which happens recurrently) the central bank effectively loses its ability to stabilize inflation around the target. Knowing this, it may be relevant to assess the central banks outcomes in terms of the effective average realized inflation. In this section, we investigate whether measuring inflation target in this alternative way matter.

Figure 8: Average realized inflation and optimal inflation

(a) US

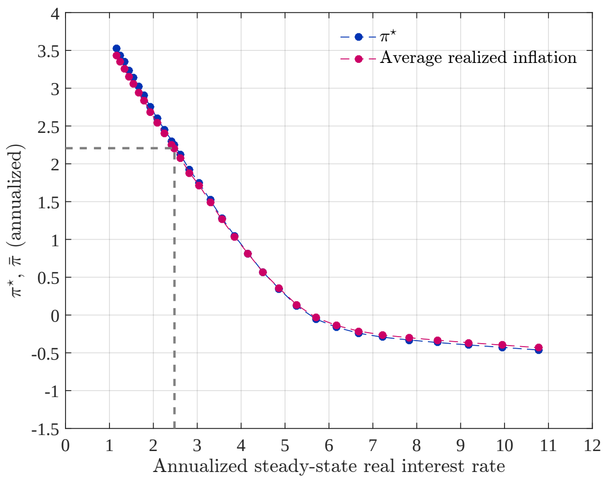

(b) EA

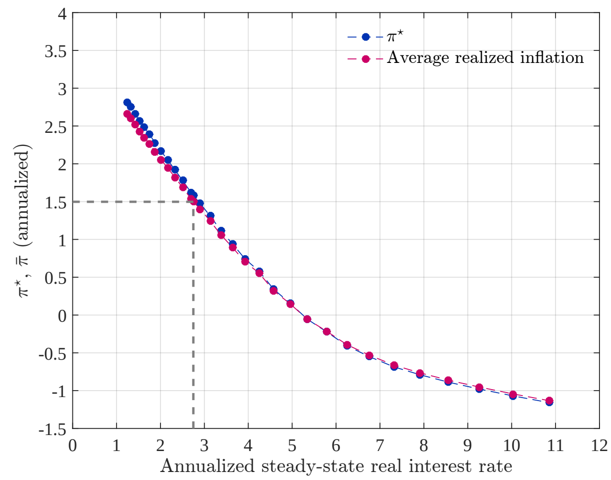

To this end, the analysis of the $(r^\star, \pi^\star)$ relation of section 4.1 is complemented here with the analysis of the relation between $r^\star$ and the average realized inflation rate $\mathbb{E}\{\pi_t\}$ obtained when simulating the model for various values of $r^\star$ and the associated optimal inflation target $\pi^\star$. In the interest of brevity, the calculations are undertaken assuming that changes in average productivity growth μ_z is the only source of variation in the natural interest rate.

Figures 8a and 8b illustrate the difference between the $(r^\star, \pi^\star)$ curve (blue dots) and the $(r^\star, \mathbb{E}\{\pi_t\})$ curve (red dots) for the US and the euro area. The overall shape of the curve is unchanged. Unsurprisingly, both curves are identical when $r^\star$ is high enough. In this case, the ZLB is (almost) not binding and average realized inflation does not differ much from $\pi^\star$. A spread between the two emerges for very low values of $r^\star$. There, for low values of the natural rate, the ZLB incidence is higher and, as a result, average realized inflation becomes indeed lower than the optimal inflation target. However, that spread remains limited, less than 10 basis points. The reason is that the implied optimal inflation target is sufficiently high to prevent the ZLB from binding too frequently, thus limiting the extent to which average realized inflation and $\pi^\star$ can differ.

Unreported simulation results show that the gap between $\pi^\star$ and average realized inflation becomes more substantial when the inflation target is below its optimal value. For instance, mean inflation is roughly zero when the central bank adopts a 1% inflation target in an economy where the optimal inflation target is $\pi^\star = 2\%$.

A Negative Effective Lower Bound The recent experience of many advanced economies (including the euro area) points to an effective lower bound (ELB) for the nominal interest rate below zero. For instance,

Figure 9: Optimal inflation with negative ELB – EA

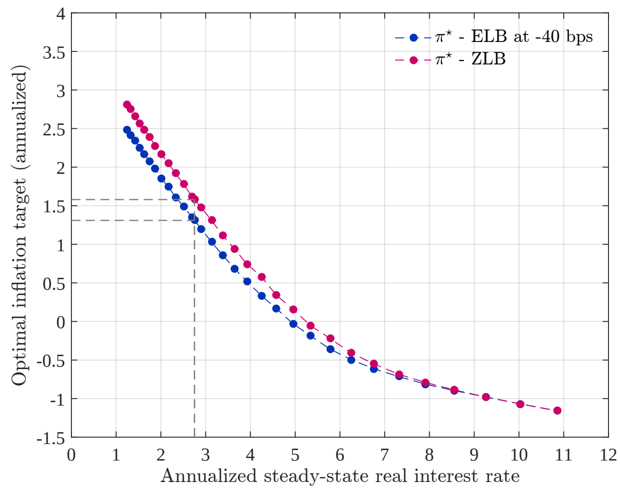

the ECB's deposit facility rate, which gears the overnight money market rate because of excess liquidity, was set at a negative value of -10 basis points in June 2014 and has been further lowered down to -40 basis points in March 2016.

We use the estimated euro-area model to evaluate the implications of a negative ELB. More precisely, we set the lower bound on the nominal rate i_t so that

$$i_t \geq e$$

and we set e to -40 basis points (in annual terms) instead of zero. Results are presented in Figure 9. As expected, the $(r^\star, \pi^\star)$ locus is shifted downwards, though by somewhat less than 40 basis points. Importantly, its slope remains identical to the baseline case: a 100 basis points downward shift in the distribution of $r^\star$ calls for a 90 basis points increase in $\pi^\star$.

A Known Reaction Function Here we study the consequences of the (plausible) assumption that the central bank actually knows the coefficients of its interest rate rule with certainty. More specifically we repeat the same simulation exercise as in subsection 4.2 but with parameters a_π, a_y and ρ_i in the reaction function 3 taken to be known with certainty. In practice we fix these three parameters at their posterior mean, instead of sampling them from their posterior distribution. This is arguably the relevant approach from the point of view of the policymaker.[21] Note, however, that all the other parameters are subject to uncertainty from the stand-point of the central bank.

[21] In practice, long-run inflation targets are seldom reconsidered while the rotation in monetary policy committees happens at a higher frequency. From this viewpoint, our baseline assumption of uncertainty on all the monetary policy rule parameters is not necessarily unwarranted.

Figure 10: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$

(a) US

(b) EA

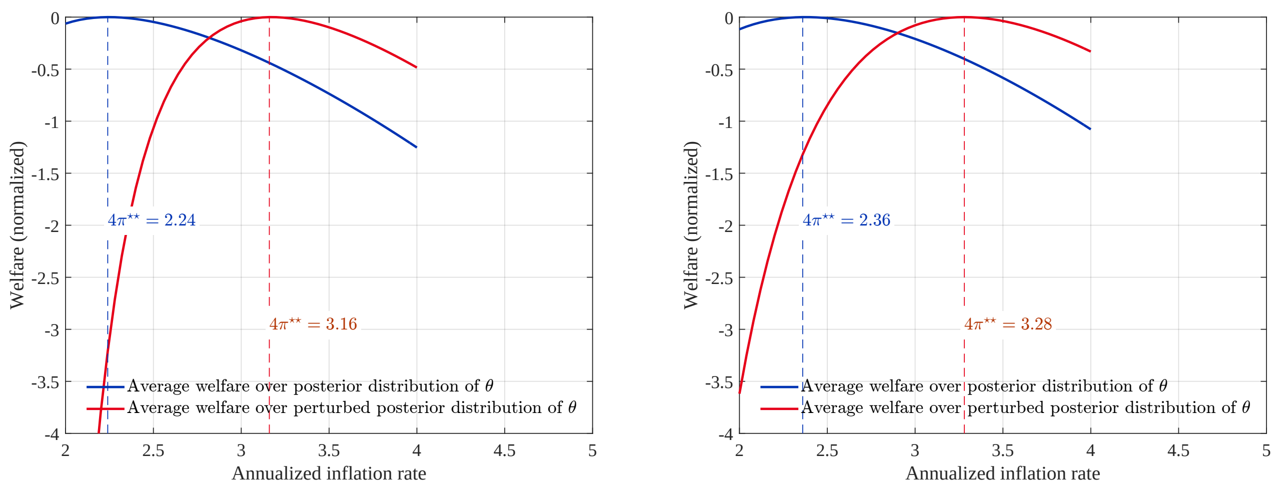

Note:: Blue curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$; Red curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$ with lower $r^\star$. In each case, ρ_{TR}, a_π, and a_y are frozen at their posterior mean values.

Figures 10a and 10b present, respectively for the US and the euro area, the Bayesian-theoretic optimal inflation targets obtained when simulating the model at the initial posteriors and after a -100 basis points level shift in the posterior distribution of the long-run growth rate μ_z and, hence, the steady-state real rate $r^\star$. According to these simulations, the inflation target should initially be $\pi^{\star\star} = 2.24\%$ in the US and $\pi^{\star\star} = 2.36\%$ in the euro area. After the counterfactual change in the distribution of $r^\star$ considered, $\pi^{\star\star}$ should be increased to 3.16% in the US and to 3.28% in the euro area, again in the ballpark of a 90 basis points increase in $\pi^\star$ to compensate for the higher probability to hit the ZLB induced by a 100 basis points downward shift in the distribution of $r^\star$.

What if interest rate smoothing is larger (or lower)? 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_y and ρ_{TR}. Among these parameters, the smoothing parameter, ρ_{TR}, has a key influence on the probability of being in a ZLB regime. A higher smoothing has 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. Such a monetary policy strategy introduces history-dependence whereby, in the instance of a ZLB episode, the central bank is committed to keep rates lower for longer. As this reaction function is known to the agents in the model, this commitment to future accommodation, through generating higher expected inflation and output, helps exiting the trap (or even not entering it).

Through both effects, a large higher degree of smoothing thus reinforces the history-dependence of monetary policy, and tends to shorten the length of ZLB episodes and the probability of hitting the ZLB constraint. Everything else 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 Figures 11a and 11b which depict the $(r^\star, \pi^\star)$ relation for respectively the US and the euro-area under three possible values of the smoothing parameter ρ_{TR}. The values used under our baseline scenario, i.e. posterior mean estimates, are .85 for the US and .87 for the euro area. We also consider two alternative setups: A higher value of $\rho_{TR} = .95$ which is close to the persistence of the central bank reaction function in Coibion et al. (2012), and a lower value of $\rho_{TR} = .8$. These two values are arguably an upper bound of the existing empirical uncertainty on the degree of smoothing, as they stand outside the 90 % probability interval of our posterior parameter estimates.

The effect of a higher interest rate smoothing is to shift downward the $(r^\star, \pi^\star)$ curve except for high values of $r^\star$ for which the probability of hitting the ZLB 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 0.5% in the US and in the euro area.[22] Conversely, a lower interest rate smoothing shifts the $(r^\star, \pi^\star)$ curve upward, even for relatively high values of $r^\star$ – because the probability of being in a ZLB regime increases under this strategy. With a lower ρ_{TR}, the pre-crisis optimal inflation rate would be close to 3.5% in the US and 4% in the euro area.

As for the slope of the $(r^\star, \pi^\star)$ 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 ZLB, somewhat alleviates the extend to which an increase in the inflation target is needed. The slope is indeed close to $-.7$ for the US and $-.6$ for the euro area in that case. For a strategy associated with a low smoothing parameter, the slope is close to -1 for both the EA and the US, so closer to the benchmark case but in induced a somewhat curve steeper. For large values of $r^\star$, the degree of smoothing is irrelevant.

What if shocks are larger? As argued before, the model is estimated using data from the Great Moderation period. One may legitimately argue that the decline in the real interest rate resulting from the secular

[22]This is not inconsistent with the result in Coibion et al. (2012) who report an optimal inflation target of 1.5% under their baseline calibration on US post WWII data. Indeed, the variance of their underlying shocks is higher than in our baseline which is based on a Great-Moderation estimates. As discussed above, a higher variance of shocks induces more frequent ZLB episodes, hence calls for a higher optimal inflation target.

Figure 11: $(r^{\star}, \pi^{\star})$ relation with alternative ρ_{TR}

(a) US

(b) EA

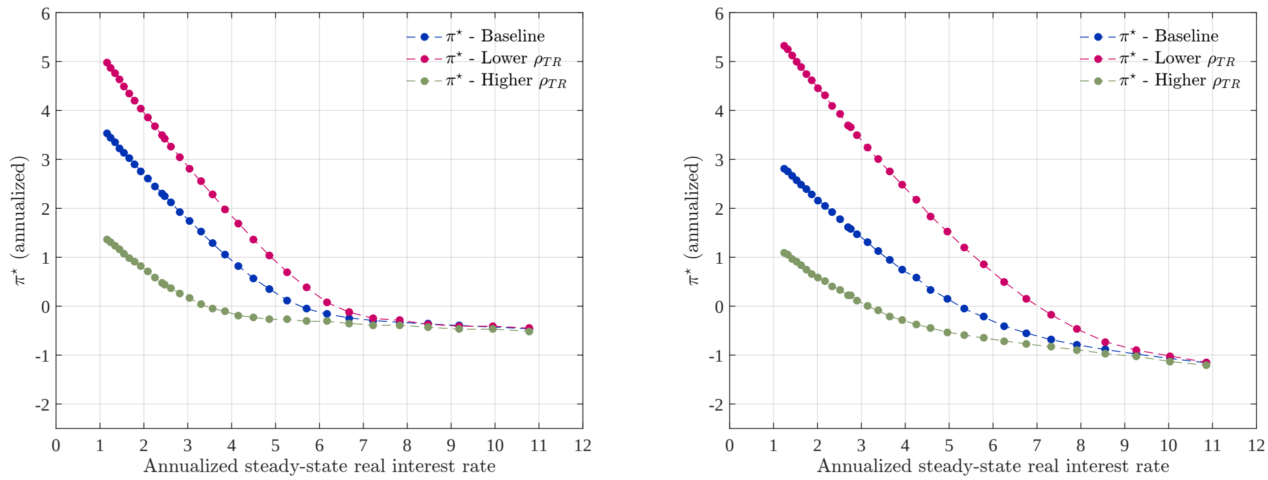

Note: the blues dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with ρ_{TR} set to 0.8. The green dots correspond to the counterfactual simulation with ρ_{TR} set to 0.95.

stagnation has come hand in hand with larger shocks, as the Great Recession suggests. To address this concern, we simulate the model assuming that demand shocks have a standard error 30 percent larger than estimated.

We conduct this exercise assuming that changes in average productivity growth μ_z are the only driver of changes in the natural rate. Apart from σ_q and σ_g, which are re-scaled, all the other parameters are frozen at their posterior mean. Given this setup, the optimal inflation target is 3.7% in the US and 2.7% in the euro area, as opposed to 2.21% and 1.58% in the baseline, respectively. Also, under the alternative shock configuration, the probability of hitting the ZLB is 5.3% in the US and 10.1% in the euro, as opposed to 5.5% and 9.4% in the baseline, respectively.

Figures 12a and 12b report the $(r^{\star}, \pi^{\star})$ relation under larger demand shocks (red dots) and compares the outcome with what obtained in the baseline (blues dots).[23] Interestingly, the $(r^{\star}, \pi^{\star})$ locus has essentially the same slope in the low $r^{\star}$ region. Here again, we find a slope close to -0.9. However, the curve is somewhat steeper in the high $r^{\star}$ region and shifted up, compared to the baseline scenario. This reflects that under larger demand shocks, even at very high levels of the natural rate, a drop in the latter is conducive to more frequent ZLB episodes. The social planner is then willing to increase the inflation target at a higher pace than in the baseline scenario and generically sets the inflation target at higher levels to hedge the economy against ZLB episodes.

[23] We obtain this figure using the same procedure as outlined before. Here again, we run several passes with successively refined inflation grids.

Figure 12: $(r^\star, \pi^\star)$ relation with larger demand shocks

(a) US (b) EA

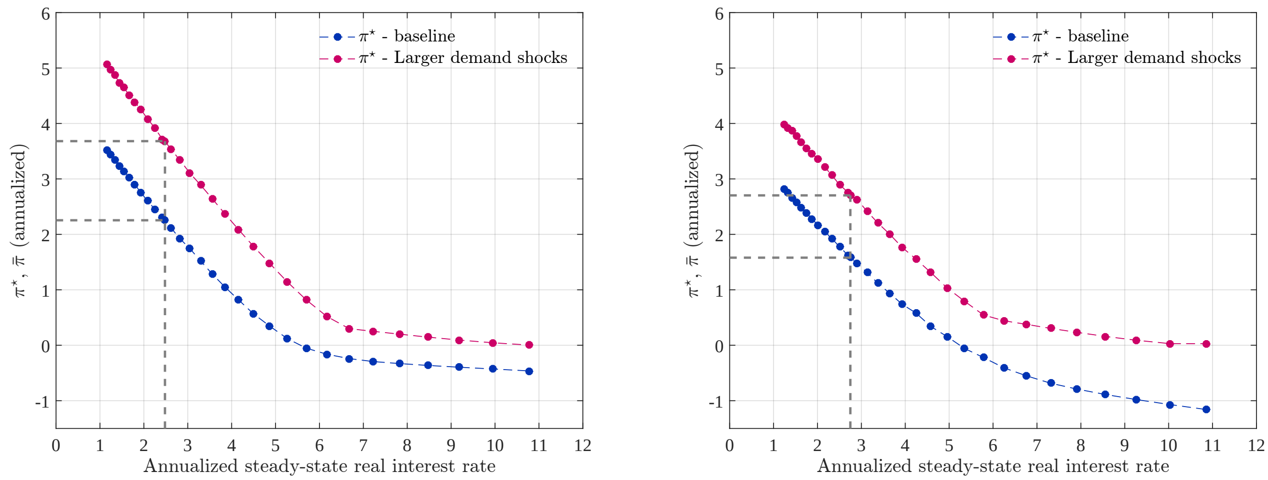

Note: the blues dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with σ_q and σ_g set to twice their baseline value.

What if mark-ups are lower (or larger)? The optimal level of inflation in our set-up depends on the elasticities of substitution between any two intermediate goods, θ_p and between any two labor types θ_w.

In our calibration, the baseline value for the elasticity of substitution θ_p is 6, leading to a price mark-up of 20%. While this value is in line with common "textbook" parameterizations (see Galí, 2015), and is close to the baseline value obtained in Hall (2018) and in Christiano et al. (2005). However, there is considerable uncertainty in the empirical literature about the level of mark-ups. For example, some estimates in Basu and Fernald (1997) and Traina (2015) point to possibly much smaller values, while Autor et al. (2017) and De Loecker and Eeckhout (2017) suggest substantially larger figures. To investigate the robustness of our results, we re-do our main simulation exercise, this time setting θ_p to a value as large as 10 or as low as θ_p set to 3. These values approximately cover the range of available empirical estimates.

Similarly, for wage mark-up, there is arguably scarcer evidence, and in any case considerable uncertainty around our baseline parameterization, given by θ_w set to 3. Here again, so as to cover a broad range of plausible estimates, we run alternatives exercises, setting in turn θ_w to 8 and θ_w to 1.5. Results are reported in Figures 13a and 13b in the case of robustness with respect to the price markup, and in Figures 14a and 14b with respect to the wage markup.

The main takeaway from these figures is that our key result is by and large preserved. That is, in the empirically relevant region (for levels of $r^\star$ lower than, say, 4 percent), the slope of the $\pi^\star - r^\star$ curve is only very mildly affected when changing the elasticity of substitution of goods or labor types, be it in the US or in the euro area. A close examination reveals, in the low $r^\star$ region, the way the $\pi^\star - r^\star$ curves moves when

Figure 13: $(r^{\star}, \pi^{\star})$ relation with alternative θ_p

(a) US (b) EA

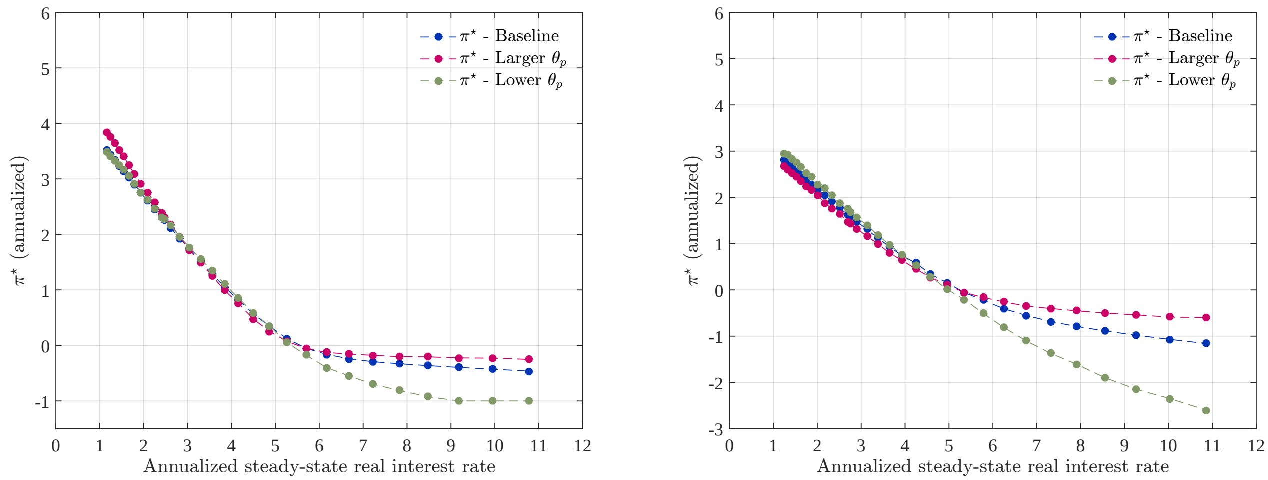

Note: the blues dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with θ_p set to 10. The green dots correspond to the counterfactual simulation with θ_p set to 3.

Figure 14: $(r^{\star}, \pi^{\star})$ relation with alternative θ_w

(a) US (b) EA

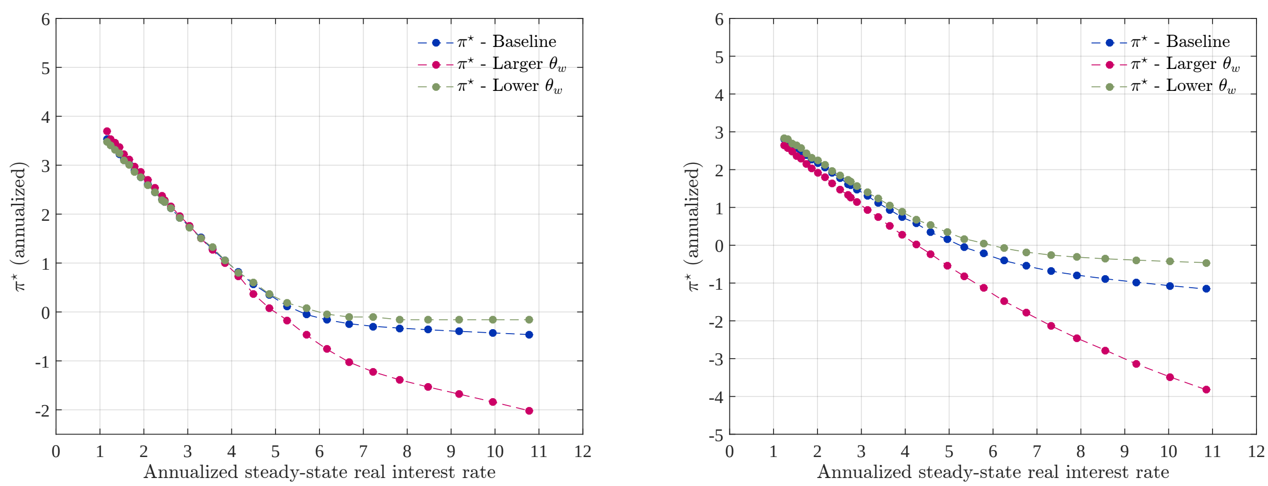

Note: the blues dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with θ_w set to 8. The green dots correspond to the counterfactual simulation with θ_w set to 1.5.

changing either θ_w or θ_p is not similar for the US and the euro area. This result turns out to depend on the degree of indexation (see Appendix for further illustration), but does not bear quantitative consequences.

Another noticeable result of this robustness exercise is that, by contrast, in the region with high steady-state real interest rates (say $r^{\star}$ larger than 5 percent) the value of optimal inflation target, and the slope of the curve of interest, are sensitive to the value of θ_p or θ_w. To see why, first notice that, in this region, the ZLB is essentially irrelevant so the standard welfare cost of inflation set-up applies. With less substitution across goods, a given level of price dispersion induced by inflation leads to smaller output dispersion (as is clear for instance in the polar case of complementary goods that leads to no output dispersion across firms at all). The effect of θ_p on output dispersion is apparent from the formulas in our appendix, or in textbook derivations of output dispersion e.g. chapter 3 in Galí (2015). Thus, with a low substitution (i.e. a low θ_p), the welfare loss due to inflation (or deflation) is smaller. Therefore a lower θ_p allows for an inflation target farther away from zero, insofar as there are motives for a non-zero steady-state inflation. Such a mechanism explains why in Figure 13a and 13b optimal inflation is more negative with lower substitution.

Interestingly, when we consider robustness with respect to parameter θ_w, the ranking of the corresponding curves is reversed. That is, a larger θ_w induces a larger inflation target in absolute value. The reason is that, with a larger sustitution across labor types, a given nominal wage growth generates dispersion of quantities across types of labor that turns out to be particularly costly. In that case, it is optimal that the burden of adjustement of real wages to growth is borne not by nominal wages, but rather on nominal prices (thus leading to a more pronounced deflation).[24]

5 Summary and conclusion

In this paper, we have assessed how changes in the steady-state natural interest rate translate into changes in the optimal inflation target in a model subject to the ZLB. Our main finding is that, starting from pre-crisis values, a 1% decline in the natural rate should be accommodated by an increase in the optimal inflation target of about 0.9%. For convenience, Table 3 recaps our results. Overall, across the different concepts of optimal inflation considered in this paper, the level of optimal inflation does vary. However it is a very robust finding that the slope of the $(r^\star, \pi^\star)$ relation is close to -0.9 in the vicinity of the pre-crisis value of steady-state real interest rates both in the US and in the euro-area.

In our analysis, we considered adjusting the inflation target as the only option at the policymaker's disposal. This is not to say that this is *the only* option in their choice set. As a matter of fact, recent discussions revolving around monetary policy in the new normal have suggested that the various non-conventional measures used in the aftermath of the Great Recession could feature permanently in policy toolbox. In particular, unconventional monetary policies could represent useful second-best instruments when the ZLB is reached, as advocated by Reifschneider (2016). An alternative would consist in a change of monetary policy strategies, e.g., adopting a price-level targeting strategy as recently advocated by Williams (2016). Beyond monetary policy measures, fiscal policies could also play a significant role, as emphasized by Cor-

[24] This can be illustrated again in the approximated welfare function, and ignoring the effects of shocks. Then the optimal inflation obeys $\pi^\star \approx -(\lambda_w(1-\gamma_z)(1-\gamma_w)/[\lambda_p(1-\gamma_p)^2+\lambda_w(1-\gamma_w)^2])\mu_z$. Inflation target is a decreasing function of λ_w, thus of θ_w.

Table 3: Effect of a decline in $r^\star$ under alternative notions of optimal inflation

	US		EA	
	Baseline	Lower $r^\star$	Baseline	Lower $r^\star$
Mean of $\pi^\star$	2.00	3.00	1.79	2.60
Median of $\pi^\star$	1.96	2.90	1.47	2.28
$\pi^\star$ at post. mean	2.21	3.20	1.58	2.39
$\pi^\star$ at post. median	2.12	3.11	1.49	2.30
$\pi^{\star\star}$	2.40	3.30	2.20	3.10
$\pi^{\star\star}$, frozen MP	2.24	3.16	2.36	3.28
$\pi^\star$ at post. mean, ELB -40 bp	—	—	1.31	2.08
Average realized inflation at post. mean	2.20	3.19	1.56	2.36
Average realized inflation at post. mean, ELB -40 bp	—	—	1.24	1.97

Note: all figures are in annualized percentage rate.

reia et al. (2013). As a result, the ZLB might be less stringent a constraint in a practical policy context than in our analysis.[25] However, the efficacy and the costs of these policies should also be part of the analysis. The complete comparison of these policy trade-offs goes beyond the scope of the present paper.

We have discussed the impact of higher inflation target, abstracting from the transition to a higher inflation target. In the current lowflation environment, increasing the inflation target in reaction to a drop in the steady-state value of the real interest rate might be challenging: because of more frequent ZLB episodes, the realizations of inflation might be on average below the initial inflation target for some time and increasing the inflation target therefore would raise some credibility issues.

Finally, our analysis has also abstracted from forces identified in the literature as warranting a small, positive inflation target, irrespective of ZLB issues, as emphasized in Bernanke et al. (1999). The first is grounded on measurement issues, following the finding from the 1996 Boskin report that the consumer price index did probably over estimate inflation in the US by over 1 percentage point in the early nineties. The second argument is rooted in downward nominal rigidities. In an economy where there are such downward rigidities (e.g. in nominal wages) a positive inflation rate can help "grease the wheel" of the labor market by facilitating relative price adjustments. Symmetrically, we also abstracted from forces calling for lower inflation targets. The most obvious is the so-called Friedman (1969) rule, according to which average inflation should equal to minus the steady state real interest rate, hence be *negative*, in order to minimize loss of resources or utility and the distortionary wedge between cash and credit goods (e.g. consumption and leisure) induced by a non-zero nominal interest rate. We conjecture that adding these elements to our

[25]See Debortoli et al. (2018) for empirical evidence in favor of this hypothesis.

setup would leave our main conclusions unchanged. A complete assessment is left for future research.

A Various long-run and optimal inflation rates considered

Table A.1: Various Notions of Long-run and optimal Inflation in the model

π	Any inflation target, used to define the "inflation gap" that enters the Taylor rule
$\mathbb{E}(\pi_t)$	Average realized inflation, might differ from π due to ZLB
$\pi^\star(\theta)$	Inflation target that minimizes the loss function given a structural parameters θ
$\pi^\star(\bar{\theta})$	$\pi^\star$ assuming parameters at post. mean
$\pi^\star(median(\theta))$	$\pi^\star$ assuming parameters at post. median
$\bar{\pi}^\star$	average of $\pi^\star(\theta)$ over the posterior distribution of θ, i.e., $\int_\theta \pi^\star(\theta)p(\theta\vert X_T)\mathrm{d}\theta$
Median($\pi^\star$)	Median of $\pi^\star(\theta)$ over the posterior distribution
$\pi^{\star\star}$	Inflation target that minimizes the average loss function over the posterior distribution of θ

B The distribution of ZLB spells duration

Figure B.1: Distribution of ZLB spells duration at the posterior mean

(a) US

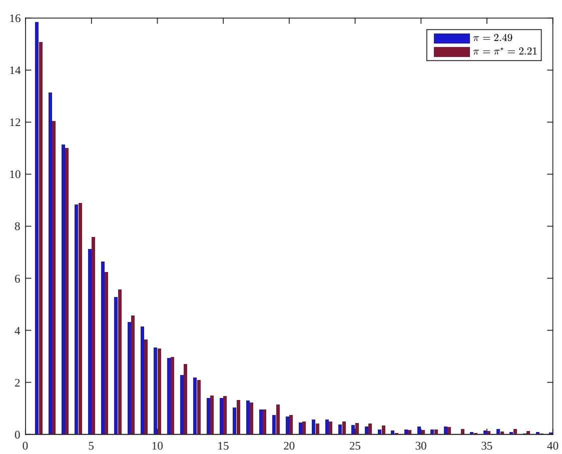

(b) EA

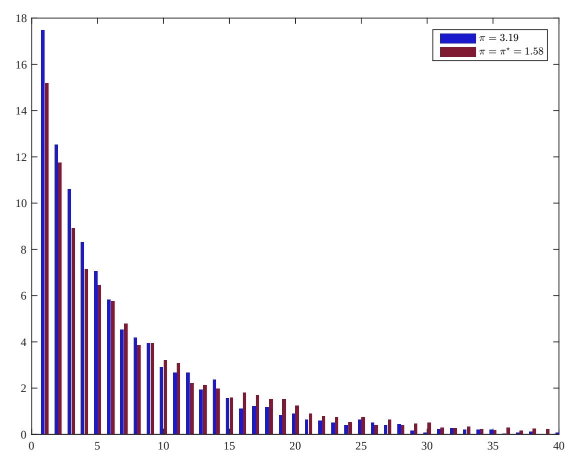

Note: Histograms are based on a simulated sample of 500,000 quarters. For both the US and the euro ares, simulations are carried out assuming in turn that the inflation target is the estimated inflation target ; and then that the inflation target is the optimal inflation target obtained using the mean of the posterior density of estimated parameters

C The welfare cost of inflation

Figure C.1: Welfare cost of inflation at the posterior mean

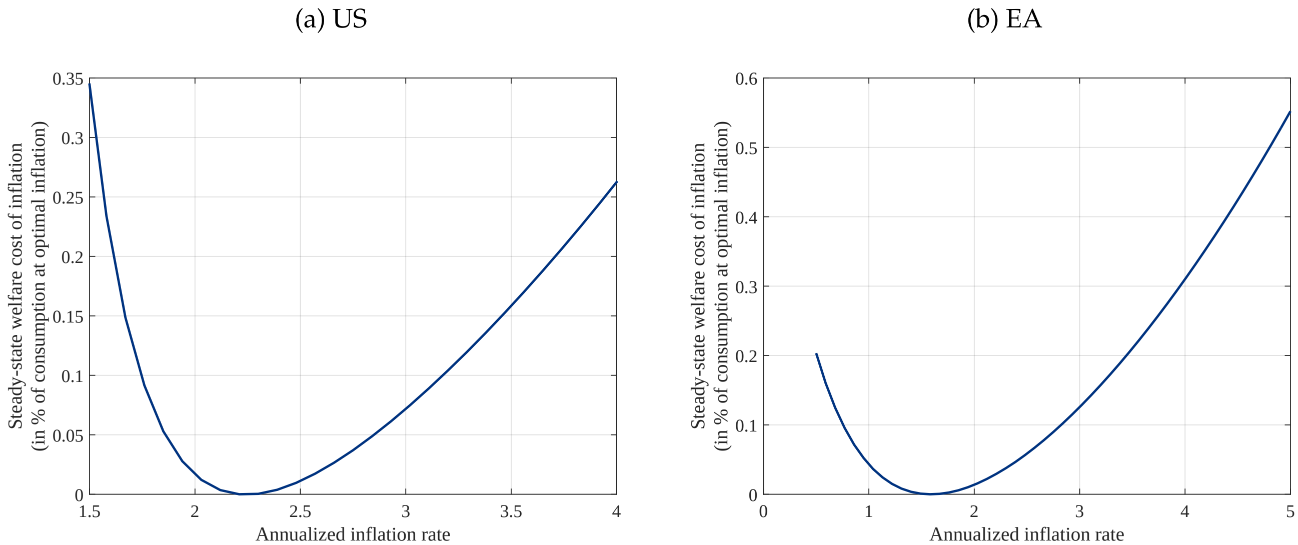

Note: The figure reports the welfare cost of inflation stated as a percentage of steady-state consumption in the optimal setting. See Appendix L for details.

D The distribution of optimal inflation targets

Figure D.1: Posterior Distribution of $\pi^\star$ - EA - Benchmark

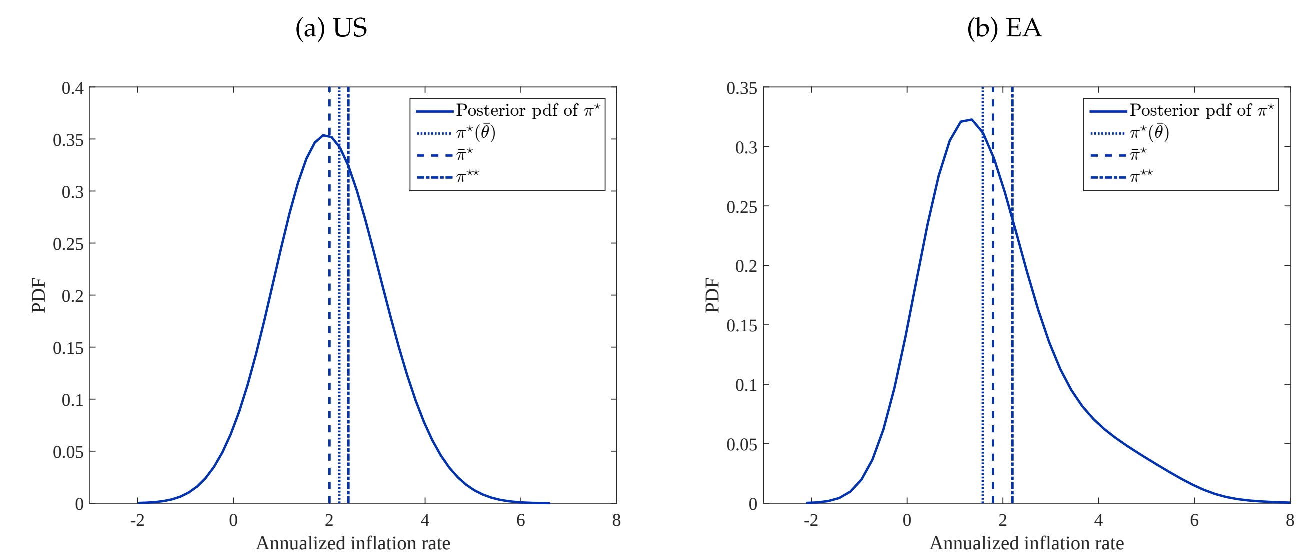

Note: Plain curve: PDF of $\pi^\star$; Dashed vertical line : Average value of $\pi^\star$ over posterior distribution; Dotted vertical line : Optimal inflation at the posterior mean of θ; Dashed-dotted vertical line : Bayesian-theoretic optimal inflation

E Impact of a decline in the natural rate on the welfare criterion

Figures E.1a and E.1b below provide a more precise sense of how $\pi^\star$ is modified following a decrease of $r^\star$.

Figure E.1: $\mathscr{W}(\pi, \mathbb{E}(\theta))$

(a) US

(b) EA

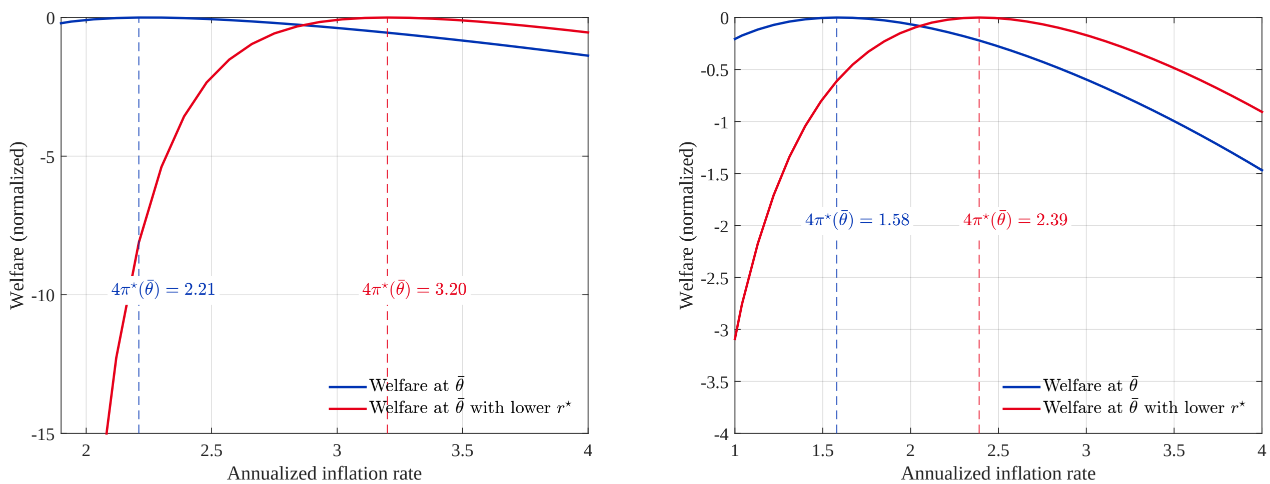

F Further illustration of the $(r^\star, \pi^\star)$ relation

F.1 When μ_z varies

Figure F.1: $(r^\star, \pi^\star)$ locus when μ_z varies

(a) US (b) EA

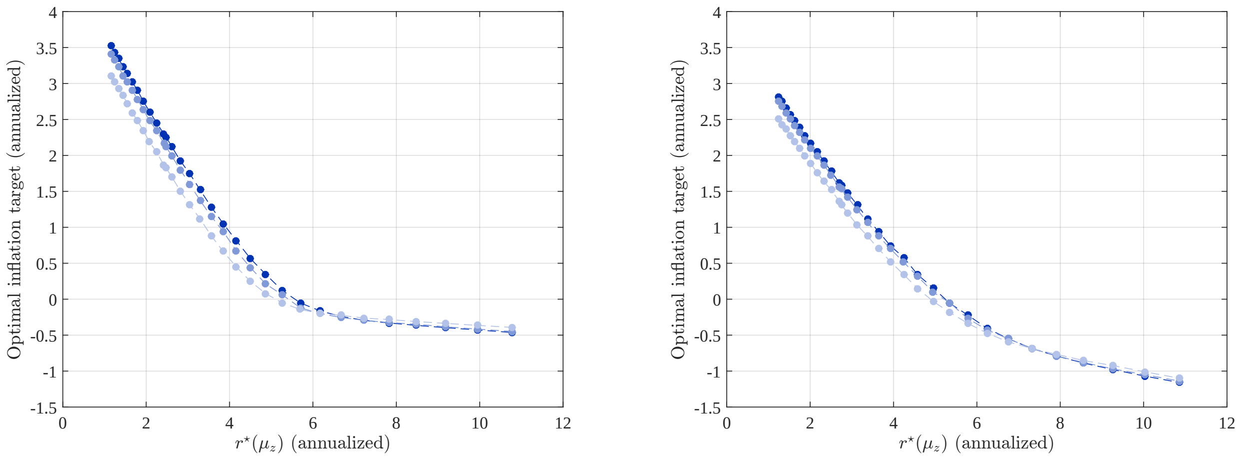

Note: Blue: parameters set at the posterior mean; light blue: parameters set at the posterior median; Lighter blue: parameters set at the posterior mode. Memo: $r^\star = \rho + \mu_z$. Range for μ_z: 0.4% to 10% (annualized) .

F.2 When ρ varies

Figure F.2: $(r^\star, \pi^\star)$ locus when ρ varies

(a) US (b) EA

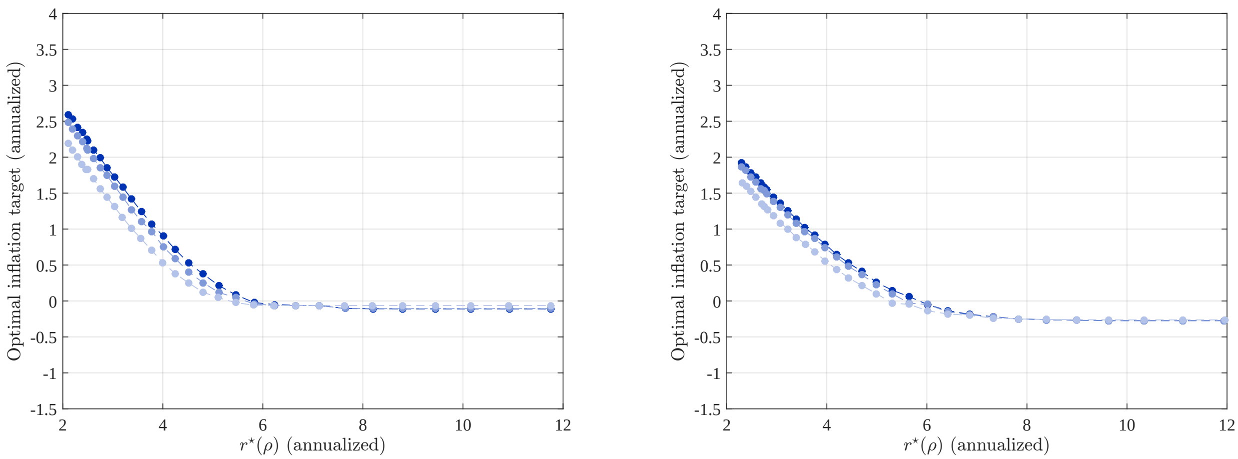

Note: Blue: parameters set at the posterior mean; light blue: parameters set at the posterior median; Lighter blue: parameters set at the posterior mode. Memo: $r^\star = \rho + \mu_z$. Range for μ_z: 0.4% to 10% (annualized) .

G Nominal and Real Interest Rates

Figure G.1: $(r^\star, i^\star)$ locus (at the posterior mean)

(a) US (b) EA

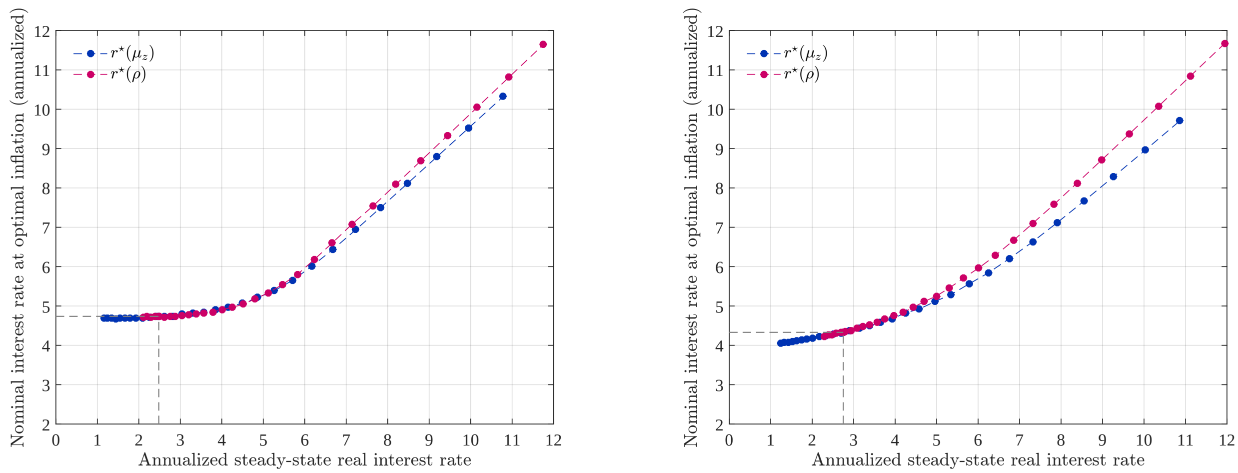

Note: the blue dots correspond to the $(r^\star, i^\star)$ locus when $r^\star$ varies with μ_z; the red dots correspond to the $(r^\star, i^\star)$ locus when $r^\star$ varies with ρ

H Distribution of $\pi^\star$ following a downward shift of the distribution of $r^\star$

Figure H.1: Counterfactual - US

(a) US (b) EA

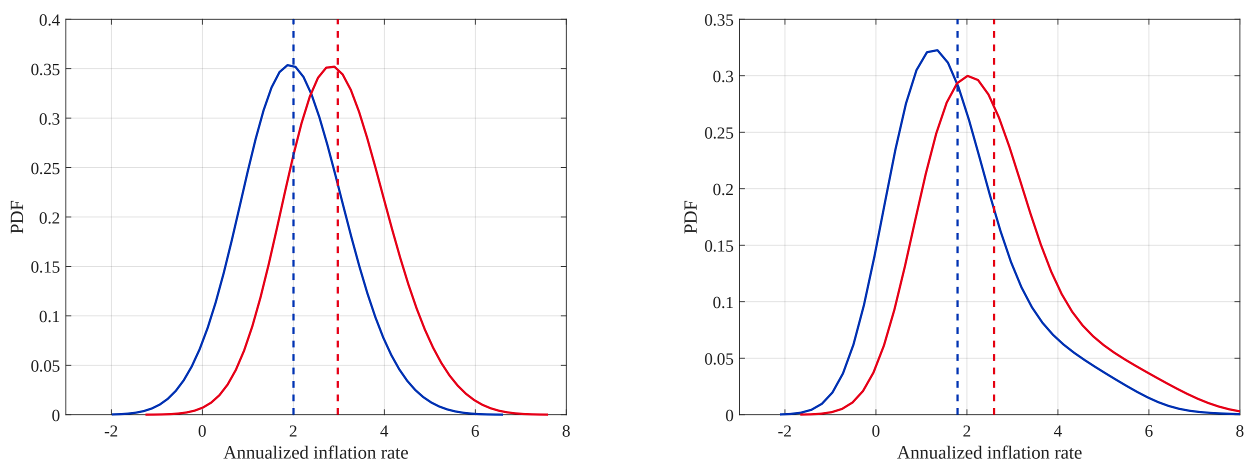

Note: The dashed vertical line indicates the mean value, i.e. $\mathbb{E}_\theta(\pi^\star(\theta))$.

I Additional results on the impact of θ_p and θ_w on the $(r^\star, \pi^\star)$ locus

Here we illustrate why depending on the configuration, a larger θ_p may lead to a larger optimal inflation target (in the US, see figure 14a) or to a lower one (in the EA, see figure 14b).

Figure I.1: $(r^\star, \pi^\star)$ relation with alternative θ_p and γ_p

(a) High γ_p (b) Low γ_p

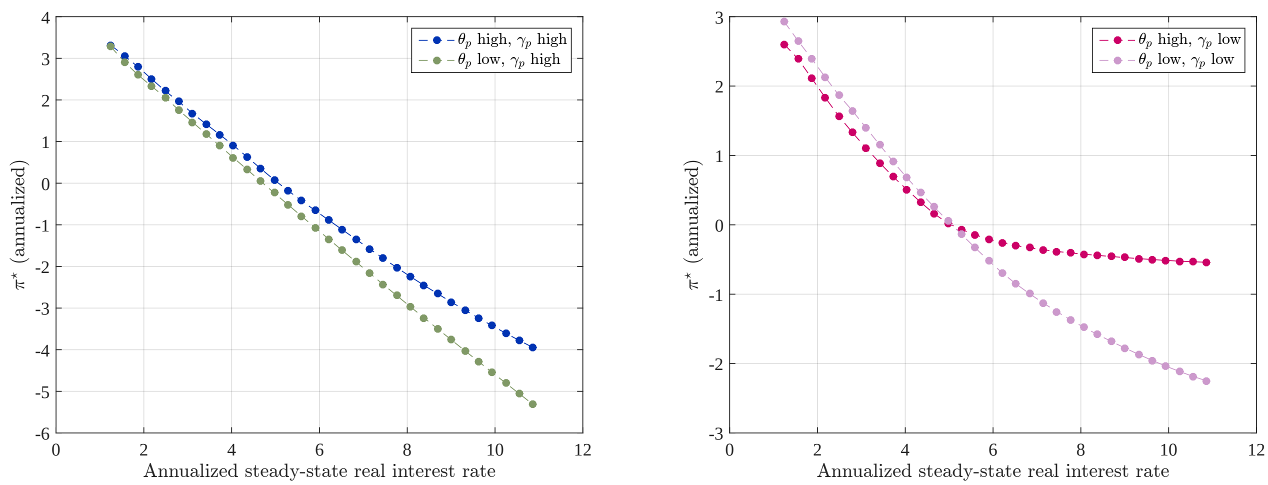

Note: the blues dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with θ_p set to 10. The green dots correspond to the counterfactual simulation with θ_p set to 3.

It turns out that there is a subtle interation between θ_p and the degree of indexation to past inflation γ_p affecting the shape of the $(r^\star, \pi^\star)$ locus. To document this, we ran additional simulations: (i) low indexation and low θ_p, (ii) low indexation, high θ_p, (iii) high indexation, low θ_p, and (iv) high indexation and high θ_p. To make as clar a point as possible, in the low indexation case, we set $\gamma_p = 0$ and in the high indexation case, we set $\gamma_p = 0.8$. Both are simulated in the case of the euro area.

To begin with, we consider the case of no indexation. See Figure I.1b. With less substitution (say complementary in the polar case), a given "cyclical" price dispersion induced by inflation, leads to smaller output dispersion. This is apparent from the formulas in our appendix, or textbook e.g. chapter 3 in Galí (2015). Thus with low subsitution there is smaller welfare loss to inflation. From this, a lower θ_p would call for an inflation target farther away from zero. This is what happens in figure I.1b.

Now with a large indexation the detrimental effect of inflation on price dispersion is mitigated. In that case, the policymaker is more willing to set an inflation target farther away from zero (see Figure I.1b). In particular, in the region with low $r^\star$ the difference between the low and high θ_p is no more material.

J Model Solution

J.1 Households

J.1.1 First Order Conditions

The associated lagrangian of program (1) under constraint (2) is

$$\mathscr{L}_t = \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \Bigg\{ e^{\zeta_{c,t+s}} \log(C_{t+s} - \hat{\eta} C_{t+s-1}) - \frac{\chi}{1+\nu} \int_0^1 e^{\zeta_{h,t+s}} (N_{t+s}(h))^{1+\nu} dh$$
$$- \frac{\Lambda_{t+s}}{P_{t+s}} \left[P_{t+s} C_{t+s} + Q_{t+s} B_{t+s} e^{-\zeta_{q,t+s}} + P_{t+s} \text{tax}_{t+s} - \int_0^1 W_{t+s}(h) N_{t+s}(h) dh - B_{t+s-1} - P_{t+s} \text{div}_{t+s} \right] \Bigg\},$$

The associated first order condition with respect to bonds is

$$\frac{\partial \mathscr{L}_t}{\partial B_t} = 0 \Leftrightarrow \Lambda_t Q_t e^{-\zeta_{q,t}} = \beta \mathbb{E}_t \left\{ \frac{\Lambda_{t+1}}{\Pi_{t+1}} \right\}, \tag{J.1}$$

and the first-order condition with respect to consumption is

$$\frac{\partial \mathscr{L}_t}{\partial C_t} = 0 \Leftrightarrow \frac{e^{\zeta_{c,t}}}{C_t - \hat{\eta} C_{t-1}} - \beta \hat{\eta} \mathbb{E}_t \left\{ \frac{e^{\zeta_{c,t+1}}}{C_{t+1} - \eta C_t} \right\} = \Lambda_t. \tag{J.2}$$

where $\Pi_t \equiv P_t / P_{t-1}$ represents the (gross) inflation rate, and

We induce stationarity by normalizing trending variables by the level of technical progress. To this end, we use the subscript z to refer to a normalized variable. For example, we define

$$C_{z,t} \equiv \frac{C_t}{Z_t}, \qquad \Lambda_{z,t} \equiv \Lambda_t Z_t,$$

where it is recalled that

$$Z_t = e^{z_t}$$

with

$$z_t = \mu_z + z_{t-1} + \zeta_{z,t}.$$

We then rewrite the first order condition in terms of the normalized variables. Equation (J.2) thus rewrites

$$\frac{e^{\zeta_{c,t}}}{C_{z,t} - \eta C_{z,t-1} e^{-\zeta_{z,t}}} - \beta \eta \mathbb{E}_t \left\{ e^{-\zeta_{z,t+1}} \frac{e^{\zeta_{c,t+1}}}{C_{z,t+1} - \eta C_{z,t} e^{-\zeta_{z,t+1}}} \right\} = \Lambda_{z,t}, \tag{J.3}$$

Similarly, equation (J.1) rewrites

$$\Lambda_{z,t} Q_t e^{-\zeta_{q,t}} = \beta e^{-\mu_z} \mathbb{E}_t \left\{ e^{-\zeta_{z,t+1}} \frac{\Lambda_{z,t+1}}{\Pi_{t+1}} \right\}, \tag{J.4}$$

where we defined

$$\eta \equiv \hat{\eta} e^{-\mu_z}.$$

Let us define $i_t \equiv -\log(Q_t)$ and for any generic variable X_t

$$x_t \equiv \log(X_t), \quad \hat{x}_t \equiv x_t - x$$

where x is the steady-state value of x. Using these definitions, log-linearizing equation (J.3) yields

$$\hat{g}_t + \beta\eta\mathbb{E}_t\{\hat{c}_{t+1}\} - (1+\beta\eta^2)\hat{c}_t + \eta\hat{c}_{t-1} - \eta(\zeta_{z,t} - \beta\mathbb{E}_t\{\zeta_{z,t+1}\}) = \varphi^{-1}\hat{\lambda}_t \tag{J.5}$$

where we defined

$$\varphi^{-1} \equiv (1-\beta\eta)(1-\eta),$$

$$\hat{g}_t = (1-\eta)(\zeta_{c,t} - \beta\eta\mathbb{E}_t\{\zeta_{c,t+1}\}).$$

Similarly, log-linearizing equation (J.4) yields

$$\hat{\lambda}_t = \hat{\imath}_t + \mathbb{E}_t\{\hat{\lambda}_{t+1} - \hat{\pi}_{t+1} - \zeta_{z,t+1}\} + \zeta_{q,t}. \tag{J.6}$$

J.2 Firms

Expressing the demand function in normalized terms yields

$$Y_{z,t}(f) = \left(\frac{P_t(f)}{P_t}\right)^{-\theta_p} Y_{z,t},$$

In the case of a firm not drawn to re-optimize, this equation specializes to (in log-linear terms)

$$\hat{y}_{t,t+s}(f) - \hat{y}_{t+s} = \theta_p(\hat{\pi}_{t,t+s} - \hat{\delta}^p_{t,t+s} - \hat{p}^\star_t(f)). \tag{J.7}$$

J.2.1 Cost Minimization

The real cost of producing $Y_t(f)$ units of good of f is

$$\frac{W_t}{P_t}L_t(f) = \frac{W_t}{P_t}\left(\frac{Y_t(f)}{Z_t}\right)^{\phi} \tag{J.8}$$

The associated real marginal cost is thus

$$S_t(f) = \phi\frac{W_t}{P_t Z_t}\left(\frac{Y_t(f)}{Z_t}\right)^{\phi-1} \tag{J.9}$$

It is useful at this stage to restate the production function in log-linearized terms:

$$\hat{y}_{z,t}(f) = \frac{1}{\phi}\hat{n}_t(f) \tag{J.10}$$

J.2.2 Price Setting of Intermediate Goods: Optimization

Firm f chooses $P_t^\star(f)$ in order to maximize

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta\alpha_p)^s \Lambda_{t+s} \left\{ (1+\tau_{p,t+s}) \frac{V_{t,t+s}^p P_t^\star(f)}{P_{t+s}} Y_{t,t+s}^\star(f) - S\left(Y_{t,t+s}(f)\right) \right\}, \tag{J.11}$$

subject to the demand function

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

and the cost schedule (J.8), where Λ_t is the representative household's marginal utility of wealth, and $\mathbb{E}_t\{\cdot\}$ is the expectation operator conditional on information available as of time t. That Λ_t appears in the above maximization program reflects the fact that the representative household is the ultimate owner of firm f.

The associated first-order condition is

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta\alpha_p)^s \Lambda_{t+s} \left\{ \left(\frac{V_{t,T}^p P_t^\star(f)}{P_{t+s}} \right)^{1-\theta_p} Y_{t+s} - \frac{\mu_p}{1+\tau_p} \mathrm{e}^{\zeta_{u,t+s}} \frac{W_{t+s}}{P_{t+s}} \phi \left(\left(\frac{V_{t,t+s}^p P_t^\star(f)}{P_{t+s}} \right)^{-\theta_p} \frac{Y_{t+s}}{Z_{t+s}} \right)^{\phi} \right\} = 0,$$

where

$$\mu_p \equiv \frac{\theta_p}{\theta_p - 1}.$$

This rewrites

$$\left(\frac{P_t^\star(f)}{P_t} \right)^{1+\theta_p(\phi-1)} = \frac{\mu_p}{1+\tau_p} \frac{K_{p,t}}{F_{p,t}}$$

where

$$K_{p,t} = \mathbb{E}_t \sum_{s=0}^{\infty} (\beta\alpha_p)^s \Lambda_{z,t+s} \mathrm{e}^{\zeta_{u,t+s}} \frac{W_{z,t+s}}{P_{t+s}} \phi \left(\left(\frac{V_{t,t+s}^p}{\Pi_{t,t+s}} \right)^{-\theta_p} Y_{z,t+s} \right)^{\phi}$$

and

$$F_{p,t} = \mathbb{E}_t \sum_{s=0}^{\infty} (\beta\alpha_p)^s \Lambda_{z,T} \left(\frac{V_{t,t+s}^p}{\Pi_{t,t+s}} \right)^{1-\theta_p} Y_{z,t+s},$$

where $\Pi_{t,t+s} \equiv P_{t+s}/P_t$.

Notice that

$$K_{p,t} = \phi \Lambda_{z,t} \mathrm{e}^{\zeta_{u,t}} \frac{W_{z,t}}{P_t} (Y_{z,t})^{\phi} + \beta\alpha_p \mathbb{E}_t \left(\frac{(\Pi_t))^{\gamma_p}}{\Pi_{t+1}} \right)^{-\phi\theta_p} K_{p,t+1},$$

and

$$F_{p,t} = \Lambda_{z,t} Y_{z,t} + \beta\alpha_p \mathbb{E}_t \left(\frac{(\Pi_t)^{\gamma_p}}{\Pi_{t+1}} \right)^{1-\theta_p} F_{p,t+1}.$$

With a slight abuse of notation, we obtain the steady-state relation

$$\left(\frac{P^\star}{P} \right)^{1+\theta_p(\phi-1)} = \frac{\mu_p}{1+\tau_p} \phi \frac{W_z}{P} Y_z^{\phi-1} \frac{1 - \beta\alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)}}{1 - \beta\alpha_p(\Pi)^{\phi\theta_p(1-\gamma_p)}}.$$

Log-linearizing yields

$$[1+\theta_p(\phi-1)](p_t^\star - p_t) = \hat{k}_{p,t} - \hat{f}_{p,t}$$

$$\hat{k}_{p,t} = (1-\omega_{K,p})[\hat{\lambda}_{z,t} + \hat{\omega}_t + \phi\hat{y}_{z,t} + \zeta_{u,t}] + \omega_{K,p}\mathbb{E}_t\{\hat{k}_{p,t+1} + \phi\theta_p(\hat{\pi}_{t+1} - \gamma_p\hat{\pi}_t)\},$$

and

$$\hat{f}_{p,t} = (1-\omega_{F,p})(\hat{\lambda}_{z,t} + \hat{y}_{z,t}) + \omega_{F,p}\mathbb{E}_t\{\hat{f}_{p,t+1} + (\theta_p - 1)(\hat{\pi}_{t+1} - \gamma_p\hat{\pi}_t)\}.$$

where we defined the de-trended real wage

$$\omega_t \equiv w_{z,t} - p_t$$

$$\omega_{K,p} \equiv \beta\alpha_p(\Pi)^{(1-\gamma_p)\phi\theta_p}$$

and

$$\omega_{F,p} \equiv \beta\alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)}$$

Finally, notice that

$$\begin{aligned} P_t^{1-\theta_p} &= \int_0^1 P_t(f)^{1-\theta_p}\mathrm{d}f \\ &= (1-\alpha_p)(P_t^\star)^{1-\theta_p} + \alpha_p\int_0^1 \left[(\Pi_{t-1}))^{\gamma_p}P_{t-1}(f)\right]^{1-\theta_p}\mathrm{d}f \end{aligned}$$

Thus

$$1 = (1-\alpha_p)\left(\frac{P_t^\star}{P_t}\right)^{1-\theta_p} + \alpha_p\left[\frac{(\Pi_{t-1})^{\gamma_p}}{\Pi_t}\right]^{1-\theta_p}$$

The steady-state relation is

$$\left(\frac{P^\star}{P}\right)^{1-\theta_p} = \frac{1-\alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)}}{1-\alpha_p}$$

Log-linearizing this yields

$$\hat{p}_t^\star = \frac{\omega_{F,p}}{\beta - \omega_{F,p}}(\hat{\pi}_t - \gamma_p\hat{\pi}_{t-1}).$$

J.3 Unions

J.3.1 Wage Setting

Union h sets $W_t^\star(h)$ so as to maximize

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

where

$$N_{t,t+s}(h) = \left(\frac{\mathrm{e}^{\gamma_z\mu_z s}V_{t,t+s}^w W_t^\star(h)}{W_{t+s}}\right)^{-\theta_w} N_{t+s}$$

The associated first-order condition is

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta\alpha_w)^s \Bigg\{ \Lambda_T \frac{W_{t+s}}{P_{t+s}} h_{t+s} \left(\frac{\mathrm{e}^{\gamma_z \mu_z s} V^w_{t,T}}{\Pi^w_{t,t+s}} \frac{W^\star_t(h)}{W_{t+s}} \right)^{1-\theta_w} - \frac{\mu_w}{1+\tau_w} \chi \mathrm{e}^{\zeta_{h,t+s}} \left(\frac{\mathrm{e}^{\gamma_z \mu_z s} V^w_{t,t+s}}{\Pi^w_{t,t+s}} \frac{W^\star_t(h)}{W_{t+s}} \right)^{-(1+\nu)\theta_w} N^{1+\nu}_{t+s} \Bigg\} = 0$$

where $\Pi^w_{t,t+s} = W_{t+s}/W_t$.

Rearranging yields

$$\left(\frac{W^\star_t(h)}{W_t} \right)^{1+\theta_w \nu} = \frac{\mu_w}{1+\tau_w} \frac{K_{w,t}}{F_{w,t}}.$$

where

$$K_{w,t} = \mathbb{E}_t \sum_{s=0}^{\infty} (\beta\alpha_w)^s \left\{ \chi \mathrm{e}^{\zeta_{h,t+s}} \left(\frac{\mathrm{e}^{\gamma_z \mu_z s} V^w_{t,t+s}}{\Pi^w_{t,t+s}} \right)^{-(1+\nu)\theta_w} N^{1+\nu}_{t+s} \right\}$$

and

$$F_{w,t} = \mathbb{E}_t \sum_{s=0}^{\infty} (\beta\alpha_w)^s \left\{ \Lambda_{t+s} \frac{W_{t+s}}{P_{t+s}} N_{t+s} \left(\frac{\mathrm{e}^{\gamma_z \mu_z s} V^w_{t,t+s}}{\Pi^w_{t,t+s}} \right)^{1-\theta_w} \right\},$$

where $\Pi^w_{t,t+s} \equiv W_{t+s}/W_t$.

Notice that

$$K_{w,t} = \chi \mathrm{e}^{\zeta_{h,t}} N^{1+\nu}_t + \beta\alpha_w \mathbb{E}_t \left\{ \left(\mathrm{e}^{\gamma_z \mu_z} \frac{(\Pi_t)^{\gamma_w}}{\Pi_{w,t+1}} \right)^{-(1+\nu)\theta_w} K_{w,t+1} \right\}$$

and

$$F_{w,t} = \Lambda_{z,t} \frac{W_{z,t}}{P_t} N_t + \beta\alpha_w \mathbb{E}_t \left\{ \left(\mathrm{e}^{\gamma_z \mu_z} \frac{(\Pi_t)^{\gamma_w}}{\Pi_{w,t+1}} \right)^{1-\theta_w} F_{w,t+1} \right\}.$$

The associated steady-state relations are

$$\left(\frac{W^\star}{W} \right)^{1+\theta_w \nu} = \frac{\mu_w}{1+\tau_w} \frac{K_w}{F_w},$$

$$K_w = \frac{\chi N^{1+\nu}}{1 - \beta\alpha_w [\mathrm{e}^{(1-\gamma_z)\mu_z} (\Pi)^{1-\gamma_w}]^{(1+\nu)\theta_w}}$$

$$F_w = \frac{\Lambda \frac{W_z}{P} H}{1 - \beta\alpha_w [\mathrm{e}^{(1-\gamma_z)\mu_z} (\Pi)^{1-\gamma_w}]^{\theta_w - 1}}.$$

Log-linearizing the above equations finally yields

$$(1+\theta_w \nu)(w^\star_t - w_t) = \hat{k}_{w,t} - \hat{f}_{w,t}.$$

$$\hat{k}_{w,t} = (1-\omega_{K,w})[(1+\nu)\hat{n}_t + \zeta_{h,t}] + \omega_{K,w} \mathbb{E}_t \{ \hat{k}_{w,t+1} + (1+\nu)\theta_w (\hat{\pi}_{w,t+1} - \gamma_w \hat{\pi}_t) \}$$

$$\hat{f}_{w,t} = (1-\omega_{F,w})(\hat{\lambda}_{z,t} + \hat{\omega}_t + \hat{n}_t) + \omega_{F,w}\mathbb{E}_t\{\hat{f}_{w,t+1} + (\theta_w - 1)(\hat{\pi}_{w,t+1} - \gamma_w \hat{\pi}_t)\},$$

where we defined

$$\omega_{K,w} = \beta\alpha_w[\mathrm{e}^{(1-\gamma_z)\mu_z}(\Pi)^{(1-\gamma_w)}]^{(1+\nu)\theta_w}$$

$$\omega_{F,w} = \beta\alpha_w[\mathrm{e}^{(1-\gamma_z)\mu_z}(\Pi)^{(1-\gamma_w)}]^{\theta_w - 1}$$

To complete this section, notice that

$$1 = (1-\alpha_w)\left(\frac{W_t^\star}{W_t}\right)^{1-\theta_w} + \alpha_w\left(\mathrm{e}^{\gamma_z\mu_z}\frac{[\Pi_{t-1}]^{\gamma_w}}{\Pi_{w,t}}\right)^{1-\theta_w}$$

and

$$w_t^\star - w_t = \frac{\omega_{F,w}}{\beta - \omega_{F,w}}(\hat{\pi}_{w,t} - \gamma_w\hat{\pi}_{t-1}).$$

J.4 Market Clearing

The clearing on the labor market implies

$$N_t = \left(\frac{Y_t}{Z_t}\right)^{\phi}\int_0^1\left(\frac{P_t(f)}{P_t}\right)^{-\phi\theta_p}\mathrm{d}f.$$

Let us define

$$\Xi_{p,t} = \left(\int_0^1\left(\frac{P_t(f)}{P_t}\right)^{-\phi\theta_p}\mathrm{d}f\right)^{-1/(\phi\theta_p)},$$

so that

$$N_t = (Y_{z,t}\Xi_{p,t}^{-\theta_p})^{\phi}.$$

Hence, expressed in log-linear terms, this equation reads

$$\hat{n}_t = \phi(\hat{y}_{z,t} - \theta_p\hat{\xi}_{p,t}).$$

Notice that

$$\Xi_{p,t}^{-\phi\theta_p} = (1-\alpha_p)\left(\frac{P_t^\star}{P_t}\right)^{-\phi\theta_p} + \alpha_p\left(\frac{[\Pi_{t-1}]^{\gamma_p}}{\Pi_t}\right)^{-\phi\theta_p}\Xi_{p,t-1}^{-\phi\theta_p}.$$

The associated steady-state relation is

$$\Xi_p^{-\phi\theta_p} = \frac{(1-\alpha_p)}{1-\alpha_p(\Pi)^{(1-\gamma_p)\phi\theta_p}}\left(\frac{P^\star}{P}\right)^{-\phi\theta_p}.$$

Log-linearizing the price dispersion yields

$$\hat{\xi}_{p,t} = (1-\omega_\Xi)(p_t^\star - p_t) + \omega_\Xi[\hat{\xi}_{p,t-1} - (\hat{\pi}_t - \gamma_p\hat{\pi}_{t-1})]$$

where we defined

$$\omega_\Xi = \alpha_p(\Pi)^{(1-\gamma_p)\phi\theta_p}.$$

J.5 Natural Rate of Output

The natural rate of output is the level of production that would prevail in an economy without nominal rigidities, i.e. $\alpha_p = \alpha_w = 0$ and without cost-push shocks (i.e., $\zeta_{u,t} = 0$). Under such circumstances, the dynamic system simplifies to

$$\hat{w}^n_{z,t} + (\phi - 1)\hat{y}^n_{z,t} = 0,$$

$$\nu\hat{n}^n_t + \zeta_{h,t} = \hat{\lambda}^n_{z,t} + \hat{w}^n_{z,t},$$

$$\hat{n}^n_t = \phi\hat{y}^n_{z,t},$$

$$\hat{g}_t + \beta\eta\mathbb{E}_t\{\hat{y}^n_{z,t+1}\} - (1+\beta\eta^2)\hat{y}^n_{z,t} + \eta\hat{y}^n_{z,t-1} - \eta(\zeta_{z,t} - \beta\mathbb{E}_t\{\zeta_{z,t+1}\}) = \varphi^{-1}\hat{\lambda}^n_{z,t},$$

where the superscript n stands for *natural*.

Combining these equations yields

$$[\varphi(1+\beta\eta^2) + \omega]\hat{y}^n_{z,t} - \varphi\beta\eta\mathbb{E}_t\{\hat{y}^n_{z,t+1}\} - \varphi\eta\hat{y}^n_{z,t-1} = \varphi\hat{g}_t - \zeta_{h,t} - \varphi\eta\zeta^*_{z,t}$$

where we defined

$$\omega \equiv \nu\phi + \phi - 1,$$

and

$$\zeta^*_{z,t} = \zeta_{z,t} - \beta\mathbb{E}_t\{\zeta_{z,t+1}\}$$

J.6 Working Out the Steady State

The steady state is defined by the following set of equations

$$\frac{1-\beta\eta}{(1-\eta)C} = \Lambda_z,$$

$$\mathrm{e}^{-i} = \beta\mathrm{e}^{-\mu_z}\Pi^{-1},$$

$$\left(\frac{P^\star}{P}\right)^{1+\theta_p(\phi-1)} = \frac{\mu_p}{1+\tau_p}\frac{K_p}{F_p},$$

$$K_p = \frac{\phi\Lambda_z\frac{W_z}{P}Y_z^\phi}{1-\beta\alpha_p(\Pi)^{\phi\theta_p(1-\gamma_p)}},$$

$$F_p = \frac{\Lambda_z Y_z}{1 - \beta\alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)}},$$

$$\left(\frac{P^\star}{P}\right)^{1-\theta_p} = \frac{1 - \alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)}}{1-\alpha_p},$$

$$\left(\frac{W^\star}{W}\right)^{1+\theta_w\nu} = \frac{\mu_w}{1+\tau_w}\frac{K_w}{F_w},$$

$$K_w = \frac{\chi N^{1+\nu}}{1 - \beta\alpha_w[\mathrm{e}^{(1-\gamma_z)\mu_z}(\Pi)^{1-\gamma_w}]^{(1+\nu)\theta_w}},$$

$$F_w = \frac{\Lambda_z \frac{W_z}{P} H}{1 - \beta\alpha_w[\mathrm{e}^{(1-\gamma_z)\mu_z}(\Pi)^{1-\gamma_w}]^{\theta_w-1}},$$

$$\left(\frac{W^\star}{W}\right)^{1-\theta_w} = \frac{1 - \alpha_w[\mathrm{e}^{(1-\gamma_z)\mu_z}(\Pi)^{(1-\gamma_w)}]^{\theta_w-1}}{1-\alpha_w},$$

$$\Pi_w = \Pi\mathrm{e}^{\mu_z}$$

We can solve for i and Π_w using

$$\Pi_w = \Pi\mathrm{e}^{\mu_z}$$

$$1 = \beta\mathrm{e}^{-\mu_z}\mathrm{e}^{i}\Pi^{-1},$$

Standard manipulations yield

$$\frac{1-\omega_{K,p}}{1-\omega_{F,p}}\left(\frac{\beta(1-\alpha_p)}{\beta-\omega_{F,p}}\right)^{\frac{1+\theta_p(\phi-1)}{\theta_p-1}} = \frac{\mu_p}{1+\tau_p}\phi\frac{W_z}{P}Y_z^{\phi-1},$$

where we used

$$\omega_{K,p} = \beta\alpha_p(\Pi)^{(1-\gamma_p)\phi\theta_p}$$

$$\omega_{F,p} = \beta\alpha_p(\Pi)^{(1-\gamma_p)(\theta_p-1)}$$

Similar manipulations yield

$$\frac{1-\omega_{K,w}}{1-\omega_{F,w}}\left(\frac{\beta(1-\alpha_w)}{\beta-\omega_{F,w}}\right)^{\frac{1+\theta_w\nu}{\theta_w-1}} = \frac{\mu_w}{1+\tau_w}\frac{\chi N^{\nu}}{\Lambda_z\frac{W_z}{P}},$$

where we used

$$\omega_{K,w} = \beta\alpha_w[\mathrm{e}^{(1-\gamma_z)\mu_z}(\Pi)^{(1-\gamma_w)}]^{(1+\nu)\theta_w}$$

$$\omega_{F,w} = \beta\alpha_w[\mathrm{e}^{(1-\gamma_z)\mu_z}(\Pi)^{(1-\gamma_w)}]^{\theta_w-1}$$

Combining these conditions yields

$$\frac{1-\omega_{K,w}}{1-\omega_{F,w}}\left(\frac{\beta(1-\alpha_w)}{\beta-\omega_{F,w}}\right)^{\frac{1+\theta_w\nu}{\theta_w-1}}\frac{1-\omega_{K,p}}{1-\omega_{F,p}}\left(\frac{\beta(1-\alpha_p)}{\beta-\omega_{F,p}}\right)^{\frac{1+\theta_p(\phi-1)}{\theta_p-1}} = \frac{\mu_w}{1+\tau_w}\frac{\mu_p}{1+\tau_p}\frac{1-\eta}{1-\beta\eta}\phi\chi N^\nu Y_z^\phi$$

Now, recall that

$$(Y_z\Xi_p^{-\theta_p})^\phi = N$$

Then, using

$$\Xi_p^{-\phi\theta_p} = \frac{1-\alpha_p}{1-\omega_\Xi}\left(\frac{P^\star}{P}\right)^{-\phi\theta_p},$$

and

$$\left(\frac{P^\star}{P}\right)^{-\phi\theta_p} = \left(\frac{\beta(1-\alpha_p)}{\beta-\omega_{F,p}}\right)^{-\phi\frac{\theta_p}{\theta_p-1}}$$

we end up with

$$N^\nu Y_z^\phi = \left(\frac{1-\alpha_p}{1-\omega_\Xi}\left(\frac{\beta(1-\alpha_p)}{\beta-\omega_{F,p}}\right)^{-\phi\frac{\theta_p}{\theta_p-1}}\right)^\nu Y_z^{(1+\nu)\phi},$$

so that

$$\Omega = \frac{\mu_w}{1+\tau_w}\frac{\mu_p}{1+\tau_p}\frac{1-\eta}{1-\beta\eta}\phi\chi Y_z^{(1+\nu)\phi},$$

where

$$\Omega = \frac{1-\omega_{K,w}}{1-\omega_{F,w}}\left(\frac{\beta(1-\alpha_w)}{\beta-\omega_{F,w}}\right)^{\frac{1+\theta_w\nu}{\theta_w-1}}\frac{1-\omega_{K,p}}{1-\omega_{F,p}}\left(\frac{\beta(1-\alpha_p)}{\beta-\omega_{F,p}}\right)^{\frac{1+\theta_p[(1+\nu)\phi-1]}{\theta_p-1}}\left(\frac{1-\omega_\Xi}{1-\alpha_p}\right)^\nu$$

We defined the natural rate of output as the level of production that would prevail in an economy without nominal rigidities, i.e. $\alpha_p = \alpha_w = 0$. Under such circumstances, the steady-state value of the (normalized) natural rate of output y^n obeys

$$1 = \frac{\mu_w}{1+\tau_w}\frac{\mu_p}{1+\tau_p}\frac{1-\eta}{1-\beta\eta}\phi\chi(Y_z^n)^{\phi(1+\nu)}.$$

It follows that the steady-state distortion due to sticky prices and wages (and less than perfect indexation) is

$$\left(\frac{Y_z}{Y_z^n}\right)^{\phi(1+\nu)} = \Omega.$$

K Welfare

Let us define for any generic variable X_t

$$\frac{X_t - X}{X} = \hat{x}_t + \frac{1}{2}\hat{x}_t^2 + \mathcal{O}(||\zeta||^3)$$

$$\frac{X_t - X^n}{X^n} = \tilde{x}_t + \frac{1}{2}\tilde{x}_t^2 + \mathcal{O}(||\zeta||^3)$$

Below, we repeatedly use the following two lemmas:

Lemma 1. *Let $g(\cdot)$ be a twice differentiable function and let X be a stationary random variable. Then*

$$\mathbb{E}\{g(X)\} = g(\mathbb{E}\{X\}) + \frac{1}{2}g''(\mathbb{E}\{X\})\mathbb{V}\{X\} + \mathcal{O}(||X||^3).$$

Lemma 2. *Let $g(\cdot)$ be a twice differentiable function and let x be a stationary random variable. Then*

$$\mathbb{V}\{g(X)\} = [g'(\mathbb{E}\{X\})]^2\mathbb{V}\{X\} + \mathcal{O}(||X||^3).$$

In the rest of this section, we take a second-order approximation of welfare, where we consider the inflation rate as an expansion parameter. It follows that we consider the welfare effects of non-zero trend inflation only up to second order.

K.1 Second-Order Approximation of Utility

Consider first the utility derived from consumption. For the sake of notational simplicity, define

$$U(C_{z,t} - \eta C_{z,t-1}\mathrm{e}^{-\zeta_{z,t}}) = \log(C_{z,t} - \eta C_{z,t-1}\mathrm{e}^{-\zeta_{z,t}})$$

We thus obtain

$$\begin{aligned}
\mathrm{e}^{\zeta_{c,t}}U(C_{z,t} - \eta C_{z,t-1}\mathrm{e}^{-\zeta_{z,t}}) = \frac{1}{1-\eta}\Bigg[&\left(\frac{C_{z,t} - C_z^n}{C_z^n}\right) - \eta\left(\frac{C_{z,t-1} - C_z^n}{C_z^n}\right) \\
&- \frac{1}{2}\frac{1}{(1-\eta)}\left(\frac{C_{z,t} - C_z^n}{C_z^n}\right)^2 + \frac{\eta}{(1-\eta)}\left(\frac{C_{z,t} - C_z^n}{C_z^n}\right)\left(\frac{C_{z,t-1} - C_z^n}{C_z^n}\right) - \frac{1}{2}\frac{\eta^2}{(1-\eta)}\left(\frac{C_{z,t-1} - C_z^n}{C_z^n}\right)^2 \\
&+ \zeta_{c,t}\left(\frac{C_{z,t} - C_z^n}{C_z^n}\right) - \eta\zeta_{c,t}\left(\frac{C_{z,t-1} - C_z^n}{C_z^n}\right) - \frac{\eta}{(1-\eta)}\zeta_{z,t}\left(\frac{C_{z,t} - C_z^n}{C_z^n}\right) + \frac{\eta}{(1-\eta)}\zeta_{z,t}\left(\frac{C_{z,t-1} - C_z^n}{C_z^n}\right)\Bigg] \\
&+ \mathrm{t.i.p} + \mathcal{O}(||\zeta||^3),
\end{aligned}$$

where t.i.p stands for terms independent of policy.

Then, using

$$\frac{C_{z,t} - C_z^n}{C_z^n} = \tilde{c}_{z,t} + \frac{1}{2}\tilde{c}_{z,t}^2 + \mathcal{O}(||\zeta||^3)$$

we obtain

$$\begin{aligned}
\mathrm{e}^{\zeta_{c,t}}U(C_{z,t} - \eta C_{z,t-1}\mathrm{e}^{-\zeta_{z,t}}) = \frac{1}{1-\eta}\Bigg[&\tilde{c}_{z,t} - \eta\tilde{c}_{z,t-1} + \frac{1}{2}(\tilde{c}_{z,t}^2 - \eta\tilde{c}_{z,t-1}^2) \\
&- \frac{1}{2}\frac{1}{1-\eta}\tilde{c}_{z,t}^2 + \frac{\eta}{1-\eta}\tilde{c}_{z,t}\tilde{c}_{z,t-1} - \frac{1}{2}\eta^2\frac{1}{1-\eta}\tilde{c}_{z,t-1}^2 \\
&+ \zeta_{c,t}(\tilde{c}_{z,t} - \eta\tilde{c}_{z,t-1}) - \frac{\eta}{1-\eta}\zeta_{z,t}(\tilde{c}_{z,t} - \tilde{c}_{z,t-1})\Bigg] + \mathrm{t.i.p} + \mathcal{O}(||\zeta||^3),
\end{aligned}$$

Using

$$\varphi^{-1} = (1-\beta\eta)(1-\eta)$$

we obtain

$$\begin{aligned}
\mathrm{e}^{\zeta_{c,t}} U(C_{z,t} - \eta C_{z,t-1}\mathrm{e}^{-\zeta_{z,t}}) = \frac{1}{1-\eta}\Big[& \tilde{y}_{z,t} - \eta\tilde{y}_{z,t-1} + \frac{1}{2}(\tilde{y}_{z,t}^2 - \eta\tilde{y}_{z,t-1}^2) \\
& - \frac{1}{2}(1-\beta\eta)\varphi\tilde{y}_{z,t}^2 + \eta(1-\beta\eta)\varphi\tilde{y}_{z,t}\tilde{y}_{z,t-1} - \frac{1}{2}\eta^2(1-\beta\eta)\varphi\tilde{y}_{z,t-1}^2 \\
& + \zeta_{c,t}(\tilde{y}_{z,t} - \eta\tilde{y}_{z,t-1}) - \eta(1-\beta\eta)\varphi\zeta_{z,t}(\tilde{y}_{z,t} - \tilde{y}_{z,t-1})\Big] + \text{t.i.p} + \mathcal{O}(||\zeta||^3),
\end{aligned}$$

where we imposed the equilibrium condition on the goods market.

Similarly, taking a second-order approximation of labor disutility in the neighborhood of the natural steady-state N^n yields

$$\begin{aligned}
\frac{\chi}{1+\nu}\mathrm{e}^{\zeta_{h,t}}(N_t(h))^{1+\nu} = {} & \chi(N^n)^{1+\nu}\left(\frac{N_t(h)-N^n}{N^n}\right) + \frac{1}{2}\chi\nu(N^n)^{1+\nu}\left(\frac{N_t(h)-N^n}{N^n}\right)^2 \\
& + \chi(N^n)^{1+\nu}\left(\frac{N_t(h)-N^n}{N^n}\right)\zeta_{h,t} + \text{t.i.p} + \mathcal{O}(||\zeta||^3).
\end{aligned}$$

Now, using

$$\frac{N_t(h)-N^n}{N^n} = \tilde{n}_t(h) + \frac{1}{2}\tilde{n}_t(h)^2 + \mathcal{O}(||\zeta||^3)$$

we get

$$\frac{\chi}{1+\nu}\mathrm{e}^{\zeta_{h,t}}(N_t(h))^{1+\nu} = \chi(N^n)^{1+\nu}\left[\tilde{n}_t(h) + \frac{1}{2}(1+\nu)\tilde{n}_t(h)^2 + \tilde{n}_t(h)\zeta_{h,t}\right] + \text{t.i.p} + \mathcal{O}(||\zeta||^3).$$

Integrating over the set of labor types, one gets

$$\int_0^1 \frac{\chi}{1+\nu}\mathrm{e}^{\zeta_{h,t}}(N_t(h))^{1+\nu}\mathrm{d}h = \chi(N^n)^{1+\nu}\left[\mathbb{E}_h\{\tilde{n}_t(h)\} + \frac{1}{2}(1+\nu)\mathbb{E}_h\{\tilde{n}_t(h)^2\} + \mathbb{E}_h\{\tilde{n}_t(h)\}\zeta_{h,t}\right] + \text{t.i.p} + \mathcal{O}(||\zeta||^3).$$

Now, since

$$\mathbb{V}_h\{\tilde{n}_t(h)\} = \mathbb{E}_h\{\tilde{n}_t(h)^2\} - \mathbb{E}_h\{\tilde{n}_t(h)\}^2$$

the above relation rewrites

$$\begin{aligned}
\int_0^1 \frac{\chi}{1+\nu}\mathrm{e}^{\zeta_{h,t}}(N_t(h))^{1+\nu}\mathrm{d}h = \chi(N^n)^{1+\nu}\Big[& \mathbb{E}_h\{\tilde{n}_t(h)\} + \frac{1}{2}(1+\nu)(\mathbb{V}_h\{\tilde{n}_t(h)\} + \mathbb{E}_h\{\tilde{n}_t(h)\}^2) \\
& + \mathbb{E}_h\{\tilde{n}_t(h)\}\zeta_{h,t}\Big] + \text{t.i.p} + \mathcal{O}(||\zeta||^3).
\end{aligned}$$

We need to express $\mathbb{E}_h\{\tilde{n}_t(h)\}$ and $\mathbb{V}_h\{\tilde{n}_t(h)\}$ in terms of the aggregate variables. To this end, we first establish a series of results, on which we draw later on.

K.2 Aggregate Labor and Aggregate Output

Notice that

$$\frac{\theta_w - 1}{\theta_w}\tilde{n}_t = \log\left(\int_0^1 \left(\frac{N_t(h)}{N^n}\right)^{(\theta_w-1)/\theta_w} \mathrm{d}h\right).$$

Then, applying lemma 1, one obtains

$$\tilde{n}_t = \mathbb{E}_h\{\tilde{n}_t(h)\} + \frac{1}{2}\frac{\theta_w}{\theta_w - 1}\mathbb{E}_h\left\{\left(\frac{N_t(h)}{N^n}\right)^{\frac{\theta_w-1}{\theta_w}}\right\}^{-2}\mathbb{V}_h\left\{\left(\frac{N_t(h)}{N^n}\right)^{\frac{\theta_w-1}{\theta_w}}\right\} + \mathcal{O}(||\zeta||^3).$$

Then, notice that

$$\mathbb{V}_h\left\{\left(\frac{N_t(h)}{N^n}\right)^{\frac{\theta_w-1}{\theta_w}}\right\} = \mathbb{V}_h\left\{\exp\left[(1-\theta_w^{-1})\log\left(\frac{N_t(h)}{N^n}\right)\right]\right\}$$

so that, by applying lemma 2, one obtains

$$\mathbb{V}_h\left\{\left(\frac{N_t(h)}{N^n}\right)^{\frac{\theta_w-1}{\theta_w}}\right\} = (1-\theta_w^{-1})^2\exp\Big((1-\theta_w^{-1})\mathbb{E}_h\{\tilde{n}_t(h)\}\Big)^2\mathbb{V}_h\{\tilde{n}_t(h)\} + \mathcal{O}(||\zeta||^3).$$

Similarly

$$\mathbb{E}_h\left\{\left(\frac{N_t(h)}{N^n}\right)^{\frac{\theta_w-1}{\theta_w}}\right\} = \mathbb{E}_h\left\{\exp\left[(1-\theta_w^{-1})\tilde{n}_t(h)\right]\right\}$$

so that, by applying lemma 1 once more, one obtains

$$\mathbb{E}_h\left\{\left(\frac{N_t(h)}{N^n}\right)^{\frac{\theta_w-1}{\theta_w}}\right\} = \exp\left[(1-\theta_w^{-1})\mathbb{E}_h\{\tilde{n}_t(h)\}\right]\left(1+\frac{1}{2}(1-\theta_w^{-1})^2\mathbb{V}_h\{\tilde{n}_t(h)\}\right) + \mathcal{O}(||\xi||^3).$$

Then combining the previous results

$$\tilde{n}_t = \mathbb{E}_h\{\tilde{n}_t(h)\} + \frac{1}{2}\frac{1}{1-\theta_w^{-1}}\frac{(1-\theta_w^{-1})^2\mathbb{V}_h\{\tilde{n}_t(h)\}}{\left(1+\frac{1}{2}(1-\theta_w^{-1})^2\mathbb{V}_h\{\tilde{n}_t(h)\}\right)^2} + \mathcal{O}(||\zeta||^3).$$

It is convenient to define

$$\Delta_{h,t} \equiv \mathbb{V}_h\{\tilde{n}_t(h)\}$$

so that finally

$$\tilde{n}_t = \mathbb{E}_h\{\tilde{n}_t(h)\} + Q_{0,h} + \frac{1-\theta_w^{-1}}{2}Q_{1,h}(\Delta_{h,t} - \Delta_n) + \mathcal{O}(||\zeta||^3).$$

where we defined

$$Q_{0,h} = \frac{\frac{1-\theta_w^{-1}}{2}\Delta_n}{\left[1+\frac{1}{2}(1-\theta_w^{-1})^2\Delta_n\right]^2}$$

and

$$Q_{1,h} = \frac{1-\frac{1}{2}(1-\theta_w^{-1})^2\Delta_n}{\left[1+\frac{1}{2}(1-\theta_w^{-1})^2\Delta_n\right]^3}$$

Applying the same logic on output and defining

$$\Delta_{y,t} \equiv \mathbb{V}_f\{\tilde{y}_t(f)\}$$

one gets

$$\tilde{y}_{z,t} = \mathbb{E}_f\{\tilde{y}_{z,t}(f)\} + Q_{0,y} + \frac{1-\theta_p^{-1}}{2} Q_{1,y}(\Delta_{y,t} - \Delta_y) + \mathcal{O}(||\zeta||^3).$$

where we defined

$$Q_{0,y} = \frac{\frac{1-\theta_p^{-1}}{2}\Delta_y}{\left[1+\frac{1}{2}(1-\theta_p^{-1})^2\Delta_y\right]^2}$$

and

$$Q_{1,y} = \frac{1-\frac{1}{2}(1-\theta_p^{-1})^2\Delta_y}{\left[1+\frac{1}{2}(1-\theta_p^{-1})^2\Delta_y\right]^3}$$

Then recall from (**??**) and from the equilibrium on the market for aggregate labor that

$$N_t = \int_0^1 L_t(f)\mathrm{d}f = \int_0^1 Y_{z,t}(f)^{\phi}\mathrm{d}f$$

which implies

$$\frac{N_t}{N^n} = \int_0^1 \left(\frac{Y_{z,t}(f)}{Y_z^n}\right)^{\phi}\mathrm{d}f$$

where we used $N^n = (Y_z^n)^{\phi}$.

This relation rewrites

$$\tilde{n}_t = \log\left(\int_0^1 \left(\frac{Y_{z,t}(f)}{Y_z^n}\right)^{\phi}\mathrm{d}f\right)$$

This expression is of the form

$$\tilde{n}_t = \log\left(\mathbb{E}_f\left\{\left(\frac{Y_{z,t}(f)}{Y_z^n}\right)^{\phi}\right\}\right).$$

Using lemmas 1 and 2, we obtain the following three approximations

$$\tilde{n}_t = \mathbb{E}_f\{\phi(\tilde{y}_{z,t}(f) - z_t)\} + \frac{1}{2}\frac{\mathbb{V}_f\left\{\left(\frac{Y_{z,t}(f)}{Y_z^n}\right)^{\phi}\right\}}{\left(\mathbb{E}_f\left\{\left(\frac{Y_{z,t}(f)}{Y_z^n}\right)^{\phi}\right\}\right)^2} + \mathcal{O}(||\zeta||^3),$$

$$\mathbb{V}_f\left\{\left(\frac{Y_{z,t}(f)}{Y_z^n}\right)^{\phi}\right\} = \phi^2\left[\exp\left[\phi\mathbb{E}\{\tilde{y}_{z,t}(f)\}\right]\right]^2\mathbb{V}_f\{\tilde{y}_{z,t}(f)\} + \mathcal{O}(||\zeta||^3),$$

$$\mathbb{E}_f\left\{\left(\frac{y_{z,t}(f)}{y_z^n}\right)^{\phi}\right\} = \exp\left[\phi\mathbb{E}\{\tilde{y}_{z,t}(f)\}\right]\left(1+\frac{1}{2}\phi^2\mathbb{V}_f\{\tilde{y}_{z,t}(f)\}\right) + \mathcal{O}(||\zeta||^3).$$

Combining these expressions as before yields

$$\tilde{n}_t = \phi\mathbb{E}_f\{\tilde{y}_{z,t}(f)\} + \frac{1}{2}\phi^2\frac{\mathbb{V}_f\{\tilde{y}_{z,t}(f)\}}{\left(1+\frac{1}{2}\phi^2\mathbb{V}_f\{\tilde{y}_{z,t}(f)\}\right)^2} + \mathcal{O}(||\zeta||^3).$$

We finally obtain

$$\tilde{n}_t = \phi \mathbb{E}_f\{\tilde{y}_{z,t}(f)\} + P_{0,y} + \frac{1}{2}\phi^2 P_{1,y}(\Delta_{y,t} - \Delta_y) + \mathcal{O}(||\zeta||^3),$$

where we used

$$\frac{\mathbb{V}_f\{\tilde{y}_{z,t}(f)\}}{\left(1+\frac{1}{2}\phi^2\mathbb{V}_f\{\tilde{y}_{z,t}(f)\}\right)^2} = \frac{\Delta_y}{\left(1+\frac{1}{2}\phi^2\Delta_y\right)^2} + \frac{1-\frac{1}{2}\phi^2\Delta_y}{\left(1+\frac{1}{2}\phi^2\Delta_y\right)^3}(\Delta_{y,t} - \Delta_y) + \mathcal{O}(||\zeta||^3)$$

and defined

$$P_{0,y} = \frac{\frac{1}{2}\phi^2\Delta_y}{\left(1+\frac{1}{2}\phi^2\Delta_y\right)^2}$$

and

$$P_{1,y} = \frac{1-\frac{1}{2}\phi^2\Delta_y}{\left(1+\frac{1}{2}\phi^2\Delta_y\right)^3}$$

K.3 Aggregate Price and Wage Levels

The aggregate price index is

$$P_t^{1-\theta_p} = \left(\int_0^1 P_t(f)^{1-\theta_p}\mathrm{d}f\right)$$

and the aggregate wage index is

$$W_t^{1-\theta_w} = \left(\int_0^1 W_t(h)^{1-\theta_w}\mathrm{d}h\right).$$

From lemma 1 and the definitions of P_t and W_t, we obtain

$$p_t = \mathbb{E}_f\{p_t(f)\} + \frac{1}{2}\frac{1}{1-\theta_p}\frac{\mathbb{V}_f\{P_t(f)^{1-\theta_p}\}}{\mathbb{E}_f\{P_t(f)^{1-\theta_p}\}^2} + \mathcal{O}(||\zeta||^3),$$

and

$$w_t = \mathbb{E}_h\{w_t(h)\} + \frac{1}{2}\frac{1}{1-\theta_w}\frac{\mathbb{V}_h\{W_t(h)^{1-\theta_w}\}}{\mathbb{E}_h\{W_t(h)^{1-\theta_w}\}^2} + \mathcal{O}(||\zeta||^3).$$

Then, from lemma 2, we obtain

$$\begin{aligned}\mathbb{V}_f\{P_t(f)^{1-\theta_p}\} &= \mathbb{V}_f\{\exp[(1-\theta_p)p_t(f)]\} \\ &= (1-\theta_p)^2\exp[(1-\theta_p)\bar{p}_t]^2\Delta_{p,t} + \mathcal{O}(||\zeta||^3),\end{aligned}$$

and

$$\begin{aligned}\mathbb{V}_h\{W_t(h)^{1-\theta_w}\} &= \mathbb{V}_h\{\exp[(1-\theta_w)w_t(h)]\} \\ &= (1-\theta_w)^2\exp[(1-\theta_w)\bar{w}_t]^2\Delta_{w,t} + \mathcal{O}(||\zeta||^3),\end{aligned}$$

where we defined

$$\bar{p}_t = \mathbb{E}_f\{p_t(f)\}, \quad \bar{w}_t = \mathbb{E}_h\{w_t(h)\},$$

$$\Delta_{p,t} = \mathbb{V}_f\{p_t(f)\}, \quad \Delta_{w,t} = \mathbb{V}_h\{w_t(h)\}.$$

Applying lemma 1 once again, we obtain

$$\begin{aligned}\mathbb{E}_f\{P_t(f)^{1-\theta_p}\} &= \mathbb{E}_f\{\exp[(1-\theta_p)p_t(f)]\} \\ &= \exp[(1-\theta_p)\bar{p}_t]\left(1+\frac{1}{2}(1-\theta_p)^2\Delta_{p,t}\right)\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}_h\{W_t(h)^{1-\theta_w}\} &= \mathbb{E}_h\{\exp[(1-\theta_w)w_t(h)]\} \\ &= \exp[(1-\theta_w)\bar{w}_t]\left(1+\frac{1}{2}(1-\theta_w)^2\Delta_{w,t}\right)\end{aligned}$$

Combining these relations, we obtain

$$p_t = \bar{p}_t + \frac{1}{2}\frac{(1-\theta_p)\Delta_{p,t}}{\left[1+\frac{1}{2}(1-\theta_p)^2\Delta_{p,t}\right]^2} + \mathcal{O}(||\zeta||^3),$$

and

$$w_t = \bar{w}_t + \frac{1}{2}\frac{(1-\theta_w)\Delta_{w,t}}{\left[1+\frac{1}{2}(1-\theta_w)^2\Delta_{w,t}\right]^2} + \mathcal{O}(||\zeta||^3).$$

Thus

$$p_t = \bar{p}_t + Q_{0,p} + \frac{1-\theta_p}{2}Q_{1,p}(\Delta_{p,t}-\Delta_p) + \mathcal{O}(||\zeta||^3),$$

and

$$w_t = \bar{w}_t + Q_{0,w} + \frac{1-\theta_w}{2}Q_{1,w}(\Delta_{w,t}-\Delta_w) + \mathcal{O}(||\zeta||^3).$$

where we defined

$$Q_{0,p} = \frac{\frac{1-\theta_p}{2}\Delta_p}{\left[1+\frac{1}{2}(1-\theta_p)^2\Delta_p\right]^2}, \quad Q_{0,w} = \frac{\frac{1-\theta_w}{2}\Delta_w}{\left[1+\frac{1}{2}(1-\theta_w)^2\Delta_w\right]^2}$$

and

$$Q_{1,p} = \frac{1-\frac{1}{2}(1-\theta_p)^2\Delta_p}{\left[1+\frac{1}{2}(1-\theta_p)^2\Delta_p\right]^3}, \quad Q_{1,w} = \frac{1-\frac{1}{2}(1-\theta_w)^2\Delta_w}{\left[1+\frac{1}{2}(1-\theta_w)^2\Delta_w\right]^3}$$

Remark that the constant terms in the second-order approximation of the log-price index can be rewritten as

$$Q_{0,p} - \frac{1-\theta_p}{2}Q_{1,p}\Delta_p = \frac{1}{2}\frac{(1-\theta_p)^3\Delta_p^2}{\left[1+\frac{1}{2}(1-\theta_p)^2\Delta_p\right]^3}.$$

Finally, using the demand functions, one obtains

$$\tilde{y}_{z,t}(f) = -\theta_p[p_t(f)-p_t] + \tilde{y}_{z,t},$$

$$\tilde{n}_t(h) = -\theta_w[w_t(h) - w_t] + \tilde{n}_t,$$

from which we deduce that

$$\Delta_{y,t} = \theta_p^2 \Delta_{p,t}$$

and

$$\Delta_{h,t} = \theta_w^2 \Delta_{w,t}.$$

K.4 Price and Wage Dispersions

We now derive the law of motion of price dispersion. Notice that

$$\Delta_{p,t} = \mathbb{V}_f\{p_t(f) - \bar{p}_{t-1}\}$$

Immediate manipulations of the definition of the cross-sectional mean of log-prices yield

$$\bar{p}_t - \bar{p}_{t-1} = \alpha_p \gamma_p \pi_{t-1} + (1 - \alpha_p)[p_t^\star - \bar{p}_{t-1}]. \tag{K.1}$$

Then, the classic variance formula yields

$$\Delta_{p,t} = \mathbb{E}_f\{[p_t(f) - \bar{p}_{t-1}]^2\} - [\mathbb{E}_f\{p_t(f) - \bar{p}_{t-1}\}]^2$$

Using this, we obtain

$$\Delta_{p,t} = \alpha_p \mathbb{E}_f\{[p_{t-1}(f) - \bar{p}_{t-1} + \gamma_p \pi_{t-1}]^2\} + (1 - \alpha_p)[p_t^\star - \bar{p}_{t-1}]^2 - [\bar{p}_t - \bar{p}_{t-1}]^2$$

Notice that

$$\begin{aligned}
&(1 - \alpha_p)[p_t^\star - \bar{p}_{t-1}]^2 - [\bar{p}_t - \bar{p}_{t-1}]^2 \\
&= (1 - \alpha_p)\left[\frac{1}{1 - \alpha_p}(\bar{p}_t - \bar{p}_{t-1}) - \frac{\alpha_p}{1 - \alpha_p}\gamma_p \pi_t)\right]^2 - [\bar{p}_t - \bar{p}_{t-1}]^2 \\
&= \frac{\alpha_p}{1 - \alpha_p}[\bar{p}_t - \bar{p}_{t-1} - \gamma_p \pi_t]^2 - \alpha_p[\gamma_p \pi_t]^2
\end{aligned}$$

Using this in the above equation yields

$$\Delta_{p,t} = \alpha_p \mathbb{E}_f\{[p_{t-1}(f) - \bar{p}_{t-1} + \gamma_p \pi_t]^2\} - \alpha_p[\gamma_p \pi_t]^2 + \frac{\alpha_p}{1 - \alpha_p}[\bar{p}_t - \bar{p}_{t-1} - \gamma_p \pi_t]^2$$

Now, notice also that

$$\alpha_p \mathbb{E}_f\{[p_{t-1}(f) - \bar{p}_{t-1}]^2\} = \alpha_p \mathbb{E}_f\{[p_{t-1}(f) - \bar{p}_{t-1} + \gamma_p \pi_t]^2\} - \alpha_p[\gamma_p \pi_t]^2$$

It then follows that

$$\Delta_{p,t} = \alpha_p \mathbb{E}_f\{[p_{t-1}(f) - \bar{p}_{t-1}]^2\} + \frac{\alpha_p}{1 - \alpha_p}[\bar{p}_t - \bar{p}_{t-1} - \gamma_p \pi_t]^2$$

Hence
$$\Delta_{p,t} = \alpha_p \Delta_{p,t-1} + \frac{\alpha_p}{1-\alpha_p}[\bar{p}_t - \bar{p}_{t-1} - \gamma_p \pi_t]^2$$

Using
$$p_t = \bar{p}_t + Q_{0,p} + \frac{1-\theta_p}{2} Q_{1,p}(\Delta_{p,t} - \Delta_p) + \mathcal{O}(||\zeta||^3),$$
we obtain
$$\bar{p}_t - \bar{p}_{t-1} = \pi_t - \frac{1-\theta_p}{2} Q_{1,p}(\Delta_{p,t} - \Delta_{p,t-1}) + \mathcal{O}(||\zeta||^3).$$
Hence
$$\Delta_{p,t} = \alpha_p \Delta_{p,t-1} + \frac{\alpha_p}{1-\alpha_p}\left[\pi_t - \frac{1-\theta_p}{2} Q_{1,p}(\Delta_{p,t} - \Delta_{p,t-1}) - \gamma_p \pi_{t-1}\right]^2 + \mathcal{O}(||\zeta||^3).$$

The steady-state value of Δ_p is thus
$$\Delta_p = \frac{(1-\gamma_p)^2 \alpha_p}{(1-\alpha_p)^2} \pi^2$$

We obtain finally
$$\Delta_{p,t} = \alpha_p \Delta_{p,t-1} + \frac{\alpha_p}{1-\alpha_p}\left[(1-\gamma_p)\pi + \hat{\pi}_t - \gamma_p \hat{\pi}_{t-1} - \frac{1-\theta_p}{2} Q_{1,p}(\Delta_{p,t} - \Delta_{p,t-1})\right]^2 + \mathcal{O}(||\zeta||^3).$$

Unless π is itself treated as an expansion variable in the above approximation, we cannot claim that $\Delta_{p,t}$ is second-order.

We now derive the law of motion of wage dispersion. Following similar steps as for price dispersion, notice that
$$\Delta_{w,t} = \mathbb{V}_h\{w_t(h) - \bar{w}_{t-1}\}$$

Immediate manipulations of the definition of the cross-sectional mean of log-wages yield
$$\bar{w}_t - \bar{w}_{t-1} = \alpha_w(\gamma_z \mu_z + \gamma_w \pi_{t-1}) + (1-\alpha_w)[w_t^\star - \bar{w}_{t-1}]. \tag{K.2}$$

Then, the classic variance formula yields
$$\Delta_{w,t} = \mathbb{E}_h\{[w_t(h) - \bar{w}_{t-1}]^2\} - [\mathbb{E}_h\{w_t(h) - \bar{w}_{t-1}\}]^2$$

Using this, we obtain
$$\Delta_{w,t} = \alpha_w \mathbb{E}_h\{[w_{t-1}(h) - \bar{w}_{t-1} + \gamma_z \mu_z + \gamma_w \pi_{t-1}]^2\} + (1-\alpha_w)[w_t^\star - \bar{w}_{t-1}]^2 - [\bar{w}_t - \bar{w}_{t-1}]^2$$

Notice that
$$w_t^\star - \bar{w}_{t-1} = \frac{1}{1-\alpha_w}(\bar{w}_t - \bar{w}_{t-1}) - \frac{\alpha_w}{1-\alpha_w}[\gamma_z \mu_z + \gamma_w \pi_t]$$

so that

$$
\begin{aligned}
&(1-\alpha_w)[w_t^\star - \bar{w}_{t-1}]^2 - [\bar{w}_t - \bar{w}_{t-1}]^2 \\
&= (1-\alpha_w)\Big[\frac{1}{1-\alpha_w}(\bar{w}_t - \bar{w}_{t-1}) - \frac{\alpha_w}{1-\alpha_w}[\gamma_z\mu_z + \gamma_w\pi_t]\Big]^2 - [\bar{w}_t - \bar{w}_{t-1}]^2 \\
&= \frac{\alpha_w}{1-\alpha_w}\Big[\bar{w}_t - \bar{w}_{t-1} - [\gamma_z\mu_z + \gamma_w\pi_t]\Big]^2 - \alpha_w[\gamma_z\mu_z + \gamma_w\pi_t]^2
\end{aligned}
$$

Using this in the above equation yields

$$
\begin{aligned}
\Delta_{w,t} = \alpha_w \mathbb{E}_h\{[w_{t-1}(h) - \bar{w}_{t-1} + \gamma_z\mu_z + \gamma_w\pi_t]^2\} \\
- \alpha_w[\gamma_w \log(1+\pi_t)]^2 + \frac{\alpha_w}{1-\alpha_w}\Big[\bar{w}_t - \bar{w}_{t-1} - [\gamma_z\mu_z + \gamma_w\pi_t]\Big]^2
\end{aligned}
$$

Now, notice also that

$$
\alpha_w \mathbb{E}_h\{[w_{t-1}(h) - \bar{w}_{t-1}]^2\} = \alpha_w \mathbb{E}_h\{[w_{t-1}(h) - \bar{w}_{t-1} + \gamma_z\mu_z + \gamma_w\pi_t]^2\} - \alpha_w[\gamma_z\mu_z + \gamma_w\pi_t]^2
$$

It then follows that

$$
\Delta_{w,t} = \alpha_w \mathbb{E}_h\{[w_{t-1}(h) - \bar{w}_{t-1}]^2\} + \frac{\alpha_w}{1-\alpha_w}\Big[\bar{w}_t - \bar{w}_{t-1} - [\gamma_z\mu_z + \gamma_p\pi_t]\Big]^2
$$

Hence

$$
\Delta_{w,t} = \alpha_w \Delta_{w,t-1} + \frac{\alpha_w}{1-\alpha_w}\Big[\bar{w}_t - \bar{w}_{t-1} - [\gamma_z\mu_z + \gamma_w\pi_t]\Big]^2
$$

which, in turn, implies

$$
\Delta_{w,t} = \alpha_w \Delta_{w,t-1} + \frac{\alpha_w}{1-\alpha_p w}\left[\bar{w}_t - \bar{w}_{t-1} - \gamma_z\mu_z - \gamma_w\pi_{t-1}\right]^2.
$$

Using

$$
w_t = \bar{w}_t + Q_{0,w} + \frac{1-\theta_w}{2}Q_{1,w}(\Delta_{w,t} - \Delta_w) + \mathcal{O}(||\zeta||^3),
$$

we obtain

$$
\bar{w}_t - \bar{w}_{t-1} = \pi_{w,t} - \frac{1-\theta_w}{2}Q_{1,w}(\Delta_{w,t} - \Delta_{w,t-1}) + \mathcal{O}(||\zeta||^3).
$$

Hence

$$
\begin{aligned}
\Delta_{w,t} = \alpha_w \Delta_{w,t-1} + \frac{\alpha_w}{1-\alpha_w}\Bigg[\pi_{w,t} - \frac{1-\theta_w}{2}Q_{1,w}(\Delta_{w,t} - \Delta_{w,t-1}) \\
- \gamma_z\mu_z - \gamma_w\pi_{t-1}\Bigg]^2 + \mathcal{O}(||\zeta||^3).
\end{aligned}
$$

The steady-state value of Δ_w is thus

$$
\Delta_w = \frac{\alpha_w}{(1-\alpha_w)^2}[(1-\gamma_z)\mu_z + (1-\gamma_w)\pi]^2
$$

We obtain finally

$$\Delta_{w,t} = \alpha_w \Delta_{w,t-1} + \frac{\alpha_w}{1-\alpha_w}\Big[(1-\gamma_z)\mu_z + (1-\gamma_w)\pi + \hat{\pi}_{w,t} - \gamma_w \hat{\pi}_{t-1} \\ - \frac{1-\theta_w}{2} Q_{1,w}(\Delta_{w,t} - \Delta_{w,t-1})\Big]^2 + \mathcal{O}(||\zeta||^3).$$

Here, treating π as an expansion parameter will not suffice to ensure that wage dispersion $\Delta_{w,t}$ is second order because an additional constant term linked to average productivity shows up in the above equations. Henceforth, we will assume that μ_z is an expansion parameter.

Then, if we treat π as an expansion variable, precisely because the steady-state value of Δ_p is of second-order, many of the expressions previously derived considerably simplify. In particular, we now obtain

$$p_t = \bar{p}_t + \frac{1-\theta_p}{2}\Delta_{p,t} + \mathcal{O}(||\zeta, \pi||^3),$$

$$w_t = \bar{w}_t + \frac{1-\theta_w}{2}\Delta_{w,t} + \mathcal{O}(||\zeta, \pi||^3),).$$

Now, because Δ_y and Δ_n are proportional to Δ_p and Δ_w, respectively, and because Δ_p and Δ_w are both proportional to π^2, we also obtain

$$\tilde{n}_t = \mathbb{E}_h\{\tilde{n}_t(h)\} + \frac{1-\theta_w^{-1}}{2}\Delta_{h,t} + \mathcal{O}(||\zeta, \pi||^3),$$

$$\tilde{n}_t = \phi(\mathbb{E}_f\{\tilde{y}_{z,t}(f)\} - z_t) + \frac{1}{2}\phi^2\Delta_{y,t} + \mathcal{O}(||\zeta, \pi||^3),$$

$$\tilde{y}_t = \mathbb{E}_f\{\tilde{y}_t(f)\} + \frac{1-\theta_p^{-1}}{2}\Delta_{y,t} + \mathcal{O}(||\zeta, \pi||^3),).$$

Thus, for sufficiently small inflation rates, we obtain formulas resembling those derived in Woodford (2003).

Finally, price and wage dispersions rewrite

$$\Delta_{p,t} = \alpha_p \Delta_{p,t-1} + \frac{\alpha_p}{1-\alpha_p}\Big[(1-\gamma_p)\pi + \hat{\pi}_t - \gamma_p \hat{\pi}_{t-1}\Big]^2 + \mathcal{O}(||\zeta, \pi||^3),).$$

$$\Delta_{w,t} = \alpha_w \Delta_{w,t-1} + \frac{\alpha_w}{1-\alpha_w}\Big[(1-\gamma_z)\mu_z + (1-\gamma_w)\pi + \hat{\pi}_{w,t} - \gamma_w \hat{\pi}_{t-1}\Big]^2 + \mathcal{O}(||\zeta, \pi||^3),).$$

K.5 Combining the Results

Combining the previous results, we obtain

$$\int_0^1 \frac{\chi}{1+\nu} e^{\zeta_{h,t}} (N_t(h))^{1+\nu} \mathrm{d}h = \chi (N^n)^{1+\nu} \Big[\tilde{n}_t + \frac{1}{2}(1+\nu)\tilde{n}_t^2 + \tilde{n}_t \zeta_{h,t} + \frac{1}{2}(1+\nu\theta_w)\theta_w \Delta_{w,t} \Big] + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),).$$

In turn, we have

$$\tilde{n}_t = \phi \tilde{y}_t + \frac{1}{2}\phi[(\phi-1)\theta_p + 1]\theta_p \Delta_{p,t} + \mathcal{O}(||\zeta,\pi||^3),$$

so that

$$\int_0^1 \frac{\chi}{1+\nu} e^{\zeta_{h,t}} (N_t(h))^{1+\nu} \mathrm{d}h = \phi\chi (N^n)^{1+\nu} \Big[(\tilde{y}_t - z_t) + \frac{1}{2}(1+\nu)\phi \tilde{y}_t^2 + \tilde{y}_t \zeta_{h,t} + \frac{1}{2}[(\phi-1)\theta_p + 1]\theta_p \Delta_{p,t} + \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w \Delta_{w,t} \Big] + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),).$$

Then, using

$$(1-\Phi)\frac{1-\beta\eta}{1-\eta} = \phi\chi(N^n)^{1+\nu},$$

where we defined

$$1-\Phi \equiv \frac{1+\tau_w}{\mu_w}\frac{1+\tau_p}{\mu_p},$$

we obtain

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \int_0^1 \frac{\chi}{1+\nu} e^{\zeta_{h,t}} (N_t(h))^{1+\nu} \mathrm{d}h \right\} = (1-\Phi)\frac{1-\beta\eta}{1-\eta} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Big[\tilde{y}_t + \frac{1}{2}(1+\nu)\phi \tilde{y}_t^2 + \tilde{y}_t \zeta_{h,t} + \frac{1}{2}[(\phi-1)\theta_p + 1]\theta_p \Delta_{p,t} + \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w \Delta_{w,t} \Big] + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),).$$

Assuming the distortions are themselves negligible, this simplifies further to

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \int_0^1 \frac{\chi}{1+\nu} e^{\zeta_{h,t}} (N_t(h))^{1+\nu} \mathrm{d}h \right\} = \frac{1-\beta\eta}{1-\eta} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Big[(1-\Phi)\tilde{y}_t + \frac{1}{2}(1+\nu)\phi \tilde{y}_t^2 + \tilde{y}_t \zeta_{h,t} + \frac{1}{2}[(\phi-1)\theta_p + 1]\theta_p \Delta_{p,t} + \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w \Delta_{w,t} \Big] + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),).$$

We now deal with the first term in the utility function. To that end, notice that

$$\sum_{t=0}^{\infty} \beta^t a_{t-1} = a_{-1} + \beta \sum_{t=0}^{\infty} \beta^{t-1} a_{t-1} = a_{-1} + \beta \sum_{t=0}^{\infty} \beta^t a_t.$$

Using this trick, we obtain

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t e^{\zeta_{c,t}} \log(C_{z,t} - \eta C_{z,t-1} e^{-\zeta_{z,t}}) = \frac{1-\beta\eta}{1-\eta} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg[\tilde{y}_{z,t} - \frac{1}{2}[\varphi(1+\beta\eta^2) - 1]\tilde{y}_{z,t}^2 + \eta\varphi\tilde{y}_{z,t}\tilde{y}_{z,t-1}$$

$$+ \varphi\hat{g}_t\tilde{y}_{z,t} - \eta\varphi\zeta^*_{z,t}\tilde{y}_{z,t} \Bigg] + \text{t.i.p} + \mathcal{O}(||\zeta||^3),$$

where we defined

$$\varphi^{-1} \equiv (1-\beta\eta)(1-\eta),$$

$$\hat{g}_t = (1-\eta)(\zeta_{c,t} - \beta\eta\mathbb{E}_t\{\zeta_{c,t+1}\}),$$

so that

$$(1-\beta\eta)\varphi\hat{g}_t \equiv (\zeta_{c,t} - \beta\eta\mathbb{E}_t\{\zeta_{c,t+1}\}).$$

and

$$\zeta^*_{z,t} = \zeta_{z,t} - \beta\mathbb{E}_t\{\zeta_{z,t+1}\}$$

Combining terms, we obtain

$$\mathsf{U}_0 = \frac{1-\beta\eta}{1-\eta} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg[\Phi\tilde{y}_{z,t} - \frac{1}{2}[\varphi(1+\beta\eta^2) + \omega]\tilde{y}_{z,t}^2 + \eta\varphi\tilde{y}_{z,t}\tilde{y}_{z,t-1}$$

$$+ (\varphi\hat{g}_t - \zeta_{h,t} - \varphi\eta\zeta^*_{z,t})\tilde{y}_{z,t}$$

$$- \frac{1}{2}[(\phi-1)\theta_p + 1]\theta_p\Delta_{p,t} - \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w\Delta_{w,t} \Bigg] + \text{t.i.p} + \mathcal{O}(||\zeta||^3),$$

where, as defined earlier

$$\omega = (1+\nu)\phi - 1$$

Now, recall that

$$[\varphi(1+\beta\eta^2) + \omega]\hat{y}^n_{z,t} - \varphi\beta\eta\mathbb{E}_t\left\{\hat{y}^n_{z,t+1}\right\} - \varphi\eta\hat{y}^n_{z,t-1} = \varphi\hat{g}_t - \zeta_{h,t} - \varphi\eta\zeta^*_{z,t}$$

Using this above yields

$$\mathsf{U}_0 = \frac{1-\beta\eta}{1-\eta} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg[\Phi\tilde{y}_{z,t} - \frac{1}{2}[\varphi(1+\beta\eta^2) + \omega]\tilde{y}_{z,t}^2 + \eta\varphi\tilde{y}_{z,t}\tilde{y}_{z,t-1}$$

$$+ [\varphi(1+\beta\eta^2) + \omega]\hat{y}^n_{z,t}\tilde{y}_{z,t} - \varphi\beta\eta\hat{y}^n_{z,t+1}\tilde{y}_{z,t} - \varphi\eta\hat{y}^n_{z,t-1}\tilde{y}_{z,t}$$

$$- \frac{1}{2}[(\phi-1)\theta_p + 1]\theta_p\Delta_{p,t} - \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w\Delta_{w,t} \Bigg] + \text{t.i.p} + \mathcal{O}(||\zeta||^3),$$

$$\mathsf{U}_0 = \frac{1-\beta\eta}{1-\eta}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg[\Phi\tilde{y}_t - \frac{1}{2}[\varphi(1+\beta\eta^2)+\omega]\tilde{y}_t^2 + \eta\varphi\tilde{y}_t\tilde{y}_{t-1}$$

$$+ [\omega + \varphi(1+\beta\eta^2)]\hat{y}_t^n\tilde{y}_t - \varphi\beta\eta\hat{y}_{t+1}^n\tilde{y}_t - \varphi\eta\hat{y}_{t-1}^n\tilde{y}_t$$

$$- \frac{1}{2}[(\phi-1)\theta_p+1]\theta_p\Delta_{p,t} - \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w\Delta_{w,t}\Bigg] + \text{t.i.p} + \mathcal{O}(||\xi,\pi||^3)$$

To simplify this expression, we seek constant terms δ_0, δ and $x^\star$ such that

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg\{ -\frac{1}{2}\delta_0[(\tilde{y}_t - \hat{y}_t^n) - \delta(\tilde{y}_{t-1} - \hat{y}_{t-1}^n) - x^\star]^2\Bigg\}$$

$$= \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg[\Phi\tilde{y}_t - \frac{1}{2}[\varphi(1+\beta\eta^2)+\omega]\tilde{y}_t^2 + \eta\varphi\tilde{y}_t\tilde{y}_{t-1}$$

$$+ [\omega + \varphi(1+\beta\eta^2)]\tilde{y}_t\hat{y}_t^n - \varphi\beta\eta\tilde{y}_t\hat{y}_{t+1}^n - \varphi\eta\tilde{y}_t\hat{y}_{t-1}^n\Bigg] + \text{t.i.p}$$

Developing yields

$$-\frac{\delta_0}{2}\Big[(\tilde{y}_t - \hat{y}_t^n) - \delta(\tilde{y}_{t-1} - \hat{y}_{t-1}^n) - x^*\Big]^2$$

$$= -\frac{1}{2}\delta_0\tilde{y}_t^2 + \delta_0\tilde{y}_t\hat{y}_t^n + \delta_0\delta\tilde{y}_t\tilde{y}_{t-1} - \delta_0\delta\tilde{y}_t\hat{y}_{t-1}^n - \delta_0\delta\tilde{y}_{t-1}\hat{y}_t^n$$

$$-\frac{1}{2}\delta_0\delta^2\tilde{y}_{t-1}^2 + \delta_0\delta^2\tilde{y}_{t-1}\hat{y}_{t-1}^n + \delta_0(\tilde{y}_t - \delta\tilde{y}_{t-1})\hat{x}^* + \text{t.i.p}$$

Thus

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg\{ -\frac{\delta_0}{2}[(\tilde{y}_t - \hat{y}_t^n) - \delta(\tilde{y}_{t-1} - \hat{y}_{t-1}^n) - x^*]^2\Bigg\}$$

$$= \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg\{\delta_0(1-\beta\delta)\hat{x}^*\tilde{y}_t - \frac{1}{2}\delta_0(1+\beta\delta^2)\tilde{y}_t^2 + \delta_0\delta\tilde{y}_t\tilde{y}_{t-1}$$

$$+ \delta_0(1+\beta\delta^2)\tilde{y}_t\hat{y}_t^n - \delta_0\delta\tilde{y}_t\hat{y}_{t-1}^n - \delta_0\delta\beta\tilde{y}_t\hat{y}_{t+1}^n\Bigg\} + \text{t.i.p}$$

Identifying term by term, we obtain

$$\delta_0(1-\beta\delta)x^* = \Phi$$

$$\delta_0(1+\beta\delta^2) = [\omega + \varphi(1+\beta\eta^2)]$$

$$\delta_0\delta = \eta\varphi$$

Combining these relations, we obtain

$$\eta\delta^2 - \frac{\omega + \varphi)(1+\beta\eta^2)}{\beta\varphi}\delta + \eta\beta^{-1} = 0,$$

or equivalently

$$\mathbb{P}(\varkappa) = \beta^{-1}\varkappa^2 - \chi\varkappa + \eta^2 = 0,$$

where

$$\varkappa = \frac{\eta}{\delta},$$

$$\chi = \frac{\omega + \varphi(1+\beta\eta^2)}{\beta\varphi} > 0.$$

Notice that

$$\mathbb{P}(0) = \eta^2 > 0,$$

$$\mathbb{P}(1) = -\frac{\omega}{\beta\varphi} < 0$$

so that the two roots of $\mathbb{P}(\varkappa) = 0$ obey

$$0 < \varkappa_1 < 1 < \varkappa_2.$$

In the sequel, we focus on the larger root and define

$$\varkappa = \varkappa_2 = \frac{\beta}{2}\left(\chi + \sqrt{\chi^2 - 4\eta^2\beta^{-1}}\right) > 1.$$

Since $\delta = \eta/\varkappa$, we have

$$0 \leq \delta \leq \eta < 1.$$

Thus, given the obtained value for $\varkappa$, we can deduce δ from which we can compute δ_0.

We thus obtain

$$\mathsf{U}_0 = -\frac{1-\beta\eta}{1-\eta}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\left\{\frac{\delta_0}{2}[(\tilde{y}_t - \hat{y}_t^n) - \delta(\tilde{y}_{t-1} - \hat{y}_{t-1}^n) - x^*]^2\right.$$

$$\left. + \frac{1}{2}[(\phi-1)\theta_p + 1]\theta_p\Delta_{p,t} + \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w\Delta_{w,t}\right\} + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),)$$

The last step consists in expressing price and wage dispersions in terms of squared price and wage inflations.

Recall that

$$\Delta_{p,t} = \alpha_p\Delta_{p,t-1} + \frac{\alpha_p}{1-\alpha_p}\left[(1-\gamma_p)\pi + \hat{\pi}_t - \gamma_p\hat{\pi}_{t-1}\right]^2 + \mathcal{O}(||\zeta,\pi||^3),).$$

Iterating backward on this formula yields

$$\Delta_{p,t} = \frac{\alpha_p}{1-\alpha_p}\sum_{s=0}^{t}\alpha_p^{t-s}[(1-\gamma_p)\pi + \hat{\pi}_s - \gamma_p\hat{\pi}_{s-1}]^2 + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),).$$

It follows that

$$\sum_{t=0}^{\infty}\beta^t\Delta_{p,t} = \frac{\alpha_p}{(1-\alpha_p)(1-\beta\alpha_p)}\sum_{t=0}^{\infty}\beta^t[(1-\gamma_p)\pi + \hat{\pi}_t - \gamma_p\hat{\pi}_{t-1}]^2 + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),).$$

and by the same line of reasoning

$$\sum_{t=0}^{\infty}\beta^t\Delta_{w,t} = \frac{\alpha_w}{(1-\alpha_w)(1-\beta\alpha_w)}\sum_{t=0}^{\infty}\beta^t[(1-\gamma_z)\mu_z + (1-\gamma_w)\pi + \hat{\pi}_{w,t} - \gamma_w\hat{\pi}_{t-1}]^2 + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),).$$

Thus, defining

$$\lambda_y \equiv \delta_0$$

$$\lambda_p \equiv \frac{\alpha_p\theta_p[(\phi-1)\theta_p+1]}{(1-\alpha_p)(1-\beta\alpha_p)}$$

$$\lambda_w \equiv \frac{\alpha_w\phi^{-1}\theta_w(1+\nu\theta_w)}{(1-\alpha_w)(1-\beta\alpha_w)}$$

The second order approximations to welfare rewrites

$$\begin{aligned}\mathsf{U}_0 = -\frac{1}{2}\frac{1-\beta\eta}{1-\eta}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\Big\{&\lambda_y[x_t - \delta x_{t-1} + (1-\delta)\bar{x} - x^*]^2 \\ &+ \lambda_p[(1-\gamma_p)\pi + \hat{\pi}_t - \gamma_p\hat{\pi}_{t-1}]^2 \\ &+ \lambda_w[(1-\gamma_z)\mu_z + (1-\gamma_w)\pi + \hat{\pi}_{w,t} - \gamma_w\hat{\pi}_{t-1}]^2\Big\} + \text{t.i.p} + \mathcal{O}(||\zeta,\pi||^3),\end{aligned}$$

where we defined

$$x_t \equiv \hat{y}_t - \hat{y}_t^n$$

$$\bar{x} \equiv \log\left(\frac{Y_z}{Y_z^n}\right).$$

L Welfare cost of Inflation

We start the analysis by restatting the welfare function

$$W(\pi) = \mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\left[\mathrm{e}^{\zeta_{g,t}}\log(\hat{C}_t(\pi) - \eta\hat{C}_{t-1}(\pi)\mathrm{e}^{-\zeta_{z,t}}) - \frac{\chi}{1+\nu}\int_0^1 N_t(\pi,h)^{1+\nu}dh\right] + \Psi_0(\mu_z,\zeta_z)$$

where we made explicit the dependence on π. Importantly, the welfare function is stated in terms of detrended consumption. The term Ψ_0 captures the part of welfare that depends exclusively on μ_z and $\zeta_{z,t}$ and is not affected by changes in the inflation target.

Let us now consider a deterministic economy in which labor supply is held constant at the undistorted steady-state level N_n and in which agents consume the constant level of detrended consumption $\hat{C}(\pi)$. We seek to find the $\hat{C}(\pi)$ such that this deterministic economy enjoys the same level of welfare as above. Thus

$$W(\pi) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\log\left((1-\eta)\hat{C}(\pi)\right) - \frac{\chi}{1+\nu} N_n^{1+\nu} \right] + \Psi_0(\mu_z, 0)$$

Direct manipulations yield

$$W(\pi) = \frac{1}{1-\beta} \left[\log\left((1-\eta)\hat{C}(\pi)\right) - \frac{\chi}{1+\nu} N_n^{1+\nu} \right] + \Psi_0(\mu_z, 0)$$

Consider now an economy with $\pi = \pi^\star$ and another one with $\pi = \tilde{\pi} \neq \pi^\star$. Imagine that in the latter, consumer are compensated in consumption units in such a way that they are as well off with $\tilde{\pi}$ as with $\pi^\star$. Let $1+\varphi$ denote this percentage increase in consumption. Thus φ is such that

$$\begin{aligned} W(\pi^\star) &= \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\log\left((1+\varphi)(1-\eta)\hat{C}(\pi)\right) - \frac{\chi}{1+\nu} N_n^{1+\nu} \right] + \Psi_0(\mu_z, 0) \\ &= \frac{\log(1+\varphi)}{1-\beta} + W(\pi) \end{aligned}$$

It then follows that

$$\varphi = \exp\{(1-\beta)[W(\pi^\star) - W(\pi)]\} - 1.$$

In practice, wefare is approximated to second order as

$$\begin{aligned} W(\pi) = -\frac{1}{2}\frac{1-\beta\eta}{1-\eta} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t & \left\{ \lambda_y [x_t - \delta x_{t-1} + (1-\delta)\bar{x}] + \lambda_p [(1-\gamma_p)\pi + \hat{\pi}_t - \gamma_p \hat{\pi}_{t-1}] \right. \\ & \left. + \lambda_w [(1-\gamma_z)\mu_z + (1-\gamma_w)\pi + \hat{\pi}_{w,t} - \gamma_p \hat{\pi}_{w,t-1}] \right\} + \Theta_0 \end{aligned}$$

where Θ_0 does not depend on π and is thus common across the economy with $\pi = \pi^\star$ and another one with $\pi = \tilde{\pi} \neq \pi^\star$.

References

Adam, K. and H. Weber (2017): "Optimal trend inflation," Discussion Papers 25/2017, Deutsche Bundesbank, Research Centre.

Amano, R., K. Moran, S. Murchison, and A. Rennison (2009): "Trend Inflation, Wage and Price Rigidities, and Productivity Growth," *Journal of Monetary Economics*, 56, 353–364.

Ascari, G., L. Phaneuf, and E. Sims (2015): "On the Welfare and Cyclical Implications of Moderate Trend Inflation," Working Paper 21392, National Bureau of Economic Research.

Ascari, G. and A. M. Sbordone (2014): "The Macroeconomics of Trend Inflation," *Journal of Economic Literature*, 52, 679–739.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. V. Reenen (2017): "The Fall of the Labor Share and the Rise of Superstar Firms," Working Paper 23396, NBER.

Ball, L. M. (2014): "The Case for a Long-Run Inflation Target of Four Percent," IMF Working Papers 14/92, International Monetary Fund.

Basu, S. and J. G. Fernald (1997): "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy*, 105, 249–283.

Bernanke, B. (2016): "Modifying the Fed's Policy Framework: Does a Higher Inflation Target Beat Negative Interest Rates?" Blog post, September 13.

Bernanke, B. S., T. Laubach, F. Mishkin, and A. Posen (1999): *Inflation Targeting. Lessons form the international experience*, Economics Books, Princeton University Press.

Bilbiie, F. O., I. Fujiwara, and F. Ghironi (2014): "Optimal Monetary Policy with Endogenous Entry and Product Variety," *Journal of Monetary Economics*, 64, 1–20.

Blanchard, O., G. Dell'Ariccia, and P. Mauro (2010): "Rethinking Macroeconomic Policy," *Journal of Money, Credit and Banking*, 42, 199–215.

Blanco, A. (2016): "Optimal Inflation Target in an Economy with Menu Costs and Zero Lower Bound," mimeo, University of Michigan.

Bodenstein, M., C. J. Erceg, and L. Guerrieri (2009): "The Effects of Foreign Shocks when Interest Rates are at Zero," International Finance Discussion Papers 983, Board of Governors of the Federal Reserve System (U.S.).

Burstein, A. and C. Hellwig (2008): "Welfare Costs of Inflation in a Menu Cost Model," *American Economic Review*, 98, 438–43.

CABALLERO, R. AND E. FARHI (2015): "The Safety Trap," Working paper, Harvard University OpenScholar.

CARLSSON, M. AND A. WESTERMARK (2016): "Labor Market Frictions and Optimal Steady-State Inflation," *Journal of Monetary Economics*, 78, 67–79.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

CHUNG, H., J.-P. LAFORTE, D. REIFSCHNEIDER, AND J. C. WILLIAMS (2012): "Have We Underestimated the Likelihood and Severity of Zero Lower Bound Events?" *Journal of Money, Credit and Banking*, 44, 47–82.

COENEN, G. (2003): "Zero Lower Bound: Is it a Problem in the Euro Area?" Working Paper Series 269, European Central Bank.

COENEN, G., A. ORPHANIDES, AND V. WIELAND (2004): "Price Stability and Monetary Policy Effectiveness when Nominal Interest Rates are Bounded at Zero," *The B.E. Journal of Macroeconomics*, 4, 1–25.

COIBION, O., Y. GORODNICHENKO, AND J. WIELAND (2012): "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?" *Review of Economic Studies*, 79, 1371–1406.

CORREIA, I., E. FARHI, J. P. NICOLINI, AND P. TELES (2013): "Unconventional Fiscal Policy at the Zero Bound," *American Economic Review*, 103, 1172–1211.

DE LOECKER, J. AND J. EECKHOUT (2017): "The Rise of Market Power and the Macroeconomic Implications," Working Papers 23687, NBER.

DEBORTOLI, D., J. GALÍ, AND L. GAMBETTI (2018): "On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint," Working Papers 1013, Barcelona GSE.

DEL NEGRO, M., D. GIANNONE, M. P. GIANNONI, AND A. TAMBALOTTI (2018): "Global Trends in Interest Rates," Working Papers 25039, National Bureau of Economic Research.

DORDAL-I-CARRERAS, M., O. COIBION, Y. GORODNICHENKO, AND J. WIELAND (2016): "Infrequent but Long-Lived Zero-Bound Episodes and the Optimal Rate of Inflation," Working Papers 22510, National Bureau of Economic Research.

EGGERTSSON, G. B., N. R. MEHROTRA, AND J. A. ROBBINS (2017): "A Model of Secular Stagnation: Theory and Quantitative Evaluation," Working Papers 23093, National Bureau of Economic Research.

ERCEG, C. J., D. W. HENDERSON, AND A. T. LEVIN (2000): "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46, 281–313.

FAGAN, G., J. HENRY, AND R. MESTRE (2001): "An Area-Wide Model (AWM) for the Euro Area," Working Paper Series 0042, European Central Bank.

FÈVE, P., J. MATHERON, AND J.-G. SAHUC (2010): "Disinflation Shocks in the Eurozone: A DSGE Perspective," *Journal of Money, Credit and Banking*, 42, 289–323.

FRIEDMAN, M. (1969): *The Optimum Quantity of Money and Other Essays*, Chicago, IL: Aldine Press, chap. The Optimum Quantity of Money.

GALÍ, J. (2015): *Monetary Policy, Inflation, and the Business Cycle. An Introduction to the New Keynesian Framework*, Princeton, NJ: Princeton University Press, chap. 5, second ed.

GOLOSOV, M. AND R. E. J. LUCAS (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199.

GORDON, R. J. (2015): "Secular Stagnation: A Supply-Side View," *American Economic Review*, 105.

GUERRIERI, L. AND M. IACOVIELLO (2015): "OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily," *Journal of Monetary Economics*, 70, 22–38.

GUST, C., E. HERBST, D. LÓPEZ-SALIDO, AND M. E. SMITH (2017): "The Empirical Implications of the Interest-Rate Lower Bound," *American Economic Review*, 107, 1971–2006.

HALL, R. E. (2018): "New Evidence on the Markup of prices over Marginal Costs and the Role of Mega-Firms in the US Economy," Working Paper 24574, NBER.

HILLS, T. S., T. NAKATA, AND S. SCHMIDT (2016): "The Risky Steady State and the Interest Rate Lower Bound," Finance and Economics Discussion Series 2016-9, Board of Governors of the Federal Reserve System.

HOLSTON, K., T. LAUBACH, AND J. C. WILLIAMS (2017): "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, 108, 59–75.

KHAN, A., R. G. KING, AND A. L. WOLMAN (2003): "Optimal Monetary Policy," *Review of Economic Studies*, 70, 825–860.

KILEY, M. T. AND J. M. ROBERTS (2017): "Monetary Policy in a Low Interest Rate World," Finance and Economics Discussion Series 2017-080, Board of Governors of the Federal Reserve System (U.S.).

LAUBACH, T. AND J. C. WILLIAMS (2016): "Measuring the Natural Rate of Interest Redux," Finance and Economics Discussion Series 2016-11, Board of Governors of the Federal Reserve System (U.S.).

LEPETIT, A. (2017): "The Optimal Inflation Rate with Discount Factor Heterogeneity," Working Papers hal-01527816, HAL.

LINDÉ, J., J. MAIH, AND R. WOUTERS (2017): "Estimation of Operational Macromodels at the Zero Lower Bound," Mimeo.

NAKAMURA, E., J. STEINSSON, P. SUN, AND D. VILLAR (2016): "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation," Working papers, NBER.

REIFSCHNEIDER, D. L. (2016): "Gauging the Ability of the FOMC to Respond to Future Recessions," Finance and Economics Discussion Series 2016-068, Board of Governors of the Federal Reserve System (U.S.).

SCHMITT-GROHÉ, S. AND M. URIBE (2010): "The Optimal Rate of Inflation," in *Handbook of Monetary Economics*, ed. by B. M. Friedman and M. Woodford, Elsevier, vol. 3 of *Handbook of Monetary Economics*, chap. 13, 653–722.

SUMMERS, L. (2014): "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound," *Business Economics*, 49, 65–73.

TRAINA, J. (2015): "Is Aggregate Market power Increasing? Production Trends Using Financial Statements," Working Paper 17, Chicago Booth, Stigler Center for the Study of the Economy and the State.

WILLIAMS, J. C. (2016): "Monetary Policy in a Low R-star World," *FRBSF Economic Letter*.

WOODFORD, M. (2003): *Interest and Prices*, Princeton, NJ: Princeton University Press.