Working Papers

FEDERAL RESERVE BANK OF BOSTON

No. 19-18

The Optimal Inflation Target and the Natural Rate of Interest

Philippe Andrade, Jordi Galí, Hervé Le Bihan, and Julien Matheron

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 lower steady-state real interest rate increases the probability of hitting the lower bound. That effect can be counteracted by an increase in the inflation target, but the resulting higher steady-state inflation has a welfare cost in and of itself. We use an estimated DSGE model to quantify that tradeoff and determine the implied optimal inflation target, conditional on the monetary policy rule in place before the financial crisis. The relation between the steady-state real interest rate and the optimal inflation target is downward sloping. While the increase in the optimal inflation rate is in general smaller than the decline in the steady-state real interest rate, in the currently empirically relevant region the slope of the relation is found to be close to –1. That slope is robust to allowing for parameter uncertainty. Under "make-up" strategies such as price level targeting, the required increase in the optimal inflation target under a lower steady-state real interest rate is, however, much smaller.

JEL Classifications: E31, E52, E58

Keywords: inflation target, effective lower bound, natural interest rate, steady-state real interest rate

---

Philippe Andrade is a senior economist in the research department at the Federal Reserve Bank of Boston; his email is Philippe.andrade@bos.frb.org. Jordi Galí is the director and senior researcher at the Center for Research in International Economics; his email is jgali@crei.cat. Hervé Le Bihan is the deputy director of the Monetary and Financial Studies Directorate at the Banque de France; his email is herve.lebihan@banque-france.fr. Julien Matheron is a senior research advisor for the directorate; his email is julien.matheron@banque-france.fr.

The authors thank Janice Eberly; James Stock; and their discussants, Mike Kiley and Luminita Stevens, for their comments and suggestions. They also thank Olivier Blanchard, Olivier Coibion, Isabel Correia, Kevin Lansing, Argia Sbordone, Raf Wouters, and Johannes Wieland as well as participants at several conferences and seminars for comments and discussions on an earlier draft of this paper.

This paper was originally prepared for presentation at a Brookings Institution conference held September 5–6, 2019, and for publication in the Fall 2019 edition of the *Brookings Papers on Economic Activity*. The views expressed herein are those of the authors and do not indicate concurrence by the Banque de France, the Eurosystem, the Federal Reserve Bank of Boston, the principals of the Board of Governors, or the Federal Reserve System. This paper, which may be revised, is available on the website of the Boston Fed at http://www.bostonfed.org/economic/wp/index.htm.

This version: October 2019 **https://doi.org/10.29412/res.wp.2019.18**

1 Introduction

A recent but sizable literature points to a permanent—or, at least very persistent—decline in the "natural" rate of interest in advanced economies (Holston et al. 2017; Laubach and Williams 2016). Various likely sources of that decline are discussed, including a lower trend growth rate of productivity (Gordon, 2015), demographic factors (Eggertsson et al. 2017), and an enhanced preference for safe and liquid assets (Caballero and Farhi 2018; Del Negro et al. 2017; 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 and bringing about more frequent (and potentially protracted) episodes of recession and below-target inflation. The low interest rate environment is a key factor behind the Federal Reserve's current review of its monetary policy framework (see Clarida 2019; Fuhrer et al. 2018).[1]

In the face of that risk, several prominent economists forcefully argue in favor of raising the inflation target (see, among others, Ball 2014; Blanchard et al. 2010; and, with qualifications, 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 as, all else being equal, a higher inflation target increases the steady-state nominal interest rate and reduces the ZLB incidence. But the answer to the practical question of how much the target should 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 Bernanke (2016), among others, argues. The answer to this question thus requires us 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 is emphasized in the recent literature, such assessment has received surprisingly little attention.

This paper contributes to this debate by asking four 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? Third, how does parameter uncertainty affect the $(r^\star, \pi^\star)$ curve? Fourth, to what extent do the strategy and rules followed by the central bank alter the relation between $r^\star$ and $\pi^\star$? We focus on the US economy, but the issues we investigate equally apply to other advanced economies—in particular the euro area—because the decline in $r^\star$ appears to be a global phenomenon (see Brand et al. 2018; Del Negro et al. 2018; Rachel and Summers 2019).[2]

We provide answers to these questions using a structural, empirically estimated macroeconomic model. Our main findings can be summarized as follows: (1) the relation between $r^\star$ and $\pi^\star$ is downward sloping, *but not necessarily, in general, one for one*; (2) in the vicinity of the pre-crisis values for $r^\star$, the slope of

[1] Note that the numerical value of the inflation target is not part of that review.

[2] We provide a comparable analysis for the euro area in a work in progress.

the $(r^\star, \pi^\star)$ locus is close to -1, though slightly less in absolute value; the relation is largely robust to the underlying source of variation in $r^\star$ for a plausible range of $r^\star$ values; (3) the slope of the $(r^\star, \pi^\star)$ locus remains close to -1 when the central bank is uncertain about the parameters of the model characterizing the economy, including $r^\star$; and (4) the slope of the curve is also robust to various alterations of the monetary policy rule, albeit not to considering rules such as price level targeting, which consist in **credibly** committing to making up for past deviations from the inflation target.

Our results are obtained from extensive simulations of a New Keynesian DSGE model estimated for the US over a Great Moderation sample.[3] The framework features (1) price stickiness and partial indexation of prices to trend inflation, (2) wage stickiness and partial indexation of wages to both inflation and productivity, and (3) a ZLB constraint on the nominal interest rate. The first two features imply the presence of potentially substantial costs associated with non-zero steady-state 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 tradeoff between the costs attached to the probability of hitting the ZLB and the costs induced by a positive steady-state inflation rate for a given monetary policy strategy. In the baseline, monetary policy follows an inertial interest rate rule estimated using pre-crisis data. Importantly, the specification of the policy rule implies that interest rates remain “low for long” after the end of a ZLB episode and that private agents expect the central bank to follow this rule. This implies that some monetary accommodation can be provided despite the ZLB constraint. This specification can thus be seen as a parsimonious way to factor in the effects of non-conventional policies that the Fed implemented during the ZLB period in our analysis.

According to our simulations, the optimal inflation target obtained when the policymaker is assumed to know the economy’s parameters with certainty (and taken to correspond to the mean of the posterior distribution) is around 2% (in annual terms). This result is obtained in an environment with a relatively low, 6% probability of hitting the ZLB when the target for the inflation rate is set at the historical mean of inflation, and given the size of the shocks estimated on our Great Moderation sample. Our simulations also show that a 100 basis point drop of $r^\star$ from its estimated 2.5% pre-crisis level 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. This optimal reaction limits the increase in the probability of hitting the ZLB to a mere half of a percentage point.

This optimal adjustment is robust to a set of alternative scenarios. It does not depend on the cause (productivity, demography, or safe assets) underlying such a structural decline. It also remains close to one for one when we consider alternative assumptions regarding key structural parameters: structural shocks with higher variance, alternative markups in the goods and labor markets, and a different degree

[3] In a work in progress, we show that very similar results are obtained in a model estimated with euro area data.

of indexation to trend inflation. Strikingly, while the level of the locus can be significantly affected by those changes—these alternative scenarios call for an optimal inflation target that would have been close to or above 2% before the crisis—overall the slope of the $(r^\star, \pi^\star)$ relation remains close to -1 in the vicinity of the pre-crisis parameter region.

More generally, one may wonder how a central bank should adjust its optimal inflation target when it is uncertain about the true values of structural parameters describing the economy. A noticeable feature of our approach is that we perform a full-blown Bayesian estimation of the model. 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 policymaker, including uncertainty with regard to the determinants of the steady-state real interest rate. When that parameter uncertainty is allowed for, the optimal inflation target value increases significantly, to 2.40%. The higher optimal target under parameter uncertainty reflects the fact that the loss function is asymmetric, so that choosing an inflation target that is below the optimal one is more costly than choosing an inflation target that is above it. In spite of the higher level, 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 $r^\star$ by 100 basis points.

Finally, we study how potential changes in the monetary policy rule or strategy affect the $(r^\star, \pi^\star)$ relation. We consider a number of different cases: (1) defining the inflation target in terms of average realized inflation as opposed to a parameter in the rule, (2) a central bank constrained by an effective lower bound on the policy rate that can be below zero, (3) a central bank with a lower or higher smoothing parameter in the interest-rate policy rule, (4) a central bank with a smoothing component that involves the lagged actual policy rate instead of the lagged shadow rate, and (5) a central bank targeting the price level rather than the inflation rate. All these changes have an impact on the level of $\pi^\star$ for any given level of $r^\star$. Yet, only in the case of a higher interest rate smoothing and price level targeting do we find a noticeable change in the slope of the $(r^\star, \pi^\star)$ relation. In these two cases, the relation is much less steep, illustrating the strength of "make-up" strategies to overturn the ZLB. However, as we discuss in the conclusion, an important caveat is that this result is obtained under the joint assumption of rational expectations, perfect information, and full credibility of the commitment.

The remainder of the paper is organized as follows. Section 1.1 reviews the related literature. Section 2 describes our baseline model. Section 3 discusses 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 under the baseline estimates as well as for a set of alternative parameters. Section 5 presents and discusses this locus under parameter uncertainty. Section 6 investigates the $(r^\star, \pi^\star)$ relation under alternative monetary policy rules and strategies. Finally, Section 7 summarizes and concludes.

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 setups that account for the ZLB. However, the analyses in Coibion et al. (2012) assume a constant steady-state natural rate of interest, so a key difference is our focus on eliciting the relation between the steady-state real interest rate and optimal inflation. Other differences are that (1) we estimate, rather than calibrate, the model, and (2) 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 assesses the probability that the US economy hit the ZLB for a given inflation target. Interestingly, our own assessment of this pre-crisis ZLB incidence falls in the ballpark of available estimates, including those in Chung et al. (2012). As we show, when the inflation target is not adjusted, but post–Great Moderation shocks are allowed, we also get "post-crisis" probabilities of hitting the ZLB that are comparable to the ones obtained in recent related studies, such as Chung et al. (2019).

In the New Keynesian setup that we consider, agents have rational expectations and make decisions that are forward looking: They fully understand that the central bank's inability to lower the policy further will lead to a deflation that magnifies the contractionary demand shocks responsible for driving the economy to the ZLB in the first place. One concern may be that this framework makes the ZLB too destabilizing, hence overweighting the benefits of a positive inflation target. However, this is partially offset by the fact that rational expectations and forward looking decisions also make the "lower for longer" monetary policies that we consider at the end of the trap very effective, which limits the length and width of ZLB episodes. Chung et al. (2019) illustrate that ZLB episodes can also be very costly in setups featuring agents that are less forward looking, such as in the FRB/US model.

Our assessment of the (welfare) cost of inflation also critically relies on our assumptions of a Calvo mechanism for price and wage setting. Among the recent papers examining the ZLB, Blanco (2016) studies optimal inflation in a state-dependent pricing model, that is, a "menu cost" model. In this setup, optimal inflation is typically positive and higher than it would be with time-dependent pricing. Indeed, as in our analysis, positive inflation edges the economy against detrimental effects of the ZLB.[4] In addition, as shown by Nakamura et al. (2018), the presence of state-dependent pricing weakens considerably the positive relationship between inflation and price dispersion, thus reducing the costs of inflation. Nakamura et al. (2018) further argue that menu costs are a more plausible mechanism for pricing frictions.[5] Two

[4]See Burstein and Hellwig (2008) for a similar exercise under menu costs without the ZLB, which leads to negative optimal inflation rate.

[5]They document that the cross-sector dispersion in the size of price changes is similar in the current, low inflation period to what it was in the high inflation period of the late 1970s. If Calvo was the relevant pricing frictions, the dispersion in size of price changes should have been much larger in the high inflation period than today.

points are, however, worth making. First, in the range of values for the inflation target that we consider, the difference between the welfare cost in a Calvo model and in a menu cost model is less dramatic than with a 10 percent or higher inflation rate, as documented, for example, by Nakamura et al. (2018). Second, most recent empirical analyses of price setting show that there is a mass of small price changes in the data that cannot be rationalized by a menu cost model. To fit the microdata, much of this recent literature typically introduces a random opportunity of price change, hence a Calvo component, in the menu cost model (see, for example, Alvarez et al. 2016). In such an augmented menu cost model, the distinction with the assessment taken from the Calvo model is bound to be attenuated.

Our paper is also connected to the voluminous literature on monetary policy under uncertainty (see, for example, Levin et al. 2006; Williams 2013), although to our knowledge this literature does not investigate the impact of uncertainty on the determination of the optimal inflation target.

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. (2018). 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 (2019) show that, even without any ZLB concern, optimal inflation might be positive in the context of a model with heterogeneous firms and systematic firm-level productivity trends. Finally, Lepetit (2018) shows that optimal inflation can be different from zero when profits and utility flows are discounted at different rates, as is generally the case in overlapping-generation models. In a parameterized example of the latter he shows the optimal steady-state inflation is significantly above zero.

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, 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). Also, all else being equal, price inflation induces (nominal) wage inflation, which in turn triggers inefficient wage dispersion in the presence of staggered wage setting. Partial indexation also magnifies the costs of non-zero price (or wage) inflation as compared with a setup in which price and wages mechanically catch up with trend inflation (Ascari and Sbordone 2014). 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 with 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 and the associated macroeconomic volatility should decline with the average rate of inflation.

Overall, the model we use, and the implied tradeoff 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.[6]

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 for 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.

[6] In their robustness analysis, Coibion et al. (2012) consider downward nominal wage rigidity, which involves mechanisms that are different from those used with Calvo-style rigidities.

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 the final good produced at t, $Y_t(f)$ is the input of intermediate good f, and θ_p is 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 setup is that the probability of price reoptimization is assumed to be invariant, with respect to the long run inflation rate. Drawing from the logic of menu cost models, the Calvo parameter of price stickiness could be expected to endogenously decrease when trend inflation rises. However, in the range of values for trend inflation that we will consider, available microeconomic 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 firm 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 subsidy paid to firms and financed via a lump-sum tax on households, and $Y_{t,t+s}$ is the demand function of a monopolistic firm that last revised its price at t faces at $t+s$; the firm obeys

$$Y_{t,t+s} = \left(\frac{V^p_{t,t+s} P^\star_t}{P_{t+s}}\right)^{-\theta_p} Y_{t+s},$$

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

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

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 in order 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.

As with 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 here is also 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 that 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 in order 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; that is, $\hat{\imath}_t \equiv i_t - i$. 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 efficient 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.

Importantly, we interpret π 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 purposes, 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 of deviations from steady state

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

The rule effectively implemented is given by:

$$\hat{\imath}_t = \max\{\hat{\imath}^n_t,\ -(\mu_z + \rho + \pi)\}, \tag{5}$$

where

$$\hat{\imath}^n_t = \rho_i \hat{\imath}^n_{t-1} + (1-\rho_i)\left(a_\pi \hat{\pi}_t + a_y \hat{x}_t\right) + \zeta_{R,t}, \tag{6}$$

with i^n_t denoting the shadow or notional rate, that is, the one that would be effective in the absence of the ZLB constraint. Thus the lagged rate that matters is the lagged notional interest rate, rather than the lagged actual rate. In making that assumption we follow Coibion et al. (2012) and a large share of the recent literature.

Before we proceed, several remarks are in order. First, note that realized inflation might be on average below the target π as a consequence of ZLB episodes; that is, $\mathbb{E}\{\pi_t\} < \pi$. 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 with no ZLB constraint.[7]

Second, we assume the central bank policy is characterized by simple interest rate rule rather than a Ramsey-type fully optimal policy of the type studied by, for example, Khan et al. (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. Moreover, two features in our setup, the inertia in the monetary policy rule, as well as the use of a lagged notional rate rather than a lagged actual rate, render the policy more persistent and thus closer to a Ramsey-like fully optimal interest rate rule. In particular, the dependence on the lagged notional rate $\hat{\imath}^n_t$ results in the nominal interest rate $\hat{\imath}_t$ being "lower for longer" in the aftermath of ZLB episodes (as $\hat{\imath}^n_t$ will stay negative for a protracted period). In Section 6, we study how alternative strategies of "lower for longer" affect the $(r^\star, \pi^\star)$ relation.

As equation (4) makes clear, μ_z, ρ, π enter symmetrically 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 declined 1 percent and the inflation target is increased 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^* must necessarily change by the same amount (with a negative sign).

The previous conjecture is, however, incorrect. The reason for this is 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 our later results imply, after a 1 percentage point decline in r^* the inflation target has to be raised more than 1 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 policymaker would also have to balance the benefits of keeping the incidence of ZLB episodes constant

[7]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.

with the additional costs of extra price dispersion and inefficient resource allocation. These costs can be substantial and may more than offset the benefits of holding the probability of hitting the 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

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.[8]

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.[9] 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 $k \in \{R, g, u, q, z\}$.

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.

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. We use a sample of quarterly data covering the period 1985Q2–2008Q3.[10] 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. The sample covers the Volcker and post-Volcker period, arguably one of relative homogeneity of monetary policy. Second, the sample coincides more or less with the so-called Great Moderation, when, as is argued in the literature, we expect smaller shocks to hit the economy. In principle, smaller shocks

[8] See Gust et al. (2017) and Lindé et al. (2017) for alternative methods that deal with the ZLB constraint at the estimation stage.

[9] See the Appendix for further details.

[10] The data are obtained from the FRED database. GDP is expressed in per capita terms.

will lead to a conservative assessment of the effects of the more stringent ZLB constraint due to lower real interest rates.

The parameters ϕ, θ_p, and θ_w are calibrated before 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 arguably large ballpark of available values used in the literature. In a robustness section, we investigate the sensitivity of our results to these parameters. The parameter ϕ is set to $1/0.7$. Given the assumed subsidy correcting the steady-state price markup distortion, this results in a steady-state labor share of 70%.

We rely on a full-system Bayesian approach to estimate the remaining model parameters. After casting 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 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 1 million from the joint posterior distribution of the parameters.

Estimation Results. Table 1 reports the parameter's postulated priors (type of distribution, mean, and standard error) and estimation results, that is, 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 that the model steady-state values match the mean values of inflation, real per capita GDP growth, and the real interest rate in our US sample. Our choice of priors for the other parameters 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.

Most of our estimated parameters are in line with the calibration adopted by Coibion et al. (2012), with important qualifications. First, we obtain a degree of price rigidity that is slightly higher 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 have only 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, whereas 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

Table 1: Estimation Results

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.

terms of half-life of convergence in the context of autoregressive model of order 1. Our value implies a half-life that is twice as short as theirs. Third, our monetary policy shock and our shocks to demand have an unconditional standard deviation that is approximately twice as small as theirs. Finally, we estimate the degree of indexation to past inflation rather than setting it to zero, as in Coibion et al. (2012). We find small though non-zero degrees of indexation to past inflation. This will translate into a higher tolerance for inflation 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. However, it turns out that, given these estimates, this effect is quantitatively small.

Properties of the estimated model, such as the response to a monetary policy shock, are standard (see the Appendix, Section B. The Appendix also illustrates the "lower for longer" property embedded in the policy rule).

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. (2017) 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

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

Here, s_t is a vector collecting all the model's variables; $\mathcal{A}^{(k)}$, $\mathcal{B}^{(k)}$, $\mathcal{C}^{(k)}$, 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, whereas in the ZLB regime, the coefficient associated with i_t is equal to one 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 (that is, 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.[11] A shortcoming of this approach is that the agents in the model are assumed to believe that the ZLB will not bind again in the future, once the current ZLB episode comes

[11]In practice we combine the implementation of the Bodenstein et al. (2017) 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.

to an end. This may bias estimates, as explained by Gust et al. (2017), even when, as in our case, estimation is performed on a pre-ZLB period. The scope of this concern is, however, dampened by evidence that in the pre-crisis environment, even experts severely underestimated the probability of the ZLB occurring; see Chung et al. (2012).[12]

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, as in Woodford (2003), assuming a small steady-state inflation rate. As detailed in the Appendix, this second-order approximation is given by

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

where t.i.p collects terms that are independent of monetary policy, and $\mathcal{O}(||\zeta,\pi||^3)$ denotes residual terms of order 3, with $||\zeta,\pi||$ denoting a bound on the amplitude of exogenous shocks and the inflation target. Parameters λ_y λ_p, λ_w are effectively weights on an output gap term, a price inflation term, and a wage inflation term. Parameter δ fulfills $0 \leq \delta \leq 1$. The parameter $\bar{x}$ is the log ratio of steady-state output to efficient output. $\bar{x}$ is zero either when trend inflation and trend productivity growth are zero, or when indexation is full, and negative otherwise (in which case, output is inefficiently low). Finally, λ_y λ_p, λ_w, δ, and $\bar{x}$ are functions of the structural parameters θ.

We let $\mathscr{W}(\pi;\theta)$ denote this welfare criterion, to emphasize 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 policymaker. In an alternative exercise, the policymaker maximizes welfare while recognizing the uncertainty associated with the model's parameters.

Using the algorithm outlined above, 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.[13] The optimal inflation rate associated with a given vector of parameters θ is then

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

[13]More precisely, a sample of size $T = 100,000$ 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.

Figure 1: Welfare and the inflation target

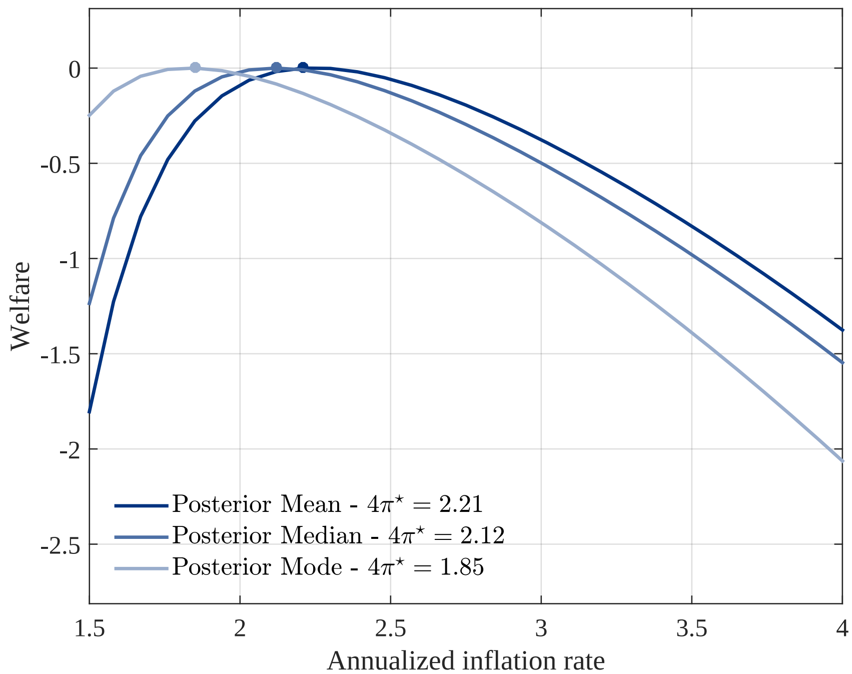

Note: Blue: parameters set at the posterior mean. Light blue: parameters set at the posterior median. Lighter blue: parameters set at the posterior mode. $\pi^\star \equiv \log(\Pi^\star)$. In all cases, the welfare functions are normalized so that they peak at 0.

obtained as the one that maximizes the welfare function; that is,

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

Given parameter estimates at the posterior mean, we can compute the weight on output and wage inflation relative to inflation, that is, λ_y/λ_p and λ_w/λ_p). These relative weights are equal to 0.22 and 0.10, respectively.[14] Note that these values are in the ballpark of values obtained in analyses of optimal inflation based on welfare criteria.

3.3 Some Properties of Loss Function and the Optimal Inflation Target in the Estimated Model

This section presents selected properties of the model related to the optimal inflation target. Figure 1 displays the welfare function—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.

As Figure 1 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

[14] The absolute value of λ_p is found to be 130.52.

Figure 2: Probability of ZLB

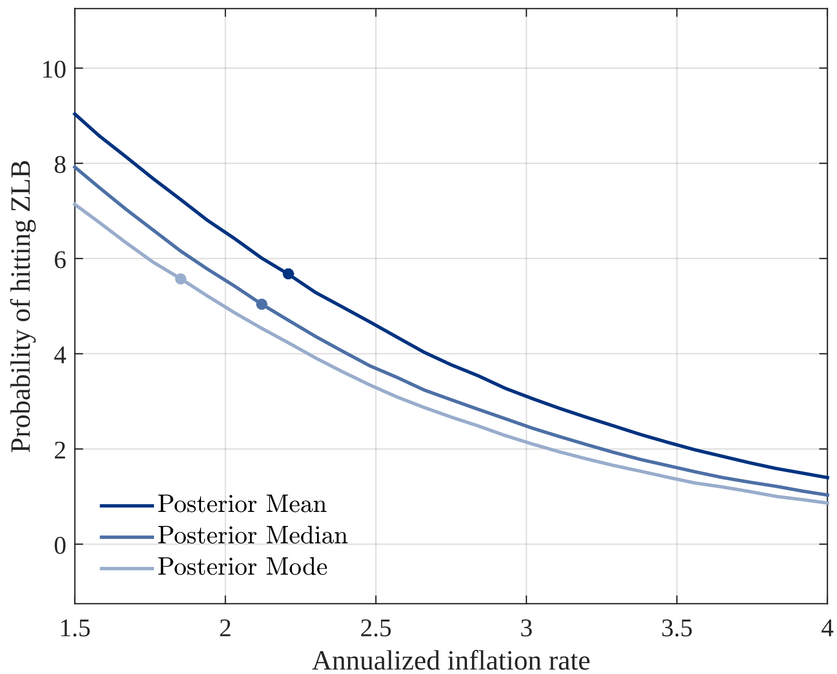

Note: Blue: parameters set at the posterior mean. Light blue: parameters set at the posterior median. Lighter blue: parameters set at the posterior mode. $\pi^\star \equiv \log(\Pi^\star)$.

of values is consistent with the ones of Coibion et al. (2012), even though in this paper it is derived from an estimated model over a much shorter sample.[15] Importantly, while the larger shocks in Coibion et al. (2012), all else being equal, induce larger inflation targets, the high degree of interest rate smoothing in the authors' analysis works in the other direction (as documented below in Section 6) .

To complement these illustrative results, Figure 2 displays 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.

The probability of hitting the ZLB associated with these positive optimal inflation targets is relatively low, at about 6%. 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 episode durations, with a significant fraction of ZLB episodes lasting more than, say, five years (see figure in Appendix D). Given the existence of a single ZLB episode in 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 further analysis in that direction).

A property of our model, as noticed by Kiley (2019), is that ZLB episodes are rather costly, compared

[15] Coibion et al. (2012) calibrate their model on a post–World War II, pre-Great Recession US sample. By contrast, we use a Great Moderation sample.

with other studies. This property reflects the absence in our model of ad hoc stabilizing devices that are present in some other papers concerned with the ZLB, such as emergency fiscal packages or exogenous caps on the maximum duration of ZLB episodes (as in Kiley and Roberts 2017; Williams 2009). Allowing for such devices would mechanically reduce the severity of ZLB episodes in our framework, resulting in a lower optimal inflation target.

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

This section investigates how the monetary authority should adjust its optimal inflation target $\pi^\star$ in response to changes in the steady-state real interest rate, $r^\star$.[16] Intuitively, with a lower $r^\star$ the ZLB is bound to bind more often, so one expects that a higher inflation target is desirable in that case. But the answer to the practical question of how much the target should 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 that induces permanent costs, as argued by, for example, 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 estimated model and ignoring parameter uncertainty. We show that the link between $\pi^\star$ and $r^\star$ depends to some extent on the factor underlying a variation in $r^\star$, that is, a change in the discount rate ρ or a change in the growth rate of technology μ_z. In our setup the first scenario roughly captures the "taste for safe asset" and "aging population" rationale for secular stagnation, while the second one captures the "decline in technological progress" rationale. Subsequently, 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 shocks or on the steady-state price and wage markups.

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 (namely 0.76% and 1.72%, respectively, in annualized terms). 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 value of μ_z or ρ and various values of inflation targets π using the procedure as before.[17] The

[16] Note that 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 such as that of Holston et al. (2017). Our exercise is arguably consistent with "secular stagnation" understood as a permanently lower real rate of interest—without having to assume a unit root process in the real rate of interest.

[17] In particular, we use the same sequence of shocks $\{\epsilon_t\}_{t=1}^T$ 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

optimal value $\pi^\star$ associated with each value of $r^\star$ is obtained as the one maximizing the welfare criterion $\mathscr{W}(\pi;\theta)$.

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 rates 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 curve 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.[18]

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

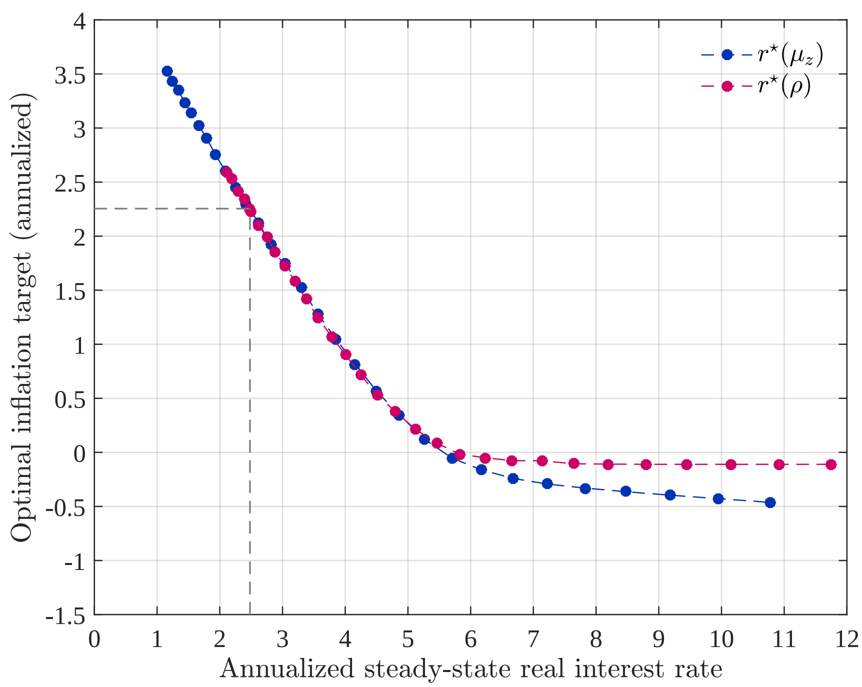

Note: The blue dots correspond to the $(r^\star, \pi^\star)$ locus when $r^\star$ varies with μ_z; the red dots correspond to the $(r^\star, \pi^\star)$ locus when $r^\star$ varies with ρ.

Figure 3 depicts the $(r^\star, \pi^\star)$ relations thus obtained. 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 gray lines indicate the benchmark result corresponding to the optimal inflation target at the posterior mean of the structural parameter distribution.

These results are complemented by Figure 4, which shows the relation between $r^\star$ and the probability of hitting the ZLB, *evaluated at the optimal inflation target*. As with Figure 3, blue dots correspond to the case

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.

[18] Figures G.1 and G.2 report similar results at the posterior mode and at the posterior median. Figure H.1 documents the relation in terms of "optimal nominal interest rate."

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

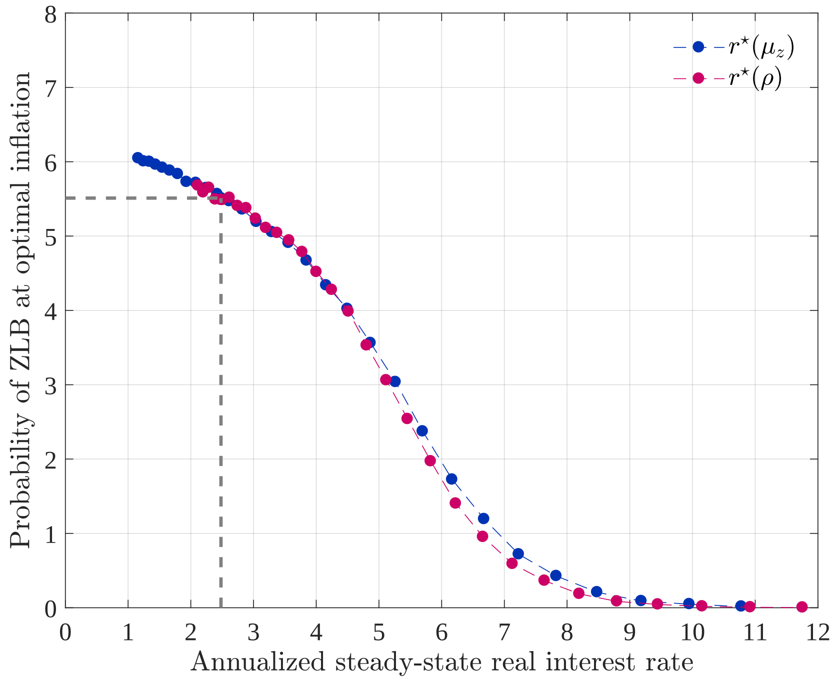

Note: The blue dots correspond to the $(r^\star, \pi^\star)$ locus when $r^\star$ varies with μ_z; the red dots correspond to the $(r^\star, \pi^\star)$ locus when $r^\star$ varies with ρ

when $r^\star$ varies with μ_z, while red dots correspond to the case when it varies with ρ.[19]

As expected, the relation in Figure 3 is decreasing. However, the slope varies with the value of $r^\star$. 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 Figure 4 shows.

At some point, the optimal inflation target becomes insensitive to changes in $r^\star$ when those changes 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 from a zero price inflation and a consequently larger nominal wage inflation.[20]

[19] Figure H.1 in the Appendix shows the relation between $r^\star$ and the nominal interest rate when the inflation target is set at its optimal value.

[20] 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 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

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

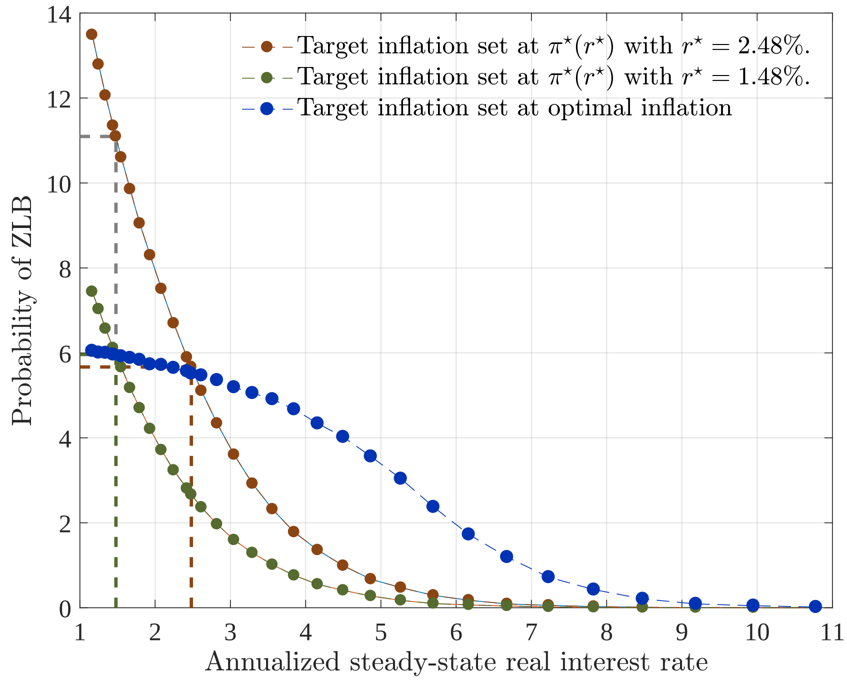

Note: The blue dots correspond to the relation linking $r^\star$ and the probability of ZLB, holding the optimal inflation target $\pi^\star$ at the baseline value. The red dots correspond to the same relation when the optimal inflation target $\pi^\star$ is set at the value consistent with a steady-state real interest rate 1 percentage point lower. The pink dots correspond to the probability of ZLB obtained under the optimal inflation target $\pi^\star$ associated with a given value of $r^\star$.

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 mildly sensitive to changes in $r^\star$ for very large values of $r^\star$, typically above 6%.

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 point decline in $r^\star$ should lead to a 99 basis point increase in $\pi^\star$. Importantly, this increase in the optimal inflation target is virtually the same regardless of whether the underlying factor causing the change in $r^\star$ is a drop in potential growth, μ_z, or a decrease in the discount factor, ρ. At the same time, the ZLB incidence evaluated at the optimal inflation rate also increases when the real rate decreases. At some point, the speed at which this probability increases diminishes, reflecting that the social planner would increase the inflation target as needed so as to avoid a higher ZLB incidence.

Figure 5 shows how the probability of ZLB changes as a function of $r^\star$ when the inflation target is held constant. We first set the inflation target at its optimal baseline value (that is, the value computed

$\pi^\star \approx -\lambda_w(1-\gamma_z)(1-\gamma_w)/[\lambda_p(1-\gamma_p)^2 + \lambda_w(1-\gamma_w)^2]\mu_z$. 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.

at the posterior mean, 2.21%). This is reported as the brown 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 1 percentage point lower (thus, inflation is set to 3.20). Here again, the other parameters are set at their posterior mean. This corresponds to the green dots in the figure. For convenience, we also report the probability of hitting the ZLB as a function of $r^{\star}$ conditional on adjusting optimally the inflation target, as in Figure 4. This corresponds to the blue dots.

Consider first the brown curve. At the level where the real interest rate prevails 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%. Imagine now that the real interest rate experiences a decline of 100 basis points. When the inflation target is kept at the same level as before the shock, the probability of reaching the ZLB climbs to approximately 11%. However, the change in the optimal inflation target brings the probability of reaching the ZLB back to approximately 6%. Thus, the social planner would almost neutralize the effects of the natural rate decline on the probability of hitting the ZLB.

Finally we investigate whether the tradeoff analyzed above translates into meaningful welfare costs, measured in terms of foregone per-period consumption. Results are reported in Appendix F. It turns out that, under sufficiently low $r^{\star}$ values, agents faced with a 1 percentage point decline in the steady-state real interest rate would require as much as a 1.5 percentage point increase in consumption to be as well-off under the former optimal inflation target (that is, 2.21%) as they were under the optimal target associated with the lower real interest rate (3.20% in this case). In other words, the welfare costs of not adjusting the target in the face of a decline in $r^{\star}$ are substantial.

4.2 Robustness to Alternative Structural Assumptions

In this section, we investigate the robustness of the $(r^{\star}, \pi^{\star})$ relation to altering (or modifying) some structural features of the environment. We consider several relevant dimensions: larger shocks, alternative calibrations for the steady-state price and wage markup, and changes in the degree of price and wage indexation.[21]

Larger Shocks. As noted previously, 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 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 deviation that is 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

[21] Robustness to altering the monetary policy rule is assessed further below.

Figure 6: $(r^\star, \pi^\star)$ relation with larger demand shocks

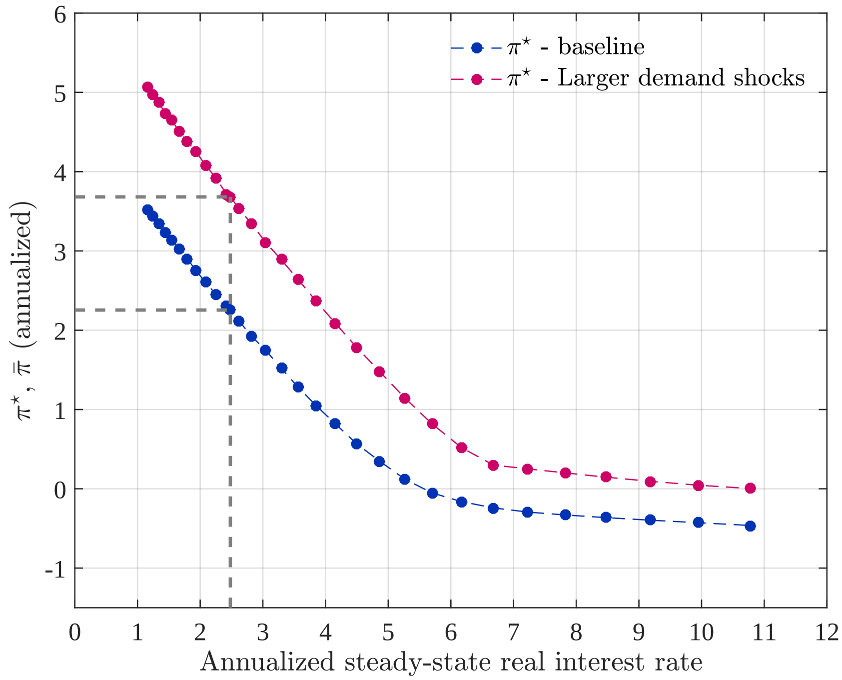

Note: The blue dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with σ_q and σ_g set to twice their baseline value.

their posterior mean. Given this setup, the optimal inflation target is 3.7% as opposed to 2.21%, conditional on the baseline value of $r^\star$. Also, under the alternative shock configuration, the probability of hitting the ZLB is 5.3%, as opposed to 5.5% in the baseline. These probabilities may seem low, especially in the case of large shocks, which, we argue, capture Great Recession-like shocks. However they are particularly low because the inflation target is chosen optimally in this setup. In particular, in the larger shocks case, the increase in the inflation target is large enough to offset the impact of larger shocks in terms of ZLB incidence. When instead we keep the inflation target unchanged, the probability of hitting the ZLB rises to 18% in the face of a 1% decline in $r^\star$.[22] In that case, these ZLB probabilities come close to the probabilities reported by Kiley and Roberts (2017) in their DSGE model (albeit, they find a higher probability of ZLB, of the order of 30% when using the FRS/US model) and by Chung et al. (2019). [23]

Figure 6 reports the $(r^\star, \pi^\star)$ relation under larger demand shocks (red dots) and compares the outcome with the baseline relation (blue dots).[24] 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 -1. However, the curve is somewhat steeper in the high $r^\star$ region and shifted up, compared with the baseline scenario. This reflects that under larger

[22]See Appendix Section I, in which such counterfactual probabilities of ZLB are reported.

[23]In addition to per period probability of ZLB, these authors also put forward and emphasize the probability that a ZLB event occurs in the next decade. By construction this number is a larger one, and the mapping between the two numbers is not fully straightforward

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

demand shocks, even at very high levels of the natural rate, a drop in the rate is conducive to more frequent ZLB episodes. The social planner is then willing to increase the inflation target faster than in the baseline scenario and generically sets the inflation target at higher levels to hedge the economy against ZLB episodes.

Alternative Markups. The optimal level of inflation in our setup depends on the elasticities of substitution among intermediate goods, θ_p, and among labor types, θ_w, since those parameters determine the extent to which the price and wage dispersion induced by inflation is translated into an inefficient allocation of resources. These parameters have been calibrated, as they cannot be identified from time-series data and a log-linearized version of the model.

In our calibration, the baseline value for the elasticity of substitution θ_p is 6, leading to a steady-state price markup 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), there is considerable uncertainty in the empirical literature about the level of markups. For example, some estimates in Basu and Fernald (1997) and Traina (2015) point to possibly much smaller values, while Autor et al. (2017), De Loecker and Eeckhout (2017), and Farhi and Gourio (2018) 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 and as low as 3. These values largely encompass the range of available empirical estimates.

Similarly, for the wage markup, there is arguably scarcer evidence, and in any case considerable uncertainty around our baseline parameterization, given by θ_w set to 3. Here again, in order 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 Figure 7 in the case of robustness with respect to the price markup, and in Figure 8 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%), the slope of the $(r^\star, \pi^\star)$ curve is only very mildly affected when the elasticity of substitution of goods or labor types is changed.

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%), the value of the optimal inflation target, and the slope of the curve of interest, are more sensitive to the value of θ_p or θ_w. To see why, first note that in this region, the ZLB is essentially irrelevant, so the standard welfare cost of the inflation setup 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 lead to no output dispersion across firms at all). The effect of θ_p on output dispersion is apparent from the formulas in our Appendix and in textbook derivations of output dispersion, such as those in chapter 3 of Galí (2015). Thus, with a low substitution (that is, a low θ_p), the welfare loss due to inflation (or deflation) is smaller. Therefore a lower

Figure 7: $(r^{\star}, \pi^{\star})$ relation with alternative θ_p

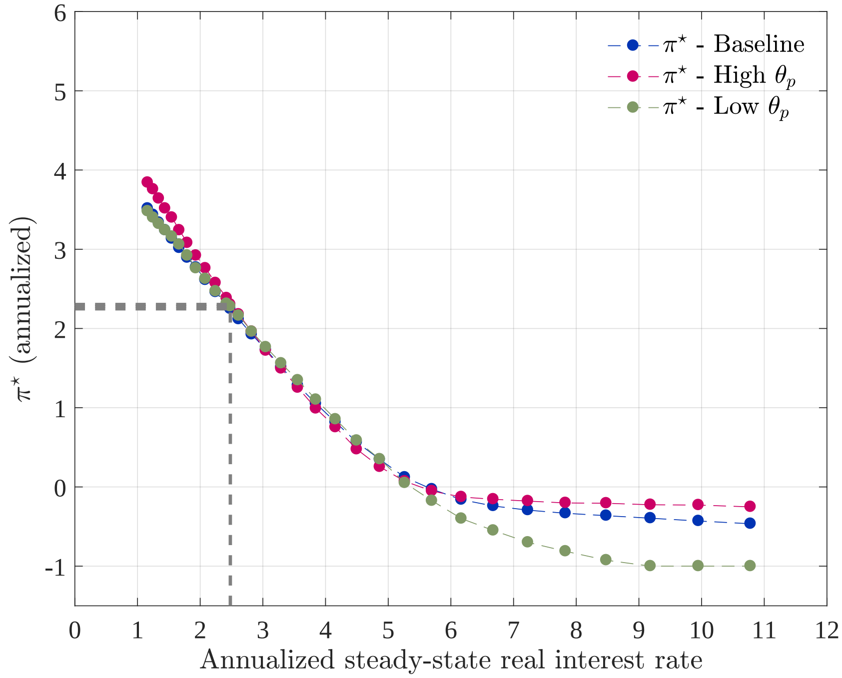

Note: The blue dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with θ_p set to 10. The green dots correspond to the counterfactual simulation with θ_p set to 3.

Figure 8: $(r^{\star}, \pi^{\star})$ relation with alternative θ_w

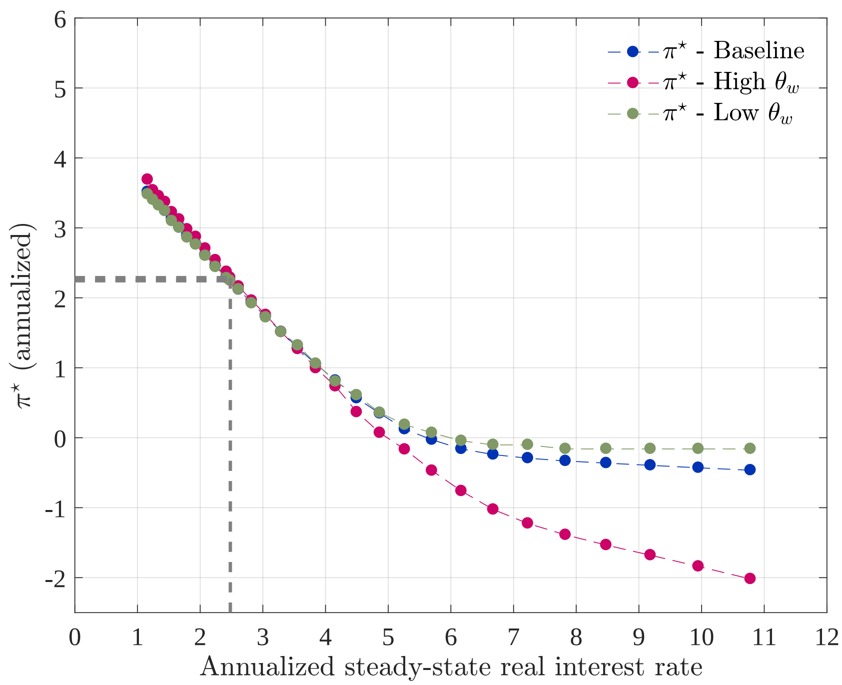

Note: The blue dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with θ_w set to 8. The green dots correspond to the counterfactual simulation with θ_w set to 1.5.

θ_p allows for an inflation target that is further from zero, insofar as there are motives for a non-zero steady-state inflation. Such a mechanism explains why, in Figure 7, 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 (see Figure 8). That is, a larger θ_w induces a larger inflation target in absolute value. The reason is that with a larger substitution across labor types, a given nominal wage growth generates dispersion of quantities across types of labor that turns out to be particularly costly. In such a case, it is optimal that the burden of adjustment of real wages to growth is borne not by nominal wages, but rather by nominal prices (thus leading to a more pronounced deflation).[25]

Alternative Degrees of Indexation. The degree of indexation of price and wage is an important determinant of the cost of inflation. In our empirical estimate, the degrees of indexation are moderate: 0.22 for prices and 0.44 for wages at the posterior mean. It is worthwhile to examine the sensitivity of our results to the degree of indexation. Indeed, some macro studies' estimates find or impose a much larger degree of indexation (Christiano et al. 2005). By contrast, micro studies hardly find any evidence of indexation. In this robustness exercise, we consider in turn a "zero indexation" case, a high indexation case (setting γ_p and γ_w to 0.7), and a very high indexation case (setting γ_p and γ_w to 0.9). The last two configurations are arguably unrealistic. Results are presented in Figure 9. In the absence of indexation, results are similar to those under our estimated indexation levels. For the high indexation case (γ_p and γ_w equal to 0.7), the results differ from the baseline only for relatively large values of the steady-state real interest rate.

In the very high indexation case, the position and shape of the curve are substantially affected. The curve is nearly a declining straight line. For a very large indexation degree the welfare cost of inflation (or deflation) is substantially reduced. Thus, it is optimal to allow for a sizable trend deflation when the natural rate is large as a result of large productivity growth. However, we can note that in the empirically relevant region, that is, for $r^\star$ below 2%, the local slope of the curve is similar whatever the degree of indexation.

5 The Effect of Parameter Uncertainty

In this section we investigate the impact of parameter uncertainty on the relation between the optimal inflation target and the steady-state real interest rate. Specifically, we analyze how a Bayesian-theoretic optimal inflation target reacts to a downward shift in the *distribution* of the steady-state real interest rate.

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

Figure 9: $(r^\star, \pi^\star)$ relation with alternative indexation degrees

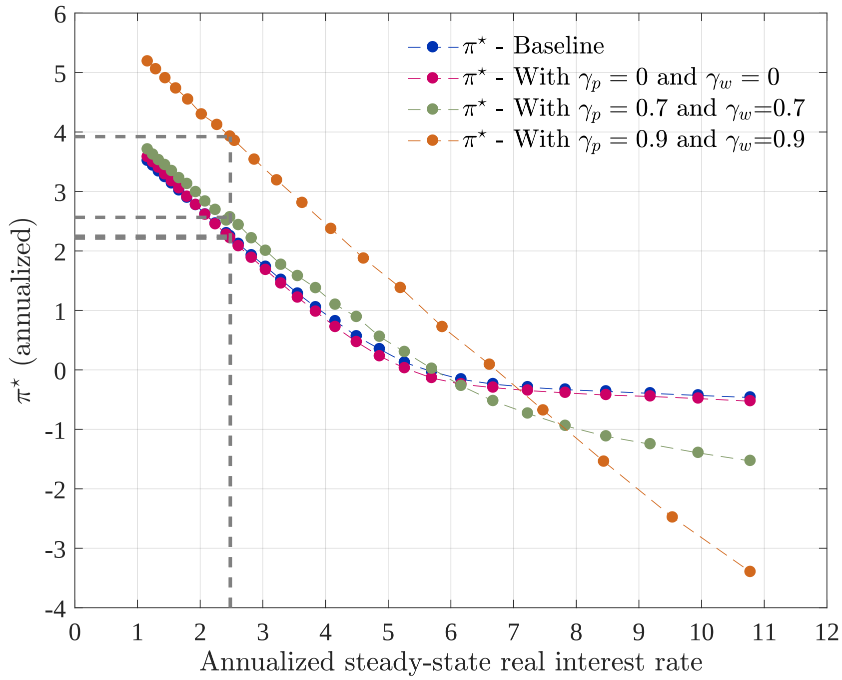

Note: The blue dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with $\gamma_p = \gamma_w = 0$. The green dots correspond to the counterfactual simulation with $\gamma_p = \gamma_w = 0.7$

A Bayesian-theoretic Optimal Inflation Target. The location of the loss function $\mathscr{W}(\pi;\theta)$ evidently depends on the vector of parameters θ describing the economy. As a result of estimation uncertainty around θ, the optimal inflation rate $\pi^\star(\theta)$ will be subject to uncertainty. Further, a policymaker may wish to take into account the uncertainty surrounding θ when determining the optimal inflation target. A relevant feature of the welfare functions in our setup is that, in general, and as shown above, they are markedly asymmetric: Adopting an inflation target 1 percentage point below the optimal value generates welfare losses larger than by setting it 1 percentage point above. As a result, the certainty equivalence does not hold. A policymaker maximizing expected welfare while recognizing the uncertainty will choose an inflation target that differs from the one 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 such a target as the inflation target $\pi^{\star\star}$ that maximizes the expected welfare not only over the realizations of shocks, but also over the realizations of parameters[26]

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

We interpret the spread between the optimal Bayesian inflation target and the "certainty-equivalent" optimal inflation target at the posterior mean $\bar{\theta}$ as a measure of how uncertainty about the parameter value affects optimal inflation. Given the nature of the asymmetry in the welfare function, the spread will turn out to be positive: A Bayesian policymaker will tend to choose a higher inflation target than a policymaker who takes θ to be known and equal to the mean of its distribution. A higher 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 $\mathrm{Spr}(\bar{\theta})$.

Results. According to the simulation exercise, $\pi^{\star\star} = 2.40\%$. This robust optimal inflation target is higher than the value obtained with θ set at its central tendency. As expected, a Bayesian policymaker chooses a higher inflation target to hedge against particularly harmful states of the world (that is, parameter draws) where the frequency of hitting the ZLB is high.[27]

Assessing how a change in $r^{\star}$ 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 somewhat smaller in absolute value) to recent estimates of the drop of the natural rate after the crisis, such as those in 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. Figure 10 depicts the counterfactual shift in the distribution of $r^{\star}$ that is considered.

The Bayesian-theoretic optimal inflation target corresponding to the counterfactual lower distribution

[26]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 $T = 100,000$ 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 of 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.

[27]Figure E.1 in the Appendix illustrates this point by showing that the posterior distribution of $\pi^{\star}(\theta)$ is broadly symmetric.

Figure 10: Posterior distributions of $r^\star$ and counterfactual $r^\star$

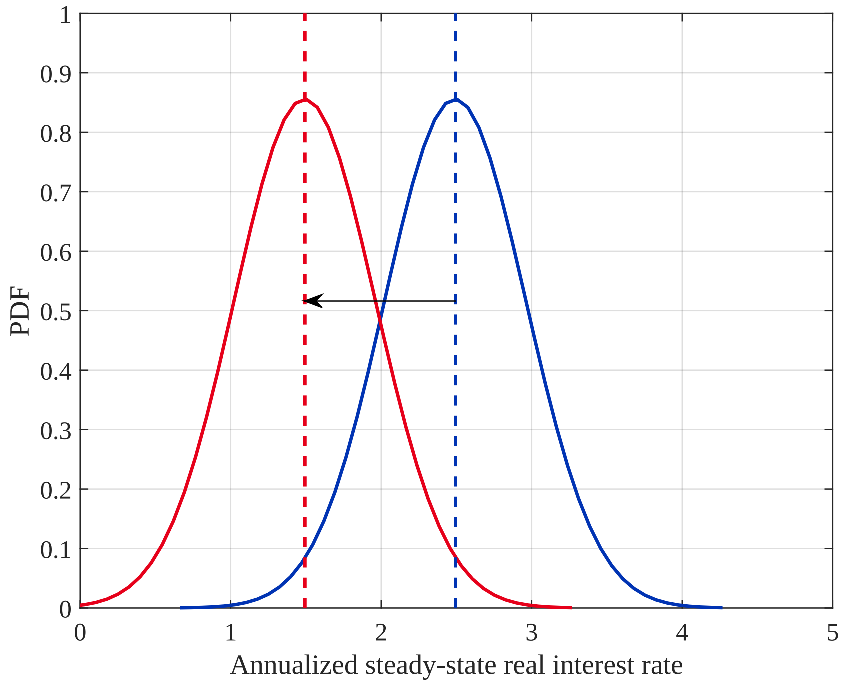

Note: Plain curve: PDF of $r^\star$. Dashed vertical line: mean value of $r^\star$.

of $r^\star$ is obtained from a simulation exercise that relies on the same procedure as before.[28] 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 downward by 1 percentage point (in annualized terms). The welfare function $\mathscr{W}(\pi;\theta_\Delta)$ is then evaluated. Since there are no changes other 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

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

Figure 11 illustrates the counterfactual change in the optimal inflation target obtained when the mean of the distribution of the steady-state real interest rate declines by 100 basis points. The simulation exercise returns a value of $\pi^{\star\star}_\Delta = 3.30\%$, that is, 90 basis points higher than the optimal value under uncertainty obtained with the posterior distribution of parameters in the pre-crisis sample $\pi^{\star\star}_\Delta = 2.40\%$.[29]

Thus, in our setup, a monetary authority that is concerned about the uncertainty surrounding the parameters driving the costs and benefits of the inflation chooses a higher optimal inflation target. However,

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

[29] Figure J.1 in Appendix shows how the posterior distribution of $\pi^\star$ is shifted after the permanent decline in the mean of $r^\star$.

Figure 11: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$ in baseline and counterfactual

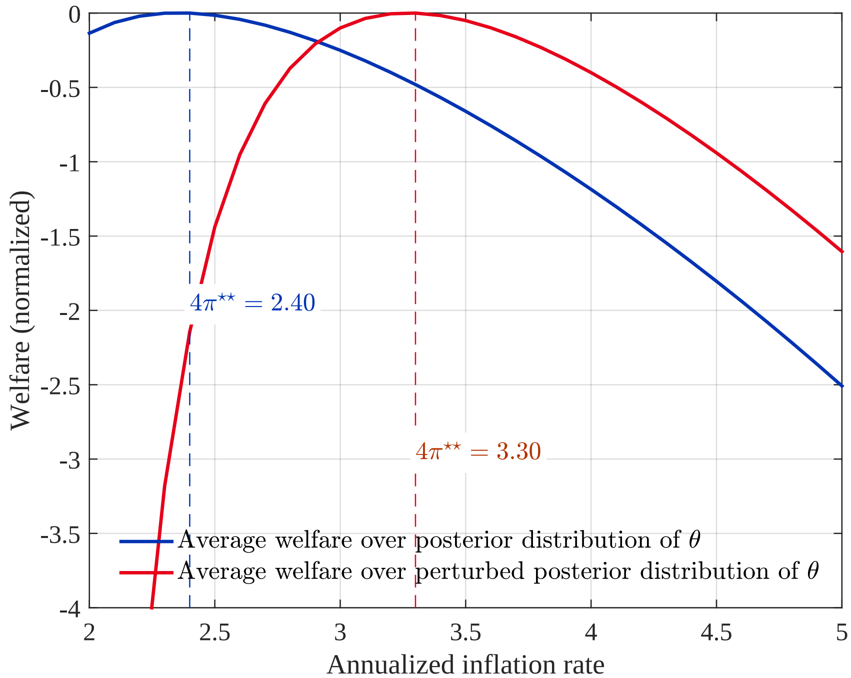

Note: Blue curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$. Red curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$ with lower $r^\star$.

the reaction of this optimal inflation target following a drop in the mean $r^\star$ is hardly altered: A 100 basis point decrease in the steady-state real interest rate calls for a roughly 90 basis point increase in the optimal inflation target, which is in the vicinity of pre-crisis parameter estimates.

A Known Reaction Function. Here we study the consequences of the (plausible) assumption that the central bank actually knows with certainty the coefficients of its interest rate rule. More specifically, we repeat the same simulation exercise as in the previous subsection 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.[30] Note, however, that all the other parameters are subject to uncertainty from the standpoint of the central bank.

Figure 12 presents the Bayesian-theoretic optimal inflation targets obtained when simulating the model at the initial posteriors and after a -100 basis point 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 $\pi^{\star\star} = 2.24\%$. After the counterfactual change in the distribution of $r^\star$ is considered, $\pi^{\star\star}$ should be increased to 3.16%, which is again in the ballpark of a 90 basis point increase in $\pi^\star$ in response

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

Figure 12: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$

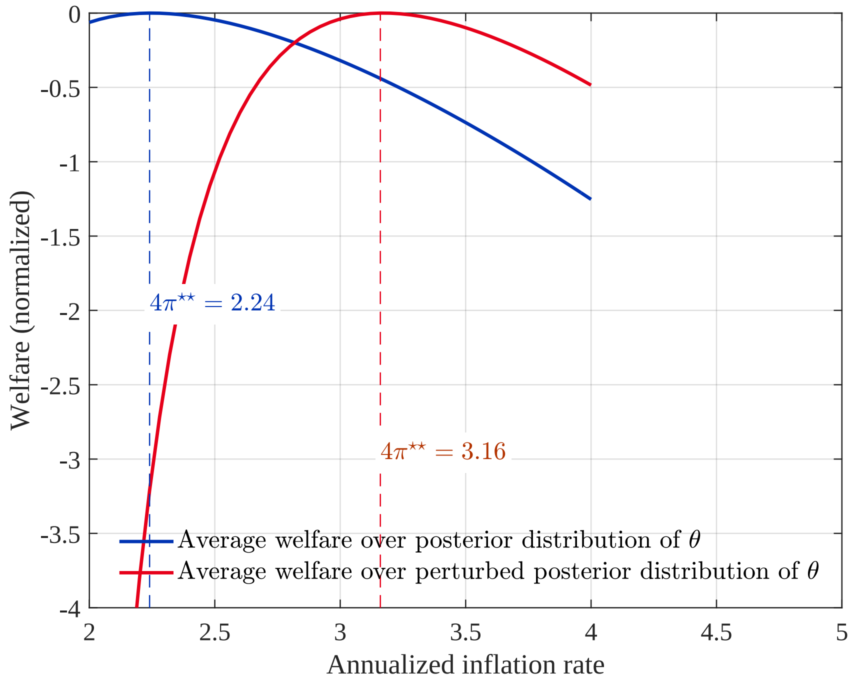

Note:: Blue curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$. Red curve: $\mathbb{E}_\theta(\mathscr{W}(\pi,\theta))$ with lower $r^\star$. In each case, ρ_i, a_π, and a_y are frozen at their posterior mean values.

to a 100 basis point downward shift in the distribution of $r^\star$.

6 Alternative Monetary Policy Rules and Environments

In this section we study the optimal adjustment of the inflation target in response to a change in the steady-state real interest rate under five alternative assumptions regarding monetary policy: (1) an inflation target that is set in terms of average realized inflation, (2) an effective lower bound on the policy rate that can be below zero, (3) alternative degrees of smoothing in the policy rule, (4) a central bank with no "lower for longer" strategy, and (5) a price level targeting rule. For simplicity, throughout this section we ignore the role of uncertainty and treat the model parameters as known.

Average versus 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 or in a linear model). This results from the fact that anytime the ZLB is binding (a recurrent event), the central bank effectively loses its ability to stabilize inflation around the target. Therefore, it may be relevant to assess the central bank's outcomes and set the corresponding target in terms of the effective average realized inflation. In this section, we investigate whether measuring the inflation target in this alternative way matters.

Figure 13: Average realized inflation and optimal inflation

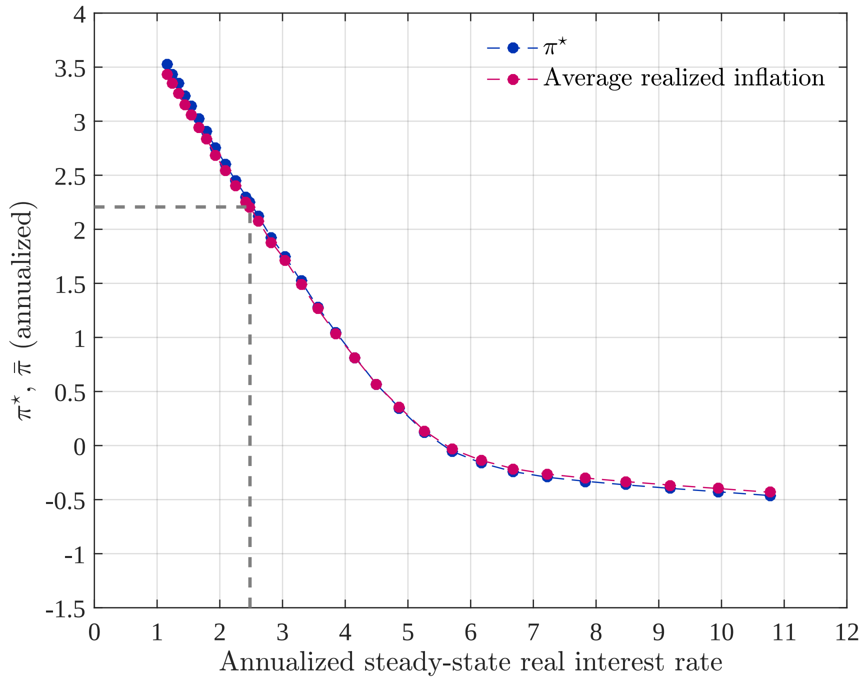

To this end, the analysis of the $(r^\star, \pi^\star)$ relation in Section 3.2 is complemented here with the analysis of the relation between $r^\star$ and the average realized inflation rate $\mathbb{E}\{\pi_t\}$ obtained when simulating the model for various values of $r^\star$ and the associated optimal inflation target $\pi^\star$. In the interest of brevity, the calculations are presented only for when the source of variation in the natural interest rate is the change in average productivity growth μ_z.

Figure 13 illustrates the difference between the $(r^\star, \pi^\star)$ curve (blue dots) and the $(r^\star, \mathbb{E}\{\pi_t\})$ curve (red dots). The overall shape of the curve is unchanged. Unsurprisingly, the 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 indeed becomes lower than the optimal inflation target. However, that spread remains limited; it is 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 14: Optimal inflation with a negative ELB

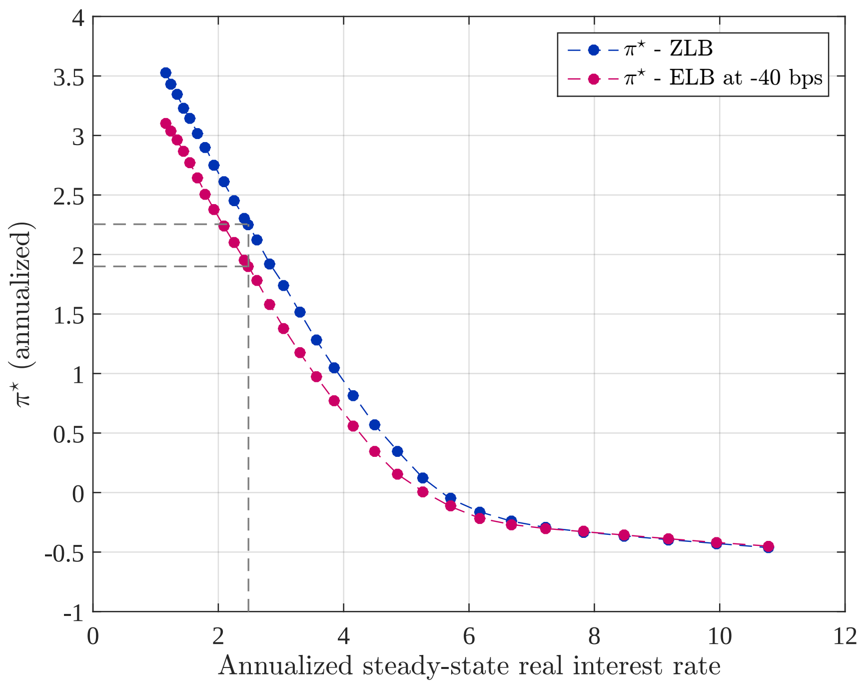

the ECB's deposit facility rate, which gears the overnight money market rate due to excess liquidity, was set at a value of -10 basis points in June 2014 and was further lowered to -40 basis points in March 2016.[31]

We use the estimated model to evaluate the implications of a negative ELB in the US. 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 14. As expected, the $(r^\star, \pi^\star)$ locus shifts downward, though by slightly fewer than 40 basis points. Importantly, its slope remains identical to the baseline case: Around the baseline value for the real interest rate, a 100 basis point downward shift in the distribution of $r^\star$ calls for an increase in $\pi^\star$ of about 90 basis points.

Alternative Degrees of Interest Rate Smoothing. Our analysis is conditional on a specific reaction function of the central bank, described in our setup by the set of parameters a_π, a_y, and ρ_i. Among these parameters, the smoothing parameter, ρ_i, 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

[31] In September 2019, the rate on the deposit facility was lowered to -0.50%

bound constraint (see equation 6), while the effective nominal interest rate is the maximum of zero and the notional rate (see equation 5). Thus the interest rate inherits the past negative values of the notional nominal rate. So, a higher smoothing results in the effective interest rate remaining at zero for an extended period of time beyond that implied by the macroeconomic shocks that initially brought the economy to 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 keeping 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 with exiting the trap (or even not entering it).

Through both effects, a 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 being equal, one should therefore expect a lower optimal inflation rate for higher values of the smoothing parameters. This property of the model is illustrated in Figure 15, which depicts the $(r^\star, \pi^\star)$ relation under three possible values of the smoothing parameter ρ_i. The value used under our baseline scenario, that is, the posterior mean estimates, are 0.85. We also consider two alternative settings: a higher value of $\rho_i = 0.95$, which is close to the inertia of the central bank reaction function in Coibion et al. (2012), and a lower value of $\rho_i = 0.8$. These two values arguably encompass 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 the $(r^\star, \pi^\star)$ curve downward, 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 have been close to 0.5% in the US.[32] 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_i, the pre-crisis optimal inflation rate would be close to 3.5%.

As for the slope of the $(r^\star, \pi^\star)$ curve, in the empirically relevant region, it is much less affected than is 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 extent to which an increase in the inflation target is needed. The slope is indeed close to -0.7 in that case. For a strategy associated with a low smoothing parameter, the slope is close to -1.0, which is closer to the benchmark case. For large values of $r^\star$, the degree of smoothing is irrelevant.

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

Figure 15: $(r^{\star}, \pi^{\star})$ relation with alternative ρ_i

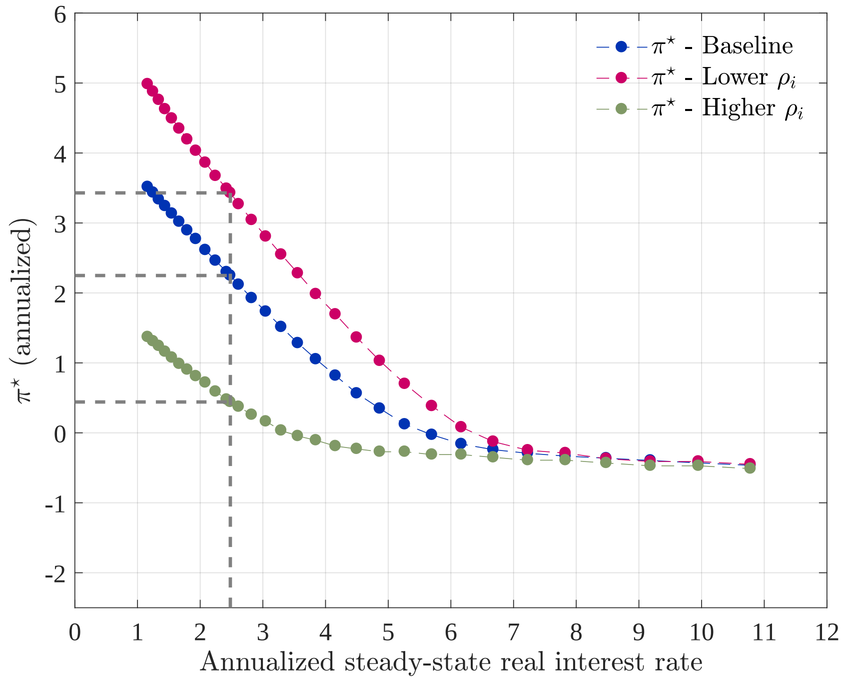

Notes: The blue dots correspond to the baseline scenario wherein all the structural parameters are set at their posterior mean $\bar{\theta}$. The red dots correspond to the counterfactual simulation with ρ_i set to 0.8. The green dots correspond to the counterfactual simulation with ρ_i set to 0.95.

More Traditional Specifications of the Policy Rule. We also considered the case of a monetary policy rule featuring no shadow rate (that is, no "lower for longer" feature) as well as the case of a simple non-inertial Taylor rule. Results are reported in Figures 16 and 17. In the first case, the lagged interest rate is the lagged actual rate. Immediately after the liftoff that follows a ZLB episode, the interest rate follows a standard path, so monetary policy does not "keep memory" that it has been constrained for some periods by the ZLB (unlike under our baseline specification). In the second case, there is no inertia at all, but we use a four-quarter inflation rate as in the standard Taylor rule (and its implementation in Kiley and Roberts 2017). In both cases, the overall degree of monetary policy inertia decreases, and so the stabilization property of the policy rule is weaker in our forward-looking model, materializing in more frequent ZLB episodes. As a result, the optimal inflation rate is in both cases larger than in the baseline, for realistic values of the real interest rate. Also the optimal inflation rate is positive for a wider range of values of $r^{\star}$. However, in both variants the slope of the $(r^{\star}, \pi^{\star})$ is similar to that of our baseline curve around the sample value of $r^{\star}$.

A Price Level Targeting Rule. Finally, we consider that the rule effectively implemented by the central bank reacts to deviations of the (log) price level $\hat{p}_t = \hat{p}_{t-1} + \hat{\pi}_t$ from a targeted path, instead of to the gap $\hat{\pi}_t$ between the inflation rate and its optimal target. Formally, we assume that the central bank sets the policy rate according to the following rule:

$$\hat{\imath}_t^{plt} = \rho_i \hat{\imath}_{t-1}^{plt} + (1 - \rho_i)\left(a_p \hat{p}_t + a_y \hat{x}_t\right) + \zeta_{R,t}, \tag{7}$$

Figure 16: $(r^{\star}, \pi^{\star})$ relation with simple standard Taylor rule

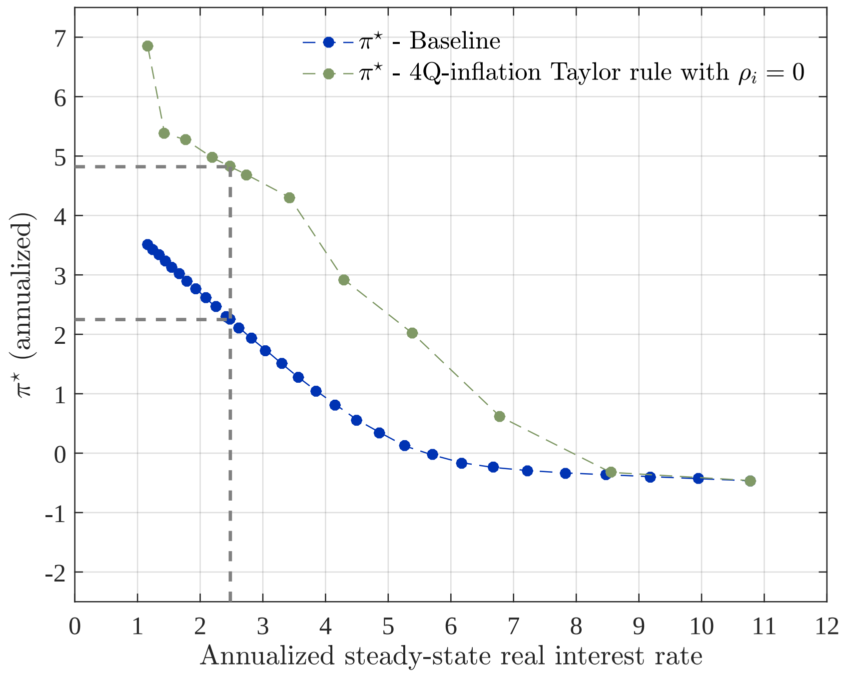

Notes: The blue dots correspond to the baseline scenario. The green dots correspond to the counterfactual simulation using the simple standard Taylor rule. Parameters of the rules on inflation and output gap are the same as in the baseline, but there is no inertia, and four-quarter inflation is used.

Figure 17: $(r^{\star}, \pi^{\star})$ relation with "'no shadow rate" rule

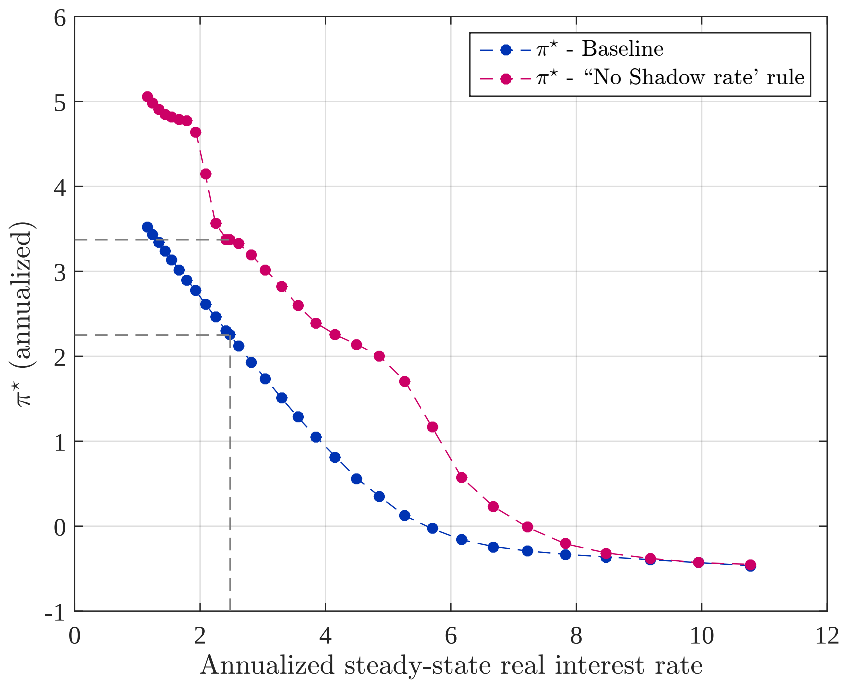

Notes: The blue dots correspond to the baseline scenario. The red dots correspond to the counterfactual non-inertial policy rule used.

Figure 18: $(r^\star, \pi^\star)$ relation with price level targeting strategy

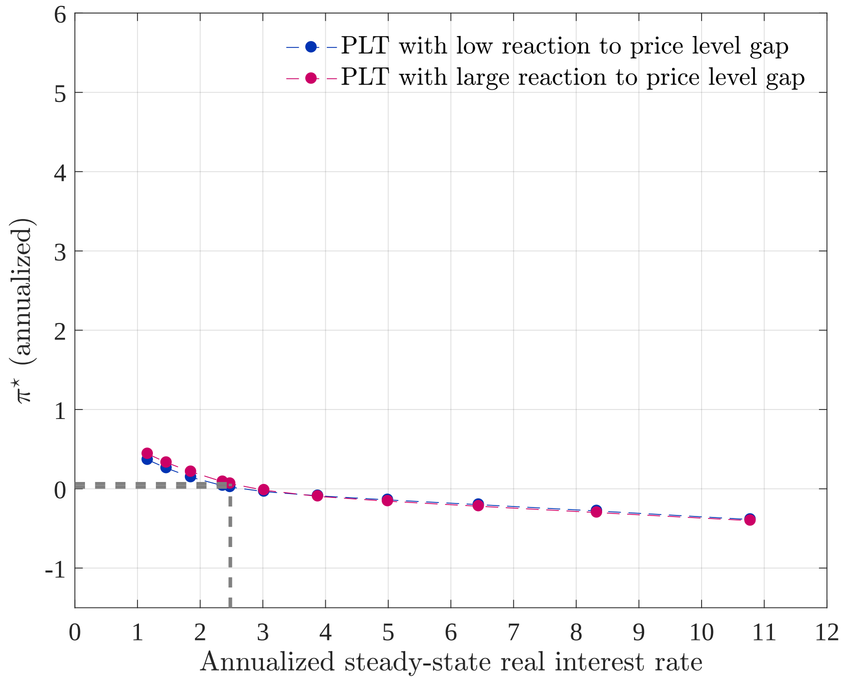

Notes: Simulations obtained under the price level targeting policy rule given in equation 7. The blue (red) dashes correspond to the scenario wherein $a_p = 0.1$ (0.5). All the other structural parameters are set at their posterior mean $\bar{\theta}$.

with $\hat{\imath}_t = \max\{\hat{\imath}_t^{plt}, -(\mu_z + \rho + \pi)\}$.

We perform the same exercises as before, focusing on the case in which average productivity growth μ_z is the driver of changes in the natural rate. We consider two values for a_p: 0.1 and 0.5. All the other parameters of the model are set to their posterior mean.

Figure 18 reports the $(r^\star, \pi^\star)$ relation obtained under these two alternative scenarios. A striking features of the new curve is that the optimal inflation target lies between 0% and 1%, as opposed to 2.21% in the baseline. Price level targeting makes a commitment to make up for past inflation undershooting (or overshooting) that is even stronger than what can be obtained when increasing the smoothing parameters in a rule that targets inflation instead. This commitment stabilizes inflation expectations by reducing both the probability of hitting the zero lower bound and the average length of such episodes. As a consequence, there is no incentive to bear the costs of a positive steady-state inflation, and the optimal inflation target is close to zero. This holds regardless of whether the central bank reacts aggressively to a deviation of the price level from its targeted path.

Another striking result is that the $(r^\star, \pi^\star)$ relation is much flatter in the vicinity of the pre-crisis level for $r^\star$ than under alternative inflation targeting monetary policy strategies. The slope is close to -0.3 instead of the (-1.0 to -0.7) range obtained previously. A price-level targeting strategy thus allows the costs of the ZLB to remain small even if the natural rate of interest drops by, say, 1% compared with the pre-crisis regime.

7 Summary and Conclusions

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 percentage point decline in the natural rate should be accommodated by an increase in the optimal inflation target of 0.9 to 1.0 percentage points. For convenience, Table 2 recaps our results. Overall, across the different concepts of optimal inflation considered in this paper, the level of optimal inflation does vary. However, our finding that the slope of the $(r^\star, \pi^\star)$ relation is close to -1 in the vicinity of the pre-crisis value of steady-state real interest rates is very robust.

Table 2: Effect of a decline in $r^\star$ under alternative notions of optimal inflation

	Baseline	Lower $r^\star$	Δ
Mean of $\pi^\star$	2.00	3.00	1.00
Median of $\pi^\star$	1.96	2.90	.94
$\pi^\star$ at post. mean	2.21	3.20	.99
$\pi^\star$ at post. median	2.12	3.11	.99
$\pi^{\star\star}$	2.40	3.30	.90
$\pi^{\star\star}$, frozen MP	2.24	3.16	.92
Average realized inflation at post. mean	2.20	3.19	.99
$\pi^\star$ at post. mean, ELB -40 bp	1.90	2.83	.93
Average realized inflation at post. mean, ELB -40 bp	1.86	2.77	.91
$\pi^\star$ at post. mean, higher ρ_i	0.44	1.13	.69
$\pi^\star$ at post. mean, price level targeting	0.06	0.32	.26

Note: All figures are annualized percentage rates.

Our analysis considers 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. Indeed, recent discussions revolving around monetary policy in the new normal suggest that the various non-conventional measures—forward guidance on interest rates and large scale asset purchases—used in the aftermath of the Great Recession could feature permanently in the policy toolbox. In particular, unconventional monetary policies could represent useful second-best instruments when the ZLB is reached, as advocated by Reifschneider (2016), Swanson (2018), and Sims and Wu (2019) (see also Eberly et al. 2019 and Chung et al. 2019 for recent work documenting the efficiency of such instruments). By implying a "low for long" interest rate at the end of the trap, the monetary policy rule that we consider in our exercise accounts, at least partially, for the effect of non-conventional policies that were implemented at the ZLB. But more aggressive non-conventional packages could be considered as alternative strategies. Beyond these monetary policy measures, fiscal policies could play a significant role, as emphasized by Correia et al. (2013). As a result, the ZLB might be less stringent a constraint in a practical policy context than in our analysis, as argued in Debortoli et al. (2019) among others. However, the efficacy and the costs of these policies should also be part of the analysis. The complete comparison of these policy tradeoffs goes beyond the scope of this paper.

An alternative would consist of a change of monetary policy strategies, for example, adopting variants of the price-level targeting strategy, as recently advocated by Williams (2016) and Bernanke et al. (2019). Our exercises emphasize that when the central bank follows such a "make-up" strategy for past inflation deviations from its target, the case for increasing the inflation target is greatly reduced. Nevertheless, these results are obtained under the assumption that private agents believe and understand the commitment of the central bank to deviating from its inflation target in order to compensate for previous deviations. This is a debatable assumption. Andrade et al. (2019) show that the lower-for-longer guidance on future interest rates that the FOMC gave during the recent ZLB episode was interpreted differently by private agents, including professional forecasters: Some viewed it as a good news of a commitment to future accommodation, and some viewed it as bad news that the lowflation would last longer. This finding emphasizes that lower-for-longer policies are much less effective in practice than is implied by theoretical models with assumptions of perfect credibility, full information, and rational expectations. They can even be detrimental if the bad signals prevail.

We discuss the potential desirability of a higher inflation target, abstracting from the challenges of implementing an eventual transition to the new objective. In the current low-inflation environment, increasing the inflation target in reaction to a drop in the steady-state value of the real interest rate might raise some credibility issues. However, a move toward make-up strategies would also raise substantial credibility issues, as these imply an arguably time-inconsistent commitment to deviate from the inflation target once it has been reached.

Finally, our analysis abstracts from forces identified in the literature as warranting a small, positive inflation target, irrespective of ZLB issues, as emphasized in, for example, Bernanke et al. (1999) and Kiley et al. (2007). The first is grounded in measurement issues, following the finding from the 1996 Boskin report that the Consumer Price Index did probably overestimate inflation in the US by more than 1 percentage point in the early 1990s. downward rigidities (for example, in nominal wages), a positive inflation rate can help "grease the wheel" of the labor market by facilitating relative price adjustments. Symmetrically, we also abstract from forces calling for lower inflation targets. The most obvious is the so-called Friedman (1969) rule, according to which average inflation should be equal to minus the steady-state real interest rate, hence *negative*, in order to minimize loss of resources or utility and the distortionary wedge between cash and credit goods (for example, consumption and leisure) induced by a non-zero nominal interest rate. Presumably, these and several other factors were taken into account when an inflation target of 2% was chosen. But an estimate of r^* was, undoubtedly, one of the key factors in that choice. Accordingly, the current re-assessment of r^* by the Federal Reserve and other central banks would seem to call for a simultaneous re-assessment of the optimal inflation target.

Appendix

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., $\int_\theta \pi^\star(\theta)p(\theta\|X_T)\mathrm{d}\theta$
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 Illustrating model properties: moments, IRF to monetary policy shock

This section illustrates basic properties of the estimated baseline model.

Table B.1: Moments of key variables

Data 1985Q2-2008Q3

Variable	Inflation	4-Quarter -Inflation	Output gap	Output growth	Interest rate
Std. dev.	0.22	0.73	—	0.54	2.20

Simulated Model (with ZLB constraint)

Variable	Inflation	4-Quarter -Inflation	Output gap	Output growth	Interest rate
Std. dev.	0.43	1.53	0.58	0.99	2.15
	(0.11)	(0.41)	(0.14)	(0.14)	(0.14)

Note: In percent. Inflation is quarterly inflation (not annualized). Interest rate is annualized. 4-Quarter inflation is the year-on-year growth rate of the price index. The model moments are based on 1000 simulations at the posterior mean. At each simulation, shocks are drawn (with replacement) from the historical shocks. The figures in parentheses are the standard deviation across bootstrap simulations.

Figure B.1: Response to a monetary policy shock

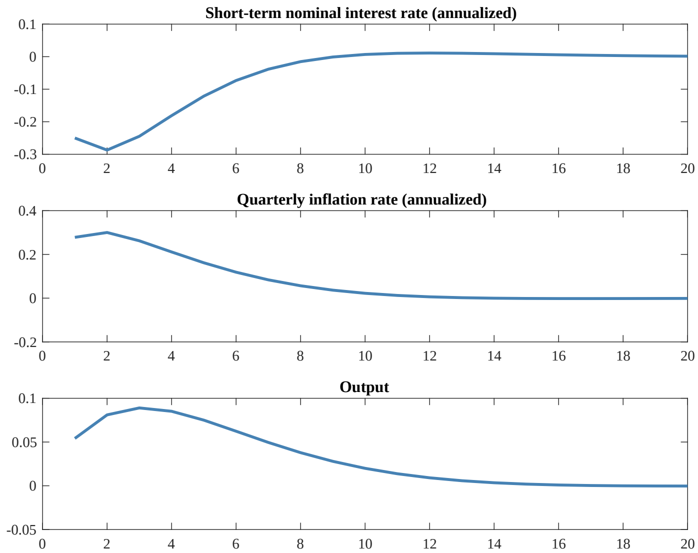

Note: Plain line : response to a monetary policy shock leading to -25 basis point cut in the nominal interest on impact. Inflation is the annualized quarterly growth rate of the price index. Interest rate is annualized.

C Illustrating the "lower for longer" property of the model policy rule

Figure C.1: Interest rate, inflation and output path in a recession with ELB scenario

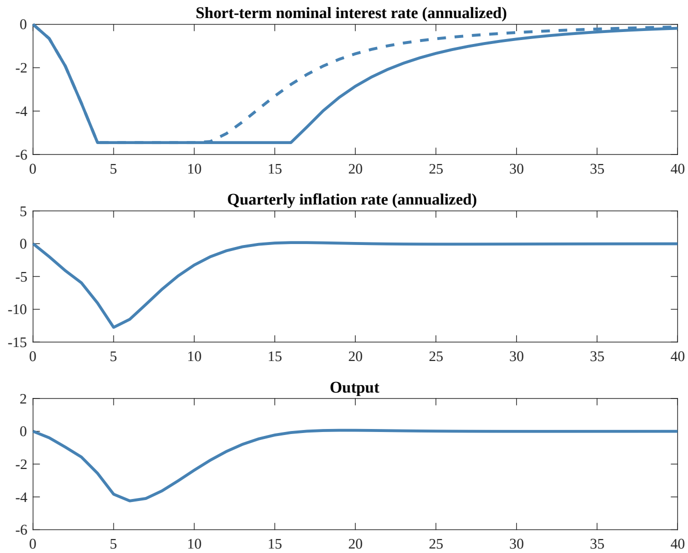

Note: Plain line : actual model policy rule. Dashed line: illustrative interest rule featuring actual rate lagged term rather than lagged "notional rate" term. The latter rule has no feedback on the model.

In this section, we illustrate how the "lower for longer" property of the model policy rule works in practice. To this end, we assume that the model starts in steady state and is hit by a series of unexpected risk-premium shocks that drive the economy to the ZLB. Given the implied path for inflation $\hat{\pi}_t$, the output gap $\hat{x}_t$, and the notional rate $\hat{\imath}^n_t$, we reconstruct the path of an alternative interest rate $\tilde{\imath}_t$ that would obey

$$\tilde{\imath}^n_t = \rho_i \tilde{\imath}_{t-1} + (1-\rho_i)(a_\pi \hat{\pi}_t + a_y \hat{x}_t) + \zeta_{R,t}$$

$$\tilde{\imath}_t = \max\{\tilde{\imath}^n_t, -(\mu_z + \rho + \pi)\}.$$

In this alternative specification, the notional rate does not depend on its lagged value but rather on the lagged value of the nominal interest rate. Away from the ZLB, this has no discernible effect. However, when the economy hits the ZLB, $\tilde{\imath}^n_t$ will mechanically increase sooner than $\hat{\imath}^n_t$. Figure C.1 reports the outcome of this simulation. The solid blue line shows the path of $\hat{\imath}_t$ while the dashed line shows the implied path for $\tilde{\imath}_t$.

D The distribution of ZLB spells duration

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

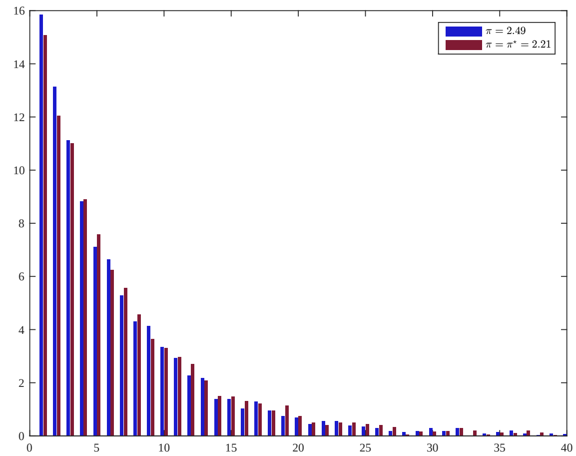

Note: Histograms are based on a simulated sample of 500,000 quarters. Simulations are carried out assuming in turn that the inflation target is the estimated inflation target; and then that the inflation target is the optimal inflation target obtained using the mean of the posterior density of estimated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

E The distribution of optimal inflation targets

Figure E.1: Posterior Distribution of $\pi^{\star}$

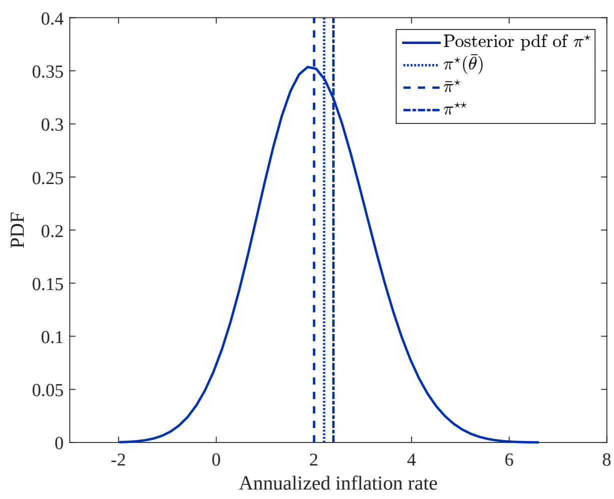

Note: Plain curve: PDF of $\pi^{\star}$; Dashed vertical line: Average value of $\pi^{\star}$ over posterior distribution; Dotted vertical line: Optimal inflation at the posterior mean of θ; Dashed-dotted vertical line : Bayesian-theoretic optimal inflation

F The welfare cost of inflation

Following a standard approach when assessing alternative policies, we complement our characterization of optimal inflation by providing measures of consumption-equivalent welfare gains/losses of choosing a suboptimal inflation target.

Let $\mathscr{W}(\pi)$ denote welfare under the inflation target π. It is defined as

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

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

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

Direct manipulations thus yield

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

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 $1 + \varphi(\pi)$ denote this percentage increase in consumption. Thus $\varphi(\pi)$ is such that

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

It then follows that

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

In practice, welfare is approximated to second order.

We compute $\varphi(\pi)$ under two alternative steady-state interest rate scenarios. In the first scenario, we set $r^\star$ to the baseline estimated value, corresponding to the posterior mean of $\mu_z + \rho$. In the second scenario, we consider a downward shift in μ_z by one percentage point (in annual terms), resulting in a lower steady-state real rate. The results are reported in Figure F.1. The blue lines show $\varphi(\pi)$ in the first scenario and the red lines show $\varphi(\pi)$ under a lower real interest rate. For ease of interpretation, the dashed, vertical lines indicate the optimal values of inflation under the two alternative interest rate scenarios.

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

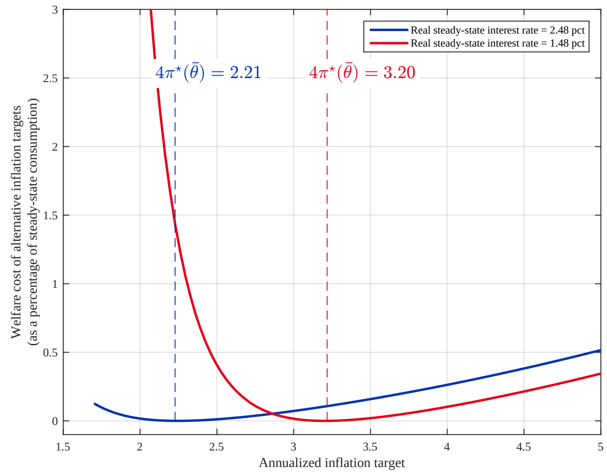

Note: The figure reports the welfare cost of inflation stated as a percentage of steady-state consumption in the optimal setting.

Figures F.1 suggests that in the baseline scenario, the welfare cost of raising or lowering the inflation target by one percentage point is relatively mild. However, this conclusion is not robust to a lower real interest rate. As the red line shows, with a one percentage rate lower $r^\star$, the welfare cost of inflation is asymmetric. It would be much costlier to lower the inflation target than to raise it in the neighborhood of the optimal target. In particular, keeping the inflation target unchanged when faced with a 1 percentage point decline in $r^\star$ gives rise to a 1.5% consumption loss.

G Further illustrations of the $(r^\star, \pi^\star)$ relation

G.1 When μ_z varies

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

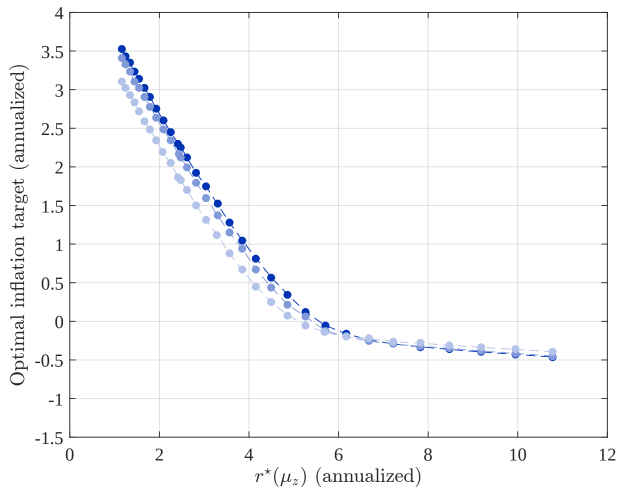

Note: Blue: parameters set at the posterior mean; light blue: parameters set at the posterior median; Lighter blue: parameters set at the posterior mode. Memo: $r^\star = \rho + \mu_z$. Range for μ_z: 0.4% to 10% (annualized) .

G.2 When ρ varies

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

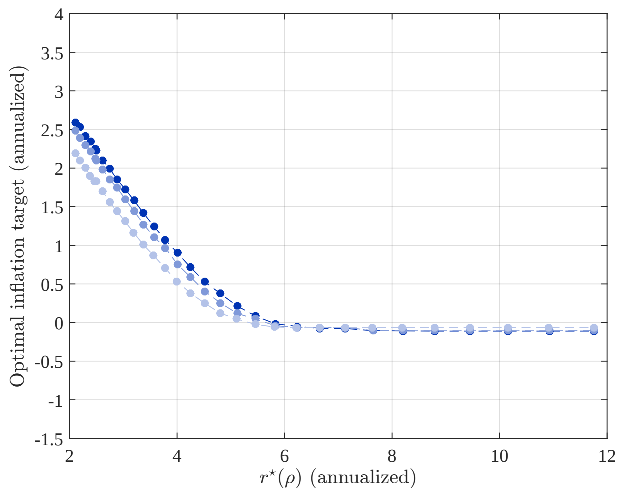

Note: Blue: parameters set at the posterior mean; light blue: parameters set at the posterior median; Lighter blue: parameters set at the posterior mode. Memo: $r^\star = \rho + \mu_z$. Range for μ_z: 0.4% to 10% (annualized) .

H Nominal and Real Interest Rates

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

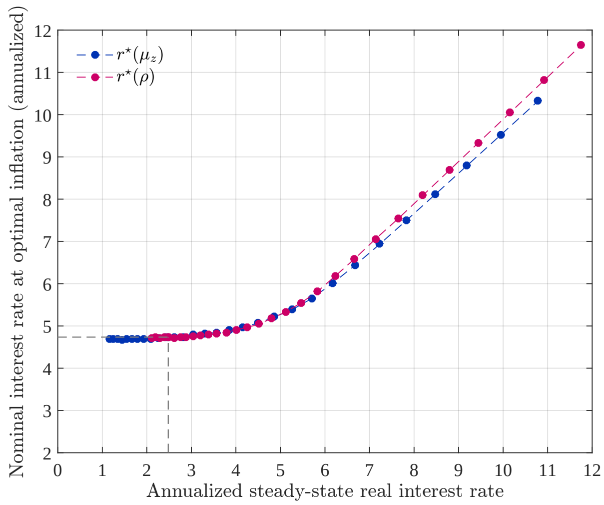

Note: the blue dots correspond to the $(r^\star, i^\star)$ locus when $r^\star$ varies with μ_z; the red dots correspond to the $(r^\star, i^\star)$ locus when $r^\star$ varies with ρ

I The probability of ZLB under large shocks

Figure I.1: Relation between probability of ZLB at optimal inflation and $r^\star$ (at the posterior mean)

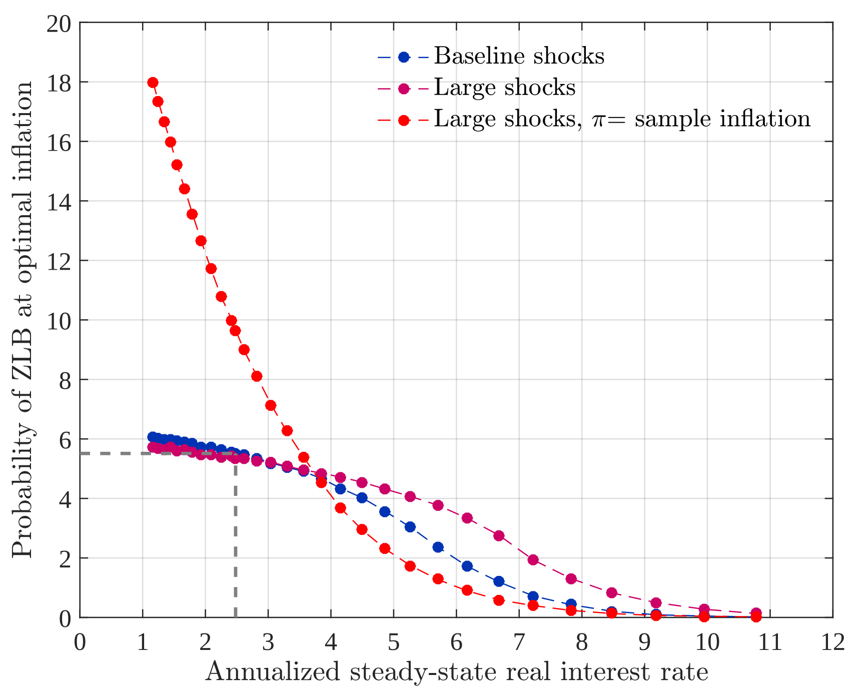

Note: the blue dots correspond to the $(r^\star, \pi^\star)$ locus when $r^\star$ varies with μ_z in the baseline; the red dots correspond to the locus in the "large shocks" case; the green dots correspond to the $(r^\star, \pi^\star)$ locus in the "large shocks" case when $\pi^\star$ is left at its sample inflation value.

J Distribution of $\pi^{\star}$ following a downward shift of the distribution of $r^{\star}$

Figure J.1: Counterfactual - US

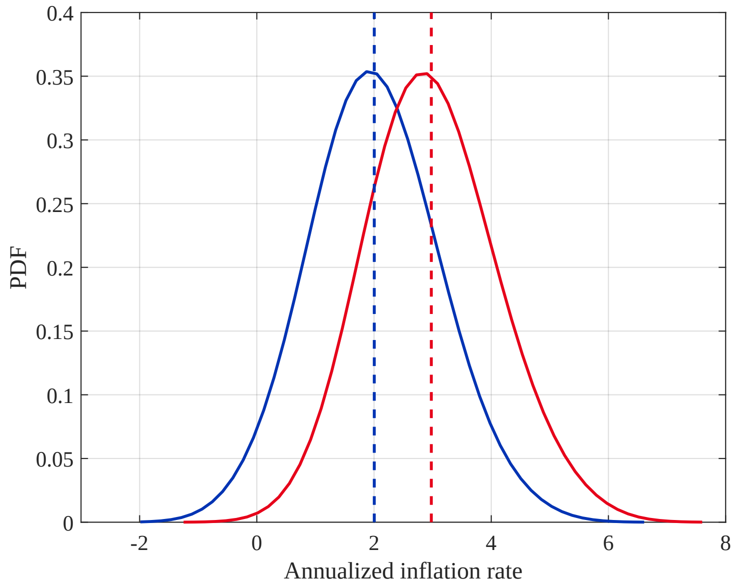

Note: The dashed vertical line indicates the mean value, i.e. $\mathbb{E}_{\theta}(\pi^{\star}(\theta))$.

K Model Solution

K.1 Households

K.1.1 First Order Conditions

The Lagrangian associated with the program (1) under constraint (2) is

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

The associated first-order condition with respect to bonds is

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

and the first-order condition with respect to consumption is

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

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

$$Z_t = \mathrm{e}^{z_t}$$

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 (K.2) thus rewrites

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

Similarly, equation (K.1) rewrites

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

where we defined

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

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 (K.3) yields

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

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 (K.4) yields

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

K.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)

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

K.2.1 Cost Minimization

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

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

The associated real marginal cost is thus

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

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

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

K.2.2 Price Setting of Intermediate Goods: Optimization

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

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

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 (K.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

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

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

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

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

Note that

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

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

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

and the auxiliary parameters

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

and

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

Finally, note 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

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

K.3 Unions

K.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

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

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

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

Note 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, note 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

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

K.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

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

Note 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

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

K.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 (that is, $\zeta_{u,t} = 0$). 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

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

K.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

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

Recall that we defined the natural rate of output as the level of production that would prevail in an economy without nominal rigidities, i.e. $\alpha_p = \alpha_w = 0$, and no cost-push shock. Under such circumstances, the steady-state value of the (normalized) natural rate of output Y_z^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

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

L 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.

L.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

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

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

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

Using

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

we obtain

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

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

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

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

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

Integrating over the set of labor types, one gets

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

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

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

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.

L.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, note 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

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

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 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

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

L.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

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

Note 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

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

L.4 Price and Wage Dispersions

We now derive the law of motion of price dispersion. Note 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

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

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

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

Note 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, note 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

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

For sufficiently small π, price dispersion $\Delta_{p,t}$ is second-order.

We now derive the law of motion of wage dispersion. Following similar steps as for price dispersion, note 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

$$\bar{w}_t - \bar{w}_{t-1} = \alpha_w(\gamma_z \mu_z + \gamma_w \pi_{t-1}) + (1-\alpha_w)[w_t^\star - \bar{w}_{t-1}]. \tag{L.2}$$

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

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

Note 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, note 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

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

For sufficiently small π and μ_z, wage dispersion $\Delta_{w,t}$ is second-order.

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 as

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

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

L.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

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

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

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

Assuming the distortions are themselves negligible, this simplifies further to

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

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

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \mathrm{e}^{\zeta_{c,t}} \log(C_{z,t} - \eta C_{z,t-1}\mathrm{e}^{-\zeta_{z,t}}) = \frac{1-\beta\eta}{1-\eta} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg[\tilde{y}_{z,t} - \frac{1}{2}[\varphi(1+\beta\eta^2) - 1]\tilde{y}_{z,t}^2 + \eta\varphi\tilde{y}_{z,t}\tilde{y}_{z,t-1}$$
$$+ \varphi\hat{g}_t\tilde{y}_{z,t} - \eta\varphi\zeta_{z,t}^*\tilde{y}_{z,t}\Bigg] + \text{t.i.p} + \mathcal{O}(||\zeta||^3),$$

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

$$\begin{aligned}\mathsf{U}_0 = \frac{1-\beta\eta}{1-\eta}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t \Bigg[& \Phi\tilde{y}_{z,t} - \frac{1}{2}[\varphi(1+\beta\eta^2)+\omega]\tilde{y}^2_{z,t} + \eta\varphi\tilde{y}_{z,t}\tilde{y}_{z,t-1} \\ & + (\varphi\hat{g}_t - \zeta_{h,t} - \varphi\eta\zeta^*_{z,t})\tilde{y}_{z,t} \\ & - \frac{1}{2}[(\phi-1)\theta_p+1]\theta_p\Delta_{p,t} - \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w\Delta_{w,t}\Bigg] + \text{t.i.p} + \mathcal{O}(||\zeta||^3),\end{aligned}$$

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

$$\begin{aligned}\mathsf{U}_0 = \frac{1-\beta\eta}{1-\eta}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t \Bigg[& \Phi\tilde{y}_{z,t} - \frac{1}{2}[\varphi(1+\beta\eta^2)+\omega]\tilde{y}^2_{z,t} + \eta\varphi\tilde{y}_{z,t}\tilde{y}_{z,t-1} \\ & + [\varphi(1+\beta\eta^2)+\omega]\hat{y}^n_{z,t}\tilde{y}_{z,t} - \varphi\beta\eta\hat{y}^n_{z,t+1}\tilde{y}_{z,t} - \varphi\eta\hat{y}^n_{z,t-1}\tilde{y}_{z,t} \\ & - \frac{1}{2}[(\phi-1)\theta_p+1]\theta_p\Delta_{p,t} - \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w\Delta_{w,t}\Bigg] + \text{t.i.p} + \mathcal{O}(||\zeta||^3),\end{aligned}$$

$$\begin{aligned}\mathsf{U}_0 = \frac{1-\beta\eta}{1-\eta}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t \Bigg[& \Phi\tilde{y}_t - \frac{1}{2}[\varphi(1+\beta\eta^2)+\omega]\tilde{y}^2_t + \eta\varphi\tilde{y}_t\tilde{y}_{t-1} \\ & + [\omega+\varphi(1+\beta\eta^2)]\hat{y}^n_t\tilde{y}_t - \varphi\beta\eta\hat{y}^n_{t+1}\tilde{y}_t - \varphi\eta\hat{y}^n_{t-1}\tilde{y}_t \\ & - \frac{1}{2}[(\phi-1)\theta_p+1]\theta_p\Delta_{p,t} - \frac{1}{2}(1+\nu\theta_w)\phi^{-1}\theta_w\Delta_{w,t}\Bigg] + \text{t.i.p} + \mathcal{O}(||\xi,\pi||^3)\end{aligned}$$

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

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

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

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

Developing yields

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

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

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

Thus

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

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

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

Identifying term by term, we obtain

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

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

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

Recall that the steady-state subsidy rates τ_p and τ_w are chosen to neutralize markups. Then, it follows that $\Phi = x^* = 0$.

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

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

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

Note 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 of expressing price and wage dispersions in terms of squared price and wage inflations.

Recall that

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

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

$$\lambda_y \equiv \delta_0$$

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

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

Using this and recalling that $x^*=0$, the second order approximations to welfare rewrites as

$$\mathsf{U}_0=-\frac{1}{2}\frac{1-\beta\eta}{1-\eta}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\Big\{\lambda_y[\hat{x}_t-\delta\hat{x}_{t-1}+(1-\delta)\bar{x}]^2+\lambda_p[(1-\gamma_p)\pi+\hat{\pi}_t-\gamma_p\hat{\pi}_{t-1}]^2$$

$$+\lambda_w[(1-\gamma_z)\mu_z+(1-\gamma_w)\pi+\hat{\pi}_{w,t}-\gamma_w\hat{\pi}_{t-1}]^2\Big\}+\text{t.i.p}+\mathcal{O}(||\zeta,\pi||^3),$$

where we defined

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

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

References

ADAM, K. AND H. WEBER (2019): "Optimal Trend Inflation," *American Economic Review*, 109, 702–737.

ALVAREZ, F., H. LE BIHAN, AND F. LIPPI (2016): "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, 106, 2817–2851.

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.

ANDRADE, P., G. GABALLO, E. MENGUS, AND B. MOJON (2019): "Forward Guidance and Heterogeneous Beliefs," *American Economic Journal: Macroeconomics*, 11, 1–29.

ASCARI, G., L. PHANEUF, AND E. R. SIMS (2018): "On the Welfare and Cyclical Implications of Moderate Trend Inflation," *Journal of Monetary Economics*, 99, 56–71.

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," 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., M. T. KILEY, AND J. M. ROBERTS (2019): "Monetary Policy Strategies for a Low-Rate Environment," *AEA Papers and Proceedings*, 109, 421–426.

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 (2017): "The Effects of Foreign Shocks when Interest Rates are at Zero," *Canadian Journal of Economics*, 50, 660–684.

BRAND, C., M. BIELECKI, AND A. PENALVER (2018): "The Natural Rate of Interest: Estimates, Drivers, and Challenges to Monetary Policy," Occasional paper 217, ECB.

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 (2018): "The Safety Trap," *Review of Economic Studies*, 85, 223–274.

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., E. GAGNON, T. NAKATA, M. PAUSTIAN, B. SCHLUSCHE, J. TREVINO, D. VILAN, AND W. ZHENG (2019): "Monetary Policy Options at the Effective Lower Bound : Assessing the Federal Reserve's Current Policy Toolkit," Finance and Economics Discussion Series 2019-003, Board of Governors of the Federal Reserve System.

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.

CLARIDA, R. H. (2019): "The Federal Reserve's Review of Its Monetary Policy Strategy, Tools, and Communication Practices," Speech at the 2019 U.S. Monetary Policy Forum, sponsored by the Initiative on Global Markets at the University of Chicago Booth School of Business, New York.

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 (2019): "On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint," in *NBER Macroeconomics Annual 2019*, vol. 34.

DEL NEGRO, M., D. GIANNONE, M. P. GIANNONI, AND A. TAMBALOTTI (2017): "Safety, Liquidity, and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, 48 (Spring).

——— (2018): "Global Trends in Interest Rates," Working Papers 25039, NBER.

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," *Annual Review of Economics*, 8, 497–520.

EBERLY, J. C., J. H. STOCK, AND J. H. WRIGHT (2019): "The Federal Reserve's Current Framework for Monetary Policy: A Review and Assessment," Working Papers 26002, NBER.

EGGERTSSON, G. B., N. R. MEHROTRA, AND J. A. ROBBINS (2017): "A Model of Secular Stagnation: Theory and Quantitative Evaluation," *American Economic Journal: Macroeconomics*, 11, 1–48.

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.

FARHI, E. AND F. GOURIO (2018): "Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia," Working Papers 25282, NBER.

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

FUHRER, J., G. OLIVEI, E. ROSENGREN, AND G. TOOTELL (2018): "Should the Fed Regularly Evaluate Its Monetary Policy Framework?" *Brookings Papers on Economic Activity*, 51 (Fall).

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., E. MAUSKOPF, AND D. WILCOX (2007): "Issues Pertaining to the Specification of a Numerical Price-Related Objective For Monetary Policy," FOMC memo.

KILEY, M. T. (2019): "The Optimal Inflation Target and the Natural Rat of Interest: Comment," *Brookings Papers on Economic Activity*, 53 (Fall).

KILEY, M. T. AND J. M. ROBERTS (2017): "Monetary Policy in a Low Interest Rate World," *Brookings Papers on Economic Activity*, 48 (Spring), 317–396.

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. (2018): "The Optimal Inflation Rate with Discount Factor Heterogeneity," Finance and Economics Discussion Series 2018-086, Board of Governors of the Federal Reserve System (US).

LEVIN, A. T., A. ONATSKI, J. WILLIAMS, AND N. M. WILLIAMS (2006): "Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models," in *NBER Macroeconomics Annual 2005, Volume 20*, National Bureau of Economic Research, Inc, NBER Chapters, 229–312.

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 (2018): "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation," *The Quarterly Journal of Economics*, 133, 1933–1980.

RACHEL, L. AND L. H. SUMMERS (2019): "On Falling Neutral Real Rates, Fiscal Policy, and the Risk of Secular Stagnation," *Brookings Papers on Economic Activity*, 52 (Spring).

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.

SIMS, E. AND C. WU (2019): "Evaluating Central Banks' Tool Kit: Past, Present, and Future," Working Papers 26040, NBER.

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

SWANSON, E. T. (2018): "The Federal Reserve Is Not Very Constrained by the Lower Bound on Nominal Interest Rates," *Brookings Papers on Economic Activity*, 51 (Fall).

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. (2009): "Heeding Daedalus: Optimal Inflation and the Zero Lower Bound," *Brookings Papers on Economic Activity*, 40 (Fall), 1–49.

——— (2013): "A Defense of Moderation in Monetary Policy," *Journal of Macroeconomics*, 38, 137–150.

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

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