

The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?

Olivier Coibion
College of William and Mary

Yuriy Gorodnichenko
U.C. Berkeley
and NBER

Johannes Wieland
U.C. Berkeley

August 21st, 2011

Abstract: We study the effects of positive steady-state inflation in New Keynesian models subject to the zero bound on interest rates. We derive the utility-based welfare loss function taking into account the effects of positive steady-state inflation and solve for the optimal level of inflation in the model. For plausible calibrations with costly but infrequent episodes at the zero-lower bound, the optimal inflation rate is low, typically less than two percent, even after considering a variety of extensions, including optimal stabilization policy, price indexation, endogenous and state-dependent price stickiness, capital formation, model-uncertainty, and downward nominal wage rigidities. On the normative side, price level targeting delivers large welfare gains and a very low optimal inflation rate consistent with price stability. These results suggest that raising the inflation target is too blunt an instrument to efficiently reduce the severe costs of zero-bound episodes.

Keywords: Optimal inflation, New Keynesian, zero bound, price level targeting.
JEL codes: E3, E4, E5.

We are grateful to anonymous referees, the Editor, Roberto Billi, Ariel Burstein, Gauti Eggertsson, Jordi Gali, Marc Gianonni, Christian Hellwig, David Romer, Eric Sims, Alex Wolman, and seminar participants at John Hopkins University, Bank of Canada, College of William and Mary, NBER Summer Institute in Monetary Economics and Economic Fluctuations and Growth, University of Wisconsin, Society for Computational Economics, Richmond Fed and UNC Chapel Hill for helpful comments.

“The crisis has shown that interest rates can actually hit the zero level, and when this happens it is a severe constraint on monetary policy that ties your hands during times of trouble. As a matter of logic, higher average inflation and thus higher average nominal interest rates before the crisis would have given more room for monetary policy to be eased during the crisis and would have resulted in less deterioration of fiscal positions. What we need to think about now is whether this could justify setting a higher inflation target in the future.”

Olivier Blanchard, February 12th, 2010

I Introduction

One of the defining features of the current economic crisis has been the zero bound on nominal interest rates. With standard monetary policy running out of ammunition in the midst of one of the sharpest downturns in post-World War II economic history, some have suggested that central banks should consider allowing for higher target inflation rates than would have been considered reasonable just a few years ago. We contribute to this question by explicitly incorporating positive steady-state (or “trend”) inflation into New Keynesian models as well as the zero lower bound (ZLB) on nominal interest rates. We derive the effects of non-zero steady-state inflation on the loss function, thereby laying the groundwork for welfare analysis. While hitting the ZLB is very costly in the model, our baseline finding is that the optimal rate of inflation is low, typically less than two percent a year, even when we allow for features that lower the costs or raise the benefits of positive steady-state inflation.

Despite the importance of quantifying the optimal inflation rate for policy-makers, modern monetary models of the business cycle, namely the New Keynesian framework, have been strikingly ill-suited to address this question because of their near exclusive reliance on the assumption of zero steady-state inflation, particularly in welfare analysis. Our first contribution is to address the implications of positive steady-state inflation for welfare analysis by solving for the micro-founded loss function in an otherwise standard New Keynesian model with labor as the only factor of production. We identify three distinct costs of positive trend inflation. The first is the steady-state effect: with staggered price setting, higher inflation leads to greater price dispersion which causes an inefficient allocation of resources among firms, thereby lowering aggregate welfare. The second is that positive steady-state inflation raises the welfare cost of a given amount of inflation volatility. This cost reflects the fact that inflation variations create distortions in relative prices given staggered price setting. Since positive trend inflation already generates some inefficient price dispersion, the additional distortion in relative prices from an inflation shock becomes more costly as firms have to compensate workers for the increasingly high marginal disutility of sector-specific labor. Thus, the increased distortion in relative prices due to an inflation shock becomes costlier as we increase the initial price dispersion which makes the variance of inflation costlier for welfare as the steady-state level of inflation rises. In addition to the two costs from relative price dispersion, a third cost of inflation in our model comes from the dynamic effect of positive inflation on pricing decisions. Greater steady-state inflation induces more forward-looking behavior when sticky-price firms are able to reset their prices because the gradual depreciation of the relative reset price can

lead to larger losses than under zero inflation. As a result, inflation becomes more volatile which lowers aggregate welfare. This cost of inflation arising from the positive relationship between the level and volatility of inflation has been well-documented empirically but is commonly ignored in quantitative analyses because of questions as to the source of the relationship.¹ As with the price-dispersion costs of inflation, this cost arises endogenously in the New Keynesian model when one incorporates positive steady-state inflation.

The key benefit of positive inflation in our model is a reduced frequency of hitting the zero bound on nominal interest rates. As emphasized in Christiano et al. (2011), hitting the zero bound induces a deflationary mechanism which leads to increased volatility and hence large welfare costs. Because a higher steady-state level of inflation implies a higher level of nominal interest rates, raising the inflation target can reduce the incidence of zero-bound episodes, as suggested by Blanchard in the opening quote. Our approach for modeling the zero bound follows Bodenstein et al. (2009) by solving for the duration of the zero bound endogenously, unlike in Christiano et al. (2011) or Eggertsson and Woodford (2004). This is important because the welfare costs of inflation are a function of the variance of inflation and output, which themselves depend on the frequency at which the zero bound is reached as well as the duration of zero bound episodes.

After calibrating the model to broadly match the moments of macroeconomic series and the historical incidence of hitting the zero lower bound in the U.S., we then solve for the rate of inflation that maximizes welfare. While the ZLB ensures that the optimal inflation rate is positive, for plausible calibrations of the structural parameters of the model and the properties of the shocks driving the economy, the optimal inflation rate is quite low: typically less than two percent per year. This result is remarkably robust to changes in parameter values, as long as these do not dramatically increase the implied frequency of being at the zero lower bound. In addition, we show that our results are robust if the central bank follows optimal stabilization policy, rather than the baseline Taylor rule. In particular, if the central bank cannot commit to a policy rule, then the optimal inflation rate remains within the range of inflation rates targeted by central banks and is of qualitatively similar magnitude as in our baseline calibration. Furthermore, we show that all three costs of inflation—the steady state effect, the increasing cost of inflation volatility, and the positive link between the level and volatility of inflation—are quantitatively important: each cost is individually large enough to bring the optimal inflation rate down to 3.6% or lower when the ZLB is present.

The key intuition behind the low optimal inflation rate is that the unconditional cost of the zero lower bound is small even though each individual ZLB event is quite costly. In our baseline calibration, an 8-quarter ZLB event at 2% trend inflation has a cost equivalent to a 6.2% permanent reduction in consumption, above and beyond the usual business cycle cost. This is, for example, significantly higher than Williams' (2009) estimate of the costs of hitting the ZLB during the current recession. However, in the model such an event is

¹ For example, Mankiw's (2007) undergraduate Macroeconomics textbook notes that "in thinking about the costs of inflation, it is important to note a widely documented but little understood fact: high inflation is variable inflation."

also rare, occurring about once every 20 years assuming that ZLB events always last 8 quarters, so that the *unconditional* cost of the ZLB at 2% trend inflation is equivalent to a 0.08% permanent reduction in consumption. This leaves little room for further improvements in welfare by raising the long-run inflation rate. Thus, even modest costs of trend inflation, *which must be borne every period*, will imply an optimal inflation rate below 2%, despite reasonable values for both the frequency and cost of the ZLB. This explains why our results are robust to a variety of settings that we further discuss below and suggests that our results are not particular to the New Keynesian model.

Furthermore, while the New Keynesian model implies that the optimal weight on the variance of the output gap in the welfare loss function is small, we show that increasing the weight on the output gap to be more than ten times as large as that on the annualized inflation variance would still leave the optimal inflation rate at less than 2.5%. Thus, it is unlikely that augmenting the baseline model with mechanisms which could raise the welfare cost of output fluctuations (such as involuntary unemployment or income disparities across agents) would significantly raise the optimal target rate of inflation. Finally, while we use historical U.S. data to calibrate the frequency of hitting the ZLB, an approach which can be problematic when applied to rare events, we show in robustness analysis that even a tripling of the frequency of being at the ZLB (such that the economy would spend 15% of the time at the ZLB for an inflation rate of 3%) would raise the optimal inflation rate only to 3% which is the upper bound of most central banks' inflation targets.

To further investigate the robustness of this result, we extend our baseline model to consider several mechanisms which might raise the optimal rate of inflation. First, in the presence of uncertainty about underlying parameter values, policy-makers might want to choose a higher target inflation rate as a buffer against the possibility that the true parameters imply more frequent and costly incidence of the zero bound. Incorporating this uncertainty only raises the optimal inflation rate from 1.5% to 1.9% per year. Second, one might be concerned that our findings hinge on modeling price stickiness as in Calvo (1983). Since this approach implies that some firms do not change prices for extended periods of time, it could overstate the cost of price dispersion and therefore understate the optimal inflation rate. To address this possibility, we reproduce our analysis using Taylor (1977) staggered price setting of fixed durations. The latter reduces price dispersion relative to the Calvo assumption but raises the optimal inflation rate to only 2.2% when prices are changed every three quarters. Another limitation of the Calvo assumption is that the rate at which prices are changed is commonly treated as a structural parameter, yet the frequency of price setting may depend on the inflation rate, even for low inflation rates like those experienced in the U.S. As a result, we consider two modifications that allow for price flexibility to vary with the trend rate of inflation. In the first specification, we let the degree of price rigidity vary systematically with the trend level of inflation. In the second specification, we employ the Dotsey et al. (1999) model of state-dependent pricing, which allows the degree of price stickiness to vary *endogenously* both in the short-run and in the long-run, and thus we address one of

the major criticisms of the previous literature on the optimal inflation rate. Both modifications yield optimal inflation rates of less than two percent per year. Finally, we incorporate downward nominal wage rigidity, which Tobin (1972) suggests might push the optimal inflation rate higher by facilitating the downward adjustment of real wages. This “greasing the wheels” effect, however, significantly *lowers* the optimal inflation rate by lowering the volatility of marginal costs and hence of inflation.

Our analysis abstracts from several other factors which might affect the optimal inflation rate. For example, Friedman (1969) argued that the optimal rate of inflation must be negative to equalize the marginal cost and benefit of holding money. Because our model is that of a cashless economy, this cost of inflation is absent, but would tend to lower the optimal rate of inflation even further, as emphasized by Khan et al. (2003), Schmitt-Grohe and Uribe (2007, 2010) and Aruoba and Schorfheide (2011). Similarly, a long literature has studied the costs and benefits of the seigniorage revenue to policymakers associated with positive inflation, a feature which we also abstract from since seigniorage revenues for countries like the U.S. are quite small, as are the deadweight losses associated with it (Cooley and Hansen 1991, Summers 1991). Feldstein (1997) emphasizes an additional cost of inflation arising from fixed nominal tax brackets, which would again lower the optimal inflation rate. Finally, because we do not model the possibility of endogenous countercyclical fiscal policy nor do we incorporate the possibility of nonstandard monetary policy actions during ZLB episodes, it is likely that we overstate the costs of hitting the ZLB and therefore again overstate the optimal rate of inflation. Nevertheless, our finding that the threat of the ZLB coupled with limited commitment on the part of the central bank implies positive but low optimal inflation rates, goes some way in resolving the “*puzzle*” pointed out by Schmitt-Grohe and Uribe (2010) that existing monetary theories routinely imply negative optimal inflation rates, and thus cannot explain the size of observed inflation targets.

This paper is closely related to recent work that has also emphasized the effects of the zero bound on interest rates for the optimal inflation rate, such as Walsh (2009), Billi (2011), and Williams (2009). A key difference between the approach taken in this paper and such previous work is that we explicitly model the effects of positive trend inflation on the steady-state, dynamics, and loss function of the model. Billi (2011) and Walsh (2009), for example, use a New Keynesian model log-linearized around zero steady-state inflation and therefore do not explicitly incorporate the positive relationship between the level and volatility of inflation, while Williams (2009) relies on a non-microfounded model. In addition, these papers do not take into account the effects of positive steady-state inflation on the approximation to the utility function and thus do not fully incorporate the costs of inflation arising from price dispersion. Schmitt-Grohe and Uribe (2010) provide an authoritative treatment of many of the costs and benefits of trend inflation in the context of New Keynesian models. However, their calibration implies that the chance of hitting the ZLB is practically zero and therefore does not quantitatively affect the optimal rate of inflation, whereas we focus on a setting where

costly ZLB events occur at their historic frequency. Furthermore, none of these papers consider the endogenous nature of price rigidity with respect to trend inflation.

An advantage of working with a micro-founded model and its implied welfare function is the ability to engage in normative analysis. In our baseline model, the endogenous response of monetary policy-makers to macroeconomic conditions is captured by a Taylor rule. Thus, we are also able to study the welfare effects of altering the systematic response of policy-makers to endogenous fluctuations (i.e. the coefficients of the Taylor rule) and determine the new optimal steady-state rate of inflation. The most striking finding from this analysis is that even modest price-level targeting (PLT) would raise welfare by non-trivial amounts for any steady-state inflation rate and come close to the Ramsey-optimal policy, consistent with the finding of Eggertsson and Woodford (2003) and Wolman (2005). In short, the optimal policy rule for the model can be closely characterized by the name of “price stability” as typically stated in the legal mandates of most central banks.

Given our results, we conclude that raising the target rate of inflation is likely too blunt an instrument to reduce the incidence and severity of zero-bound episodes. In all of the New Keynesian models we consider, even the small costs associated with higher trend inflation rates, which must be borne every period, more than offset the welfare benefits of fewer and less severe ZLB events. Instead, changes in the policy rule, such as PLT, may be more effective both in avoiding and minimizing the costs associated with these crises. In the absence of such changes to the interest rate rule, our results suggest that addressing the large welfare losses associated with the ZLB is likely to best be pursued through policies targeted specifically to these episodes, such as countercyclical fiscal policy or the use of non-standard monetary policy tools.

Section 2 presents the baseline New Keynesian model and derivations when allowing for positive steady-state inflation, including the associated loss function. Section 3 includes our calibration of the model as well as the results for the optimal rate of inflation while section 4 investigates the robustness of our results to parameter values. Section 5 then considers extensions of the baseline model which could potentially lead to higher estimates of the optimal inflation target. Section 6 considers additional normative implications of the model, including optimal stabilization policy and price level targeting. Section 7 concludes.

II A New Keynesian Model with Positive Steady-State Inflation

We consider a standard New Keynesian model with a representative consumer, a continuum of monopolistic producers of intermediate goods, a fiscal authority and a central bank.

2.1 Model

The representative consumer maximizes the present discounted value of the utility stream from consumption and leisure

$$\max E_t \sum_{j=0}^{\infty} \beta^j \left\{ \log(C_{t+j} - hGA_{t+j}C_{t+j-1}) - \frac{\eta}{\eta+1} \int_0^1 N_{t+j}(i)^{1+1/\eta} di \right\} \quad (1)$$

where C is consumption of the final good, $N(i)$ is labor supplied to individual industry i , GA is the gross growth rate of technology, η is the Frisch labor supply elasticity, h the internal habit parameter and β is the discount factor.² The budget constraint each period t is given by

$$\xi_t: C_t + S_t/P_t \leq \int_0^1 (N_t(i)W_t(i)/P_t)di + S_{t-1}q_{t-1}R_{t-1}/P_t + T_t \quad (2)$$

where S is the stock of one-period bonds held by the consumer, R is the gross nominal interest rate, P is the price of the final good, $W(i)$ is the nominal wage earned from labor in industry i , T is real transfers and profits from ownership of firms, q is a risk premium shock, and ξ_t is the shadow value of wealth.³ The first order conditions from this utility-maximization problem are then:

$$(C_t - hGA_t C_{t-1})^{-1} - \beta hE_t GA_{t+1} (C_{t+1} - hGA_{t+1} C_t)^{-1} = \xi_t, \quad (3)$$

$$N_t(i)^{1/\eta} = \xi_t W_t(i)/P_t, \quad (4)$$

$$\xi_t/P_t = \beta E_t [\xi_{t+1} q_t R_t/P_{t+1}]. \quad (5)$$

Production of the final good is done by a perfectly competitive sector which combines a continuum of intermediate goods into a final good per the following aggregator

$$Y_t = \left[\int_0^1 Y_t(i)^{(\theta-1)/\theta} di \right]^{\theta/(\theta-1)} \quad (6)$$

where Y is the final good and $Y(i)$ is intermediate good i , while θ denotes the elasticity of substitution across intermediate goods, yielding the following demand curve for goods of intermediate sector i

$$Y_t(i) = Y_t (P_t(i)/P_t)^{-\theta} \quad (7)$$

and the following expression for the aggregate price level

$$P_t = \left[\int_0^1 P_t(i)^{(1-\theta)} di \right]^{1/(1-\theta)}. \quad (8)$$

The production of each intermediate good is done by a monopolist facing a production function linear in labor

$$Y_t(i) = A_t N_t(i) \quad (9)$$

where A denotes the level of technology, common across firms. Each intermediate good producer has sticky prices, modeled as in Calvo (1983) where $1 - \lambda$ is the probability that each firm will be able to reoptimize its price each period. We allow for indexation of prices to steady-state inflation by firms who do not reoptimize their prices each period, with ω representing the degree of indexation (0 for no indexation to 1 for full indexation). Denoting the optimal reset price of firm i by $B(i)$, re-optimizing firms solve the following profit-maximization problem

$$\max E_t \sum_{j=0}^{\infty} \lambda^j Q_{t,t+j} [Y_{t+j}(i) B_t(i) \bar{\pi}^{j\omega} - W_{t+j}(i) N_{t+j}(i)] \quad (10)$$

² We use internal habits rather than external habits because they more closely match the (lack of) persistence in consumption growth in the data. The gross growth rate of technology enters the habit term to simplify derivations.

³ As discussed in Smets and Wouters (2007), a positive shock to q , which is the wedge between the interest rate controlled by the central bank and the return on assets held by the households, increases the required return on assets and reduces current consumption. The shock q has similar effects as net-worth shocks in models with financial accelerators. Amano and Shukayev (2010) document that shocks like q are essential for generating a binding zero lower bound.

where Q is the stochastic discount factor and $\bar{\Pi}$ is the gross steady-state level of inflation. The optimal relative reset price is then given by

$$\frac{B_t(i)}{P_t} = \frac{\theta}{\theta-1} \frac{E_t \sum_{j=0}^{\infty} \lambda^j Q_{t,t+j} Y_{t+j} (P_{t+j}/P_t)^{\theta+1} \bar{\Pi}^{-j\omega\theta} (MC(i)_{t+j}/P_{t+j})}{E_t \sum_{j=0}^{\infty} \lambda^j Q_{t,t+j} Y_{t+j} (P_{t+j}/P_t)^{\theta} \bar{\Pi}^{-j\omega(\theta-1)}} \quad (11)$$

where firm-specific marginal costs can be related to aggregate variables using

$$\frac{MC_{t+j}(i)}{P_{t+j}} = \left(\frac{\xi_{t+j}^{-1}}{A_{t+j}} \right) \left(\frac{Y_{t+j}}{A_{t+j}} \right)^{1/\eta} \left(\frac{B_t(i)}{P_t} \right)^{-\theta/\eta} \left(\frac{P_{t+j}}{\bar{\Pi}^{j\omega P_t}} \right)^{\theta/\eta}. \quad (12)$$

Given these price-setting assumptions, the dynamics of the price level are governed by

$$P_t^{1-\theta} = (1-\lambda)B_t^{1-\theta} + \lambda P_{t-1}^{1-\theta} \bar{\Pi}^{\omega(1-\theta)}. \quad (13)$$

We allow for government consumption of final goods (G), so the goods market clearing condition for the economy is

$$Y_t = C_t + G_t. \quad (14)$$

We define the aggregate labor input as

$$N_t = \left[\int_0^1 N_t(i)^{(\theta-1)/\theta} di \right]^{\theta/(\theta-1)}. \quad (15)$$

2.2 Steady-state and log-linearization

Following Coibion and Gorodnichenko (2011), we log-linearize the model around the steady-state in which inflation need not be zero. Since positive trend inflation may imply that the steady state and the flexible price level of output differ, we adopt the following notational convention. Variables with a bar denote steady state values, e.g. \bar{Y}_t is the steady state level of output. Lower-case letters denote the log of a variable, e.g. $y_t = \log Y_t$ is the log of current output. We assume that technology is a random walk and hence we normalize all non-stationary real variables by the level of technology. We let hats on lower case letters denote log deviations from steady state, e.g. $\hat{y}_t = y_t - \bar{y}_t$ is the approximate percentage deviation of output from steady state. Since we define the steady state as embodying the current level of technology, deviations from the steady state are stationary. Finally, we denote deviations from the flexible price level steady state with a tilde, e.g. $\tilde{y}_t = y_t - \bar{y}_t^F$ is the approximate percentage deviation of output from its flexible price steady state, where the superscript F denotes a flexible price level quantity. Define the net steady-state level of inflation as $\bar{\pi} = \log(\bar{\Pi})$. The log-linearized consumption Euler equation is

$$\hat{\xi}_t = E_t[\hat{\xi}_{t+1} + \hat{r}_t - \hat{\pi}_{t+1} + \hat{q}_t] \quad (16)$$

where the marginal utility of consumption is given by

$$\hat{\xi}_t = \frac{h}{(1-h)(1-\beta h)} \hat{c}_{t-1} - \frac{1-\beta h^2}{(1-h)(1-\beta h)} \hat{c}_t + \frac{\beta h}{(1-h)(1-\beta h)} E_t \hat{c}_{t+1}$$

and the goods market clearing condition becomes

$$\hat{y}_t = \bar{c}_y \hat{c}_t + \bar{g}_y \hat{g}_t \quad (17)$$

where \bar{c}_y and \bar{g}_y are the steady-state ratios of consumption and government to output respectively. Also, integrating over firm-specific production functions and log-linearizing yields

$$\hat{y}_t = \hat{n}_t. \quad (18)$$

Allowing for positive steady-state inflation (i.e., $\bar{\pi} > 0$) primarily affects the steady-state and price-setting components of the model. For example, the steady-state level of the output gap (which is defined as the deviation of steady state output from its flexible price level counterpart $\bar{X}_t = \bar{Y}_t/\bar{Y}_t^F$) is given by

$$\bar{X}^{(\eta+1)/\eta} = \frac{1-\lambda\beta\bar{\pi}^{(1-\omega)\theta(\eta+1)/\eta}}{1-\lambda\beta\bar{\pi}^{(1-\omega)(\theta-1)}} \left(\frac{1-\lambda}{1-\lambda\bar{\pi}^{(1-\omega)(\theta-1)}} \right)^{(\eta+\theta)/(\eta(\theta-1))}. \quad (19)$$

Note that the steady-state level of the gap is equal to one when steady-state inflation is zero (i.e., $\bar{\pi} = 1$) or when the degree of price indexation is exactly equal to one. As emphasized by Ascari and Ropele (2007), there is a non-linear relationship between the steady-state levels of inflation and output. For very low but positive trend inflation, \bar{X} is increasing in trend inflation but the sign is quickly reversed so that \bar{X} is falling with trend inflation for most positive levels of trend inflation.

Secondly, positive steady-state inflation affects the relationship between aggregate inflation and the re-optimizing price. Specifically, the relationship between the two in the steady state is now given by

$$\overline{(B/P)} = \left(\frac{1-\lambda}{1-\lambda\bar{\pi}^{(1-\omega)(\theta-1)}} \right)^{1/(\theta-1)} \quad (20)$$

and the log-linearized equation is described by

$$\hat{\pi}_t = \left(\frac{1-\lambda\bar{\pi}^{(1-\omega)(\theta-1)}}{\lambda\bar{\pi}^{(1-\omega)(\theta-1)}} \right) \hat{b}_t \Rightarrow \hat{b}_t = M\hat{\pi}_t \quad (21)$$

so that inflation is less sensitive to changes in the re-optimizing price as steady-state inflation rises because goods with high relative prices receive a smaller share of expenditures.

Similarly, positive steady-state inflation has important effects on the log-linearized optimal reset price equation, which is given by

$$\begin{aligned} \left(1 + \frac{\theta}{\eta}\right) \hat{b}_t &= (1 - \gamma_2) \sum_{j=0}^{\infty} \gamma_2^j \left(\frac{1}{\eta} E_t \hat{y}_{t+j} - E_t \hat{\xi}_{t+j} \right) + E_t \sum_{j=1}^{\infty} (\gamma_2^j - \gamma_1^j) (\widehat{\mathcal{G}}_{t+j} + \hat{r}_{t+j-1}) \\ &+ \sum_{j=1}^{\infty} \left[\gamma_2^j \left(1 + \frac{\theta(\eta+1)}{\eta} \right) - \gamma_1^j \theta \right] E_t \hat{\pi}_{t+j} + \hat{m}_t \end{aligned} \quad (22)$$

where \hat{m}_t is a cost-push shock, $\gamma_1 = \lambda\beta\bar{\pi}^{(1-\omega)(\theta-1)}$ and $\gamma_2 = \gamma_1\bar{\pi}^{(1-\omega)(1+\theta/\eta)}$ so that without steady-state inflation or full indexation we have $\gamma_1 = \gamma_2$. When $\omega < 1$, a higher $\bar{\pi}$ increases the coefficients on future output and inflation but also leads to the inclusion of a new term composed of future differences between output growth and interest rates. Each of these effects makes price-setting decisions more forward-looking. The increased coefficient on expectations of future inflation, which reflects the expected future depreciation of the reset price and the losses associated with it, plays a particularly important role. In response to an inflationary shock, a firm which can reset its price will expect higher inflation today and in the future as other firms update their prices in response to the shock. Given this expectation, the more forward looking a firm is (the higher is $\bar{\pi}$), the greater the optimal reset price must be in anticipation of other firms raising their prices in

the future. Thus, reset prices become more responsive to current shocks with higher $\bar{\pi}$. We confirm numerically that this effect dominates the reduced sensitivity of inflation to the reset price in equation (21), thereby endogenously generating a positive relationship between the level and the volatility of inflation.

To close the model, we assume that the log deviation of the desired gross interest rate from its steady state value (\hat{r}_t^*) follows a Taylor rule

$$\hat{r}_t^* = \rho_1 \hat{r}_{t-1}^* + \rho_2 \hat{r}_{t-2}^* + (1 - \rho_1 - \rho_2) [\phi_\pi(\pi_t - \pi^*) + \phi_y(y_t - y^*) + \phi_{gy}(gy_t - gy^*) + \phi_p(p_t - p_t^*)] + \varepsilon_t^r$$

where $\phi_\pi, \phi_y, \phi_{gy}, \phi_p$ capture the strength of the policy response to deviations of inflation, the output gap, the output growth rate and the price level from their respective targets, parameters ρ_1 and ρ_2 reflect interest rate smoothing, while ε_t^r is a policy shock. We set $\pi^* = \bar{\pi}, p_t^* = \pi^* t = \bar{\pi} t, y^* = \bar{y}$ and $gy^* = \bar{g}\bar{y}$ so that the central bank has no inflationary or output bias. The growth rate of output is related to the output gap by

$$\widehat{gy}_t = \hat{y}_t - \hat{y}_{t-1} \quad (23)$$

Since the actual level of the net interest rate is bounded by zero, the log deviation of the gross interest rate is bounded by $\hat{r}_t = \log(R_t) - \log(\bar{R}) \geq -\log(\bar{R}) = -\bar{r}$ with the dynamics of the actual interest rate given by

$$\hat{r}_t = \max \{ \hat{r}_t^*, -\bar{r} \}. \quad (24)$$

We consider the Taylor rule a reasonable benchmark, because it is likely to be the closest description of the current policy process, and because suggestions to raise the optimal inflation rate are not commonly associated with simultaneous changes in the way that stabilization policy is conducted. However, in section 6.1, we also derive the optimal $\bar{\pi}$ given optimal stabilization policy under discretion and commitment.

2.3 Shocks

We assume that technology follows a random walk process with drift:

$$a_t = a_{t-1} + \mu + \varepsilon_t^a. \quad (25)$$

Each of the risk premium, government, and Phillips Curve shocks follow AR(1) processes

$$\hat{q}_t = \rho_q \hat{q}_{t-1} + \varepsilon_t^q, \quad (26)$$

$$\hat{g}_t = \rho_g \hat{g}_{t-1} + \bar{g}_y^{-1} \varepsilon_t^g, \quad (27)$$

$$\hat{m}_t = \rho_m \hat{m}_{t-1} + \varepsilon_t^m. \quad (28)$$

We assume that $\varepsilon_t^a, \varepsilon_t^q, \varepsilon_t^g, \varepsilon_t^m, \varepsilon_t^r$ are mutually and serially uncorrelated.

2.4 Welfare function

To quantify welfare for different levels of steady-state inflation, we use a second-order approximation to the household utility function as in Woodford (2003).⁴ The main result can be summarized by the following proposition, with all proofs in Appendix A.

⁴ In our welfare calculations, we use the 2nd order approximation to the consumer utility function while the structural relationships in the economy are approximated to first order. As discussed in Woodford (2010), this approach is valid if distortions to the steady state are small so that the first order terms in the utility approximation are premultiplied by coefficients that can also be treated as first order terms. Since given our parameterization the distortions from imperfect competition and inflation are small (as in Woodford 2003), this condition is satisfied in our analysis. Furthermore, we show in Appendix F that the log-linear solution closely approximates the nonlinear solution, which implies that second order effects on the moments of inflation and output are small and can be ignored in the welfare calculations.

Proposition 1. The 2nd order approximation to expected per period utility in eq. (1) is⁵

$$\Theta_0 + \Theta_1 \text{var}(\hat{y}_t) + \Theta_2 \text{var}(\hat{\pi}_t) + \Theta_3 \text{var}(\hat{c}_t) \quad (29)$$

where parameters $\Theta_i, i = \{0,1,2\}$ depend on the steady state inflation $\bar{\pi}$ and are given by

$$\begin{aligned} \Theta_0 &= \left[1 - \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} (1 - (1+\eta^{-1})Q_y^0) \right] \log \bar{X} - \left[\frac{(1-\Phi)(1+\eta^{-1})(1-\beta h)}{2(1-\bar{g}_y)(1-h)} \right] [\log \bar{X}]^2 \\ &\quad - \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} \left\{ (1+\eta^{-1})[Q_y^0]^2 - Q_y^0 + \frac{\theta^2(1+\eta^{-1})}{2} \bar{\Delta} \right\} \\ \Theta_1 &= -\frac{(1+\eta^{-1})(1-\beta h)}{2(1-\bar{g}_y)(1-h)}, \\ \Theta_2 &= -\frac{\theta^2}{2(1-\bar{g}_y)(1-h)} \Gamma_3 \left\{ [Q_y^1(\theta^{-1} - 1) + (1+\eta^{-1})(1 + \frac{\theta^{-1}}{\theta} Q_y^0 Q_p^1)] - (1+\eta^{-1}) \frac{\theta^{-1}}{\theta} Q_y^1 \log \bar{X} \right\}, \\ \Theta_3 &= -\frac{h(1-\rho_c)}{(1-h)^2} \\ \Gamma_0 &= \left\{ 1 + (\theta - 1)Q_p^1[(1-\lambda)(\bar{b} + Q_p^0) - \lambda(1-\omega)\bar{\pi}] \right\}^{-1}, \Gamma_1 = \{1 - (\theta - 1)(1-\omega)\bar{\pi}Q_p^1\}\Gamma_0, \\ \Gamma_3 &= \frac{\Gamma_0}{1-\lambda\Gamma_1} \{(1-\lambda)M^2 + \lambda\}, \\ Q_y^0 &= \frac{\theta^{-1}}{2\theta} \bar{Y} / [1 + \frac{1}{2}(\frac{1-\theta}{\theta})^2 \bar{Y}]^2, Q_y^1 = [1 - \frac{1}{2}(\frac{1-\theta}{\theta})^2 \bar{Y}] / [1 + \frac{1}{2}(\frac{1-\theta}{\theta})^2 \bar{Y}]^3, Q_p^0 = \frac{1-\theta}{2} \bar{\Delta} / [1 + \frac{1}{2}(1-\theta)^2 \bar{\Delta}]^2, \\ Q_p^1 &= [1 - \frac{1}{2}(1-\theta)^2 \bar{\Delta}] / [1 + \frac{1}{2}(1-\theta)^2 \bar{\Delta}]^3, \bar{\Delta} = \bar{\pi}^{2\lambda} \frac{(1-\omega)^2}{(1-\lambda)^2}, \Phi = -\log(\frac{\theta-1}{\theta}), \text{ and } \rho_c = \text{corr}(\hat{c}_t, \hat{c}_{t-1}). \end{aligned}$$

This approximation of the household utility places no restrictions on the path of nominal interest rates and thus is invariant to stabilization policies chosen by the central bank.

The loss function in Proposition 1 illustrates the three mechanisms via which trend inflation affects welfare: the steady-state effects, the effects on the coefficients of the utility-function approximation, and the dynamics of the economy via the second moments of macroeconomic variables.⁶ First, the term Θ_0 captures the steady-state effects from positive trend inflation, which hinge on the increase in the cross-sectional steady-state dispersion in prices (and therefore in inefficient allocations of resources across sectors) associated with positive trend inflation.⁷ Note that as $\bar{\pi}$ approaches zero, Θ_0 converges to zero. As shown by Ascari and Ropele (2007), when $\bar{\pi} = 0$, $\partial\Theta_0/\partial\bar{\pi} > 0$, but the sign of the slope quickly reverses at marginally positive inflation rates. In our baseline calibration, Θ_0 is strictly negative and $\partial\Theta_0/\partial\bar{\pi} < 0$ when trend inflation exceeds 0.12% per annum. Thus for quantitatively relevant inflation rates, the welfare loss from steady-state effects is increasing in the steady-state level of inflation. This is intuitive since, except for very small levels of inflation, the steady state level of output declines with higher $\bar{\pi}$ because the steady state cross-sectional price dispersion rises. The steady-state cost of inflation from price dispersion is one of the best-known costs of

⁵ The complete approximation also contains three linear terms, the expected output gap, expected consumption and expected inflation. Since the distortions to the steady state are small for the levels of trend inflation we consider, the coefficients that multiply these terms can be considered as first order so we can evaluate these terms using the first order approximation to the laws of motion as in Woodford (2003). We confirmed in numeric simulations that they can be ignored. Furthermore, second order effects on the expected output gap and expected inflation are likely to be quantitatively small since the linear solution closely approximates the nonlinear solution to the model (see Appendix F).

⁶ When $\bar{\pi} = 0$, equation (41) reduces to the standard second-order approximation of the utility function as in Proposition 6.4 of Woodford (2003). There is a slight difference between our approximation and the approximation in Woodford (2003) since we focus on the per-period utility while Woodford calculated the present value.

⁷ The parameter Φ measures the deviation of the flexible-price level of output from the flexible-price perfect-competition level of output. See Woodford (2003) for derivation.

inflation and arises naturally from the integration of positive trend inflation into the New Keynesian model. Consistent with this effect being driven by the increase in dispersion, one can show that the steady-state effect is eliminated with full indexation of prices and mitigated with partial indexation.

Second, the coefficient on the variance of output around its steady state $\Theta_1 < 0$ does not depend on trend inflation. This term is directly related to the increasing disutility of labor supply. With a convex cost of labor supply, the expected disutility rises with the variance of output around its steady state. However, even though Θ_1 is independent of $\bar{\pi}$, this does not imply that a positive $\bar{\pi}$ does not impose any output cost. Rather, trend inflation reduces the steady state level of output, which is already captured by Θ_0 . Once this is taken into account, then log utility implies that a given level of output variance around the (new) steady state is as costly as it was before. Furthermore, the variance of output around its steady state depends on the dynamic properties of the model which are affected by the level of trend inflation.

The coefficient on the variance of inflation $\Theta_2 < 0$ captures the sensitivity of the welfare loss due to the cross-sectional dispersion of prices. One can also show analytically that for $\bar{\pi} \approx 0$, $\partial\Theta_2/\partial\bar{\pi} < 0$ so that the cross-sectional dispersion of prices becomes *ceteris paribus* costlier in terms of welfare. Because an inflationary shock creates distortions in relative prices and positive trend inflation already generates some price dispersion and an inefficient allocation of resources, firms operating at an inefficient level have to compensate workers for the increasingly high marginal disutility of sector-specific labor. With this rising marginal disutility, the increased distortion in relative prices due to an inflation shock becomes costlier due to the higher initial price dispersion making the variance of inflation costlier for welfare as the trend level of inflation rises. This is a second, and to the best of our knowledge previously unidentified, channel through which the price dispersion from staggered price setting under positive inflation reduces welfare.

Finally, the coefficient on the variance of consumption $\Theta_3 < 0$ captures the desire of habit-driven consumers to smooth consumption. The greater the degree of habit formation, the more costly a given amount of consumption volatility becomes. Conversely, the greater the autocorrelation of consumption, the smaller are period-by-period changes in consumption, and the less costly consumption volatility becomes. Trend inflation changes this coefficient only by affecting the persistence of consumption.

III Calibration and Optimal Inflation

Having derived the approximation to the utility function, we now turn to solving for the optimal inflation rate. Because utility depends on the volatility of macroeconomic variables, this will be a function of the structural parameters and shock processes. Therefore, we first discuss our parameter selection and then consider the implications for the optimal inflation rate in the model.

3.1 Parameters

Our baseline parameter values are illustrated in Table 1. For the utility function, we set η , the Frisch labor supply elasticity, equal to one. The steady-state discount factor β is set to 0.998 to match the real rate of

2.3% per year on 6-month commercial paper or assets with similar short-term maturities given that we set the steady-state growth rate of real GDP per capita to be 1.5% per year ($\overline{G\dot{Y}} = 1.015^{0.25}$), as in Coibion and Gorodnichenko (2011). We set the elasticity of substitution across intermediate goods θ to 7, so that steady-state markups are equal to 17%. This size of the markup is consistent with estimates presented in Burnside (1996) and Basu and Fernald (1997). The degree of price stickiness (λ) is set to 0.55, which amounts to firms resetting prices approximately every 7 months on average. This is midway between the micro estimates of Bils and Klenow (2004), who find that firms change prices every 4 to 5 months, and those of Nakamura and Steinsson (2008), who find that firms change prices every 9 to 11 months.

The degree of price indexation is assumed to be zero in the baseline for three reasons. First, the workhorse New Keynesian model is based only on price stickiness, making this the most natural benchmark (Woodford 2003). Second, any price indexation implies that firms are constantly changing prices, a feature strongly at odds with the empirical findings of Bils and Klenow (2004) and more recently Nakamura and Steinsson (2008), among many others. Third, while indexation is often included to replicate the apparent role for lagged inflation in empirical estimates of the New Keynesian Phillips Curve (NKPC; see Gali and Gertler 1999), Cogley and Sbordone (2008) show that once one controls for steady-state inflation, estimates of the NKPC reject the presence of indexation in price setting decisions. However, we relax the assumption of no indexation in the robustness checks.

The coefficients for the Taylor rule are taken from Coibion and Gorodnichenko (2011). These estimates point to strong long-run responses by the central bank to inflation and output growth (2.5 and 1.5 respectively) and a moderate response to the output gap (0.43).⁸ The steady-state share of consumption is set to 0.80 so that the share of government spending is twenty percent. The calibration of the shocks is primarily taken from the estimated DSGE model of Smets and Wouters (2007) with the exception of the persistence of the risk premium shocks for which we consider a larger value calibrated at 0.947 to match the historical frequency of hitting the ZLB and the routinely high persistence of risk premia in financial time series.⁹

In our baseline model, positive trend inflation is costly because it leads to more price dispersion and therefore less efficient allocations, more volatile inflation, and a greater welfare cost for a given amount of inflation volatility. On the other hand, positive trend inflation gives policy-makers more room to avoid the ZLB on interest rates. Therefore, a key determinant of the tradeoff between the two depends on how frequently the ZLB is binding for different levels of trend inflation. To illustrate the implications of our parameter calibration for how often we hit the ZLB, Figure 1 plots the fraction of time spent at the ZLB from simulating our model for different steady-state levels of the inflation rate. In addition, we plot the steady-state

⁸ Because empirical Taylor rules are estimated using annualized rates while the Taylor rule in the model is expressed at quarterly rates, we rescale the coefficient on the output gap in the model such that $\phi_y = 0.43/4 = 0.11$.

⁹ This calibration is, e.g., consistent with the persistence of the spread between Baa and Aaa bonds which we estimate to be 0.945 between 1920:1 and 2009:2 and 0.941 between 1950:1 and 2009:2 at the quarterly frequency.

level of the nominal interest rate associated with each inflation rate, where the steady-state nominal rate in the model is determined by $\bar{R} = \bar{\pi} \cdot \bar{G}\bar{Y}/\bar{\beta}$. Our calibration implies that with a steady-state inflation rate of approximately 3.5%, the average rate for the U.S. since the early 1950's, the economy should be at the ZLB approximately 5 percent of the time. This is consistent with the post-WWII experience of the U.S.: with U.S. interest rates at the ZLB since late 2008 and expected to remain so until the end of 2011, this yields a historical frequency of being at the ZLB of 5 percent (i.e. around 3 years out of 60).¹⁰

In addition, this calibration agrees with the historical changes in interest rates associated with post-WWII U.S. recessions. For example, starting with the 1958 recession and excluding the current recession, the average decline in the Federal Funds Rate during a recession has been 4.76 percentage points.¹¹ The model predicts that the average nominal interest rate with 3.5% steady-state inflation is around 6%, so the ZLB would not have been binding during the average recession, consistent with the historical experience. Only the 1981-82 recession led to a decline in nominal interest rates that would have been sufficiently large to reach the ZLB (8.66% drop in interest rates), but did not because nominal interest rates and estimates of trend inflation over this period were much higher than their average values. Thus, with 3-3.5% inflation, our calibration (dotted line in Figure 1) implies that it would take unusually large recessions for the ZLB to become binding. In addition, our calibration indicates that at much lower levels of $\bar{\pi}$, the ZLB would be binding much more frequently: e.g. at $\bar{\pi} = 0$, the ZLB would be binding 27% of the time. This seems conservative since it exceeds the historical frequency of U.S. recessions. The model predicts a steady-state level of interest rates of less than 2.5% when $\bar{\pi} = 0$, and six out the last eight recessions (again excluding the current episode) were associated with decreases in interest rates that exceeded this value (specifically the 1969, 1973, 1980, 1981, 1990 and 2001 recessions). Our calibration is also largely in line with the frequency of the ZLB we would have observed given historical declines in nominal interest rates during recessions and counterfactual levels of trend inflation (broken line in Figure 1). Thus, our parameterization provides a reasonable representation of the likelihood of hitting the ZLB for different inflation rates given the historical experience of the U.S.

Our calibration also accounts for the key moments of output, inflation, interest rates and consumption. Table 2 presents the variance and autocorrelation of each HP-filtered variable in the model and U.S. data from 1947Q1 to 2011Q1. The model reproduces both the absolute and relative volatilities of these variables as well as their persistence, although the persistence of consumption and output are slightly lower in the model than the data. The model also replicates the strong comovement of consumption with output and the much lower comovement of inflation and interest rates with output.

¹⁰ Of possible concern may be that this calculation includes the high-inflation environment from 1970-85. Excluding those years generates a historical frequency at the ZLB of 3/45=6.66% but now at a lower trend inflation rate of 3% per year. Our baseline calibration generates approximately that frequency at 3% trend inflation.

¹¹ This value is calculated by taking the average level of the Federal Funds rate (FFR) over the last 6 months prior to the start of each NBER recession and subtracting the minimum level of the FFR in the aftermath of that recession.

3.2 Optimal Inflation

Having derived the dynamics of the model, parameterized the shocks, and obtained the second-order approximation to the utility function, we now simulate the model for different levels of trend inflation $\bar{\pi}$ and compute the expected utility for each $\bar{\pi}$. We use the Bodenstein et al. (2009) algorithm to solve the non-linear model and verify in Appendix F that this algorithm has very high accuracy, even after large shocks leading to a binding ZLB. The results taking into account the ZLB and in the case when we ignore the ZLB are plotted in Panel A of Figure 2. When the ZLB is not taken into account, the optimal rate of inflation is zero because there are only costs to inflation and no benefits. Figure 2 also plots the other extreme when we include the ZLB but do not take into account the effects of positive steady-state inflation on the loss function or the dynamics of the model. In this case, there are no costs to inflation so utility is strictly increasing as steady-state inflation rises and the frequency of the ZLB diminishes. Our key result is the specification which incorporates both the costs and benefits of inflation. As a result of the ZLB constraint, we find that utility is increasing at very low levels of inflation so that zero inflation is not optimal when the zero bound is present. Second, *the peak level of utility is reached when the inflation rate is 1.5% at an annualized rate*. This magnitude is close to the bottom end of the target range of most central banks, which are commonly between 1% and 3%. Thus, our baseline results imply that taking into account the zero bound on interest rates raises the optimal level of inflation, but with no additional benefits to inflation included in the model, the optimal inflation rate is within the standard range of inflation targets. Third, the costs of even moderate inflation can be nontrivial: a 5% inflation rate would lower utility by approximately 1% relative to the optimal level, which given log utility in consumption is equivalent to a permanent 1% decrease in the level of consumption. As we show later, the magnitude of the welfare costs of inflation varies with the calibration and price setting assumptions, but the optimal rate of inflation is remarkably insensitive to these modifications.

Panel B of Figure 2 quantifies the importance of each of the three costs of inflation – the steady state effect, the increasing cost of inflation volatility, and the positive link between the level and volatility of inflation – by calculating the optimal inflation rate subject to the ZLB when only one of these costs, in turn, is included. The first finding to note is that allowing for any of the three inflation costs is sufficient to bring the optimal inflation rate to 3.6% or below. Thus, all three inflation costs incorporated in the model are individually large enough to prevent the ZLB from pushing the optimal inflation rate much above the current target range of most central banks. Second, the steady-state cost is the largest cost of inflation out of the three, bringing the optimal inflation rate down to 1.6% by itself. However, even if we omit steady-state costs and include only the other two channels, the optimal inflation rate would be less than 3%.

To get a sense of which factors drive these results, the top row of Figure 3 plots the coefficients of the second-order approximation to the utility function from Proposition 1. First, higher $\bar{\pi}$ has important negative steady-state effects on utility, as the increasing price dispersion inefficiently lowers the steady-

state level of production and consumption. Second, the coefficient on the variance of consumption becomes slightly smaller in absolute value for low levels of inflation then rises moderately at higher levels of inflation. Third, the coefficient on inflation variance is decreasing in $\bar{\pi}$, i.e., holding the inflation variance constant, higher $\bar{\pi}$ raises the utility cost of the variance in inflation. This reflects the fact that when the steady state level of price dispersion is already high then a temporary increase in price dispersion due to an inflation shock is even more costly. Moving from zero inflation to six percent inflation raises the coefficient on the inflation variance by almost 30% in absolute value. Thus, as $\bar{\pi}$ rises, policy-makers should place an increasing weight on the variance of inflation relative to the variance of the output gap.

The middle row of Figure 3 plots the effects of $\bar{\pi}$ on the variance of inflation, consumption and the output gap, i.e. the dynamic effects of steady-state inflation and the ZLB. In addition, we plot the corresponding moments in the absence of the zero-bound on interest rates to characterize the contribution of the zero-bound on macroeconomic dynamics. A notable feature of the figure is how rapidly consumption, output and inflation volatility rise as $\bar{\pi}$ falls when the ZLB is present. Intuitively, the ZLB is hit more often at a low $\bar{\pi}$. With the nominal rate fixed at zero, the central bank cannot stabilize the economy by cutting interest rates further and thus macroeconomic volatility increases. As we increase $\bar{\pi}$, macroeconomic volatility diminishes. This is the benefit of higher $\bar{\pi}$ in the model. The effect of changes in $\bar{\pi}$, however, is non-linear for the variance of inflation when we take into account the zero-bound on interest rates. At low levels of inflation, increasing $\bar{\pi}$ reduces the volatility of inflation for the same reason as for output: the reduced frequency of hitting the zero bound. On the other hand, higher $\bar{\pi}$ also tends to make pricing decisions more forward-looking, so that, absent the zero bound, inflation volatility is consistently rising with $\bar{\pi}$, a feature emphasized in Kiley (2007) and consistent with a long literature documenting a positive relationship between the level and variance of inflation (Okun 1971, Taylor 1981). When $\bar{\pi}$ rises past a specific value, the latter effect dominates and the variance of inflation rises with $\bar{\pi}$. Given our baseline values, this switch occurs at an annualized trend inflation rate of approximately 3.5%. These results show the importance of modeling both the ZLB and the effects of $\bar{\pi}$ on the dynamics of the model.

The bottom row of Figure 3 then plots the contribution of these different effects on the welfare costs of inflation, i.e. each of the terms in Proposition 1. These include the steady-state effects of $\bar{\pi}$ as well as the interaction of the effects of $\bar{\pi}$ on the coefficients of the utility function approximation and the dynamics of the economy. The most striking result is that the welfare costs and benefits of positive $\bar{\pi}$ are essentially driven by only two components: the steady-state effect and the contribution of inflation variance to utility. In particular, the U-shape pattern of the inflation variance combined with decreasing θ_2 plays the key role in delivering a positive level of the optimal inflation rate, while the effects of the ZLB on the contribution of the output gap and consumption variability are an order of magnitude smaller and therefore play a limited role in determining the optimal inflation rate.

3.3 Are the costs of business cycles and the ZLB too small in the model?

The minor contribution of output gap volatility to the optimal inflation rate might be interpreted as an indication that the model understates the costs of business cycles in general and the ZLB in particular. For the former, the implied welfare costs of business cycles in our model are approximately 0.5% of steady-state consumption at the historical trend inflation rate, in line with many of the estimates surveyed in Barlevy (2004) and much larger than in Lucas (1987). To assess the cost of hitting the ZLB, we compute the average welfare loss net of steady-state effects from simulating the model under different inflation rates both with and without the zero bound. The difference between the two provides a measure of the additional welfare cost of business cycles due to the presence of the ZLB. We can then divide this cost by the average frequency of being at the zero bound from our simulations, for each level of steady-state inflation, to get a per-quarter average welfare loss measure conditional on being at the ZLB which is plotted in Panel A of Figure 4. As $\bar{\pi}$ rises, this per-period cost declines because the average duration of ZLB episodes gets shorter and the output losses during the ZLB are increasing non-linearly with the duration of the ZLB (see Christiano et al. 2011). For example, the average cost of a quarter spent at the ZLB is approximately equivalent to a permanent 1.4% reduction in consumption when inflation is 1% but declines to 0.4% at a 3.5% rate of inflation. The latter implies that the additional cost of being restrained by the zero bound for 8 quarters is equivalent to a 3.2% *permanent* reduction in consumption, or approximately \$320 billion per year based on 2008 consumption data. For comparison, Williams (2009) uses the Federal Reserve's FRB/US model to estimate that the ZLB between 2009 and 2010 cost \$1.8 trillion in lost output over four years, or roughly \$300 billion per year in lost consumption over four years if one assumes that the decline in consumption was proportional to the decline in output. Thus, the costs of both business cycles and the ZLB in the model cannot be described as being uncharacteristically small.

However, while the conditional costs of long ZLB events are quite large, they also occur relatively infrequently. For example, if we assume that all ZLB episodes are 8 quarters long, then at 3.5% trend inflation an 8-quarter episode at the ZLB occurs with probability 0.007 each quarter, or about 3 times every 100 years. This implies that the expected cost of the ZLB is a 0.02% permanent reduction of consumption. Similar calculations for 2% trend inflation reveal that while the conditional cost of an 8-quarter ZLB event is about a 6.2% permanent reduction of consumption, the unconditional cost of the ZLB is only a 0.08% permanent reduction in consumption. Thus, while the model implies that a higher inflation target can significantly reduce the cost of a given ZLB event, as suggested by Blanchard, taken over a long horizon the expected gain in mitigating the ZLB from such a policy is small. As a result, even modest steady-state costs of inflation, because they must be borne every period, are sufficient to push the optimal inflation rate below 2%.

3.4 How does optimal inflation depend on the coefficient on the variance of the output gap?

Even though the costs of business cycles are significant and ZLB episodes are both very costly and occurring with reasonable probability, one may be concerned that these costs are incorrectly measured due to the small relative weight assigned to output gap fluctuations in the utility function. At $\bar{\pi} = 0$, the

coefficient on the output gap variance in the loss function is less than one-hundredth that on the quarterly inflation variance (or one-tenth for the annualized inflation variance), and this difference becomes even more pronounced as $\bar{\pi}$ rises. The low weight on output gap volatility is standard in New Keynesian models and could reflect the lack of involuntary unemployment, which inflicts substantial hardship to a fraction of the population and whose welfare effects may be poorly approximated by changes in aggregate consumption and employment, or the absence of distribution effects, as business cycles disproportionately impact low income/wealth agents with higher marginal utilities of consumption than the average consumer.

To assess how sensitive the optimal inflation rate is to the coefficient on the output gap variance, we increase this coefficient by a factor ranging from 1 to 100 and reproduce our results for the optimal $\bar{\pi}$ for each factor (see Panel B of Figure 4). Raising the coefficient on the variance of the output gap pushes the optimal inflation rate higher, but the coefficient on the output gap variance needs to be very large to qualitatively affect our findings. For example, much of the traditional literature on optimal monetary policy assumed equal weights on output and annualized inflation variances in the loss function. With inflation being measured at an annualized rate, this equal weighting obtains at zero steady-state inflation when Θ_1 is multiplied by a factor of approximately 10. Yet this weighting would push the optimal inflation rate up only modestly to 1.6% per year. Even if one increased the coefficient on output gap volatility by a factor of 100, the optimal inflation rate would rise only to 2.4%. Placing such weight on output volatility would raise the implied per-quarter cost of having the ZLB bind to an equivalent of a 3% permanent reduction in consumption, such that an episode of 8 consecutive quarters at the ZLB would deliver welfare losses equivalent to roughly 24% of steady-state consumption, above and beyond the costs of the shock in the absence of the ZLB. Thus, while one can mechanically raise the optimal inflation rate via larger weights on output fluctuations than implied by the model, weighting schemes which meaningfully raise the optimal inflation rate point to welfare costs of business cycles, and particularly episodes at the ZLB, that depart from the conventional wisdom.

IV Robustness of the Optimal Inflation Rate to Alternative Parameter Values

In this section, we investigate the robustness of the optimal inflation rate to our parameterization of the model. We focus particularly on pricing and utility parameters, the discount factor, and the risk premium shock.

4.1 Pricing and Utility Parameters

Figure 5 plots the optimal inflation rates and associated welfare losses for optimal levels of $\bar{\pi}$ for alternative pricing and utility parameters. First, we consider the role of the elasticity of substitution θ (Panel A), which plays a critical role in determining the cost of price dispersion, and therefore costs of inflation, in the model. Note that the welfare costs of inflation are larger when θ is high. This result captures the fact that a higher elasticity of substitution generates more steady-state output dispersion and, through greater real rigidity, higher welfare cost of fluctuations for any $\bar{\pi}$. However, the effects of this parameter on the optimal $\bar{\pi}$ are

relatively small: even a value of $\theta = 3$ yields an optimal $\bar{\pi}$ of less than 2%. Thus, the optimal inflation rate is robust to the range of values of θ commonly considered in the macroeconomics literature.

However, microeconomic estimates of demand elasticities commonly find much lower values than those employed in macroeconomic models. Hausman et al. (1994), for example, document that while the elasticity of demand for individual beer brands may be around 5, the overall elasticity of demand for beer is closer to 1.5. To the extent that goods across industries may be less substitutable than goods within industries, our baseline model could overstate the amount of price and output dispersion and therefore the steady-state costs of inflation. This effect would likely be particularly pronounced if price changes within industries were synchronized so that the relevant amount of price dispersion for welfare arises primarily from cross-industry substitution at low elasticities.

We can assess the implications of these possibilities by introducing a few modifications to our baseline model. Specifically, we can consider a two-tier model in which the elasticity of substitution within industries is given by θ and the elasticity of substitution across industries is given by μ where we assume $\theta > \mu$. Prices are sticky at the firm level but not at the industry level – industries simply combine individual goods and sell them off to consumers. In the absence of industry-specific shocks and price-synchronization within industries, average price levels across industries must be equal since industries are symmetric and firm-level Calvo shocks which allow firms to reset prices wash out at the industry level. As a result, there would be no price dispersion across industries. In addition, if we assume that there is a continuum of firms that aggregate for a given industry i , then there is no industry mark-up. Thus the steady state mark-up on goods would be $\frac{\theta}{\theta-1}$ and this model would have exactly the same implications as our baseline model with all of the results being driven by θ rather than μ . This suggests that, absent price synchronization, it is the elasticity of substitution within industries which is most relevant for our results. Intuitively, since all industries have the same price index and the same intra-industry price dispersion, we can combine them into a “representative” industry, thereby reducing this model to our baseline setup.

An important potential caveat to this result, however, is if price changes are synchronized within industries, as in Bhaskar (2002). Price synchronization within industries could significantly reduce price dispersion within industries while the remaining price dispersion across industries would matter little for welfare under low cross-industry elasticities of substitution. To quantify this possibility, we used a model with Taylor pricing (as in section 5.4) with four industries, each of which had staggered pricing over four quarters, but in which disproportionate shares of firms within each industry reset their prices in the same quarter. We allowed for the elasticity of substitution within industries ($\theta=7$) to exceed that across industries ($\mu=3$). Even high levels of synchronization (i.e. 70% of firms in an industry resetting their prices in the same quarter) had only modest effects on steady-state levels of price dispersion and thus on our welfare

calculations.¹² While these modifications suggest that our results are robust to a number of alternative assumptions about industrial structure and pricing assumptions, one should nonetheless bear in mind that there is significant uncertainty about how best to model the substitution and pricing of goods across industries and that alternative specifications may quantitatively affect the predictions about the costs of price dispersion.

Second, we also investigate the role of price indexation. In our baseline, we assumed $\omega = 0$, based on the fact that firms do not change prices every period in the data, as documented by Bils and Klenow (2004) and Nakamura and Steinsson (2008), as well as the results of Cogley and Sbordone (2008) who argue that once one controls for time-varying trend inflation, we cannot reject the null that $\omega = 0$ for the US. However, because price-indexation is such a common component of New Keynesian models, we consider the effects of price indexation on our results. Panel B of Figure 5 indicates that higher levels of indexation lead to higher optimal rates of inflation because indexation reduces the dispersion of prices. Yet with $\omega = 0.5$, which is most likely an upper bound for an empirically plausible degree of indexation in low-inflation economies like the U.S., the optimal $\bar{\pi}$ remains less than 2.4%.¹³

Third, we examine the effects of price stickiness. Our baseline calibration, $\lambda = 0.55$, is midway between the findings of Bils and Klenow (2004) of median price durations of 4-5 months and those of Nakamura and Steinsson (2008) of median price durations of 9-11 months. We now consider values of λ ranging from 0.50 to 0.75 (Panel C). With more price stickiness, price dispersion is greater, and this effect is amplified at higher levels of steady-state inflation, thereby generating much larger welfare losses. Nonetheless, this has only minor effects on the optimal inflation rate.

Finally, we consider variation in the levels of habit formation and the Frisch labor supply elasticity in Panels D and E. The specific values applied to both parameters can significantly affect the level of welfare losses at the optimal inflation rate but neither qualitatively alters the optimal inflation rate in the model.

4.2 Discount Factor and Risk Premium Shocks

We also consider the sensitivity of our results to the discount factor and the parameters governing the risk premium shocks (Figure 6). First, we reproduce our baseline welfare figure for different levels of the persistence to risk premium shocks. The results are quite sensitive to this parameter, which reflects the fact that these shocks play a crucial role in hitting the zero lower bound. For example, Figure 6 illustrates that

¹² These results are available upon request. The empirical evidence on price synchronization is mixed. Dhyne and Konieczny (2007), for example, document that price synchronization within industries is larger than across industries but also document a remarkable degree of staggering of price changes within industries. For example, they observe (p. 11), "...price changes at the individual product category level are neither perfectly staggered nor synchronized, but their behavior is much closer to perfect staggering." Likewise, a more recent study, Klenow and Malin (2011) survey the literature which uses micro-level price data and conclude that the timing of price changes is little synchronized across sellers even within very narrow product categories.

¹³ Standard estimates of the Phillips Curve without trend inflation (e.g., Gali and Gertler 1999 and Levin et al. 2005) suggest that the fraction of indexing firms is at most 0.5 and more likely between 0.25-0.35. We also solved the model with dynamic indexation and found nearly identical results.

when we raise the persistence of the shock from 0.947 to 0.96, the optimal inflation rate rises from 1.5% to 3% because this increase in the persistence of the shock has a large effect on the frequency and duration of being at the ZLB. At 3.5% inflation, this frequency more than doubles relative to our baseline scenario, thereby raising the benefit of higher steady-state inflation. The reverse occurs with lower persistence of risk premium shocks: the frequency of being at the ZLB declines sharply as does the optimal inflation rate. Second, similar results obtain when we vary the volatility of the risk premium shock. When we increase the standard deviation of these shocks to 0.0035 from our baseline of 0.0024, the optimal inflation rate again rises to slightly over 3 percent. As with the persistence of the shocks, this is driven by a higher frequency of being at the ZLB: at 3.5% inflation, this alternative shock volatility implies the economy would be at the ZLB three times as often as under our baseline calibration.

Third, we consider the sensitivity of our results to the steady-state level of the discount factor β . This parameter is also important in determining the frequency at which the economy is at the ZLB since it affects the steady-state level of nominal interest rates. As with the risk premium shock variables, a higher value of β is associated with a lower steady-state level of nominal interest rates, so that the ZLB will be binding more frequently. For example, with $\beta=0.9999$ (which corresponds to a real rate of 1.54% per year), the ZLB is binding approximately 7% of the time when steady-state inflation is 3.5%. At the maximum, however, the optimal $\bar{\pi}$ is only 0.6% higher than implied by our baseline results.

These robustness checks clearly illustrate how important the frequency at which the economy hits the ZLB is for our results. Naturally, parameter changes which make the ZLB binding more often raise the optimal rate of inflation because a higher $\bar{\pi}$ lowers the frequency of hitting the ZLB. Thus, the key point is not the specific values chosen for these parameters but rather having a combination of them that closely reproduces the historical frequency of hitting the ZLB for the U.S. Nonetheless, even if we consider parameter values that double or even triple the frequency of hitting the ZLB at the historical average rate of inflation for the U.S., the optimal inflation rate rises only to about 3%. This suggests that the evidence for an inflation target in the neighborhood of 2% is robust to a wide range of plausible calibrations of hitting the ZLB.

4.3 Summary

These results indicate that the optimal inflation rate in the baseline model is robust to reasonable variations in the parameters as calibrated to the U.S. economy. However, the fact that more persistent and volatile shocks could potentially raise the frequency of the ZLB suggests that the optimal inflation rate is likely to vary across countries. For example, smaller open economies are typically more subject to volatile terms of trade shocks than the U.S. and this increased volatility could justify higher target rates of inflation if this increased volatility increases the incidence of hitting the ZLB. Similarly, economies such as Japan in which the real interest rate has historically been very low might also find it optimal to pursue higher inflation targets. More broadly, our

results highlight the importance of carefully calibrating the model parameters to the specifics of the economy before drawing general conclusions about optimal policies.

V What Could Raise the Optimal Inflation Rate?

While our baseline model emphasizes the tradeoff between higher $\bar{\pi}$ to insure against the zero-bound on nominal interest rates versus the utility costs associated with higher trend inflation, previous research has identified additional factors beyond the lower-bound on nominal interest rates which might lead to higher levels of optimal inflation. In this section, we extend our analysis to assess their quantitative importance. First, we include capital formation in the model. Second, we allow for uncertainty about parameter values on the part of policy-makers. Third, we integrate downward nominal wage rigidity, i.e. “greasing the wheels,” into the model. Fourth, we explore whether our results are sensitive to using Taylor pricing. Finally, we consider the possibility that the degree of price stickiness varies with $\bar{\pi}$.

5.1 Capital

First, we consider how sensitive our results are to the introduction of capital. We present a detailed model in Appendix B and only provide a verbal description in this section. In this model, firms produce output with a Cobb-Douglas technology (capital share $\alpha = 0.33$). All capital goods are homogeneous and can be equally well employed by all firms. Capital is accumulated by the representative consumer subject to a quadratic adjustment cost to capital ($\psi = 3$ as in Woodford 2003) and rented out in a perfectly competitive rental market. The aggregate capital stock depreciates at rate $\delta = 0.02$ per quarter. We calculate the new steady state level of output relative to the flexible price level output and derive the analogue of Proposition 1 in Appendix B with proofs in Appendix C.

By allowing capital to freely move between firms we reduce the steady state welfare cost from trend inflation since firms that have a relatively low price can now hire additional capital rather than sector-specific workers to boost their output. Capital also increases the likelihood of hitting the ZLB and therefore the benefits of inflation because including capital permits disinvestment when agents prefer storing wealth in safe bonds rather than capital, so we are more likely to be in a situation where an increase in q_t pushes interest rates to zero. We isolate the first channel by setting $\rho_q = 0.943$ to match the historical frequency at the ZLB. As shown in Panel A of Figure 7, utility peaks at a trend inflation rate of 2.1% suggesting that capital does not lower the cost of inflation substantially.

5.2 Model Uncertainty

An additional feature that could potentially lead to higher rates of optimal inflation is uncertainty about the model on the part of policy-makers. If some plausible parameter values lead to much higher frequencies of hitting the ZLB or raise the output costs of being at the ZLB, then policy-makers might want to insure against these outcomes by allowing for a higher $\bar{\pi}$. We quantify this uncertainty via the variance-

covariance matrix of the estimated parameters from Smets and Wouters (2007), placing an upper bound on parameter values to eliminate draws where the ZLB binds unrealistically often, in excess of 10% at the historical average of annual trend inflation.

To assess the optimal inflation rate given uncertainty about parameter values, we compute the expected utility associated with each level of steady-state inflation by repeatedly drawing from the distribution of parameter values. Panel B of Figure 7 plots the implied levels of expected utility associated with each steady-state level of inflation. Maximum utility is achieved with an inflation rate of 1.9% per year. As expected, this is higher than our baseline result, which reflects the fact that some parameter draws lead to much larger costs of being at the zero-bound, a feature which also leads to a much more pronounced inverted U-shape of the welfare losses from steady-state inflation. Nonetheless, this optimal rate of inflation remains well within the bounds of current inflation targets of modern central banks.

We also consider another exercise in which we repetitively draw from the parameter space and solve for the optimal inflation rate associated with each draw. The distribution of optimal inflation rates has a 90% confidence interval ranging from 0.3% to 2.9% per year, which again is very close to the target range for inflation of most central banks.

5.3 Downward Wage Rigidity

A common motivation for positive trend inflation, aside from the zero-lower bound, is the “greasing the wheels” effect raised by Tobin (1972). If wages are downwardly rigid, as usually found in the data (e.g., Dickens et al. 2007), then positive trend inflation will facilitate the downward-adjustment of real wages required to adjust to negative shocks. To quantify the effects of downward nominal wage rigidity in our model, we integrate it in a manner analogous to the zero-bound on interest rates by imposing that changes in the aggregate nominal wage index be above a minimum bound $\Delta\widehat{w}_t = \max\{\Delta\widehat{w}_t^m, \Delta\widehat{w}_t^*\}$ where $\Delta\widehat{w}_t^*$ is the change in wages that would occur in the absence of the zero-bound on nominal wages and $\Delta\widehat{w}_t^m$ is the lower-bound on nominal wage changes. Note that even with zero steady-state inflation, steady-state nominal wages grow at the rate of technological progress. Thus, we set $\Delta\widehat{w}_t^m$ to be equal to minus the sum of the growth rate of technology and the steady state rate of inflation.

Panels C and D of Figure 7 present the utility associated with different $\bar{\pi}$ under both the zero-bound on interest rates and downward-wage rigidity. The result is striking: the optimal inflation rate falls to 0.3% per year with downward wage rigidity. With downward wage rigidity, marginal costs are less volatile, so the variance of inflation is reduced relative to the case with flexible wages. In addition, the fact that marginal costs are downwardly rigid means that, in the face of a negative demand shock, inflation will decline by less and therefore interest rates will fall less, reducing the frequency of the ZLB. With $\bar{\pi} = 0$, the ZLB binds less than 20 percent of the time with downward wage rigidity but over 27 percent of the time

with flexible wages. This decrease in the frequency of binding ZLB at low $\bar{\pi}$ reduces the benefit of higher trend inflation and leads to lower estimates of the optimal inflation rate.¹⁴

5.4 Taylor pricing

Our baseline model relies on the Calvo (1983) price-setting framework. While analytically convenient, this approach suffers from several drawbacks that could affect our estimates of the optimal inflation rate. One such factor is that with exogenous probabilities of changing prices each period, there is always a fraction of firms using very outdated prices. With positive steady-state inflation, these prices will be very low in relative terms, leading to high estimates of the cost of price dispersion even at moderate inflation rates. To assess how important this factor is for the optimal rate of inflation, we consider an alternative approach to Calvo pricing, namely the staggered contracts approach of Taylor (1977), in which firms set prices for a pre-determined duration of time. With fixed durations of price stickiness, price dispersion should be smaller than under Calvo pricing for sufficiently high inflation rates. The derivation of the utility approximation as well as the structural log-linearized equations of the model is similar to Calvo pricing (Appendix D contains details for the utility approximation when the duration of price spells is equal to three quarters).

Panels E and F of Figure 7 compare the results under Taylor pricing (using durations of price contracts equal to 3 and 4 quarters) with the results from Calvo pricing. The optimal inflation rates for the Taylor model are 2.2 and 1.8 percent per year for price durations of 3 and 4 quarters respectively, which is close to the 1.5 percent per year found for the baseline Calvo model. The models yield similar volatilities of inflation and the output gap and frequencies of hitting ZLB (see Coibion et al. 2011 for more details) are approximately the same in all models. Instead, the main difference is that the Taylor model has smaller welfare losses as $\bar{\pi}$ increases above 2 percent because of the reduced steady-state price dispersion. Intuitively, since firms under Calvo pricing may be stuck with a suboptimal price for a long time, the cost of positive steady state inflation is larger than in the Taylor model where firms are guaranteed to change prices in a fixed number of periods. The quantitative implications for the optimal inflation rate, however, are quite small.

5.5 Endogenous and State-Dependent Price Stickiness

Because theory implies that the cost to firms of not changing prices should increase as $\bar{\pi}$ rises (Romer 1990), higher levels of inflation should be associated with lower levels of price stickiness, which would tend to lower the welfare costs of positive trend inflation. Thus, by ignoring this endogeneity, we might be overstating the costs of positive inflation and thereby underestimating the optimal rate of inflation. To address this possibility, we consider the sensitivity of our baseline results to a possible systematic link between the inflation rate and the degree of price stickiness in two ways. As a first approach, we follow Nakamura and Steinsson's empirical approach and posit a linear relationship between the (monthly) frequency of price changes and the steady state

¹⁴ Kim and Ruge-Murcia (2009) similarly find that downward wage rigidity, by itself, has little positive effect on the optimal inflation rate in an estimated DSGE model.

annual rate of inflation, with the coefficient on inflation denoted by β_π . The average estimate of Nakamura and Steinsson across price measures and time periods is approximately $\beta_\pi = 0.5$, and the upper bound of their confidence intervals is approximately $\beta_\pi = 1$. We calibrate the degree of price rigidity such that $\lambda=0.55$ (our baseline value) at a steady-state level of annual inflation of 3.5%. In this setting, the degree of price stickiness varies with the steady-state rate of inflation but not with business cycle conditions.

As a second approach, we consider state-dependent pricing in the spirit of Dotsey et al. (1999). Using the same model as in section II, we replace the exogenous probability of changing prices with an explicit optimizing decision based on comparison of menu costs and the gains from price adjustment. Specifically, we assume that each period, firms draw from a uniform distribution of costs to changing prices and, conditional on their draw, decide whether or not to reset their price. Because all firms are identical, every firm that chooses to reset its price picks the same price. As in the Taylor (1977) staggered contracts model, the distribution of prices in the economy depends only on past reset prices and the share of firms which changed their price in previous quarters, but, in contrast to time-dependent models, these shares are time-varying. Derivations of the model and the associated welfare loss function are provided in Appendix E. Using the same parameter values as the baseline model, we calibrate the average size of menu costs to yield the same degree of price stickiness at a steady-state inflation rate of 3.5% as in the baseline case. Because this calibration implies strong strategic complementarity in price setting, large menu costs (approximately 7% of output) are necessary to match the duration of price spells in the data.¹⁵ We therefore quantify the sensitivity of our results to this parameterization by varying the elasticity of substitution across intermediate goods from $\theta = 10$ to $\theta = 5$ while simultaneously varying the menu costs from 11% to 4% of revenue to maintain the same average price duration at 3.5% inflation as in our baseline calibration.

Panels A and D in Figure 8 illustrate the implied variation in the degree of price stickiness for both methods. In each case, higher steady state levels of inflation lead to more frequent updating of prices. Panels B and E plot the welfare associated with different $\bar{\pi}$ and different parameters for each approach. In the first approach, more endogeneity in the degree of price stickiness leads to slightly lower optimal levels of inflation. The fact that the optimal inflation rate declines when the sensitivity of price stickiness to inflation rises may seem counterintuitive: higher inflation should lead to faster updating of prices and therefore less price dispersion. But the endogenous rate of price stickiness also implies that inflation volatility rises more rapidly with average inflation than in the baseline case which raises the cost of inflation. Given our parameter values and the range of inflation rates that we consider, the latter effect approximately offsets the former at low levels of inflation so endogenous price stickiness actually leads to slightly lower optimal inflation rates

¹⁵ The size of the menu costs is in line with those reported in Zbaracki et al. (2004) of approximately 6% of operating expenses. However, Zbaracki et al. (2004) emphasize that these costs appear to be largely associated with multiproduct pricing decisions, whereas in our model firms produce a single good. Unfortunately, there is little other direct empirical evidence on the size of menu costs.

than the standard Calvo model. Second, the welfare costs of inflation at the optimal rate are rising with endogenous price setting. This reflects the fact that, given the same low optimal rate of inflation, more endogeneity is associated with higher degrees of price stickiness and therefore a higher cost of inflation when $\bar{\pi} < 3.5\%$. Third, despite the fact that the optimal rate of inflation varies little with endogenous price stickiness, the costs of much higher $\bar{\pi}$ are significantly lower relative to our baseline, because higher inflation leads to more frequent price changes and therefore price dispersion rises less rapidly with steady state inflation than under constant price stickiness.

Similarly, allowing for state-dependent pricing lowers the optimal inflation rate relative to our baseline model as well as the welfare losses from higher inflation rates, with the latter reflecting the reduced sensitivity of price dispersion to inflation due to endogenous pricing decisions, as in Burstein and Hellwig (2008).¹⁶ Reducing the size of menu costs raises the optimal inflation rate, but the effects are quantitatively small. In short, the state-dependent model confirms the results reached using the endogenous Calvo price durations: allowing for price stickiness to fall with inflation lowers the optimal inflation rate relative our baseline model, even as the costs of higher inflation are substantially reduced.

VI Normative Implications

We have so far been treating the question of the optimal inflation rate independently of the systematic response of the central bank to macroeconomic fluctuations. In this section, we investigate the implications of optimal stabilization policy, both under commitment and discretion, for the optimal rate of inflation. We then consider whether altering the parameters of the Taylor rule can lead to outcomes that approach those achieved using optimal stabilization policy with commitment.

6.1 Optimal Stabilization Policy

In the baseline model, stabilization policy follows a Taylor rule calibrated to match the historical behavior of the Federal Reserve. Following Giannoni and Woodford (2010) and Woodford (2010), we derive the optimal policy rules under commitment and discretion, and simulate the model with our baseline parameter values.¹⁷ Panel A of Figure 9 plots expected utility under both policies. When the central bank can commit to a particular policy rule then the optimal inflation rate is practically zero (0.2%). This is not because the ZLB binds less frequently—it actually binds much *more* frequently at the optimal inflation rate—but rather because the cost of the ZLB with commitment is negligible. This reflects the fact that in the event of a large shock, the central bank will promise to keep interest rates low for an extended period, which significantly reduces the impact of the shock and thus the cost of the ZLB. Under discretionary policy, the reverse is true. The inability of the

¹⁶ For easier comparison with the other models, we present welfare results net of menu costs. Including the menu costs pushes the optimal inflation rate down by approximately 0.2%-0.5% per year.

¹⁷ We assume that the central bank can always commit to a long-run inflation target, even if it cannot commit to a particular stabilization policy rule. Therefore, there is no inflation bias under this policy.

central bank to promise to keep interest rates low after the ZLB constraint ceases to bind yields very large costs of the ZLB. As a result, optimal policy entails a higher level of trend inflation, reaching 2.7% in our baseline calibration, which is still within the target range of most central banks today.¹⁸ Thus, a positive but low optimal inflation rate does not hinge on the assumption that the central bank follows a Taylor rule, but also obtains with optimal policy under discretion. The commitment case, while likely of limited practical relevance, suggests that an improved monetary policy could deliver important welfare gains by reducing the costs of the ZLB when they occur rather than trying to avoid these episodes.

6.2 Taylor Rule Parameters, Price-Level Targeting and the Optimal Inflation Rate

Given the size of the welfare differential between optimal policy with discretion/Taylor rule versus commitment, we investigate to what extent alternative monetary policy rules can improve outcomes in the face of the ZLB. Thus, we consider the implications of alternative parameter values in the Taylor rule, illustrated in Panels C-F of Figure 9. Varying the long-run response to inflation from 2 to 5 has unambiguous effects on welfare: stronger long-run responses to inflation ϕ_π raise welfare in the model for all inflation rates. Intuitively, this stronger systematic response reduces inflation and output volatility, thereby leading to a lower frequency of being at the ZLB and higher utility. However, this has little effect on the optimal inflation rate, which ranges from 1.7% when $\phi_\pi = 2$ to 1.2% when $\phi_\pi = 5$. In terms of responses to the real side of the economy which are captured by coefficients ϕ_y and ϕ_{gy} , stronger responses to output growth generally lower welfare while higher responses to the output gap are welfare-improving. However, the quantitative changes in welfare are small and neither measure plays a significant role in determining the optimal inflation rate within the determinacy region of the parameter space.

While our baseline specification of the Taylor rule restricts the endogenous response of the central bank to inflation and the real side of the economy, an additional factor sometimes considered is price-level targeting (PLT). While the evidence for central banks actually following PLT remains scarce, PLT has nonetheless received substantial attention in the literature for several reasons. First, as emphasized in Woodford (2003), PLT guarantees determinacy under zero trend inflation for any positive response to the price level gap. Second, Coibion and Gorodnichenko (2011) show that PLT ensures determinacy for $\bar{\pi} > 0$ as well, and is not subject to the deterioration of the Taylor principle as a result of positive trend inflation which occurs when the central bank responds only to inflation. Third, Gorodnichenko and Shapiro (2007) show that PLT robustly helps stabilize inflation expectations, thereby yielding smaller inflation and output volatility than would occur in inflation-targeting regimes.

¹⁸ Billi (2011) reports an optimal inflation rate of 13% under discretion in the presence of the ZLB. The main differences between his approach and ours are that a) he assumes higher levels of indexation (0.90) and b) he log-linearizes the New Keynesian model and performs welfare approximations under the assumption of zero steady-state inflation so none of the three costs of inflation which we emphasize here are included in his model.

We extend our baseline model to include PLT in the central bank's reaction function ($\phi_p > 0$). Panel F of Figure 9 shows the effects of PLT on welfare for different $\bar{\pi}$ as well as its implications for the optimal rate of inflation. First, PLT strictly increases welfare for any $\bar{\pi}$, especially at low levels of inflation. Second, PLT leads to much lower levels of optimal inflation than inflation-targeting regimes. Even for moderate responses to the price level gap, the optimal level of inflation is less than 0.3 percent per year, which is close to the optimal 0.2 percent per year inflation rate in the Ramsey problem under commitment. This magnitude practically implies *price level stability* (rather than inflation stability) which is, in fact, the mandated objective for most central banks.

The intuition for why PLT delivers such a small optimal inflation rate is straightforward. First, as observed in Gorodnichenko and Shapiro (2007), PLT stabilizes expectations and has a profound effect on output and inflation volatility. In our simulations, the reduction in inflation and output volatility is so substantial that the welfare costs of inflation are almost exclusively driven by the steady-state effects. As a result of reduced volatility, the ZLB binds less frequently. For example, with $\phi_p = 0.3$, the ZLB binds less than 1.4 percent of the time at a steady-state level of inflation of 3.5%. Second, even if the nominal rate hits zero, the policy rule remains a potent factor in stimulating the economy despite the ZLB because agents know that the deflationary pressures during the ZLB will have to be offset by above-average inflation in the future. This limits the downward movement in inflationary expectations and therefore the associated increase in real interest rates. In short, PLT limits the extent of deflationary spirals so that the exit from a ZLB episode occurs more rapidly and the welfare costs of the ZLB are substantially reduced. To give a sense of the magnitude of the associated welfare change, we note that by increasing ϕ_p from zero to 0.25 (combined with the appropriate change in the optimal rate of inflation), a policymaker could raise welfare by the equivalent to a permanent increase in consumption of 0.5 percent and approach the welfare gains for the optimal policy in the Ramsey problem under commitment. Thus, these results provide a new justification for the consideration of PLT by monetary policymakers.

VII Concluding remarks

If nothing else, the Great Recession has taught monetary economists one lesson: the zero lower bound (ZLB) is not a theoretical curiosity of interest only to historians of the Great Depression or as a precautionary tale against overly cautious policy-makers such as the Japanese monetary and fiscal authorities in the early 1990s. Instead, the pervasiveness of the zero bound constraint among major industrial countries has demonstrated the necessity of incorporating this feature into modern macroeconomic models. Indeed, the recent interest in raising the inflation targets of central banks has resurrected a basic question for macroeconomists: what is the optimal inflation rate? Strikingly, New Keynesian models, with their pervasive reliance on the assumption of zero steady-state inflation, have been ill-equipped to answer this key question for central bankers.

We provide an integrated treatment of the effects of non-zero steady-state inflation in New Keynesian models. Most importantly, we derive an approximation to the utility function of the representative agent which incorporates the various dimensions along which steady-state inflation matters: the steady-state, the dynamics of the model, and the coefficients of the utility-function approximation. This allows us to study the optimal rate of inflation using a welfare criterion derived explicitly from the microfoundations of the model. Combining this with the zero-bound on nominal interest rates, we are then able to study the costs and benefits of steady-state inflation and quantify the optimal rate of inflation in models with time and state-dependent pricing. Our baseline result is that this optimal rate of inflation is fairly low: typically less than two percent a year. We show that this result is robust to a variety of parameter specifications and modifications of the model.

Given that most central banks are targeting inflation rates between 1% and 3% a year, our results can be interpreted as supporting the current regimes, while providing little evidence in favor of raising these targets to provide additional insurance against the zero-bound constraint on interest rates. Furthermore, our results go some way in resolving the apparent disconnect between observed inflation targets and prescriptions from standard monetary models. From a normative point of view, we also show that welfare could be substantially improved by introducing price-level targeting. The latter helps to stabilize economic fluctuations and reduces the probability of hitting the zero-lower bound. As a result, the optimal inflation rate under a price level targeting regime would be close to zero. In other words, optimal monetary policy, characterized as a combination of a low inflation target and a systematic response of nominal interest rates to deviations of the price-level from its target, can be interpreted as being very close to the “price stability” enshrined in the legal mandates of most central banks.

In the absence of such a change in the interest rate rule, our results suggest that higher inflation targets are likely too blunt an instrument to address the ZLB in a way that significantly increases aggregate welfare. This is not because ZLB episodes are not costly, but rather because the perpetual costs of higher inflation outweigh the benefits of more infrequent ZLB episodes. Addressing the very large costs associated with ZLB episodes is therefore likely to require alternative policies more explicitly focused on these specific episodes, such as countercyclical fiscal policies or the use of non-standard monetary policy tools. The lack of consensus on the efficacy of these policy tools, however, suggests that they should be high on the research agenda of macroeconomists.

References

- Amano, Robert, and Malik Shukayev, 2010. “Risk Premium Shocks and the Zero Bound on Nominal Interest Rates,” forthcoming in *Journal of Money, Credit and Banking*.
- Aruoba, S. Boragan, and Frank Schorfheide, 2011. “Sticky Prices versus Monetary Frictions: An Estimation of Policy Trade-offs,” *AEJ Macroeconomics* 3(1), 60-90.
- Ascari, Guido, and Tiziano Ropele, 2007. “Trend Inflation, Taylor Principle, and Indeterminacy,” *Journal of Money, Credit and Banking* 48(1), 1557-1584.

- Basu, Susanto, and John G. Fernald, 1997. "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy* 105, 249-283.
- Billi, Roberto M., 2011, "Optimal Inflation for the US Economy?" *American Economic Journal: Macroeconomics* 3(3), 29-52.
- Bils, Mark, and Peter J. Klenow, 2004. "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy* 112(5), 947-985.
- Barlevy, Gadi, 2004. "The Costs of Business Cycles and the Benefits of Stabilization: A Survey" NBER Working Paper 10926.
- Bhaskar, Venkataraman, 2002. "On Endogenously Staggered Prices," *Review of Economic Studies* 69(1), 97-116.
- Bodenstein, Martin, Christopher J. Erceg, and Luca Guerrieri, 2009. "The effects of foreign shocks when interest rates are at zero," International Finance Discussion Paper 983, Board of Governors.
- Burnside, Craig, 1996. "Production Function Regressions, Returns to Scale, and Externalities," *Journal of Monetary Economics* 37, 177-201.
- Burstein, Ariel and Christian Hellwig, 2008. "Welfare Costs of Inflation in a Menu Cost Model," *American Economic Review* 98(2), 438-443.
- Christiano, Lawrence J., Martin Eichenbaum, and Sergio Rebelo, 2011. "When is the government spending multiplier large?" *Journal of Political Economy* 119(1), 78-121.
- Cogley, Timothy, and Argia Sbordone, 2008. "Trend Inflation, Indexation and Inflation Persistence in the New Keynesian Phillips Curve," *American Economic Review* 98(5), 2101-2026.
- Coibion, Olivier, and Yuriy Gorodnichenko, 2011. "Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation," *American Economic Review* 101(1), 341-370.
- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland, 2011. "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Interest Rates in Light of the ZLB? NBER Working Paper 16093.
- Cooley, Thomas F. and Gary D. Hansen, 1991. "The Welfare Costs of Moderate Inflation," *Journal of Money, Credit and Banking* 23(3) 483-503.
- Dickens, W., L. Goette, E. Groshen, S. Holden, J. Messina, M. Schweitzer, J. Turunen, M. Ward, 2007. "How wages change: Micro evidence from the International Wage Flexibility Project," *Journal of Economic Perspectives* 21(2), 195-214.
- Dhyne, Emmanuel, and Jerzy D. Konieczny, 2007. "Temporal Distribution of Price Changes: Staggering in the Large and Synchronization in the Small" National Bank of Belgium Working Paper 116.
- Dotsey, Michael, Robert G. King, and Alexander L. Wolman, 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics* 114, 655-690.
- Eggertsson, Gauti, and Michael Woodford, 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 1, 139-211.
- Eggertsson, Gauti, and Michael Woodford, 2004. "Policy Options in a Liquidity Trap," *American Economic Review* 94(2), 76-79.
- Fair, Ray C., and John B. Taylor, 1983. "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models," *Econometrica* 51(4), 1169-1185.
- Feldstein, Martin S., 1997. "The Costs and Benefits of Going from Low Inflation to Price Stability," in *Reducing Inflation: Motivation and Strategy*, 123-166. NBER, Inc.
- Friedman, Milton, 1969. *The Optimum Quantity of Money*, Macmillan.
- Gali, Jordi, and Mark Gertler, 1999. "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* 44(2), 195-222.
- Giannoni, Mark P., and Michael Woodford, 2010. "Optimal Target Criterion for Stabilization Policy" NBER Working Paper 15757.
- Gorodnichenko, Yuriy, and Serena Ng, 2010. "Estimation of DSGE models when data are persistent," *Journal of Monetary Economics* 57, 325-340.
- Gorodnichenko, Yuriy, and Matthew D. Shapiro, 2007. "Monetary Policy When Potential Output Is Uncertain: Understanding the Growth Gamble of the 1990s," *Journal of Monetary Economics* 54(4), 1132-1162.

- Hausman, Jerry, Gregory Leonard and J. Douglas Zona, 1994. "Competitive Analysis with Differentiated Products," *Annales d'Economie et de Statistique* 34, 159-180.
- Kiley, Michael E., 2007. "Is Moderate-to-High Inflation Inherently Unstable?" *International Journal of Central Banking* 3(2), 173-201.
- Khan, A., Robert G. King and Alexander L. Wolman, 2003. "Optimal Monetary Policy," *Review of Economic Studies* 70(4), 825-860.
- Kim, Jinill, and Francisco J. Ruge-Murcia, 2009. "How much inflation is necessary to grease the wheels?" *Journal of Monetary Economics* 56(3), 365-377.
- Peter J. Klenow and Benjamin A. Malin, 2011. "Microeconomic Evidence on Price-Setting," *Handbook of Monetary Economics* 3A, B. Friedman and M. Woodford ed.: Elsevier, 231-284.
- Levin, Andrew, Alexei Onatski, John Williams, Noah Williams, 2005. "Monetary Policy under Uncertainty in Micro-Founded Macroeconometric Models," *NBER Macroeconomics Annual* 20, 229-287.
- Lubik, Thomas A., and Frank Schorfheide, 2003. "Computing Sunspot Equilibria in Linear Rational Expectations Models," *Journal of Economic Dynamics and Control* 28(2), 273-285.
- Lucas, Robert, 1987. *Models of Business Cycles*. Oxford: Basil Blackwell.
- Mankiw, N. Gregory, 2007. *Macroeconomics*, Sixth Edition. New York: Worth Publishers.
- Nakamura, Emi, and Jón Steinsson, 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics* 123(4), 1415-1464.
- Okun, Arthur, 1971. "The Mirage of Steady Inflation," *Brookings Papers on Economic Activity* 2, 485-498.
- Orphanides, Athanasios, 2003. "Historical Monetary Policy Analysis and the Taylor Rule," *Journal of Monetary Economics* 50(5), 983-1022.
- Romer, David, 1990, "Staggered Price Setting with Endogenous Frequency of Adjustment," *Economic Letters* 32(3), 205-210.
- Smets, Frank, and Rafael Wouters, 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review* 97(3), 586-606.
- Schmitt-Grohe, Stephanie and Martin Uribe, 2007, "Optimal Inflation Stabilization in a Medium-Scale Macroeconomic Model" in *Monetary Policy Under Inflation Targeting*, edited by Klaus Schmidt-Hebbel and Rick Mishkin, Central Bank of Chile, Santiago, Chile, 2007, p. 125-186.
- Schmitt-Grohe, Stephanie, and Martin Uribe, 2010. "The Optimal Rate of Inflation" in *Handbook of Monetary Economics*, vol. 3B (forthcoming, Elsevier).
- Summers, Lawrence, 1991. "Price Stability: How Should Long-Term Monetary Policy Be Determined?" *Journal of Money, Credit and Banking* 23(3), 625-631.
- Taylor, John, 1977. "Staggered Wage Setting in a Macro Model," *American Economic Review* 69(2), 108-13.
- Taylor, John B., 1981. "On the Relationship Between the Variability of Inflation and the Average Inflation Rate," *Carnegie-Rochester Conference Series on Public Policy* 15, 57-85.
- Tobin, James, 1972. "Inflation and Unemployment," *American Economic Review* 62(1), 1-18.
- Walsh, Carl E., 2009. "Using Monetary Policy to Stabilize Economic Activity," in Federal Reserve Bank of Kansas City *Financial Stability and Macroeconomic Policy*, Jackson Hole Symposium 2010: 245-96.
- Williams, John C., 2009. "Heeding Deдалus: Optimal Inflation and the Zero Bound" *Brookings Papers on Economic Activity* 2009(2), 1-37.
- Wolman, Alexander L., 2005. "Real Implications of the Zero Bound on Interest Rates," *Journal of Money, Credit and Banking* 37(2) 273-296.
- Woodford, Michael, 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- Woodford, Michael, 2010. "Optimal Monetary Stabilization Policy," *Handbook of Monetary Economics*, vol. 3B (forthcoming, Elsevier).
- Zbaracki, M.J., M. Ritson M, D. Levy, S. Dutta, M. Bergen, 2004. "Managerial and customer costs of price adjustment: Direct evidence from industrial markets," *Review of Economics and Statistics* 86(2), 514-33.

Table 1: Baseline Parameter Values

<u>Parameters of Utility Function</u>		<u>Steady-State Values</u>	
η : Frisch Labor Elasticity	1.00	$\overline{g_y}$: Growth Rate of RGDP/cap	1.5% p.a.
β : Discount factor	0.998	$\overline{c_y}$: Consumption Share of GDP	0.80
h : Habit in consumption	0.7	$\overline{g_y}$: Government Share of GDP	0.20
<u>Pricing Parameters</u>		<u>Shock Persistence</u>	
θ : Elasticity of substitution	7	ρ_g : Government Spending Shocks	0.97
λ : Degree of Price Stickiness	0.55	ρ_m : Cost-Push Shocks	0.90
ω : Price indexation	0.00	ρ_q : Risk Premium Shocks	0.947
<u>Taylor Rule Parameters</u>		<u>Shock Volatility</u>	
ϕ_π : Long run response to inflation	2.50	σ_g : Government Spending Shocks	0.0052
ϕ_{g_y} : Long run response to output growth	1.50	σ_m : Cost-Push Shocks	0.0014
ϕ_x : Long run response to output gap	0.11	σ_q : Risk Premium Shocks	0.0024
ρ_1 : Interest smoothing	1.05	σ_a : Technology Shocks	0.0090
ρ_2 : Interest smoothing	-0.13	σ_r : Monetary Policy Shocks	0.0024

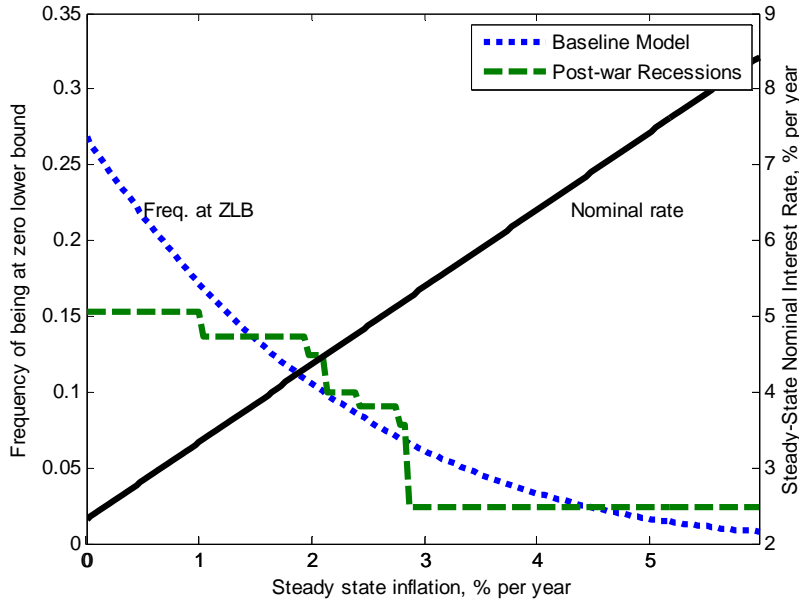
Note: The table presents the baseline parameter values assigned to the model in section 3.1 and used for solving for the optimal inflation rate in section 3.2. “p.a.” means per annum.

Table 2: Model Fit at the Baseline Parameterization

	Standard deviation		Serial correlation, AR(1) coefficient		Corr(x,output)	
	Model	Data	Model	Data	Model	Data
Consumption	0.0121	0.0128	0.75	0.81	0.82	0.78
Output	0.0124	0.0167	0.71	0.85	1.00	1.00
Inflation	0.0067	0.0065	0.77	0.76	0.10	0.19
Interest Rates	0.0081	0.0070	0.96	0.97	0.10	0.17

Notes: The table presents moments of the data and simulated series from the model parameterized at baseline values. Consistent with Gorodnichenko and Ng (2010), both model and data series are detrended with the Hodrick-Prescott ($\lambda = 1600$) filter to remove the trending component in the series.

Figure 1. Frequency of being in the Zero Lower Bound and Steady-State Nominal Interest Rate

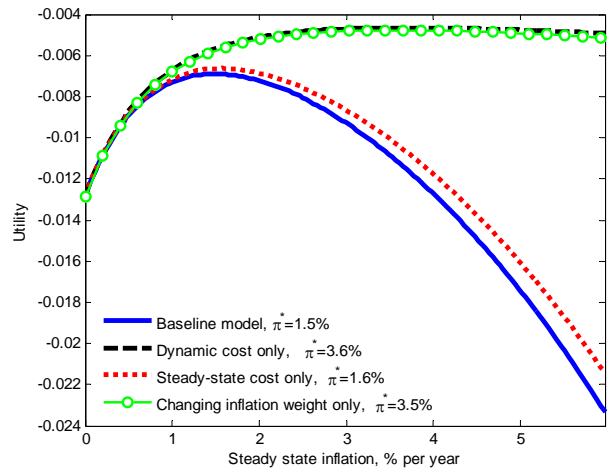
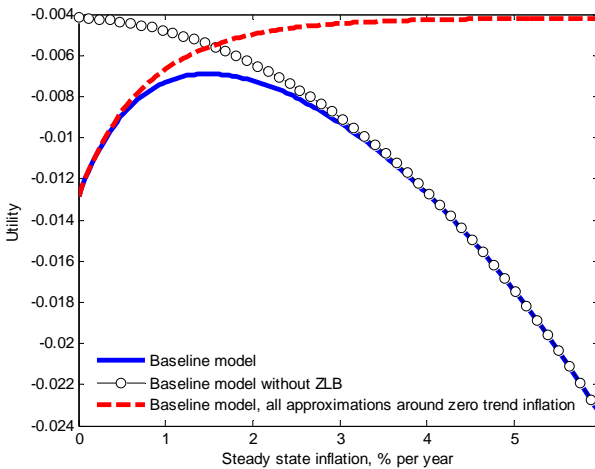


Note: The figure plots the steady-state level of the annualized nominal interest rate (right axis) implied by the baseline model of section 3 for different steady-state inflation rates. In addition, the figure plots the frequency of hitting the zero bound on nominal interest rates (left axis) from simulating the baseline model at different steady-state inflation rates as well as historical frequencies of ZLB for the U.S. for counterfactual inflation rates. See section 3.1 for details.

Figure 2. Utility at Different Levels of Steady-State Inflation

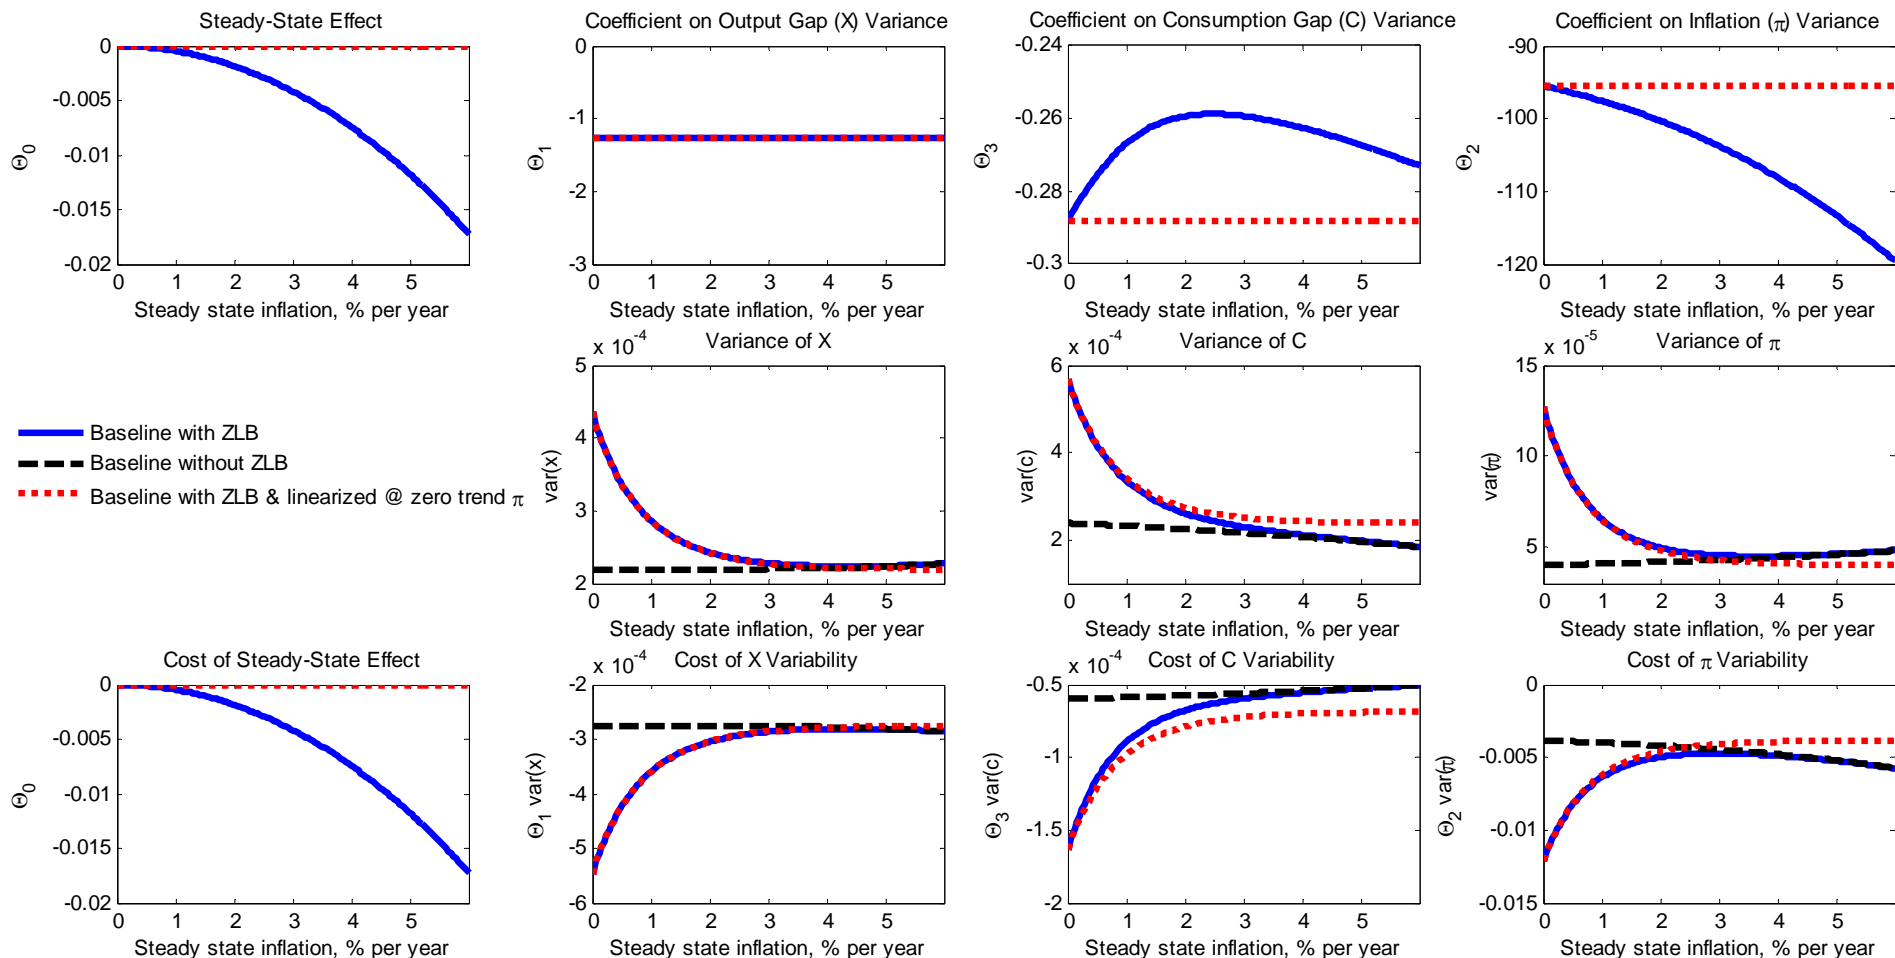
Panel A: Effect of ZLB

Panel B: Effects of positive trend inflation



Note: The figures plot the approximation to the utility function in Proposition 1 from simulating the model for different levels of steady-state inflation. Panel A includes results for the baseline model, the baseline model without the ZLB, as well as the model with the ZLB but omitting the three cost channels of inflation: steady-state effects, the changing coefficient on inflation variance in utility and the dynamic effects. Panel B reproduces our baseline with ZLB, then presents results when we restrict the model to include only one cost of inflation and the ZLB. “Dynamic cost only” includes only the dynamic effects of positive inflation and keeps the rest of the model being approximated around zero trend inflation, “Steady-state cost only” includes only the steady-state cost of inflation and keeps the rest of the model being approximated around zero trend inflation, while “Changing inflation weight only” includes only the changing coefficients on inflation variance in the loss function and keeps the rest of the model being approximated around zero trend inflation. See section 3.2 for details.


Figure 3. The Sources of Utility Costs of Inflation




Note: The first row of the figure plots the coefficients of the approximation to the utility function from Proposition 1 for different levels of trend inflation. The second row plots the variance of macroeconomic variables that enter the approximation to the utility function in Proposition 1 from simulating the model subject to the zero bound on nominal interest rates for different levels of steady-state inflation using the baseline parameter values of the model. The dashed black lines are the corresponding moments without the zero-bound on nominal interest rates, while the dotted lines correspond to the moments of the model when the dynamics are approximated at zero trend inflation. The third row plots the contribution of the different components of the approximation to the utility function in Proposition 1. See section 3.2 for details.

Figure 4. The Costs of Business Cycles

Panel A: The costs of the ZLB

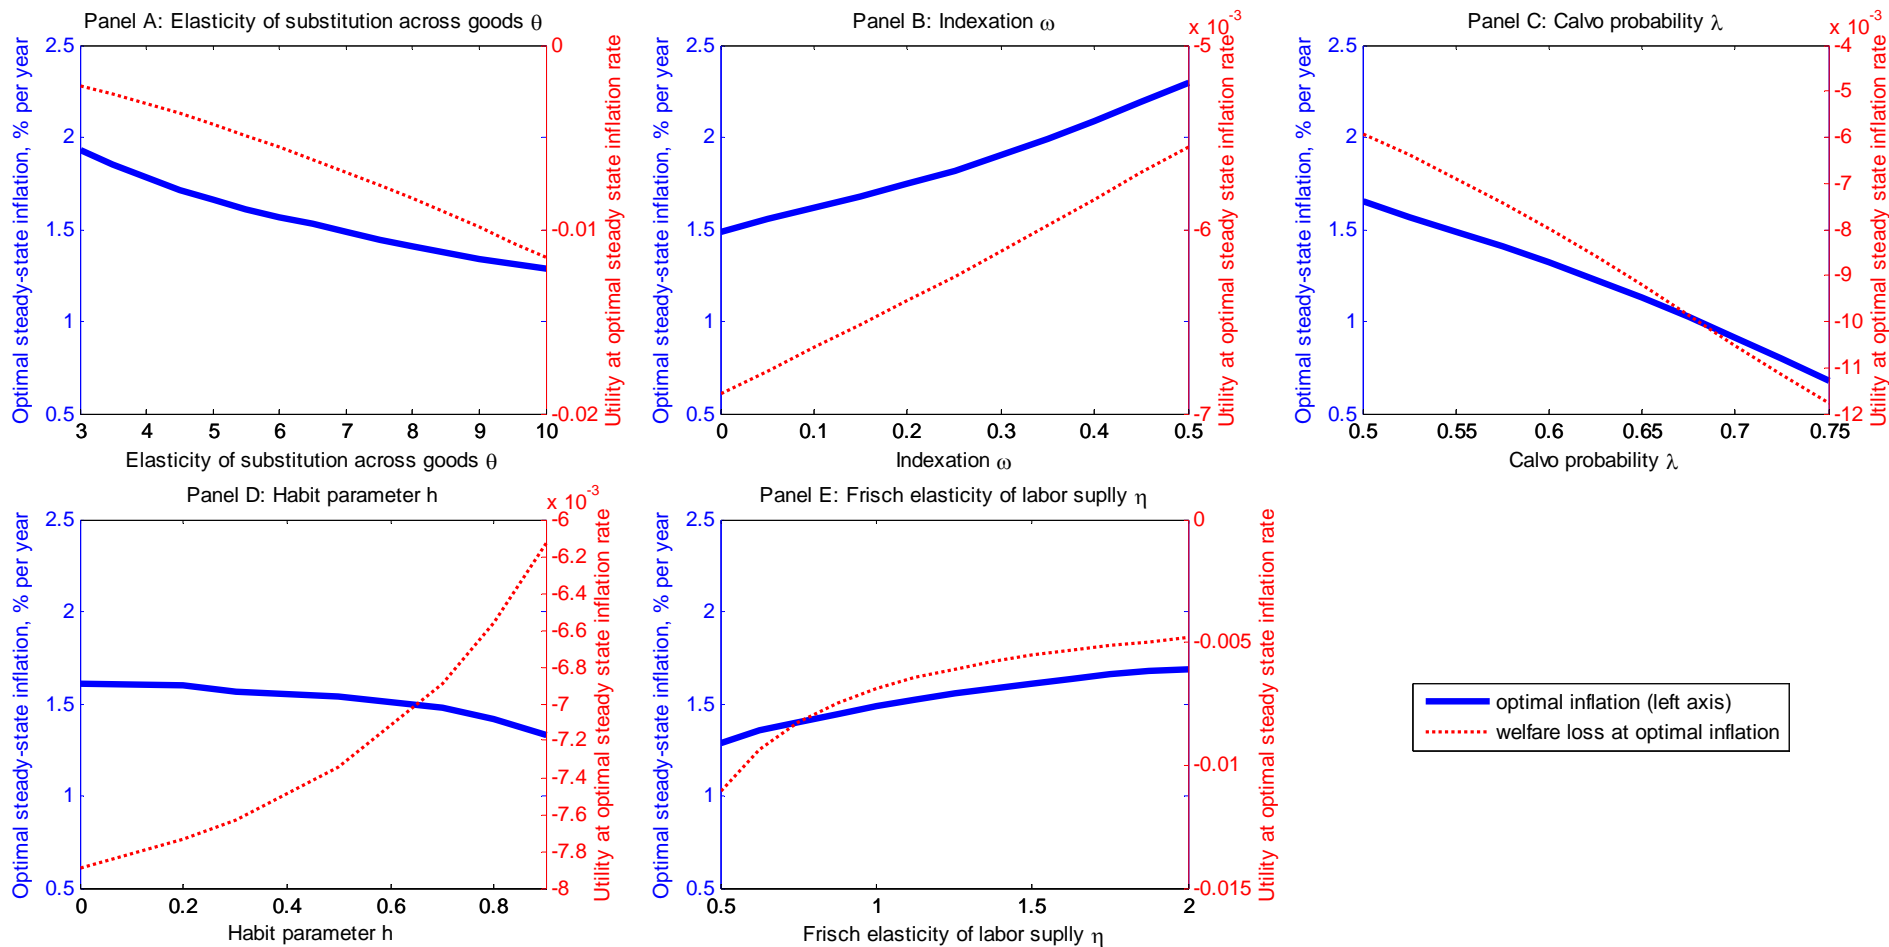


Panel B: Changing the weight on output gap volatility in the loss function



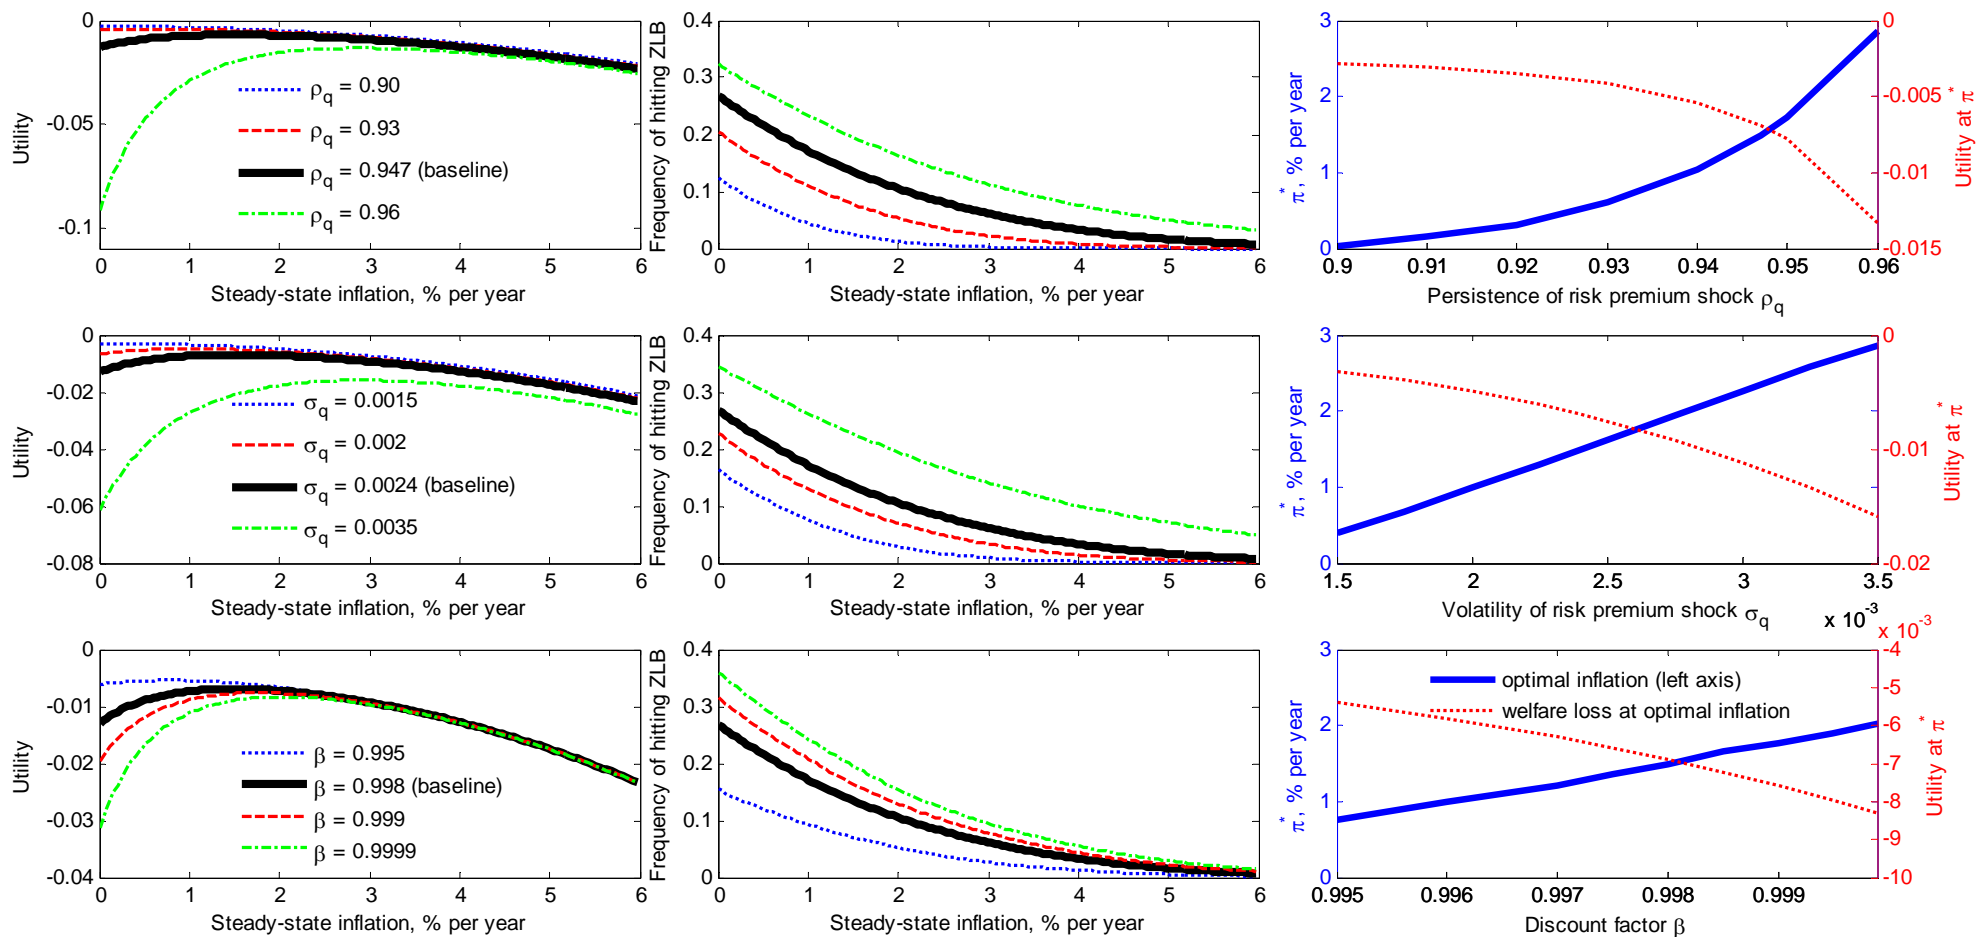
Note: The top panel plots the average duration of ZLB episodes in the baseline calibration and the implied average welfare cost per quarter of being at the ZLB for different levels of trend inflation. The bottom panel plots the effects of changing the coefficient on the variance of the output gap in the utility function approximation of Proposition 1 on the optimal inflation rate (left graph), the welfare costs of business cycle fluctuations (middle graph), and the average welfare costs of hitting the ZLB (right graph). The latter two are measured using Proposition 1 net of the steady-state effects of trend inflation. See sections 3.3 and 3.4 for details.

Figure 5. Sensitivity: Pricing Parameters, Habit, and Labor Supply Elasticity.



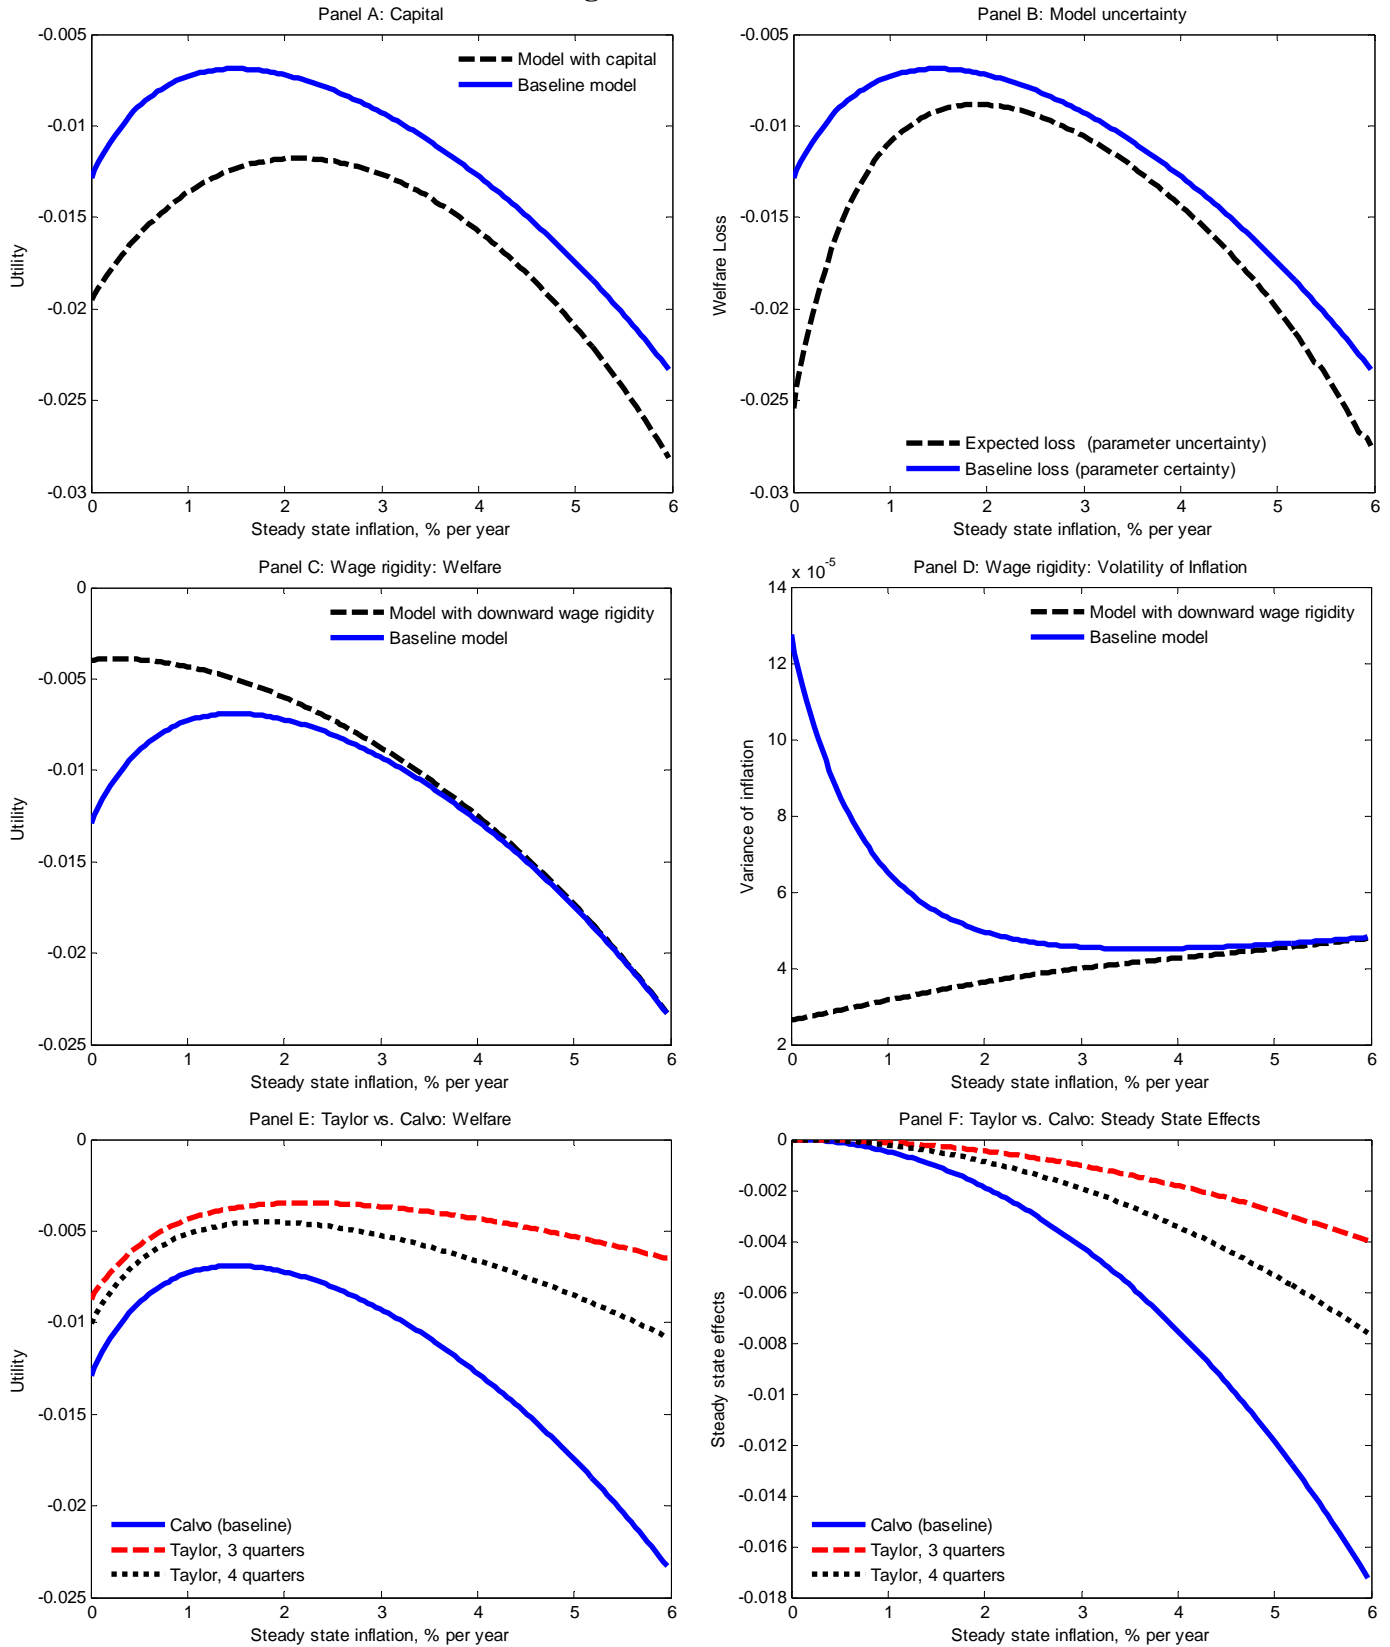
Note: Figures plot the optimal level of steady-state inflation and the welfare loss at the optimal steady state level of inflation as a function of a structural parameter. See section 4.1 for details.

Figure 6. Sensitivity: Risk Premium Shocks and the Discount Factor



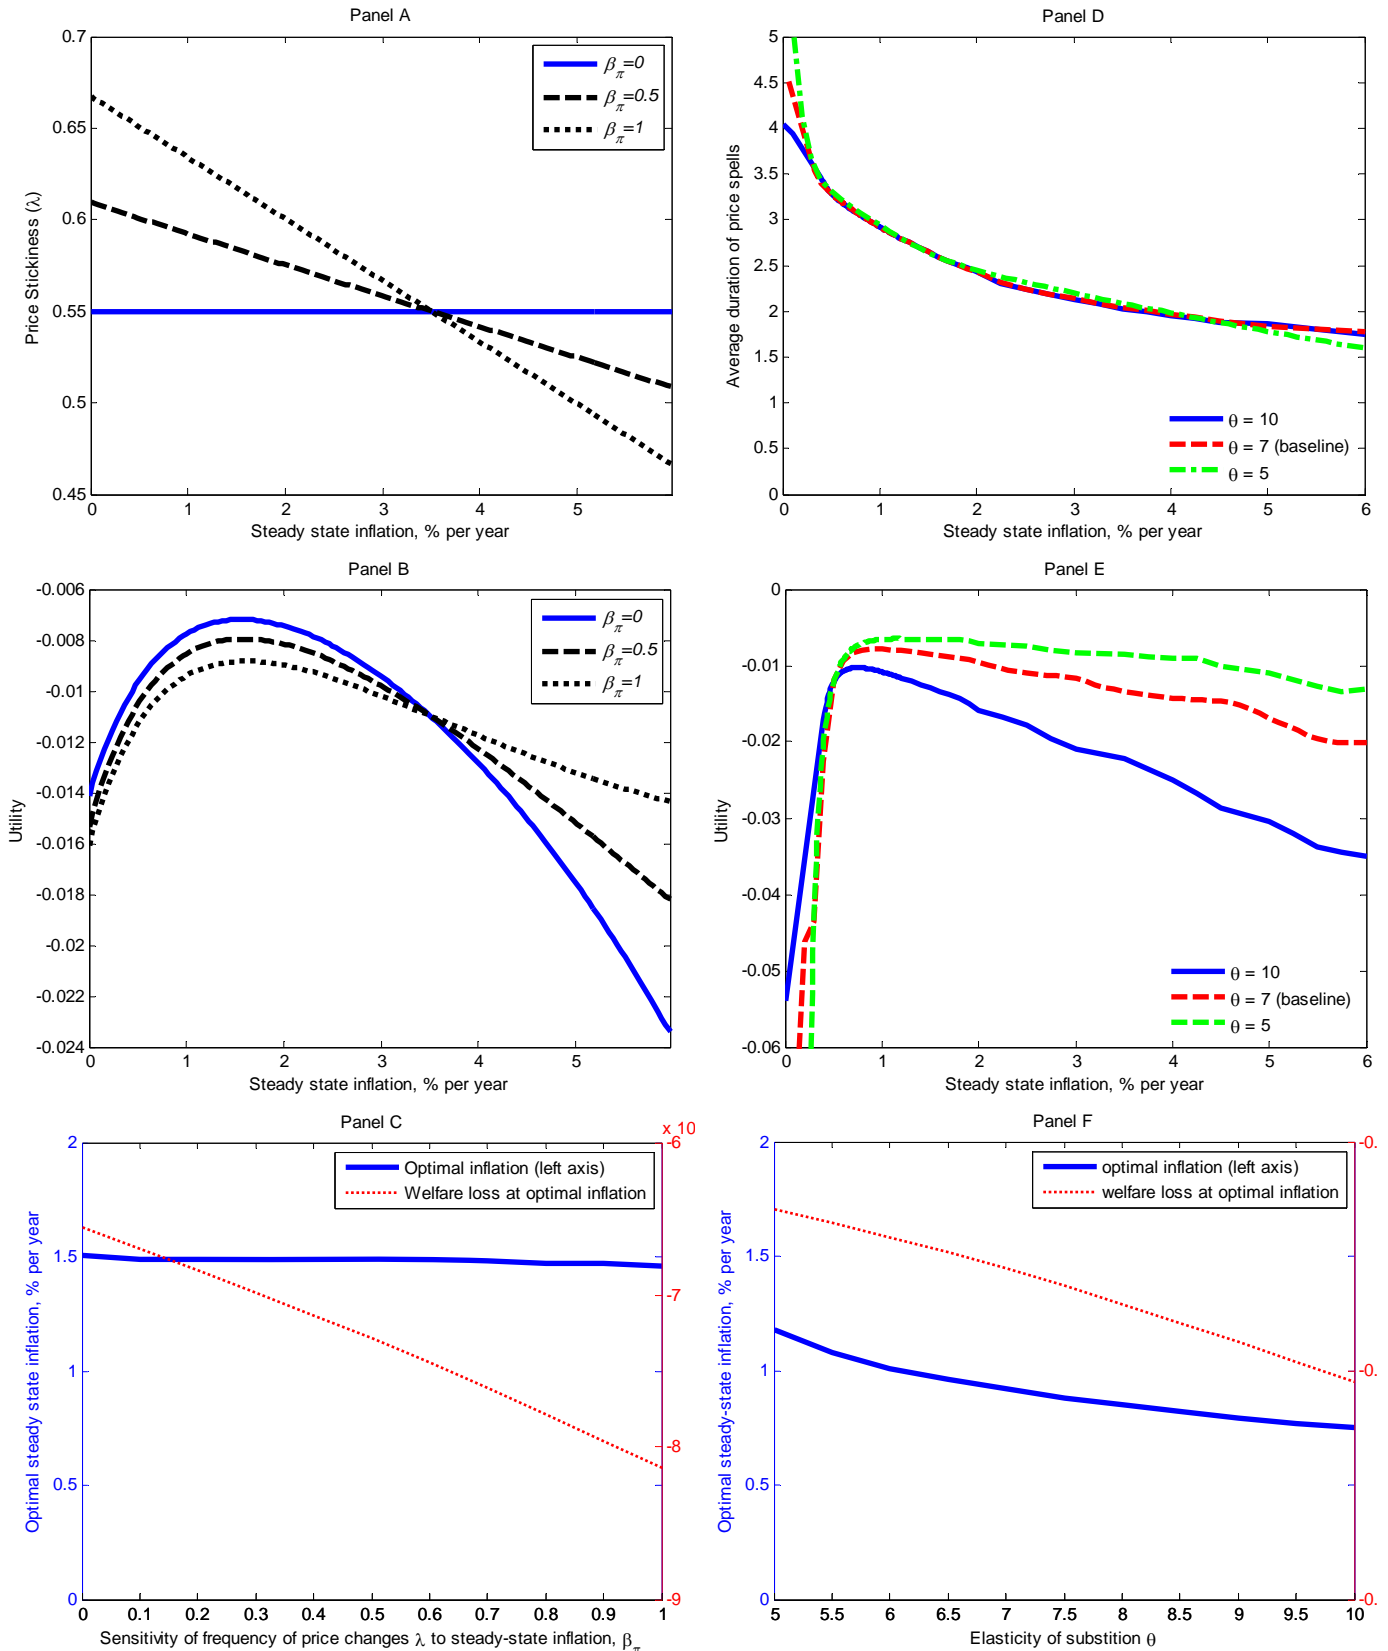
Note: Figures in the left column plot the welfare loss as a function of steady state inflation for alternative values of structural parameters. The solid thick back line corresponds to the baseline parameterization. Figures in the middle column plot the frequency of hitting the ZLB for different parameter values. Figures in the right column plot the optimal level of steady-state inflation and the welfare loss at the optimal steady state level of inflation as a function of a structural parameter. See section 4.2 for details.

Figure 7. Robustness.



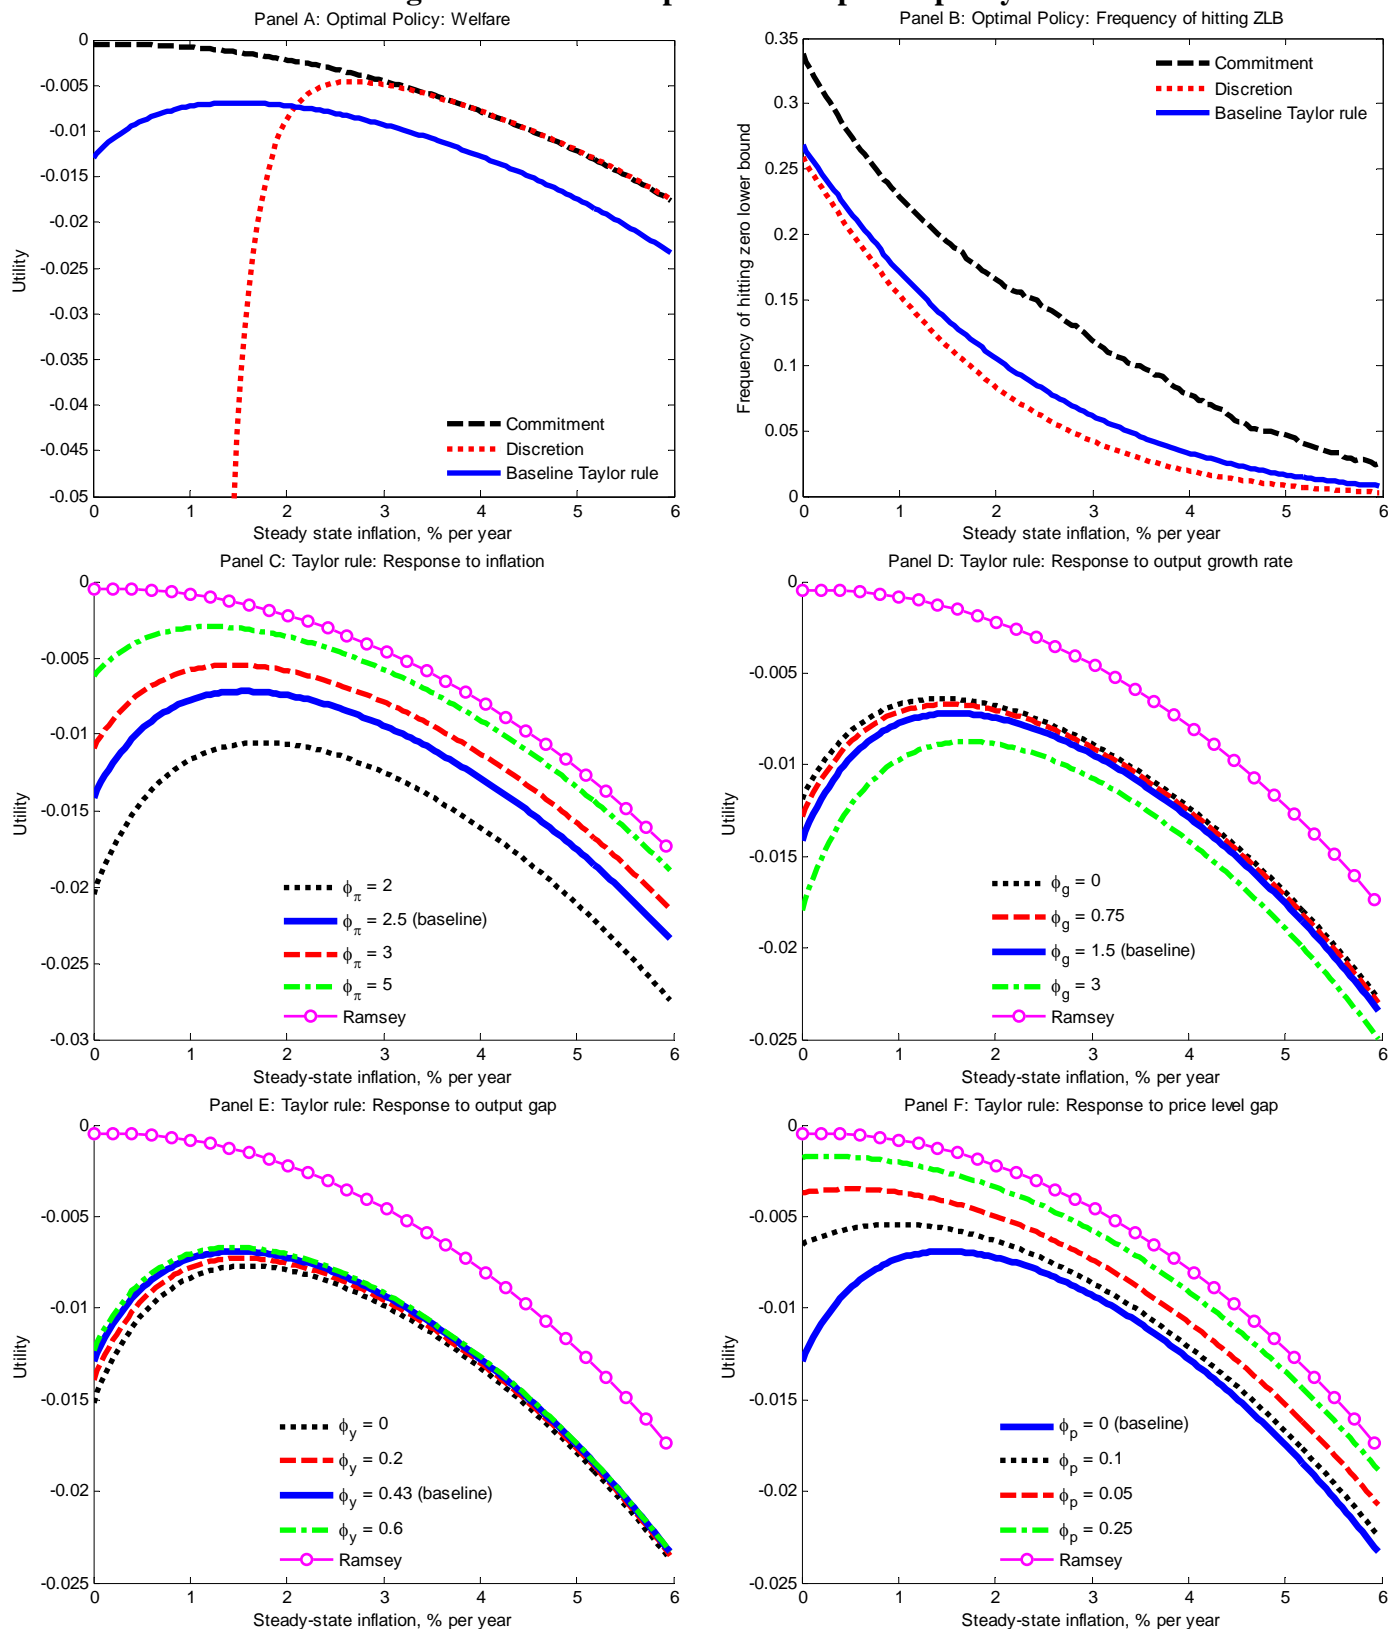
Notes: Panel A plots welfare in the baseline model and the model with capital. See section 5.1 for details. Panel B plots the expected utility in the baseline model and model-uncertainty. See section 5.2 for details. Panel C plots the utility associated with different steady-state inflation rates under the baseline model as well as the model with downward nominal wage rigidity. Panel D plots the variance of inflation for different steady-state inflation rates using our baseline model and the model with downward nominal wage rigidity. See section 5.3 for details. Panels E and F plot the implications of Calvo vs. Taylor price setting for welfare. *Taylor, X quarters* corresponds to the duration of price contracts equal to *X* quarters. See section 5.4 for details.

Figure 8. Endogenous and State-Dependent Price Stickiness
Endogenous Calvo Parameter **State-Dependent Pricing (Dotsey et al. 1999)**



Note: The figures plot the implications of endogenous and state-dependent price stickiness on the model. β_π is the effect of steady state inflation on the frequency of price changes $1 - \lambda$. $\beta_\pi = 0$ is our baseline case of exogenous price stickiness. See section 5.5 for details.

Figure 9. Positive Implications: Optimal policy.



Notes: Panels A and B plot outcomes for three scenarios: optimal policy with commitment (Ramsey); optimal policy without commitment (discretion); the baseline Taylor rule described in section 3.1. Optimal policies with and without discretion are described in section 6.1. Panels C- F plot the welfare loss as a function of $\bar{\pi}$ for alternative values of the monetary policy rule parameters. The solid thick blue line corresponds to the baseline parameterization. ϕ_π is the long-run response of interest rates to inflation, ϕ_g is the response to output growth, ϕ_y is the response to the output gap, and ϕ_p is the response to the price-level gap. See section 6.2 for details.

Not for publication material

Appendix A. Proofs for section 2.4.

Derivation of Proposition 1:

Proposition 1. The 2nd order approximation to expected per period utility in eq. (1) is¹⁹

$$\Theta_0 + \Theta_1 \text{var}(\hat{y}_t) + \Theta_2 \text{var}(\hat{\pi}_t) + \Theta_3 \text{var}(\hat{c}_t) \quad (\text{A.1})$$

where parameters $\Theta_i, i = \{0,1,2\}$ depend on the steady state inflation $\bar{\pi}$ and are given by

$$\begin{aligned} \Theta_0 &= \left[1 - \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} (1 - (1 + \eta^{-1})Q_y^0) \right] \log \bar{X} - \frac{(1-\Phi)(1+\eta^{-1})(1-\beta h)}{2(1-\bar{g}_y)(1-h)} [\log \bar{X}]^2 \\ &\quad - \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} \left\{ (1 + \eta^{-1})[Q_y^0]^2 - Q_y^0 + \frac{\theta^2(1+\eta^{-1})}{2} \bar{\Delta} \right\}, \\ \Theta_1 &= -\frac{\frac{1}{2}(1+\eta^{-1})(1-\beta h)}{1-\bar{g}_y(1-h)}, \\ \Theta_2 &= -\frac{\theta^2}{2(1-\bar{g}_y)(1-h)} \Gamma_3 \left\{ \left[Q_y^1(\theta^{-1} - 1) + (1 + \eta^{-1}) \left(1 + \frac{\theta^{-1}}{\theta} Q_y^0 Q_y^1 \right) \right] - (1 + \eta^{-1}) \frac{\theta^{-1}}{\theta} Q_y^1 \log \bar{X} \right\}, \\ \Theta_3 &= -\frac{h(1-\rho_c)}{(1-h)^2} \\ \Gamma_3 &= \frac{\Gamma_0}{1-\lambda\Gamma_1} \{(1-\lambda)M^2 + \lambda\}, \\ \Phi &= -\log\left(\frac{\theta-1}{\theta}\right). \end{aligned}$$

The proof of Proposition 1 follows from the following 5 Lemmas. First of all, we decompose utility described in equation (1) in the text into utility due to consumption and (dis)utility due to labor supply. Lemmas 1 and 2 provide second order approximations for each component.

Lemma 1. Utility from consumption in equation (1) is given by

$$u(C_t) = \frac{1}{1-h} (\tilde{c}_t - h\tilde{c}_{t-1}) - \frac{1}{2} \frac{h}{(1-h)^2} (\tilde{c}_t^2 - 2\tilde{c}_t\tilde{c}_{t-1} + \tilde{c}_{t-1}^2) + hot + tip \quad (\text{A.2})$$

where $\tilde{c}_t = \log(C_t/C_t^F)$ is the percent deviation of consumption from its flexible-price level *t.i.p.* stands for terms independent from policy, and *h.o.t.* means higher order terms.

Proof: The expansion of the utility derived from consumption around flexible price steady state is

$$\begin{aligned} u(C_t) &= u(\bar{C}_t^F(1-h)) + u'_{\bar{C}_t^F}(C_t - \bar{C}_t^F) + u'_{\bar{C}_{t-1}^F}(C_{t-1} - \bar{C}_{t-1}^F) + \frac{1}{2} u''_{\bar{C}_t^F}(C_t - \bar{C}_t^F)^2 \\ &\quad + u''_{\bar{C}_t^F \bar{C}_{t-1}^F}(C_t - \bar{C}_t^F)(C_{t-1} - \bar{C}_{t-1}^F) + \frac{1}{2} u''_{\bar{C}_{t-1}^F}(C_{t-1} - \bar{C}_{t-1}^F)^2 + hot \\ &= \bar{C}_t^F u'_{\bar{C}_t^F} \left(\tilde{c}_t + \frac{1}{2} \tilde{c}_t^2 \right) + \bar{C}_{t-1}^F u'_{\bar{C}_{t-1}^F} \left(\tilde{c}_{t-1} + \frac{1}{2} \tilde{c}_{t-1}^2 \right) + \frac{1}{2} (\bar{C}_t^F)^2 u''_{\bar{C}_t^F} \tilde{c}_t^2 + \frac{1}{2} \bar{C}_t^F \bar{C}_{t-1}^F u''_{\bar{C}_t^F \bar{C}_{t-1}^F} \tilde{c}_t \tilde{c}_{t-1} + \\ &\quad \frac{1}{2} (\bar{C}_{t-1}^F)^2 u''_{\bar{C}_{t-1}^F} \tilde{c}_{t-1}^2 + hot + tip \\ &= \frac{1}{1-h} \left(\tilde{c}_t + \frac{1}{2} \tilde{c}_t^2 \right) - \frac{h}{1-h} \left(\tilde{c}_{t-1} + \frac{1}{2} \tilde{c}_{t-1}^2 \right) - \frac{1}{2} \left(\frac{1}{1-h} \right)^2 \tilde{c}_t^2 + \frac{h}{(1-h)^2} \tilde{c}_t \tilde{c}_{t-1} - \frac{1}{2} \left(\frac{h}{1-h} \right)^2 \tilde{c}_{t-1}^2 + hot + tip \\ &= \frac{1}{1-h} (\tilde{c}_t - h\tilde{c}_{t-1}) - \frac{1}{2} \frac{h}{(1-h)^2} (\tilde{c}_t^2 - 2\tilde{c}_t\tilde{c}_{t-1} + \tilde{c}_{t-1}^2) + hot + tip \end{aligned}$$

where \bar{C}_t^F is the flexible price level of consumption in the steady state. Tildes denote percent deviations from the flexible price level, *tip* stands for terms independent from policy, *hot* means higher order terms. ■

Lemma 2. Using production function (9), define $\tilde{v}(Y_t(i)) \equiv v(N_t(i)) \equiv N_t(i)^{1+1/\eta}$. Then

$$\tilde{v}(Y_t(i)) \approx \bar{Y}_t^F \tilde{v}'_{\bar{Y}_t^F} \{ \tilde{y}_t(i) + \frac{1}{2} (1 + \eta^{-1}) \tilde{y}_t^2(i) \} + t. i. p. + h. o. t. \quad (\text{A.3})$$

where $\tilde{y}_t(i) = \log(Y_t(i)/\bar{Y}_t^F)$ is the deviation of firm *i*'s output from the flexible-price level of output \bar{Y}_t^F . Correspondingly, the total disutility from labor supply is

¹⁹ The approximation also contains three linear terms, the expected output gap, expected consumption and expected inflation, which are omitted here. Since the distortions to the steady state are small for the levels of trend inflation we consider, the coefficients that multiply these terms can be considered as first order so we can evaluate these terms using the first order approximation to the laws of motion as in Woodford (2003). We confirmed in numeric simulations that they can be ignored. Furthermore, second order effects on the expected output gap and expected inflation are likely to be quantitatively small since the linear solution closely approximates the nonlinear solution to the model (see Appendix F).

$$\int_0^1 \tilde{v}(Y_t(i)) di = \bar{Y}_t^F \tilde{v}'_{Y_t^F} \left\{ E_i \tilde{y}_t(i) + \frac{1}{2}(1 + \eta^{-1}) \int_0^1 \tilde{y}_t^2(i) di \right\} + t. i. p. + h. o. t. \quad (\text{A.4})$$

Proof: See Proposition 6.1 in Woodford (2003).

The key insight from Lemmas 1 and 2 is that welfare is diminished when consumption is low relative to its flexible-price level and when the cross-sectional dispersion of output is large. To understand and assess the implications of cross-sectional output dispersion, we need to examine the cross-sectional dispersion of prices.

Denote the cross-sectional dispersion of prices at time t with $\Delta_t = \text{var}_i(\log(P_t(i)))$ and let $\bar{\Delta}$ be the cross-sectional dispersion of prices in the non-stochastic steady state. It is straightforward to show that $\bar{\Delta} = \bar{\pi}^{\frac{2\lambda(1-\omega)^2}{(1-\lambda)^2}}$ where $\bar{\Delta}$ is increasing in price stickiness λ and steady-state inflation $\bar{\pi}$ and decreasing in the degree of indexation ω . Define $\bar{P}_t = E_i \log P_t(i)$ as the average (across firms) log price of goods.

Lemma 3. The difference between the log price index P_t and the average log price across firms \bar{P}_t is given by

$$\log P_t - \bar{P}_t = Q_p^0 + \frac{1-\theta}{2} Q_p^1 (\Delta_t - \bar{\Delta}) + h. o. t. \quad (\text{A.5})$$

where $Q_p^0 = \frac{1-\theta}{2} \bar{\Delta} / [1 + \frac{1}{2}(1-\theta)^2 \bar{\Delta}]^2$ and $Q_p^1 = [1 - \frac{1}{2}(1-\theta)^2 \bar{\Delta}] / [1 + \frac{1}{2}(1-\theta)^2 \bar{\Delta}]^3$.

Proof: Note that

$$\log P_t = \frac{1}{1-\theta} \log \left(\int_0^1 P_t^{1-\theta}(i) di \right) = \frac{1}{1-\theta} \left(\int_0^1 \log P_t^{1-\theta}(i) di \right) + \frac{1}{2} \frac{1}{1-\theta} \frac{1}{[EP_t^{1-\theta}(i)]} \text{var}_i \left(P_t^{1-\theta}(i) \right) + hot \quad (\text{A.6})$$

By the delta method, we have

$$\begin{aligned} \text{var}_i \left(P_t^{1-\theta}(i) \right) &= \text{var}_i \left(\exp\{(1-\theta)E \log P_t(i)\} \right) \\ &\approx (1-\theta)^2 [\exp\{(1-\theta)E \log P_t(i)\}]^2 \text{var}_i(\log P_t(i)) = (1-\theta)^2 [\exp\{(1-\theta)\bar{P}_t\}]^2 \Delta_t \end{aligned} \quad (\text{A.7})$$

Now observe that

$$\begin{aligned} EP_t^{1-\theta}(i) &= E[\exp\{(1-\theta) \log P_t(i)\}] \\ &\approx \exp\{(1-\theta)E \log P_t(i)\} + \frac{1}{2}(1-\theta)^2 [\exp\{(1-\theta)E \log P_t(i)\}] \text{var}_i(\log P_t(i)) \\ &= \exp\{(1-\theta)\bar{P}_t\} \left[1 + \frac{1}{2}(1-\theta)^2 \Delta_t \right] \end{aligned} \quad (\text{A.8})$$

Hence it follows that

$$\begin{aligned} \log P_t &= \bar{P}_t + \frac{1}{2} \frac{1}{1-\theta} \frac{(1-\theta)^2 [\exp\{(1-\theta)\bar{P}_t\}]^2 \Delta_t}{[\exp\{(1-\theta)\bar{P}_t\}]^2 [1 + \frac{1}{2}(1-\theta)^2 \Delta_t]^2} + hot = \bar{P}_t + \frac{1}{2} \frac{(1-\theta)\Delta_t}{[1 + \frac{1}{2}(1-\theta)^2 \Delta_t]^2} + hot \\ &= \bar{P}_t + Q_p^0 + \frac{1-\theta}{2} Q_p^1 (\Delta_t - \bar{\Delta}) + hot \end{aligned} \quad (\text{A.9})$$

where $Q_p^0 = \frac{\frac{(1-\theta)\bar{\Delta}}{2}}{[1 + \frac{1}{2}(1-\theta)^2 \bar{\Delta}]^2}$ and $Q_p^1 = \frac{1 - \frac{(1-\theta)^2 \bar{\Delta}^2}{2}}{[1 + \frac{1}{2}(1-\theta)^2 \bar{\Delta}]^3}$. ■

Lemma 3 is a manifestation of Jensen's inequality. Note that since $\bar{\Delta}$ is quadratic in $\bar{\pi}$, the dispersion of prices $\bar{\Delta}$ is approximately zero when $\bar{\pi} \approx 0$ and therefore $Q_p^0 \approx 0$, $Q_p^1 \approx 1$ so that $\log P_t - \bar{P}_t \approx \frac{1-\theta}{2} \Delta_t$, which is the standard result. Again, since $\bar{\Delta}$ is quadratic in $\bar{\pi}$, one can show that $\partial Q_p^0 / \partial \bar{\pi} \approx 0$, $\partial Q_p^1 / \partial \bar{\pi} \approx 0$ when $\bar{\pi} \approx 0$.

Using Lemma 3, we describe the dynamic properties of the price dispersion in Lemma 4.

Lemma 4. Let $\Xi_t = \Delta_t - \bar{\Delta}$ be the deviation of cross-section price dispersion from its non-stochastic steady state level $\bar{\Delta}$. Then

$$\begin{aligned} \Xi_t &= (\pi_t - \bar{\pi})^2 \Gamma_0 \left\{ (1-\lambda)M^2 + \lambda + \frac{(1-\theta)^2}{4} [Q_p^1]^2 \Gamma_2^2 + (1-\lambda)(1-\theta)MQ_p^1 \Gamma_2 - \lambda(1-\theta)Q_p^1 \Gamma_2 \right\} \\ &\quad + \Gamma_2 (\pi_t - \bar{\pi}) + \lambda \Gamma_1 \Xi_{t-1} + h. o. t. \end{aligned} \quad (\text{A.10})$$

where

$$\begin{aligned} \Gamma_0 &= \left\{ 1 + (\theta-1)Q_p^1 [(1-\lambda)(\bar{b} + Q_p^0) - \lambda(1-\omega)\bar{\pi}] \right\}^{-1}, \\ \Gamma_1 &= \left\{ 1 - (\theta-1)(1-\omega)\bar{\pi}Q_p^1 \right\} \Gamma_0, \end{aligned}$$

$$\Gamma_2 = 2\{(1-\lambda)M(\bar{b} + Q_p^0) + \lambda(1-\omega)\bar{\pi}\}\Gamma_0,$$

\bar{b} is the log of the optimal reset price in the non-stochastic steady state.

Proof: Recall that $b_t = \log(B_t/P_t)$ (note that in our notation b_t is the price relative to *current* price level) and $\bar{P}_t = E_i \log P_t(i)$, and $P_t = \left[\int_0^1 P_t(i)^{1-\theta} di \right]^{1/(1-\theta)}$. Also keep in mind that the fraction of firms that do not change their price index it to a fraction $0 \leq \omega \leq 1$ of steady-state inflation, $P_t(i) = P_{t-1}(i)\bar{\Pi}^\omega$. Then

$$\begin{aligned} \Delta_t &= \text{var}_i(\log P_t(i)) = \int_0^1 [\log P_t(i) - \bar{P}_t]^2 di = \int_0^{1-\lambda} [\log P_t(i) - \bar{P}_t]^2 di + \int_{1-\lambda}^1 [\log P_t(i) - \bar{P}_t]^2 di \\ &= \int_0^{1-\lambda} [\log P_t(i) - \log P_t]^2 di + 2 \int_0^{1-\lambda} [\log P_t(i) - \log P_t][\log P_t - \bar{P}_t] di + \int_0^{1-\lambda} [\log P_t - \bar{P}_t]^2 di + \\ &\quad \int_{1-\lambda}^1 [(\log P_{t-1}(i) - \bar{P}_{t-1}) + (\bar{P}_{t-1} - \bar{P}_t + \omega\bar{\pi})]^2 di \\ &= (1-\lambda)b_t^2 + 2(1-\lambda)b_t[\log P_t - \bar{P}_t] + (1-\lambda)[\log P_t - \bar{P}_t]^2 + \lambda\Delta_{t-1} + \lambda[\bar{P}_{t-1} - \bar{P}_t + \omega\bar{\pi}]^2 \end{aligned} \quad (\text{A.11})$$

From Lemma 3 we have

$$\log P_t - \bar{P}_t \approx Q_p^0 + \frac{1-\theta}{2} Q_p^1 (\Delta_t - \bar{\Delta}) \quad (\text{A.12})$$

Also from equation (21) in the paper we have that

$$\pi_t - \bar{\pi} = \frac{1-\lambda\bar{\Pi}^{(\theta-1)(1-\omega)}}{\lambda\bar{\Pi}^{(\theta-1)(1-\omega)}} (b_t - \bar{b}) \Rightarrow b_t = \bar{b} + \frac{\lambda\bar{\Pi}^{(\theta-1)(1-\omega)}}{1-\lambda\bar{\Pi}^{(\theta-1)(1-\omega)}} (\pi_t - \bar{\pi}) = \bar{b} + M(\pi_t - \bar{\pi}). \quad (\text{A.13})$$

We use a guess and verify procedure to determine the dispersion of prices. In particular, we posit that if the deviation of cross-sectional price dispersion from its non-stochastic steady state level $\bar{\Delta}$ is $\Xi_t = (\Delta_t - \bar{\Delta})$ then

$$\Xi_t = K_Q(\pi_t - \bar{\pi}) + Z_{1,Q}(\pi_t - \bar{\pi})^2 + F_Q\Xi_{t-1} \quad (\text{A.14})$$

where we assume that K_Q is of the same order as the shock processes, so that the first term is of second order. Assuming that the initial degree of price dispersion is also of second order then the degree of price dispersion is of second order at all times.

Plugging (A.13) and (A.11) into (A.10) and using the guess (A.13) yields,

$$\begin{aligned} \Delta_t &= (1-\lambda) \left(\bar{b} + M(\pi_t - \bar{\pi}) \right)^2 + 2(1-\lambda) \left[\bar{b} + M(\pi_t - \bar{\pi}) \right] \left[Q_p^0 + \frac{1-\theta}{2} Q_p^1 \Xi_t \right] \\ &\quad + (1-\lambda) \left[Q_p^0 + \frac{1-\theta}{2} Q_p^1 \Xi_t \right]^2 + \lambda\Delta_{t-1} \\ &\quad + \lambda \{ [\log \bar{P}_t - \log \bar{P}_{t-1} - \omega\bar{\pi}] - [\log \bar{P}_t - \bar{P}_t] + [\log \bar{P}_{t-1} - \bar{P}_{t-1}] \}^2 \\ &= (1-\lambda)\bar{b}^2 + (1-\lambda)M^2(\pi_t - \bar{\pi})^2 + 2(1-\lambda)\bar{b}M(\pi_t - \bar{\pi}) \\ &\quad + 2(1-\lambda)\bar{b}Q_p^0 + (1-\lambda)(1-\theta)\bar{b}Q_p^1\Xi_t + 2(1-\lambda)Q_p^0M(\pi_t - \bar{\pi}) + \\ &\quad (1-\lambda)(1-\theta)MQ_p^1K_Q(\pi_t - \bar{\pi})^2 \\ &\quad + (1-\lambda)(Q_p^0)^2 + (1-\lambda)(1-\theta)Q_p^0Q_p^1\Xi_t + (1-\lambda) \left(\frac{1-\theta}{2} \right)^2 (Q_p^1)^2 K_Q^2(\pi_t - \bar{\pi})^2 + \lambda\Delta_{t-1} \\ &\quad + \lambda(\pi_t - \bar{\pi})^2 + \lambda(1-\omega)^2\bar{\pi}^2 + 2\lambda(1-\omega)\bar{\pi}(\pi_t - \bar{\pi}) + \lambda \left(\frac{1-\theta}{2} \right)^2 (Q_p^1)^2 K_Q^2(\pi_t - \bar{\pi})^2 + \\ &\quad \lambda(Q_p^0)^2 + \lambda(1-\theta)Q_p^0Q_p^1\Xi_t \\ &\quad + \lambda \left(\frac{1-\theta}{2} \right)^2 (Q_p^1)^2 \Xi_{t-1}^2 + \lambda(Q_p^0)^2 + \lambda(1-\theta)Q_p^0Q_p^1\Xi_{t-1} \\ &\quad - \lambda(1-\theta)Q_p^1K_Q(\pi_t - \bar{\pi})^2 - 2\lambda Q_p^0(\pi_t - \bar{\pi}) - 2\lambda(1-\omega)\bar{\pi}Q_p^0 - \lambda(1-\theta)(1-\omega)\bar{\pi}Q_p^1\Xi_t \\ &\quad + 2\lambda Q_p^0(\pi_t - \bar{\pi}) + 2\lambda(1-\omega)\bar{\pi}Q_p^0 + \lambda(1-\theta)(1-\omega)\bar{\pi}Q_p^1\Xi_{t-1} - 2\lambda(Q_p^0)^2 \\ &\quad - \lambda(1-\theta)Q_p^0Q_p^1\Xi_{t-1} - \lambda(1-\theta)Q_p^0Q_p^1\Xi_t + \text{hot}. \end{aligned}$$

The non-stochastic steady state value is given by

$$\bar{\Delta} = (1-\lambda)\bar{b}^2 + 2(1-\lambda)\bar{b}Q_p^0 + (1-\lambda)(Q_p^0)^2 + \lambda\bar{\Delta} + \lambda(1-\omega)^2\bar{\pi}^2 = (\bar{b} + Q_p^0)^2 + \frac{\lambda(1-\omega)^2}{(1-\lambda)}\bar{\pi}^2 \quad (\text{A.15})$$

We can solve for deviations from steady state dispersion $\Xi_t = (\Delta_t - \bar{\Delta})$

$$\begin{aligned} \Xi_t &= (\pi_t - \bar{\pi}) \{ 2(1-\lambda)\bar{b}M + 2(1-\lambda)Q_p^0M + 2\lambda(1-\omega)\bar{\pi} \} \\ &\quad + (\pi_t - \bar{\pi})^2 \left\{ (1-\lambda)M^2 + \lambda + \left(\frac{1-\theta}{2} \right)^2 (Q_p^1)^2 K_Q^2 + (1-\lambda)(1-\theta)MQ_p^1K_Q - \lambda(1-\theta)Q_p^1K_Q \right\} \\ &\quad + \Xi_t \{ (1-\lambda)(1-\theta)\bar{b}Q_p^1 + (1-\lambda)(1-\theta)Q_p^0Q_p^1 - \lambda(1-\theta)(1-\omega)\bar{\pi}Q_p^1 \} \\ &\quad + \Xi_{t-1} \{ \lambda + \lambda(1-\theta)(1-\omega)\bar{\pi}Q_p^1 \} + \text{hot}. \end{aligned} \quad (\text{A.16})$$

Note that the form of (A.15) is the same as the form of our guess in (A.13). Let

$$\begin{aligned}\Gamma_0 &= \{1 + (\theta - 1)Q_p^1[(1 - \lambda)(\bar{b} + Q_p^0) - \lambda(1 - \omega)\bar{\pi}]\}^{-1}, \\ \Gamma_1 &= [1 - (\theta - 1)(1 - \omega)Q_p^1\bar{\pi}]\Gamma_0.\end{aligned}$$

Then matching coefficients with the guess and verify equation (A.13), we have

$\Gamma_2 = 2\{(1 - \lambda)M(\bar{b} + Q_p^0) + \lambda(1 - \omega)\bar{\pi}\}\Gamma_0$ which confirms the guess for the functional form of $(\Delta_t - \bar{\Delta})$ with $F_Q = \lambda\Gamma_1$. Hence

$$\begin{aligned}\Xi_t &= \Gamma_2(\pi_t - \bar{\pi}) + \lambda\Gamma_1\Xi_{t-1} + \\ &\quad + (\pi_t - \bar{\pi})^2\Gamma_0 \left\{ (1 - \lambda)M^2 + \lambda + \left(\frac{1-\theta}{2}\right)^2 (Q_p^1)^2\Gamma_2^2 + (1 - \lambda)(1 - \theta)MQ_p^1\Gamma_2 - \lambda(1 - \theta)Q_p^1\Gamma_2 \right\}\end{aligned}$$

which completes the proof. ■

This lemma shows that the cross-sectional price dispersion is a function of its past values as well as the deviation of inflation from its steady state level. In the vicinity of $\bar{\pi} = 0$, $\Gamma_0 \approx 1$, $\Gamma_1 \approx 1$, $\Gamma_2 \approx 0$ and thus cross-sectional price dispersion varies very little over time since it is only a function of the variance of inflation. This is the standard result for welfare calculations in a zero steady state inflation environment (see e.g. Proposition 6.3 in Woodford (2003)). However, $\partial\Gamma_0/\partial\bar{\pi} > 0$, $\partial\Gamma_1/\partial\bar{\pi} < 0$, $\partial\Gamma_2/\partial\bar{\pi} > 0$ locally at $\bar{\pi} \approx 0$. Hence, deviations of inflation from its steady state level have an increasingly strong effect on the cross-sectional price dispersion as $\bar{\pi}$ rises and, as a result, the dynamics of price dispersion can become first-order when $\bar{\pi}$ is sufficiently high. However, given our parameter values and for the levels of trend inflation that we consider, Γ_2 remains tiny so that price dispersion effectively remains of second order as in Woodford (2003) and thus we can use a linear approximation of the structural relationships in the economy and a second order approximation of consumer utility for welfare calculations.²⁰ Henceforth, we will treat Γ_2 as negligibly small.

Using the demand condition (7), we can link the cross-sectional dispersion of output to the cross-sectional dispersion of prices:

$$\tilde{y}_t(i) = \log Y_t(i) - \log \bar{Y}_t^F = \log Y_t - \log \bar{Y}_t^F - \theta\{\log P_t(i) - \log P_t\} \quad (\text{A.17})$$

and hence

$$Y_t \equiv \text{var}_i \hat{y}_t(i) = \text{var}_i \tilde{y}_t(i) = \theta^2 \text{var}_i(\log(P_t(i))) = \theta^2 \Delta_t. \quad (\text{A.18})$$

Let \bar{Y} be the cross-sectional dispersion of output in the non-stochastic steady state. The remaining piece in the second-order approximation of household's utility is $E_i \tilde{y}_t(i)$, which is the average deviation of output from flexible price level at the firm level. Using the insight of Lemma 3, we can relate $E_i \tilde{y}_t(i)$ to the deviation of output from its flexible-price level at the aggregate level.

Lemma 5. If the deviation of output from its flexible-price level at the aggregate level is defined as $\tilde{y}_t = \log(Y_t/\bar{Y}_t^F)$, then

$$E_i \tilde{y}_t(i) = \tilde{y}_t - Q_y^0 - \frac{\theta-1}{2\theta} Q_y^1 (Y_t - \bar{Y}) + h. o. t. \quad (\text{A.19})$$

where $Q_y^0 = \frac{\theta-1}{2\theta} \bar{Y} / [1 + \frac{1}{2}(\frac{1-\theta}{\theta})^2 \bar{Y}]^2$ and $Q_y^1 = [1 - \frac{1}{2}(\frac{1-\theta}{\theta})^2 \bar{Y}] / [1 + \frac{1}{2}(\frac{1-\theta}{\theta})^2 \bar{Y}]^3$.

Proof: Consider the term $E \hat{y}_t(i)$. From the Dixit-Stiglitz aggregator we have

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}} \Rightarrow Y_t/\bar{Y}_t^F = \left[\int_0^1 (Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}} \quad (\text{A.20})$$

Note that

$$\begin{aligned}\log(Y_t/\bar{Y}_t^F) &= \frac{\theta}{\theta-1} \log \left(\int_0^1 (Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} di \right) \\ &= \frac{\theta}{\theta-1} \left(\int_0^1 \frac{\theta-1}{\theta} \log(Y_t(i)/\bar{Y}_t^F) di \right) + \frac{1}{2} \frac{\theta}{\theta-1} \left[E \left((Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} \right) \right]^{-2} \text{var}_i \left((Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} \right) + hot \\ &= E_i \hat{y}_t(i) + \frac{1}{2} \frac{\theta}{\theta-1} \left[E \left((Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} \right) \right]^{-2} \text{var}_i \left((Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} \right) + hot\end{aligned} \quad (\text{A.21})$$

Using the delta method, we have

$$\begin{aligned}\text{var}_i \left((Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} \right) &= \text{var}_i \left(\exp \left\{ \frac{\theta-1}{\theta} \log(Y_t(i)/\bar{Y}_t^F) \right\} \right) \\ &\approx \left(\frac{\theta-1}{\theta} \right)^2 \left[\exp \left\{ \frac{\theta-1}{\theta} E \log(Y_t(i)/\bar{Y}_t^F) \right\} \right]^2 \text{var}_i(\log(Y_t(i)/\bar{Y}_t^F))\end{aligned} \quad (\text{A.22})$$

²⁰ In our baseline calibration the highest value for Γ_2 is 0.04 which is reached at 6% annual inflation.

Also observe that

$$\begin{aligned}
E(Y_t(i)/\bar{Y}_t^F)^{\frac{\theta-1}{\theta}} &= E \left[\exp \left\{ \frac{\theta-1}{\theta} \log(Y_t(i)/\bar{Y}_t^F) \right\} \right] \\
&\approx \exp \left\{ \frac{\theta-1}{\theta} \log E(Y_t(i)/\bar{Y}_t^F) \right\} + \frac{1}{2} \left(\frac{\theta-1}{\theta} \right)^2 \left[\exp \left\{ \left(\frac{\theta-1}{\theta} \right) E \log(Y_t(i)/\bar{Y}_t^F) \right\} \right] \text{var}_i(\log(Y_t(i)/\bar{Y}_t^F)) \\
&= \exp \left\{ \frac{\theta-1}{\theta} \log E(Y_t(i)/\bar{Y}_t^F) \right\} \left[1 + \frac{1}{2} \left(\frac{\theta-1}{\theta} \right)^2 \text{var}_i(\log(Y_t(i)/\bar{Y}_t^F)) \right]
\end{aligned} \tag{A.23}$$

Combine (A.22) and (A.23) to re-write (A.21) as follows

$$\begin{aligned}
\tilde{y}_t &= \log(Y_t/\bar{Y}_t^F) = E_i \tilde{y}_t(i) + \frac{1}{2} \frac{\theta}{\theta-1} \frac{\left(\frac{\theta-1}{\theta} \right)^2 \text{var}_i(\log(Y_t(i)/\bar{Y}_t^F))}{\left[1 + \frac{1}{2} \left(\frac{\theta-1}{\theta} \right)^2 \text{var}_i(\log(Y_t(i)/\bar{Y}_t^F)) \right]} + hot \\
&= E_i \tilde{y}_t(i) + Q_y^0 + \frac{1}{2} \frac{\theta-1}{\theta} Q_y^1 (Y_t - \bar{Y}) + hot
\end{aligned} \tag{A.24}$$

$$\text{when } Q_y^0 = \frac{\frac{1}{2} \frac{\theta-1}{\theta} \bar{Y}}{\left[1 + \frac{1}{2} \left(\frac{\theta-1}{\theta} \right)^2 \bar{Y} \right]^2} \text{ and } Q_y^1 = \frac{1 - \frac{1}{2} \left(\frac{\theta-1}{\theta} \right)^2 \bar{Y}^2}{\left[1 + \frac{1}{2} \left(\frac{\theta-1}{\theta} \right)^2 \bar{Y} \right]^3}. \blacksquare$$

Similar to the cross-sectional price dispersion, one can show that, since \bar{Y} is quadratic in $\bar{\pi}$, $Q_y^0 \approx 0$, $Q_y^1 \approx 1$ and $\partial Q_y^0 / \partial \bar{\pi} \approx 0$, $\partial Q_y^1 / \partial \bar{\pi} \approx 0$ when $\bar{\pi} \approx 0$.

The central result can be summarized with Proposition 1, whose proof given Lemmas 1-5 is as follows.

Proof of Proposition 1:

From Lemma 5, we have $E_i \tilde{y}_t(i) = \tilde{y}_t - Q_y^0 - \frac{1}{2} \frac{\theta-1}{\theta} Q_y^1 (Y_t - \bar{Y})$. Therefore, we can express the approximation from the disutility of labor supply given in Lemma 2 as follows:

$$\begin{aligned}
\int_0^1 \tilde{v}(Y_t(i)) di &\approx \bar{Y}_t^F \tilde{v}'_{\bar{Y}_t^F} \left\{ E_i \tilde{y}_t(i) + \frac{1}{2} (1 + \eta^{-1}) \int_0^1 \tilde{y}_t^2(i) di \right\} \\
&= \bar{Y}_t^F \tilde{v}'_{\bar{Y}_t^F} \left\{ E_i \tilde{y}_t(i) + \frac{1}{2} (1 + \eta^{-1}) \left(\text{var}_i \hat{y}_t(i) + [E_i \tilde{y}_t(i)]^2 \right) \right\} \\
&= \bar{Y}_t^F \tilde{v}'_{\bar{Y}_t^F} \left\{ \left(1 - (1 + \eta^{-1}) Q_y^0 \right) \tilde{y}_t + \frac{1}{2} \left[Q_y^1 \theta^{-1} - Q_y^1 + (1 + \eta^{-1}) \left(1 + \frac{\theta-1}{\theta} Q_y^0 Q_y^1 \right) \right] \text{var}_i \hat{y}_t(i) \right. \\
&\quad \left. + \frac{1}{2} (1 + \eta^{-1}) \left(\left(\frac{1}{2} \frac{\theta-1}{\theta} \right)^2 [Q_y^1 (\text{var}_i \hat{y}_t(i) - \bar{Y})]^2 - \frac{\theta-1}{\theta} \tilde{y}_t Q_y^1 [\text{var}_i \hat{y}_t(i) - \bar{Y}] \right) + \frac{1}{2} (1 + \eta^{-1}) \tilde{y}_t^2 \right. \\
&\quad \left. + (1 + \eta^{-1}) (Q_y^0)^2 - Q_y^0 + \frac{\theta-1}{2\theta} Q_y^1 \left[1 - (1 + \eta^{-1}) Q_y^0 \right] \bar{Y} \right\} + hot + tip
\end{aligned} \tag{A.25}$$

As Woodford (2003) shows, $\tilde{v}'_{\bar{Y}_t^F} = \tilde{v}'_{\bar{C}_t^F} (1 - \Phi)(1 - \beta h)$, where (with log utility) $\Phi = -\log \left(\frac{\theta-1}{\theta} \right)$ which follows from the intra-temporal condition for labor supply. As in Lemma 1, $\bar{C}_t^F u'_{\bar{C}_t^F} = 1/(1 - h)$. We also have $\bar{C}_t^F = (1 - \bar{g}_y) \bar{Y}_t^F$. Hence,

$$\begin{aligned}
\left(\frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} \right)^{-1} \int_0^1 \tilde{v}(Y_t(i)) di &= (1 - (1 + \eta^{-1}) Q_y^0) \tilde{y}_t \\
&\quad + \frac{1}{2} \left[Q_y^1 \theta^{-1} - Q_y^1 + (1 + \eta^{-1}) \left(1 + \frac{\theta-1}{\theta} Q_y^0 Q_y^1 \right) \right] \text{var}_i \hat{y}_t(i) \\
&\quad + \frac{1}{2} (1 + \eta^{-1}) \left(\left(\frac{1}{2} \frac{\theta-1}{\theta} \right)^2 [Q_y^1 (\text{var}_i \hat{y}_t(i) - \bar{Y})]^2 - \frac{\theta-1}{\theta} \tilde{y}_t Q_y^1 [\text{var}_i \hat{y}_t(i) - \bar{Y}] \right) \\
&\quad + \frac{1}{2} (1 + \eta^{-1}) \tilde{y}_t^2 + (1 + \eta^{-1}) (Q_y^0)^2 - Q_y^0 + \frac{\theta-1}{2\theta} Q_y^1 \left[1 - (1 + \eta^{-1}) Q_y^0 \right] \bar{Y} + hot + tip
\end{aligned} \tag{A.26}$$

Using equation (A.18) from the derivations above and Lemma 4, we replace the cross-section dispersion of output in (A.26) with the cross-sectional dispersion of prices.

Given the low levels of inflation we consider in our analysis, Γ_2 is close to zero and $\Gamma_2(\pi_t - \bar{\pi})$ has (in expectation) negligible effects on Ξ_t so that

$$E(\Xi_t) = \Gamma_3 E(\hat{\pi}_t^2) \tag{A.27}$$

where

$$\Gamma_3 = \Gamma_0 (1 - \lambda \Gamma_1)^{-1} \{ (1 - \lambda) M^2 + \lambda \}. \tag{A.28}$$

Note that from the definition of Ξ_t , we have

$$E(\Delta_t) = E(\Xi_t) + \bar{\Delta}. \tag{A.29}$$

Since all variables are stationary and $\hat{y}_t = \tilde{y}_t - \log(\bar{X})$ and using (A.27) and (A.29), we can compute the expected per period utility as follows

$$\begin{aligned}
EU_t &\approx E \left(\frac{1}{1-h} (\tilde{c}_t - h\tilde{c}_{t-1}) - \frac{1}{2} \frac{h}{(1-h)^2} (\tilde{c}_t^2 - 2\tilde{c}_t\tilde{c}_{t-1} + \tilde{c}_{t-1}^2) \right) \\
&\quad - \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} (1 - (1+\eta^{-1})Q_y^0) E\tilde{y}_t - \frac{1}{2} \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} (1+\eta^{-1}) E\tilde{y}_t^2 \\
&\quad - \frac{1}{2} \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} [Q_y^1 \theta^{-1} - Q_y^1 + (1+\eta^{-1})(1 + \frac{\theta-1}{\theta} Q_y^0 Q_y^1)] \theta^2 E\Delta_t \\
&\quad + E \frac{1}{2} \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} (1+\eta^{-1}) \frac{\theta-1}{\theta} Q_y^1 E(\tilde{y}_t \Xi_t) \\
&\quad - \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} \left[(1+\eta^{-1})(Q_y^0)^2 - Q_y^0 + \frac{\theta-1}{2\theta} Q_y^1 [1 - (1+\eta^{-1})Q_y^0] \bar{Y} \right] \\
&= \left[1 - \frac{(1-\Phi)(1-\beta h)}{(1-\bar{g}_y)(1-h)} (1 - (1+\eta^{-1})Q_y^0) \right] \log \bar{X} - \left[\frac{(1-\beta h)(1-\Phi)(1+\eta^{-1})}{(1-h)2(1-\bar{g}_y)} \right] [\log \bar{X}]^2 \\
&\quad - \frac{(1-\beta h)(1-\Phi)}{(1-h)(1-\bar{g}_y)} \left\{ (1+\eta^{-1})[Q_y^0]^2 - Q_y^0 \right\} \\
&\quad - \frac{(1-\beta h)\theta^2(1-\Phi)}{(1-h)2(1-\bar{g}_y)} (1+\eta^{-1}) \bar{\Delta} \\
&\quad - \frac{1}{2} \frac{(1+\eta^{-1})(1-\beta h)}{(1-\bar{g}_y)(1-h)} \text{var}(\hat{y}_t) \\
&\quad - \frac{h(1-\rho_c)}{(1-h)^2} \text{var}(\hat{c}_t) \\
&\quad - \frac{(1-\beta h)}{(1-h)} \left[\frac{\theta^2}{2(1-\bar{g}_y)} \Gamma_3 \left\{ \left[Q_y^1 \left(\frac{1}{\theta} - 1 \right) + \left(1 + \frac{1}{\eta} \right) \left(1 + \frac{\theta-1}{\theta} Q_y^0 Q_y^1 \right) \right] - \left(1 + \frac{1}{\eta} \right) \frac{\theta-1}{\theta} Q_y^1 \log \bar{X} \right\} \right] \text{var}(\hat{\pi}_t). \blacksquare
\end{aligned}$$

Appendix B: Model with Capital

In this modification of model, we introduce homogenous capital owned by households. The model is otherwise identical to the baseline model and thus we focus only on elements which are novel. The budget constraint each period is now given by

$$\xi_t: C_t + \frac{S_t}{P_t} + I_t + \frac{\psi}{2} \left(\frac{I_t}{K_{t-1}} - \delta - \overline{GY} + 1 \right)^2 K_{t-1} \leq \int_0^1 \left(\frac{N_t(i)W_t(i)}{P_t} \right) di + \frac{S_{t-1}q_{t-1}R_{t-1}}{P_t} + \frac{K_{t-1}R_{t-1}^k}{P_t} + T_t \quad (\text{B.1})$$

Where, in addition to standard variables, K is capital owned by the representative consumer and R^k is the nominal rental rate on capital. The fourth term on the left hand side is a quadratic adjustment cost to the stock of capital held by the consumer, where ψ measures the strength of the adjustment cost. δ is the rate of depreciation on capital. Capital can be accumulated according to the law of motion

$$\mu_t: K_t = (1 - \delta)K_{t-1} + I_t \quad (\text{B.2})$$

The additional first order conditions from this utility-maximization problem are:

$$q_t = 1 + \psi \left(\frac{I_t}{K_{t-1}} - \delta - \overline{GY} + 1 \right) \quad (\text{B.3})$$

$$0 = E_t \beta \lambda_{t+1} \left[\tilde{R}_t^k - \frac{\psi}{2} \left(\frac{I_{t+1}}{K_t} - \delta - \overline{GY} + 1 \right)^2 + \psi \frac{I_{t+1}}{K_t} \left(\frac{I_{t+1}}{K_t} - \delta - \overline{GY} + 1 \right) \right] + (1 - \delta) E_t \beta \mu_{t+1} - \mu_t \quad (\text{B.4})$$

where λ_t/μ_t is Tobin's q and $\tilde{R}_t^k = R_t^k/P_t$ is the real rental rate of capital.

The production of each intermediate good is done by a monopolist facing a Cobb-Douglas production function with capital share α

$$Y_t(i) = K_{t-1}(i)^\alpha (A_t N_t(i))^{1-\alpha} \quad (\text{B.5})$$

We assume that capital is perfectly mobile across firms

$$K_t = \int_0^1 K_t(i) di. \quad (\text{B.6})$$

Each intermediate good producer has sticky prices, modeled as in Calvo (1983) where $1-\lambda$ is the probability that each firm will be able to reoptimize its price each period. We allow for indexation of prices to steady-state inflation by firms who do not reoptimize their prices each period, with ω representing the degree of indexation (0 for no indexation to 1 for full indexation). Denoting the optimal reset price of firm i by $B(i)$, re-optimizing firms solve the following profit-maximization problem

$$\max E_t \sum_{j=0}^{\infty} \lambda^j Q_{t,t+j} [Y_{t+j}(i) B_t(i) \bar{\Pi}^{j\omega} - W_{t+j}(i) N_{t+j}(i) - R_{t+j}^k K_{t+j-1}(i)] \quad (\text{B.7})$$

where Q is the stochastic discount factor and $\bar{\Pi}$ is the gross steady-state level of inflation. The optimal relative reset price is given by equation (11) with firm-specific marginal costs

$$\frac{MC_{t+j}(i)}{P_{t+j}} = \left(\frac{z_{t+j}^{-1}}{A_{t+j}} \right)^{\frac{(1-\alpha)\eta}{\eta+\alpha}} \left(\frac{Y_{t+j}}{A_{t+j}} \right)^{\frac{1-\alpha}{\eta+\alpha}} \left(\frac{R_{t+j}^k}{P_{t+j}} \right)^{\frac{\alpha(\eta+1)}{\eta+\alpha}} \left(\frac{B_t(i)}{P_t} \right)^{-\frac{\theta(1-\alpha)}{\eta+\alpha}} \left(\frac{P_{t+j}}{\bar{\Pi}^{\omega} P_t} \right)^{\frac{\theta(1-\alpha)}{\eta+\alpha}} [\alpha^\alpha (1-\alpha)^{1-\alpha}]^{\frac{\eta+1}{\eta+\alpha}} \quad (\text{B.8})$$

Firm-specific marginal cost remain a function of aggregate variables and the reset price only because full capital mobility implies that all firms have the same (adjusted) capital labor ratio. The first order condition for the firms' input usage is $\frac{\alpha}{1-\alpha} (N_t(i)/K_{t-1}(i)) = R_t^k/W_t(i)$ which we combine with the first order condition for labor supply to get $\frac{\alpha}{1-\alpha} (N_t(i)^{1+1/\eta}/K_{t-1}(i)) = \xi_t R_t^k/P_t(i)$. Since the right hand side consists of aggregate variables only, the left hand side must be identical across firms.

The goods market clearing condition for the economy is $Y_t = C_t + G_t + I_t$. We define the aggregate labor input as $N_t = [\int_0^1 N_t(i)^{1+1/\eta} di]^{\eta/(1+\eta)}$. Note that this definition of labor input differs from the labor-only model. We make different assumptions to simplify our aggregation problems.

Log-linearization

In the model with capital the goods market clearing condition becomes $\hat{y}_t = \bar{c}_y \hat{c}_t + \bar{g}_y \hat{g}_t + \bar{l}_y \hat{l}_t$ where $\bar{c}_y, \bar{g}_y, \bar{l}_y$ are the steady-state ratios of consumption, government and investment to output respectively.

Lemma 6. Let $\bar{c}_y, \bar{g}_y, \bar{l}_y$ be the steady-state ratios of consumption, government and investment to output respectively. Then $\bar{c}_y = (1 - \bar{g}_y) - \frac{\alpha\beta(\delta + \overline{GY} - 1)}{\overline{GY} - \beta(1 - \delta)}$, and $\bar{l}_y = \frac{\alpha\beta(\delta + \overline{GY} - 1)}{\overline{GY} - \beta(1 - \delta)}$.

Proof: See Appendix C.

Corollary 1. Since $\bar{c}_y, \bar{g}_y, \bar{l}_y$ are independent of the degree of price stickiness λ , they are equal to their flexible price level counterparts, $\bar{c}_y = \bar{c}_y^F, \bar{g}_y = \bar{g}_y^F, \bar{l}_y = \bar{l}_y^F$.

Also, integrating over firm-specific production functions and log-linearizing yields the aggregate production function:

Lemma 7. Let lower case letters denote the deviations from steady state values. Then the aggregate production function up to a first order approximation is given by $\hat{y}_t = \alpha \hat{k}_{t-1} + (1 - \alpha) \hat{n}_t$.

Proof: See Appendix C.

Allowing for capital also changes the steady state effects of positive trend inflation. For example, the steady-state level of the output gap (which is defined as the deviation of output from its flexible price

level counterpart) is given by $\bar{X}^{\frac{(1-\alpha)(\eta+1)}{\eta+\alpha}} = \frac{1-\lambda\beta\bar{\Pi}^{\frac{(1-\omega)\theta(\eta+1)}{\eta+\alpha}}}{1-\lambda\beta\bar{\Pi}^{(1-\omega)\theta}} \left(\frac{1-\lambda}{1-\lambda\bar{\Pi}^{(1-\omega)(\theta-1)}} \right)^{\frac{(1-\alpha)\theta+\eta+\alpha}{(\eta+\alpha)(\theta-1)}}$. Note that the steady-state level of the gap is again equal to one when steady-state inflation is zero (i.e., $\bar{\Pi} = 1$) or when the degree of price indexation is exactly equal to one.

Lemma 8. The expansion of the utility part that corresponds to disutility from labor supply in the capital model is given by

$$\int_0^1 v(N_t(i)) di = \bar{N}_t^F v'_{\bar{N}_t^F} \left\{ (1 - (1 + \eta^{-1})Q_n^0) \hat{n}_t + \frac{1}{2}(1 + \eta^{-1}) [1 - Q_n^1 + (1 + \eta^{-1})Q_n^0 Q_n^1] \text{var}_i \hat{n}_t(i) \right. \\ \left. + \frac{1}{8}(1 + \eta^{-1})^3 [Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y})^2 - \frac{1}{2}(1 + \eta^{-1})^2 Q_n^1 \hat{n}_t [\text{var}_i \hat{n}_t(i) - \bar{Y}] + \frac{1}{2}(1 + \eta^{-1}) \hat{n}_t^2 \right. \\ \left. + (1 + \eta^{-1})(Q_n^0)^2 - Q_n^0 + \frac{1}{2}(1 + \eta^{-1})Q_n^1 [1 - (1 + \eta^{-1})Q_n^0] \bar{Y} \right\} + \text{hot} + \text{tip},$$

where $Q_n^0 = \frac{1}{2}(1 + \eta^{-1})\bar{Y} / [1 + \frac{1}{2}(1 + \eta^{-1})^2\bar{Y}]^2$ and $Q_n^1 = \{1 - \frac{1}{2}(1 + \eta^{-1})^2\bar{Y}\} / [1 + \frac{1}{2}(1 + \eta^{-1})^2\bar{Y}]^3$ and \bar{Y} is the steady state dispersion of labor supply.

Proof: See Appendix C.

Lemma 9. The cross sectional variations in labor input is $\text{var}_i \hat{n}_t(i) = (\frac{\theta}{1+\alpha/\eta})^2 \text{var}_i \log p_t(i)$.

Proof: See Appendix C.

This implies that the steady state dispersion of labor supply is given by $\bar{Y} = (\frac{\theta}{1+\alpha/\eta})^2 \bar{\Delta}$

Before we proceed with Proposition 2 we write the inter-temporal condition for labor supply as $(\bar{N}_t^F v'_{\bar{N}_t^F} / u'_{\bar{c}_t^F}) = \bar{N}_t^F (W_t / P_t)$. The right hand side of this expression is simply total labor income.

Because of the distortion due to monopolistic competition it is equal to $(\bar{N}_t^F v'_{\bar{N}_t^F} / u'_{\bar{c}_t^F}) = (1 - \beta h)(1 - \Phi)(1 - \alpha) \bar{Y}_t^F$ where $\Phi = -\log(\frac{\theta-1}{\theta})$. Since we assume log utility this expression further simplifies to $\bar{N}_t^F v'_{\bar{N}_t^F} = (1 - \beta h) / (1 - h)(1 - \Phi)(1 - \alpha) (\bar{c}_y^F)^{-1}$ where \bar{c}_y^F is the share of consumption in output at the flexible price steady state.

Proposition 2. In the model for capital, the second order approximation to utility in equation (1) is $\Omega_0 + \Omega_1 \text{var}(\hat{n}_t) + \Omega_2 \text{var}(\hat{\pi}_t) + \Omega_3 \text{var}(\hat{c}_t)$ where parameters $\Omega_i, i = 1, 2, 3$ depend on the steady state inflation $\bar{\pi}$ and are given by

$$\Omega_0 = \log \bar{X} - \frac{1(1-\beta h)}{2(1-h)} (1 - \Phi)(1 - \alpha) (\bar{c}_y^F)^{-1} [2(1 + \eta^{-1})(Q_n^0)^2 - 2Q_n^0 + (1 + \eta^{-1})Q_n^1 (1 - (1 + \eta^{-1})Q_n^0) \bar{Y}], \\ \Omega_1 = -\frac{1(1-\beta h)}{2(1-h)} (1 - \alpha) (\bar{c}_y^F)^{-1} (1 + \eta^{-1}), \\ \Omega_2 = -\frac{1(1-\beta h)}{2(1-h)} (1 - \alpha) (\bar{c}_y^F)^{-1} (1 + \eta^{-1}) [1 - Q_n^1 + (1 + \eta^{-1})Q_n^1 Q_n^0] (\frac{\theta}{1+\alpha/\eta})^2 \Gamma_3 \\ \Omega_3 = \Theta_3.$$

and Γ_3, Θ_3 are defined in Proposition 1.

Proof: See Appendix C.

Appendix C: Proofs for model with capital

Lemma 6

From the first order condition for the use of capital we know that capital for a firm that reset its price at time t evolves as follows

$$k_{t+j-1}(i) = GY_t^j \bar{\Pi}^{j(1-\omega)(\theta-1)} \left\{ Y_t \alpha (\bar{R}_t^k)^{-1} (\bar{B}/P)^{1-\theta} \right\} = Y_t^j \bar{\Pi}^{j(1-\omega)(\theta-1)} k_{t-1}^R \quad (C.1)$$

where k_t^R is the optimal reset stock of capital at time t . Due to Calvo pricing the aggregate capital stock today is an aggregate of an infinite sum of past capital stocks, adjusted for their growth rate in the interim

$$\begin{aligned} K_{t-1} &= \int_0^1 k_{t-1}(i) di = \sum_{j=0}^{\infty} (1-\lambda) \lambda^j k_{t-j-1}^R GY_t^j \bar{\Pi}^{j(1-\omega)(\theta-1)} \\ &= \sum_{j=0}^{\infty} (1-\lambda) \lambda^j k_{t-j-1}^R Y_{t-j} \alpha (\bar{R}_t^k)^{-1} (\bar{B}/P)^{1-\theta} GY_t^j \bar{\Pi}^{j(1-\omega)(\theta-1)} \\ &= Y_t \alpha (\bar{R}_t^k)^{-1} \frac{1-\lambda}{1-\lambda \bar{\Pi}^{(1-\omega)(\theta-1)}} (\bar{B}/P)^{1-\theta} = Y_t \frac{\alpha \beta}{\bar{G}\bar{Y} - \beta(1-\delta)} = Y_t \bar{k}_y \end{aligned} \quad (C.2)$$

Thus, the ratio of capital to output \bar{k}_y . From the capital accumulation equation we know that in the steady state the share of output going to investment is equal to

$$\bar{y} = \frac{\alpha \beta (\delta + \bar{G}\bar{Y} - 1)}{\bar{G}\bar{Y} - \beta(1-\delta)}, \quad (C.3)$$

The conclusion for \bar{c}_y then follows directly from the resource constraint. ■

Lemma 7

We first rewrite the individual production function as

$$y_t(i) = A_t^{1-\alpha} n_t(i)^{1+\alpha/\eta} \left\{ k_{t-1}(i)/n_t(i)^{1+1/\eta} \right\}^\alpha \quad (C.4)$$

From the first order condition for labor supply and the optimal capital labor ratio, we know that the term in parentheses is the same across firms. Integrating this equation then yields,

$$Y_t = A_t^{1-\alpha} \left[\int_0^1 n_t(i)^{\frac{(1+\alpha\eta^{-1})(\theta-1)}{\theta}} di \right]^{\frac{\theta}{\theta-1}} \left(\frac{K_{t-1}}{N_t^{1+\eta^{-1}}} \right)^\alpha \quad (C.5)$$

We then take logs of this expression and consider a first order approximation to the integral,

$$\begin{aligned} \frac{\theta}{\theta-1} \ln \left[\int_0^1 n_t(i)^{\frac{(1+\alpha\eta^{-1})(\theta-1)}{\theta}} di \right] &\approx \\ &\approx \frac{\eta+\alpha}{\eta+1} \left[\int_0^1 \ln n_t(i)^{(1+\eta^{-1})} di \right] + \frac{1}{2} \frac{\theta}{\theta-1} \frac{[\frac{\eta+\alpha}{\eta} \frac{\theta-1}{\theta}]^2 \text{var}_i(\ln n_t(i))}{\{1 + \frac{1}{2} [\frac{\eta+\alpha}{\eta} \frac{\theta-1}{\theta}]^2 \text{var}_i(\ln n_t(i))\}^2} \\ &\approx \ln N_t^{1+\alpha\eta^{-1}} + \underbrace{\frac{1}{2} \frac{\theta}{\theta-1} \frac{[\frac{\eta+\alpha}{\eta} \frac{\theta-1}{\theta}]^2 \bar{Y}}{\{1 + \frac{1}{2} [\frac{\eta+\alpha}{\eta} \frac{\theta-1}{\theta}]^2 \bar{Y}\}^2}}_{Z_n^0} + \underbrace{\frac{1}{2} \frac{\theta}{\theta-1} \frac{[\frac{\eta+\alpha}{\eta} \frac{\theta-1}{\theta}]^2 \{1 - \frac{1}{2} [\frac{\eta+\alpha}{\eta} \frac{\theta-1}{\theta}]^2 \bar{Y}\}}{\{1 + \frac{1}{2} [\frac{\eta+\alpha}{\eta} \frac{\theta-1}{\theta}]^2 \bar{Y}\}^3}}_{Z_n^1} (Y_t - \bar{Y}) \\ &= \ln N_t^{1+\alpha\eta^{-1}} + Z_n^0 + Z_n^1 \frac{\theta^2}{(1+\alpha\eta^{-1})^2} \underbrace{(\bar{\Delta}_t - \bar{\Delta})}_{hot} + hot \\ &= \ln N_t^{1+\alpha\eta^{-1}} + Z_n^0 + hot \end{aligned} \quad (C.6)$$

where the last line follows from price dispersion being of second order (Lemma 4, Appendix A). ■

Lemma 8

Note that

$$v(n_t(i)) \approx \bar{N}_t^F v'_{\bar{N}_t^F} \left\{ \hat{n}_t(i) + \frac{1}{2} (1 + \eta^{-1}) \hat{n}_t(i)^2 \right\} + tip \quad (C.7)$$

where $\hat{n}_t(i) = \log(n_t(i)/\bar{N}_t^F)$ is the deviation of firm i 's labor input from flexible price level of labor input \bar{N}_t^F and η is the Frisch elasticity of labor supply. Because of our definition of labor input in the capital model the flexible price steady state labor input and the Calvo-pricing steady state labor input are the same, $\bar{N}_t = \bar{N}_t^F \Rightarrow \hat{n}_t = \tilde{n}_t$. Compute the average across firms

$$\int_0^1 v(n_t(i)) di = \bar{N}_t^F v'_{\bar{N}_t^F} \left\{ E_i \hat{n}_t(i) + \frac{1}{2} (1 + \eta^{-1}) \int_0^1 \hat{n}_t(i)^2 di \right\} + hot + tip \quad (C.8)$$

Consider the first term $E_i \hat{n}_t(i)$. We define aggregate labor input as

$$N_t = \left[\int_0^1 n_t(i)^{1+\eta^{-1}} di \right]^{1/(1+\eta^{-1})} \Rightarrow \frac{N_t}{\bar{N}_t^F} = \left[\int_0^1 \left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{1+\eta^{-1}} di \right]^{1/(1+\eta^{-1})} \quad (\text{C.9})$$

Note that

$$\begin{aligned} \log \left(\frac{N_t}{\bar{N}_t^F} \right) &= \frac{\eta}{\eta+1} \log \left(\int_0^1 \left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{1+\eta^{-1}} di \right) \\ &= \frac{\eta}{\eta+1} \int_0^1 \frac{\eta+1}{\eta} \log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) di + \frac{1}{2} \frac{\eta}{\eta+1} \left[E \left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{\frac{1+\eta}{\eta}} \right]^{-2} \text{var}_i \left(\left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{\frac{1+\eta}{\eta}} \right) + \text{hot} \\ &= E_i \hat{n}_t(i) + \frac{1}{2} \frac{\eta}{\eta+1} \left[E \left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{\frac{1+\eta}{\eta}} \right]^{-2} \text{var}_i \left(\left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{\frac{1+\eta}{\eta}} \right) + \text{hot} \end{aligned} \quad (\text{C.10})$$

Using the delta method, we have

$$\text{var}_i \left(\left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{\frac{\eta+1}{\eta}} \right) = \text{var}_i \left(\exp \left\{ \frac{\eta+1}{\eta} \log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right\} \right) \approx \left(\frac{\eta+1}{\eta} \right)^2 \left[\exp \left\{ \frac{\eta+1}{\eta} E \log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right\} \right]^2 \text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right) \quad (\text{C.11})$$

Also observe that

$$\begin{aligned} E \left(\left(\frac{n_t(i)}{\bar{N}_t^F} \right)^{\frac{\eta+1}{\eta}} \right) &= E \left[\exp \left\{ \frac{\eta+1}{\eta} \log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right\} \right] \approx \\ &\approx \exp \left\{ \frac{\eta+1}{\eta} E \log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right\} + \frac{1}{2} \left(\frac{\eta+1}{\eta} \right)^2 \exp \left\{ \frac{\eta+1}{\eta} E \log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right\} \text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right) \\ &= \exp \left\{ \frac{\eta+1}{\eta} E \log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right\} \left[1 + \frac{1}{2} \left(\frac{\eta+1}{\eta} \right)^2 \text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right) \right] \end{aligned} \quad (\text{C.12})$$

Combine (C.11) and (C.12) to re-write (C.10) as follows

$$\log N_t / \bar{N}_t^F = E_i \hat{n}_t(i) + \frac{1}{2} \frac{\eta}{\eta+1} \frac{\left(\frac{\eta+1}{\eta} \right)^2 \text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right)}{\left[1 + \frac{1}{2} \left(\frac{\eta+1}{\eta} \right)^2 \text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right) \right]^2} + \text{hot} \quad (\text{C.13})$$

We use Taylor series expansion of the second term $\log \left(\frac{N_t}{\bar{N}_t^F} \right) = E_i \hat{n}_t(i) + Q_n^0 + \frac{1}{2} \frac{\eta+1}{\eta} Q_n^1 (Y_t - \bar{Y}) + \text{hot}$

where $Q_n^1 = \left[1 - \frac{1}{2} \left(\frac{\eta+1}{\eta} \right)^2 \bar{Y} \right] / \left[1 + \frac{1}{2} \left(\frac{\eta+1}{\eta} \right)^2 \bar{Y} \right]^3 \leq 1$ and $Q_n^0 = \frac{1}{2} \frac{\eta+1}{\eta} \bar{Y} / \left[1 + \frac{1}{2} \left(\frac{\eta+1}{\eta} \right)^2 \bar{Y} \right]^2$, $Y_t = \text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right)$, and $\bar{Y} = E Y_t$. We can consider a modification where

$$\log \left(\frac{N_t}{\bar{N}_t^F} \right) = E_i \hat{n}_t(i) + Q_n^0 + \frac{1}{2} \frac{\eta+1}{\eta} Q_n^1 (Y_t - \bar{Y}) + \text{hot} \quad (\text{C.14})$$

It then follows from (C.14),

$$\hat{n}_t = \log \left(\frac{N_t}{\bar{N}_t^F} \right) = \log \left(\frac{N_t}{\bar{N}_t^F} \right) = E_i \hat{n}_t(i) + Q_n^0 + \frac{1}{2} \frac{\eta+1}{\eta} Q_n^1 \left[\text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right) - \bar{Y} \right] + \text{hot} \quad (\text{C.15})$$

which implies that the first order term in (C.18) is equal to

$$\begin{aligned} E_i \hat{n}_t(i) &= \hat{n}_t - Q_n^0 - \frac{1}{2} \frac{\eta+1}{\eta} Q_n^1 \left[\text{var}_i \left(\log \left(\frac{n_t(i)}{\bar{N}_t^F} \right) \right) - \bar{Y} \right] + \text{hot} \\ &= \hat{n}_t - Q_n^0 - \frac{1}{2} \frac{\eta+1}{\eta} Q_n^1 \left[\text{var}_i \left(\hat{n}_t(i) \right) - \bar{Y} \right] + \text{hot} \end{aligned} \quad (\text{C.16})$$

Now consider the second term in (C.8):

$$\int_0^1 \hat{n}_t(i)^2 di = \int_0^1 \left[\hat{n}_t(i) - E_i(\hat{n}_t(i)) + E_i(\hat{n}_t(i)) \right]^2 di = \text{var}_i \hat{n}_t(i) + \left[E_i(\hat{n}_t(i)) \right]^2 \quad (\text{C.17})$$

Using (C.16), we can re-write (C.17) as follows

$$\begin{aligned} \int_0^1 \hat{n}_t(i)^2 di &= \text{var}_i \hat{n}_t(i) + \left[\hat{n}_t - Q_n^0 - \frac{1}{2} \frac{\eta+1}{\eta} Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y}) \right]^2 \\ &= \text{var}_i \hat{n}_t(i) + \hat{n}_t^2 + \frac{1}{4} \left(\frac{\eta+1}{\eta} \right)^2 \left[Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y}) \right]^2 - \frac{\eta+1}{\eta} \hat{n}_t Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y}) + \\ &\quad + \left[Q_n^0 \right]^2 - 2 Q_n^0 \hat{n}_t + \frac{\eta+1}{\eta} Q_n^0 Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y}) \end{aligned} \quad (\text{C.18})$$

Now we combine (C.18) and (C.16) to finally get the part which corresponds to the second-order approximation of utility due to the disutility of labor supply (C.8):

$$\int_0^1 \nu(n_t(i)) di = \bar{N}_t^F \nu'_{\bar{N}_t^F} \left\{ E_i \hat{n}_t(i) + \frac{1}{2} (1 + \eta^{-1}) \int_0^1 \hat{n}_t(i)^2 di \right\} + \text{hot}$$

$$\begin{aligned}
&= \bar{N}_t^F \mathbf{v}'_{\bar{N}_t^F} \left\{ \hat{n}_t - Q_n^0 - \frac{1}{2} \frac{\eta+1}{\eta} Q_n^1 [\text{var}_i \hat{n}_t(i) - \bar{Y}] \right. \\
&\quad + \frac{1}{2} (1 + \eta^{-1}) \left([1 + \frac{\eta+1}{\eta} Q_n^0 Q_n^1] \text{var}_i \hat{n}_t(i) + \hat{n}_t^2 - 2Q_n^0 \hat{n}_t \right) + \frac{1}{8} (1 + \eta^{-1}) \left(\frac{\eta+1}{\eta} \right)^2 [Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y})]^2 \\
&\quad \left. - \frac{1}{2} (1 + \eta^{-1}) \frac{\eta+1}{\eta} \hat{n}_t Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y}) + (1 + \eta^{-1}) [Q_n^0]^2 - (1 + \eta^{-1}) \frac{\eta+1}{\eta} Q_n^0 Q_n^1 \bar{Y} \right\} + \text{hot} + \text{tip} \\
&= \bar{N}_t^F \mathbf{v}'_{\bar{N}_t^F} \left\{ \left(1 - \left(\frac{\eta+1}{\eta} \right) Q_n^0 \right) \hat{n}_t + \frac{1}{2} \left[\left(\frac{\eta+1}{\eta} \right) (1 - Q_n^1 + \frac{\eta+1}{\eta} Q_n^0 Q_n^1) \right] \text{var}_i \hat{n}_t(i) + \frac{1}{8} \left(\frac{\eta+1}{\eta} \right)^3 [Q_n^1 (\text{var}_i \hat{n}_t(i) - \bar{Y})]^2 \right. \\
&\quad \left. + \frac{1}{2} \left(\frac{\eta+1}{\eta} \right) \hat{n}_t^2 - Q_n^0 + \left(\frac{\eta+1}{\eta} \right) [Q_n^0]^2 + \frac{\eta+1}{\eta} Q_n^1 \left[\frac{1}{2} - \left(\frac{\eta+1}{\eta} \right) Q_n^0 \right] \bar{Y} \right\} + \text{hot} + \text{tip}. \blacksquare
\end{aligned} \tag{C.19}$$

Lemma 9

We first note that $\text{var}_i \hat{n}_t(i) = \text{var}_i (n_t(i) - \ln N_t^F) = \text{var}_i (n_t(i))$. To replace variability of log labor with variability of log prices we will use the market firms' first order conditions for inputs and combine it with the agents first order condition for labor supply,

$$\frac{1-\alpha}{\alpha} \frac{k_t(i)}{n_t(i)} = \frac{W_t(i)}{R_t^k} = \frac{P_t \xi_t^{-1} n_t(i)^{\eta-1}}{R_t^k}.$$

This implies that the idiosyncrasy in the capital-labor ratio is proportional to log labor input across firms,

$$\ln \{k_t(i)/n_t(i)\} = \eta^{-1} \ln n_t(i) + \text{aggregate terms}.$$

This then allows us to calculate the cross sectional variation in output,

$$\text{var}_i \ln y_t(i) = \text{var}_i (\ln n_t(i)^{1+\alpha\eta^{-1}} + \alpha \ln \{k_t(i)/n_t(i)\} + \ln A_t) = (1 + \alpha\eta^{-1})^2 \text{var}_i \ln n_t(i)$$

Finally, using the demand curve for each variety, we have

$$\text{var}_i \ln y_t(i) = \text{var}_i (\ln Y_t - \theta (\ln p_t(i) - \ln P_t)) = \theta^2 \text{var}_i \ln p_t(i) \tag{C.20}$$

and it follows that the cross-sectional variation in labor input is proportional to the cross-sectional variation in prices.

$$\text{var}_i \hat{n}_t(i) = \left(\frac{\theta}{1+\alpha\eta^{-1}} \right)^2 \text{var}_i \ln p_t(i) \tag{C.21}$$

which completes the proof. \blacksquare

Proposition 2.

Combine **Lemmas 8-9** and **Corollary 1** with **Lemma 1** and **Lemma 4** to get

$$\begin{aligned}
EU_t &= \log \bar{X} - \frac{h(1-\rho_c)}{(1-h)^2} \text{var}(\hat{c}_t) - \frac{1}{2} \frac{(1-\beta h)}{1-h} (1-\alpha) (c_y^F)^{-1} \left(\frac{\eta+1}{\eta} \right) E(\hat{n}_t^2) \\
&\quad - \frac{1}{2} \frac{(1-\beta h)}{1-h} (1-\alpha) (c_y^F)^{-1} \left[\left(\frac{\eta+1}{\eta} \right) \left(1 - Q_n^1 + \left(\frac{\eta+1}{\eta} \right) Q_n^0 Q_n^1 \right) \frac{\theta^2}{(1+\alpha\eta^{-1})^2} \Gamma_3 \text{var}(\hat{n}_t) \right. \\
&\quad \left. - \frac{1}{2} \frac{(1-\beta h)}{1-h} (1-\alpha) (c_y^F)^{-1} (1-\Phi) \left[2 \frac{\eta+1}{\eta} (Q_n^0)^2 - 2Q_n^0 + Q_n^1 \frac{\eta+1}{\eta} \left[1 - \frac{\eta+1}{\eta} Q_n^0 \right] \frac{\theta^2}{(1+\alpha\eta^{-1})^2} \bar{\Delta} \right] \right. \\
&\quad \left. - \frac{1}{2} \frac{(1-\beta h)}{1-h} (1-\alpha) (c_y^F)^{-1} (1-\Phi) \left[\frac{\eta+1}{\eta} \left(1 - Q_n^1 + \frac{\eta+1}{\eta} Q_n^0 Q_n^1 \right) \right] \frac{\theta^2}{(1+\alpha\eta^{-1})^2} \bar{\Delta} + \text{hot} + \text{tip}. \blacksquare
\end{aligned}$$

Appendix D: Taylor-pricing model with 3-quarter length contracts

In the Taylor-pricing model with 3-quarter length contracts, in each period one third of all firms can change their price. This price will then have to last for 3 quarters before it can be reset. The model is otherwise identical to the baseline model.

The firms maximization problem is now given by

$$\max E_t \sum_{j=0}^2 Q_{t,t+j} [Y_{t+j}(i) B_t(i) \bar{\Pi}^{j\omega} - W_{t+j}(i) N_{t+j}(i)] \quad (D.1)$$

since the price is only fixed for 3 quarters. The first order condition from this optimization problem is

$$\frac{B_t(i)}{P_t} = \frac{\theta}{\theta-1} \frac{E_t \sum_{j=0}^2 Q_{t,t+j} Y_{t+j}(P_{t+j}/P_t)^{\theta+1} \bar{\Pi}^{-j\omega\theta} (MC(i)_{t+j}/P_{t+j})}{E_t \sum_{j=0}^2 Q_{t,t+j} Y_{t+j}(P_{t+j}/P_t)^\theta \bar{\Pi}^{-j\omega(\theta-1)}} \quad (D.2)$$

Given these price-setting assumptions, the dynamics of the price level are governed by

$$P_t^{1-\theta} = \frac{1}{3} [B_t^{1-\theta} + B_{t-1}^{1-\theta} + B_{t-2}^{1-\theta}] \quad (D.3)$$

The remaining equations are as in the baseline model.

Steady State:

The steady state real reset price is then equal to $\frac{\bar{P}}{P} = \left[\frac{1}{3} (1 + \bar{\Pi}^{\theta-1} + \bar{\Pi}^{2(\theta-1)}) \right]^{1/(\theta-1)}$ which is greater than one for positive trend inflation. This implies that the output gap is equal to

$$\bar{X} = \left\{ \left(\frac{\bar{P}}{P} \right)^{1+\theta\eta^{-1}} \frac{\left\{ \sum_{j=0}^2 [\beta \bar{\Pi}^{\theta-1}]^j \right\}}{\left\{ \sum_{j=0}^2 [\beta \bar{\Pi}^{\theta(1+\eta^{-1})}]^j \right\}} \right\}^{1/(1+\eta^{-1})} \quad (D.4)$$

In the steady state, the cross sectional mean of prices is given by

$$\bar{P}_t = E_i \log P_t(i) = \frac{1}{3} [B_t + B_{t-1} + B_{t-2}] = \bar{b} + \log P_t - \log \bar{\Pi} \quad (D.5)$$

Hence, the cross sectional variance of prices in the steady state is equal to $\bar{\Delta} = \text{var}_i(\log P_t(i)) = \frac{2}{3} \bar{\pi}^2$.

Dynamic Equations:

With staggered pricing the log-linearized equations are now as follows. For equation (D.3),

$$\hat{b}_t = \bar{\Pi}^{\theta-1} (\hat{\pi}_t - \hat{b}_{t-1}) + \bar{\Pi}^{2(\theta-1)} (\hat{\pi}_t + \hat{\pi}_{t-1} - \hat{b}_{t-2}) \quad (D.6)$$

and for equation (D.2)

$$\begin{aligned} \left(1 + \frac{\theta}{\eta}\right) \hat{b}_t = & -G_2^{-1} (\hat{\xi}_t + \gamma_2 \hat{\xi}_{t+1} + \gamma_2^2 \hat{\xi}_{t+2}) + G_2^{-1} \eta^{-1} (\hat{y}_t + \gamma_2 \hat{y}_{t+1} + \gamma_2^2 \hat{y}_{t+2}) \\ & + (\widehat{g}_{y_{t+1}} - \hat{r}_t) [G_2^{-1} (\gamma_2 + \gamma_2^2) - G_1^{-1} (\gamma_1 + \gamma_1^2)] + (\widehat{g}_{y_{t+2}} - \hat{r}_{t+1}) (G_2^{-1} \gamma_2^2 - G_1^{-1} \gamma_1^2) \\ & + \hat{\pi}_{t+1} \left[\left(1 + \theta \left(1 + \frac{\theta}{\eta}\right)\right) G_2^{-1} (\gamma_2 + \gamma_2^2) - \theta G_1^{-1} (\gamma_1 + \gamma_1^2) \right] \\ & + \hat{\pi}_{t+2} \left[\left(1 + \theta \left(1 + \frac{\theta}{\eta}\right)\right) G_2^{-1} \gamma_2^2 - \theta G_1^{-1} \gamma_1^2 \right], \end{aligned} \quad (D.7)$$

where $\gamma_1 = \lambda \beta \bar{\Pi}^{(1-\omega)(\theta-1)}$ and $\gamma_2 = \gamma_1 \bar{\Pi}^{(1-\omega)(1+\theta/\eta)}$, and $G_t = 1 + \gamma_2 + \gamma_2^2$, $G_1 = 1 + \gamma_1 + \gamma_1^2$.

Lemma 4 in Taylor model:

In the baseline model, we derived the cross-sectional price dispersion up to a second order approximation. Because, the pricing contracts are different in the Taylor model, this section derives an equivalent Lemma for the Taylor model.

Lemma 4 (Taylor):

Let $\Xi_t = \Delta_t - \bar{\Delta}$ be the deviation of cross-section price dispersion from its non-stochastic steady state level $\bar{\Delta}$. Then

$$\begin{aligned} \Delta_t - \bar{\Delta} = & \frac{1}{3} \Gamma_1 [\hat{b}_t^2 + 2\bar{b}\hat{b}_t + \hat{b}_{t-1}^2 + \hat{\pi}_t^2 - 2\hat{b}_{t-1}\hat{\pi}_t + 2\hat{b}_{t-1}(\bar{b} - \bar{\pi}) - 2\hat{\pi}_t(\bar{b} - \bar{\pi}) \\ & + \hat{b}_{t-2}^2 + \hat{\pi}_t^2 + \hat{\pi}_{t-1}^2 - 2\hat{b}_{t-2}\hat{\pi}_t - 2\hat{b}_{t-2}\hat{\pi}_{t-1} + 2\hat{\pi}_t\hat{\pi}_{t-1} + 2\hat{b}_{t-2}(\bar{b} - 2\bar{\pi}) - 2\hat{\pi}_t(\bar{b} - 2\bar{\pi}) \\ & - 2\hat{\pi}_{t-1}(\bar{b} - 2\bar{\pi})] + \frac{2}{3} \Gamma_1 [\hat{b}_t + \hat{b}_{t-1} + \hat{b}_{t-2} - 2\hat{\pi}_t - \hat{\pi}_{t-1}] [Q_p^0] + h.o.t \end{aligned} \quad (D.8)$$

where $\Gamma_1 = [1 + (\theta - 1)Q_p^1(\bar{b} - \bar{\pi} + Q_p^0)]^{-1}$.

Proof:

At any given time the cross sectional dispersion is given by

Appendix E: A model with state dependent pricing

The state dependent pricing model follows the model developed in Dotsey et al. (1999). In the beginning of each period, every firm draws a cost of price adjustment. The draws are independent across firms and time. The drawn cost is a fraction of labor costs. Objective functions and constraints are the same as in the baseline model unless specified otherwise.

Let ζ_{jt} be the fraction of firms at the beginning of time t with price spell equal to j periods. The end of period share of firms at time t with price spell equal to j periods is denoted with ω_{jt} . Define α_{jt} as the probability of resetting price for firms with price spells equal to j periods. We assume, as in Dotsey et al. (1999), that the maximum duration of price spells is J periods and all firms must adjust prices after J periods. Given these definitions, we have $\omega_{jt} = (1 - \alpha_{jt})\zeta_{jt}$, $j = 1, \dots, J - 1$.

The expected cost of price adjustment (after normalization by the level of technology) for firms with price spells equal to j periods is then equal to $\Xi_{jt} = w_{jt} \int_0^{G^{-1}(\alpha_{jt})} xg(x)dx$ where G and g are the c.d.f. and p.d.f. of the random cost of price adjustment and w_{jt} is the wage rate paid by firms with price spells equal to j periods.

Similar to derivations for the Taylor model, the optimal reset price B is given by

$$\left(\frac{B_t}{P_t}\right)^{1+\frac{\theta}{\eta}} = \frac{\theta}{\theta-1} \frac{E_t \sum_{j=0}^{J-1} [\prod_{i=1}^j R_{t+i}^{-1} G Y_{t+i+1} \Pi_{t+i}^{1+\theta+\theta/\eta} (1-\alpha_{i,t+i})] \left(\frac{\xi_{t+j}^{-1}}{A_{t+j}}\right) \left(\frac{Y_{t+j}}{A_{t+j}}\right)^{1/\eta}}{E_t \sum_{j=0}^{J-1} [\prod_{i=1}^j R_{t+i}^{-1} G Y_{t+i+1} \Pi_{t+i}^{\theta} (1-\alpha_{i,t+i})]}$$

where R is the gross nominal interest rate, Π is the gross rate of inflation, GY is the gross rate of output growth, C is consumption, Y is output, and A is the level of technology.

The value of a firm at time t with price spell equal to $j = 0, \dots, J - 1$ periods is given by

$$v_{jt} = profit_{jt} + E_t R_t^{-1} \Pi_{t+1} G A_{t+1} [(1 - \alpha_{j,t+1}) v_{j+1,t+1} + \alpha_{j,t+1} v_{0,t+1} - \Xi_{j+1,t+1}],$$

$$profit_{jt} = C_t \left[\left(\frac{B_{t-j}}{P_{t-j}} \{\prod_{i=0}^{j-1} \Pi_{t-i}^{-1}\}\right)^{1-\theta} - \left(\frac{\xi_t^{-1}}{A_t}\right) \left(\frac{Y_t}{A_t}\right)^{1/\eta} \left(\frac{B_{t-j}}{P_{t-j}} \{\prod_{i=0}^{j-1} \Pi_{t-i}^{-1}\}\right)^{-\theta(1+\frac{1}{\eta})} \right].$$

For firms with price spell equal to $J - 1$ the value is given by

$$v_{J-1} = profit_{J-1,t} + E_t R_t^{-1} \Pi_{t+1} G A_{t+1} [v_{0,t+1} - \Xi_J]$$

Firms adjust prices if the drawn cost of changing prices is smaller than the gain in the value of the firm. Hence the probability of price adjustment is $\alpha_{jt} = G([v_{0,t} - v_{jt}]/w_{jt})$. Given that firms have to pay the cost of changing prices, the resource constraint is now given by $Y_t = C_t + G_t + \sum_{j=1}^J \Xi_{jt}$. The price level P is given by a weighted sum of current and past reset prices $P_t^{1-\theta} = \sum_{j=0}^{J-1} \omega_{jt} B_{t-j}^{1-\theta}$.

Steady state

The new elements in the steady of the system are given by:

$$\text{Reset price: } \left(\frac{\bar{B}}{\bar{P}}\right) = \left[\sum_{j=0}^{J-1} \bar{\omega}_j \bar{\Pi}^{j(\theta-1)}\right]^{1/(\theta-1)}, \bar{b} = \log\left(\frac{\bar{B}}{\bar{P}}\right)$$

$$\text{Average price: } \bar{P} = E_i(\log(P_t(i))) = \bar{b} + \log P_t - \bar{\pi} \sum_{j=0}^{J-1} \bar{\omega}_j j.$$

Price dispersion:

$$\bar{\Delta} = \text{var}_i(\log(P_t(i))) = \int_0^1 [\log(P_t(i)) - \bar{P}_t]^2 di = \sum_{j=0}^{J-1} \bar{\omega}_j [\bar{b} + \log P_{t-j} - \bar{P}_t]^2$$

$$= \sum_{j=0}^{J-1} \bar{\omega}_j [\bar{b} + \log P_t - j\bar{\pi} - \bar{P}_t]^2 = \bar{\pi}^2 \left\{ \sum_{j=0}^{J-1} \bar{\omega}_j j^2 - \left[\sum_{j=0}^{J-1} \bar{\omega}_j j\right]^2 \right\}.$$

$$\text{Output is determined from } \left(\frac{\bar{B}}{\bar{P}}\right)^{1+\theta/\eta} = \frac{(1-h)}{1-\beta h} \left[(1-g) - \frac{\sum_{j=1}^J \bar{\Xi}_t}{\bar{Y}} \right] \bar{Y}^{1+1/\eta} \frac{\theta}{\theta-1} \frac{\sum_{j=0}^{J-1} [\prod_{i=1}^j \bar{R}^{-1} \bar{G} \bar{Y} \bar{\Pi}^{1+\theta+\theta/\eta} (1-\bar{\alpha}_i)]}{E_t \sum_{j=0}^{J-1} [\prod_{i=1}^j \bar{R}^{-1} \bar{G} \bar{Y} \bar{\Pi}^{\theta} (1-\bar{\alpha}_i)]}$$

while potential output (flexible prices) is given by $\bar{Y}^* = \left(\frac{\theta-1}{\theta(1-g)} \frac{(1-\beta h)}{(1-h)}\right)^{\eta/(\eta+1)}$ so that output gap in the steady state is $\bar{X} = \bar{Y}/\bar{Y}^*$.

Linearized system of equations

After solving for the steady state values of the variables (which we denote with a bar above a given variable) and utilizing previous results, we log-linearize the new elements in the system as follows:

Reset price:

$$\left(1 + \frac{\theta}{\eta}\right) \hat{b}_t = \sum_{j=0}^{J-1} \delta_{2,j} \left(-\hat{\xi}_{t+j} + \frac{1}{\eta} \hat{y}_{t+j}\right) + \sum_{j=1}^{J-1} \left[\left(\sum_{k=j}^{J-1} \delta_{2,k}\right) \left(1 + \theta + \frac{\theta}{\eta}\right) - \left(\sum_{k=j}^{J-1} \delta_{1,k}\right) \theta \right] \hat{\pi}_{t+j} + \sum_{j=1}^{J-1} \left[\left(\sum_{k=j}^{J-1} \delta_{1,k}\right) \left(\widehat{g}_{t+j} - \hat{r}_{t+j-1} - \frac{\bar{\alpha}_j}{1-\bar{\alpha}_j} \hat{\alpha}_{j,t+j}\right) \theta \right]$$

where $\delta_{2,j} = \gamma_{2,j} / \sum_{k=0}^{J-1} \gamma_{2,k}$, $\delta_{1,j} = \gamma_{1,j} / \sum_{k=0}^{J-1} \gamma_{1,k}$, $\gamma_{2,j} = \bar{R}^{-j} \bar{G} \bar{Y}^j \Pi^{j(1+\theta+\theta/\eta)}$, $\prod_{i=1}^j (1 - \bar{\alpha}_i)$, and $\gamma_{2,j} = \bar{R}^{-j} \bar{G} \bar{Y}^j \Pi^{j\theta} \prod_{i=1}^j (1 - \bar{\alpha}_i)$.

Price level:

$$0 = \sum_{j=0}^{J-1} \bar{\omega}_j \left(\frac{\bar{B}}{\bar{P}}\right)^{1-\theta} \bar{\Pi}^{-j(1-\theta)} \hat{b}_{t-j} - \sum_{j=1}^{J-1} \left[\sum_{k=j}^{J-1} \bar{\omega}_k \left(\frac{\bar{B}}{\bar{P}}\right)^{1-\theta} \bar{\Pi}^{-k(1-\theta)} \right] \hat{\pi}_{t-j+1} + \frac{1}{1-\theta} \sum_{j=0}^{J-1} \left[\bar{\omega}_j \left(\frac{\bar{B}}{\bar{P}}\right)^{1-\theta} \bar{\Pi}^{-j(1-\theta)} \right] \hat{\omega}_{jt}$$

Distribution:

$$\hat{\zeta}_{j+1,t+1} = -\frac{\bar{\alpha}_j}{1-\bar{\alpha}_j} \hat{\alpha}_{jt} + \hat{\zeta}_{jt}, \quad j = 1, \dots, J-1 \text{ and } \sum_{j=0}^{J-1} \bar{\zeta}_j \hat{\zeta}_{jt} = 0$$

$$\hat{\omega}_{jt} = -\frac{\bar{\alpha}_j}{1-\bar{\alpha}_j} \hat{\alpha}_{jt} + \hat{\zeta}_{jt}, \quad j = 1, \dots, J-1 \text{ and } \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} = 0.$$

$$\text{Reset probabilities: } \hat{\alpha}_{jt} = g \left(\frac{[\bar{v}_0 - \bar{v}_j]}{\bar{w}_j} \right) \left[\frac{\bar{v}_0}{\bar{\alpha}_j \bar{w}_j} \hat{v}_{0,t} - \frac{\bar{v}_j}{\bar{\alpha}_j \bar{w}_j} \hat{v}_{jt} - \frac{[\bar{v}_0 - \bar{v}_j]}{\bar{\alpha}_j \bar{w}_j} \hat{\omega}_{jt} \right], \quad j = 1, \dots, J-1$$

$$\text{Wages: } \hat{w}_{jt} = -\hat{\xi}_t + \frac{1}{\eta} \hat{y}_t - \frac{\theta}{\eta} \hat{b}_{t-j} + \frac{\theta}{\eta} \bar{\Pi}^j \sum_{k=1}^j \hat{\pi}_{t-k+1}$$

Value functions:

$$\hat{v}_{jt} = \frac{\overline{\text{profit}}_j}{\bar{v}_j} \widehat{\text{profit}}_{jt} + \beta (1 - \bar{\alpha}_{j+1}) \frac{\bar{v}_{j+1}}{\bar{v}_j} E_t \left[\hat{v}_{j+1,t+1} + \hat{\pi}_{t+1} - \hat{r}_t + \widehat{g} \hat{a}_{t+1} - \frac{\bar{\alpha}_{j+1}}{1-\bar{\alpha}_{j+1}} \hat{\alpha}_{j+1,t+1} \right] + \beta \frac{\bar{v}_0}{\bar{v}_j} E_t \left[\hat{v}_{0,t+1} + \hat{\pi}_{t+1} - \hat{r}_t + \widehat{g} \hat{a}_{t+1} + \hat{\alpha}_{j+1,t+1} \right] - \beta \frac{\bar{\Xi}_{j+1}}{\bar{v}_j} E_t \left[\bar{\Xi}_{j+1,t+1} + \hat{\pi}_{t+1} - \hat{r}_t + \widehat{g} \hat{a}_{t+1} \right], \quad j=1, \dots, J-1$$

$$\hat{v}_{j-1,t} = \frac{\overline{\text{profit}}_{j-1}}{\bar{v}_{j-1}} \widehat{\text{profit}}_{j-1,t} + \beta \frac{\bar{v}_0}{\bar{v}_{j-1}} E_t \left[\hat{v}_{0,t+1} + \hat{\pi}_{t+1} - \hat{r}_t + \widehat{g} \hat{a}_{t+1} \right] - \beta \frac{\bar{\Xi}_j}{\bar{v}_{j-1}} E_t \left[\hat{\pi}_{t+1} - \hat{r}_t + \widehat{g} \hat{a}_{t+1} \right]$$

Profits:

$$\widehat{\text{profit}}_{jt} = \hat{c}_t + \left[\frac{\bar{c} \bar{\xi}^{-1} \bar{Y}^{\frac{1}{\eta}}}{\overline{\text{profit}}_j} \left(\frac{\bar{B}}{\bar{P}}\right)^{-\theta(1+\frac{1}{\eta})} \bar{\Pi}^{j\theta(1+\frac{1}{\eta})} \right] \hat{\xi}_t - \frac{1}{\eta} \left[\frac{\bar{c} \bar{\xi}^{-1} \bar{Y}^{\frac{1}{\eta}}}{\overline{\text{profit}}_j} \left(\frac{\bar{B}}{\bar{P}}\right)^{-\theta(1+\frac{1}{\eta})} \bar{\Pi}^{j\theta(1+\frac{1}{\eta})} \right] \hat{y}_t + \frac{\bar{c}}{\overline{\text{profit}}_j} \left[\left(\frac{\bar{B}}{\bar{P}}\right)^{1-\theta} \bar{\Pi}^{-j(1-\theta)} (1-\theta) + \theta \bar{C} \bar{Y}^{\frac{1}{\eta}} \left(\frac{\bar{B}}{\bar{P}}\right)^{-\theta(1+\frac{1}{\eta})} \bar{\Pi}^{j\theta(1+\frac{1}{\eta})} \right] \hat{b}_{t-j} - \sum_{k=1}^j \frac{\bar{c}}{\overline{\text{profit}}_j} \left[\left(\frac{\bar{B}}{\bar{P}}\right)^{1-\theta} \bar{\Pi}^{-j(1-\theta)} (1-\theta) + \theta \bar{C} \bar{Y}^{\frac{1}{\eta}} \left(\frac{\bar{B}}{\bar{P}}\right)^{-\theta(1+\frac{1}{\eta})} \bar{\Pi}^{j\theta(1+\frac{1}{\eta})} \right] \pi_{t-k+1}$$

Adjustment costs:

$$\hat{\Xi}_{jt} = \left[\frac{\partial G^{-1}(\bar{\alpha}_j)}{\partial \bar{\alpha}_j} \times G^{-1}(\bar{\alpha}_j) \times g(G^{-1}(\bar{\alpha}_j)) \times \frac{\bar{\alpha}_j}{\bar{\Xi}_j} \right] \hat{\alpha}_{jt} + \hat{\omega}_{jt}, \quad j = 1, \dots, J-1$$

$$\hat{\Xi}_{jt} = 0$$

$$\text{Resource constraint: } \hat{y}_t = s_c \hat{c}_t + s_g \hat{g}_t + \sum_{j=1}^J s_{\Xi,j} \hat{\Xi}_{jt}.$$

Approximation to utility

Lemmas 1, 2, 3 and 5 derived for the baseline model apply to the current setup. Modified Lemma 4 is given below.

Lemma 4E. Let $\Xi_t = \Delta_t - \bar{\Delta}$ be the deviation of cross-section price dispersion from its non-stochastic steady state level $\bar{\Delta}$. Then

$$\begin{aligned} \Gamma_1^{-1}(\Delta_t - \bar{\Delta}) &= \sum_{j=0}^{J-1} \bar{\omega}_j \hat{b}_{t-j}^2 + \sum_{j=0}^{J-1} \bar{\omega}_j [\sum_{k=0}^{j-1} \hat{\pi}_{t-k}]^2 - 2 \sum_{j=0}^{J-1} \bar{\omega}_j [\sum_{k=0}^{j-1} \hat{b}_{t-j} \hat{\pi}_{t-k}] + \\ & 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \hat{b}_{t-j} - 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) [\sum_{k=0}^{j-1} \hat{\pi}_{t-k}] + \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} (\bar{b} - j\bar{\pi})^2 + \\ & 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \hat{\omega}_{jt} \hat{b}_{t-j} - 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) [\sum_{k=0}^{j-1} \hat{\omega}_{jt} \hat{\pi}_{t-k}] + 2Q_p^0 \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} \hat{b}_{t-j} - \\ & 2Q_p^0 \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} [\sum_{k=0}^{j-1} \hat{\pi}_{t-k}] + 2Q_p^0 \sum_{j=0}^{J-1} \bar{\omega}_j \{ \hat{b}_{t-j} - \sum_{k=0}^{j-1} \hat{\pi}_{t-k} \} + 2Q_p^0 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \hat{\omega}_{jt}, \end{aligned}$$

where $\Gamma_1 = \left\{ 1 + (\theta - 1)Q_p^1 [Q_p^0 + \sum_{j=0}^{J-1} \omega_j (\bar{b} - j\bar{\pi})] \right\}^{-1}$.

Proof:

$$\begin{aligned} \Delta_t \equiv \text{var}_i(\log(P_t(i))) &= \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} + \log P_{t-j} - \bar{P}_t]^2 = \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} + (\log P_t - \bar{P}_t) - \sum_{k=0}^{j-1} \pi_{t-k}]^2 \\ &= \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}]^2 + (\log P_t - \bar{P}_t)^2 + 2 \sum_{j=0}^{J-1} \omega_j [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}] (\log P_t - \bar{P}_t) \\ &= \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}]^2 + (Q_p^0)^2 + (1 - \theta)Q_p^0 Q_p^1 (\Delta_t - \bar{\Delta}) \\ & \quad + 2Q_p^0 \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}] + (1 - \theta)Q_p^1 \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}] (\Delta_t - \bar{\Delta}) \\ & \quad + \text{hot} + \text{tip} \end{aligned}$$

Consider each term in turn. The first term is

$$\begin{aligned} \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}]^2 &= \sum_{j=0}^{J-1} (\omega_{jt} - \bar{\omega}_j + \bar{\omega}_j) [(b_{t-j} - \bar{b}) + (\bar{b} - j\bar{\pi}) - \sum_{k=0}^{j-1} (\pi_{t-k} - \bar{\pi})]^2 \\ &= \sum_{j=0}^{J-1} \bar{\omega}_j [\hat{b}_{t-j} + (\bar{b} - j\bar{\pi}) - \sum_{k=0}^{j-1} \hat{\pi}_{t-k}]^2 + \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} [\hat{b}_{t-j} + (\bar{b} - j\bar{\pi}) - \sum_{k=0}^{j-1} \hat{\pi}_{t-k}]^2 \\ &= \sum_{j=0}^{J-1} \bar{\omega}_j \hat{b}_{t-j}^2 + \sum_{j=0}^{J-1} \bar{\omega}_j [\sum_{k=0}^{j-1} \hat{\pi}_{t-k}]^2 + \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi})^2 - 2 \sum_{j=0}^{J-1} \bar{\omega}_j \hat{b}_{t-j} \sum_{k=0}^{j-1} \hat{\pi}_{t-k} + \\ & \quad \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \hat{b}_{t-j} - 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \sum_{k=0}^{j-1} \hat{\pi}_{t-k} + \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \hat{\omega}_{jt} + \\ & \quad 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \hat{\omega}_{jt} \hat{b}_{t-j} - 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \sum_{k=0}^{j-1} \hat{\omega}_{jt} \hat{\pi}_{t-k} + \text{hot} + \text{tip}. \end{aligned}$$

The second to last term is

$$\begin{aligned} \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}] &= \sum_{j=0}^{J-1} (\omega_{jt} - \bar{\omega}_j + \bar{\omega}_j) [(b_{t-j} - \bar{b}) + (\bar{b} - j\bar{\pi}) - \sum_{k=0}^{j-1} (\pi_{t-k} - \bar{\pi})] \\ &= \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} \hat{b}_{t-j} - \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} \sum_{k=0}^{j-1} \hat{\pi}_{t-k} + \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) + \sum_{j=0}^{J-1} \bar{\omega}_j [\hat{b}_{t-j} - \sum_{k=0}^{j-1} \hat{\pi}_{t-k}] + \\ & \quad \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \hat{\omega}_{jt}. \end{aligned}$$

The last term is

$$\begin{aligned} \sum_{j=0}^{J-1} \omega_{jt} [b_{t-j} - \sum_{k=0}^{j-1} \pi_{t-k}] (\Delta_t - \bar{\Delta}) &= \\ &= \sum_{j=0}^{J-1} (\omega_{jt} - \bar{\omega}_j + \bar{\omega}_j) [(b_{t-j} - \bar{b}) + (\bar{b} - j\bar{\pi}) - \sum_{k=0}^{j-1} (\pi_{t-k} - \bar{\pi})] (\Delta_t - \bar{\Delta}) \\ &= \sum_{j=0}^{J-1} (\omega_{jt} - \bar{\omega}_j) (b_{t-j} - \bar{b}) (\Delta_t - \bar{\Delta}) + \sum_{j=0}^{J-1} (\omega_{jt} - \bar{\omega}_j) (\bar{b} - j\bar{\pi}) (\Delta_t - \bar{\Delta}) \\ & \quad - \sum_{j=0}^{J-1} \sum_{k=0}^{j-1} (\pi_{t-k} - \bar{\pi}) (\omega_{jt} - \bar{\omega}_j) (\Delta_t - \bar{\Delta}) \\ & \quad + \sum_{j=0}^{J-1} \bar{\omega}_j [(b_{t-j} - \bar{b}) + (\bar{b} - j\bar{\pi}) - \sum_{k=0}^{j-1} (\pi_{t-k} - \bar{\pi})] (\Delta_t - \bar{\Delta}) \\ &= \sum_{j=0}^{J-1} \bar{\omega}_j \hat{\omega}_{jt} (\bar{b} - j\bar{\pi}) (\Delta_t - \bar{\Delta}) + \sum_{j=0}^{J-1} \bar{\omega}_j [\hat{b}_{t-j} + (\bar{b} - j\bar{\pi}) - \sum_{k=0}^{j-1} \hat{\pi}_{t-k}] (\Delta_t - \bar{\Delta}) + \text{hot} + \text{tip} \\ &= (\Delta_t - \bar{\Delta}) \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) + \text{hot} + \text{tip} \end{aligned}$$

Regroup terms and take $(\Delta_t - \bar{\Delta})$ to get the expression in the lemma. ■

Similar to the derivations for the Taylor and Calvo models, we consider low levels of trend inflation so that $(\bar{b} - j\bar{\pi})$ is close to zero. Note that as j increases, the sum of steady state levels of inflation increases but this increase in the sum is discounted more heavily since $\bar{\omega}_j$ is decreasing in j . Hence, terms which are linear in \hat{b}_{t-j} and $\hat{\pi}_{t-k}$ can be ignored. The new terms relative to the Taylor model capture the covariance between shares of firms resetting prices and the size of the reset price or inflation.

Using Lemma 4E, we can also find

$$\begin{aligned} \Gamma_1^{-1} E(\Xi_t) &= \Gamma_1^{-1} E(\Delta_t - \bar{\Delta}) = \\ & \text{var}(\hat{b}_{t-j}) + \sum_{j=0}^{J-1} \bar{\omega}_j \sum_{k=0}^{j-1} \sum_{s=0}^{j-1} \text{cov}(\hat{\pi}_{t-k}, \hat{\pi}_{t-s}) - 2 \sum_{j=0}^{J-1} \bar{\omega}_j [\sum_{k=0}^{j-1} \text{cov}(\hat{b}_{t-j}, \hat{\pi}_{t-k})] \\ & + 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) \text{cov}(\hat{\omega}_{jt}, \hat{b}_{t-j}) - 2 \sum_{j=0}^{J-1} \bar{\omega}_j (\bar{b} - j\bar{\pi}) [\sum_{k=0}^{j-1} \text{cov}(\hat{\omega}_{jt}, \hat{\pi}_{t-k})] + \\ & 2Q_p^0 \sum_{j=0}^{J-1} \bar{\omega}_j \text{cov}(\hat{\omega}_{jt}, \hat{b}_{t-j}) - 2Q_p^0 \sum_{j=0}^{J-1} \bar{\omega}_j [\sum_{k=0}^{j-1} \text{cov}(\hat{\omega}_{jt}, \hat{\pi}_{t-k})]. \end{aligned}$$

Proposition 1E. Given Lemmas 1, 2, 3, 4E and 5, the second order approximation to per period utility in eq. (1) is

$$\begin{aligned}
EU_t \approx & \log\left\{\frac{(1-s_c)}{(1-\bar{g}_y)}\right\} + \left\{1 - \left(\frac{1-\beta h}{1-h}\right) \frac{1-\beta h}{1-h} \frac{1-\Phi}{1-\bar{g}_y} (1 - (1 + \eta^{-1})Q_y^0)\right\} \log \bar{X} \\
& - \left(\frac{1}{2} \frac{1-\beta h}{1-h} \frac{1-\Phi}{1-\bar{g}_y} (1 + \eta^{-1})\right) (\log \bar{X})^2 \\
& - \frac{1}{2} \frac{(1-\beta h)}{(1-h)} \frac{1}{1-\bar{g}_y} [Q_y^1 \theta^{-1} - Q_y^1 + (1 + \eta^{-1})(1 + \frac{\theta-1}{\theta} Q_y^1 Q_y^0)] \theta^2 \bar{\Delta} \\
& - \frac{1-\beta h}{1-h} \frac{1-\Phi}{1-\bar{g}_y} \left\{ (1 + \eta^{-1})(Q_y^0)^2 - Q_y^0 + \frac{\theta-1}{2\theta} Q_y^1 [1 - (1 + \eta^{-1})Q_y^0] \bar{Y} \right\} \\
& - \frac{h(1-\rho_e)}{(1-h)^2} \text{var}(\hat{c}_t) - \frac{1}{2} \frac{1-\beta h}{1-h} \frac{1-\Phi}{1-\bar{g}_y} (1 + \eta^{-1}) \text{var}(\hat{y}_t) \\
& - \frac{1}{2} \frac{(1-\beta h)}{(1-h)} \frac{1}{1-\bar{g}_y} \theta^2 \left\{ [Q_y^1 \theta^{-1} - Q_y^1 + (1 + \eta^{-1})(1 + \frac{\theta-1}{\theta} Q_y^1 Q_y^0)] - \frac{\theta-1}{\theta} (1 + \eta^{-1}) Q_y^1 \log \bar{X} \right\} E(\Xi_t) \\
& + \text{hot} + \text{tip}
\end{aligned}$$

Proof:

The derivation of the dispersion term follows the derivation in Proposition 1. Note that because costs of changing prices enter the resource constraint, we use $\hat{y}_t = \tilde{y}_t - \log(\bar{X})$ to express the deviation of consumption from its flexible prices level as follows:

$$\begin{aligned}
\tilde{c}_t = c_t - c_t^F &= (c_t - \bar{c}_t) + (\bar{c}_t - c_t^F) = \hat{c}_t + \log\{(1 - s_c)\bar{Y}_t\} - \log\{(1 - \bar{g}_y)Y_t^F\} \\
&= \hat{c}_t + \log\{(1 - s_c)/(1 - \bar{g}_y)\} + \log(\bar{X})
\end{aligned}$$

Following the steps in Proposition 1, we derive the expected value of the per period utility. ■

Appendix F: Non-linear solution to the baseline model

In the log-linearized model, the unconditional variances of output and inflation can also be calculated from the moving average representation. Since this moving average representation is equivalent to the impulse response function (IRF), we check the accuracy of our solution by comparing the IRF of the approximated model with the IRF from a non-linear solution method.

First, we rewrite the system of nonlinear equations so that the price setting equations of the baseline model are in recursive form as in Woodford (2010):

$$\frac{1-\lambda\Pi_t^{\theta-1}}{1-\lambda} = \left(\frac{K_t}{H_t}\right)^{\frac{\theta-1}{1+\theta\eta^{-1}}},$$

$$K_t = \frac{\theta}{\theta-1} Y_t^{1/\eta} \xi_t^{-1} + \beta\lambda E_t \Pi_{t+1}^{\theta(1+1/\eta)} K_{t+1},$$

$$H_t = 1 + \beta\lambda E_t (\Pi_{t+1}^{\theta-1} H_{t+1})$$

where K and H are the numerator and denominator in equation (11).

Since computing a solution to the nonlinear model is computationally costly, we focus only on the worst case scenario: zero trend inflation and a large shock which leads to binding zero lower bound. In particular, we consider a 3 standard deviation shock to the risk premium q , which will cause the ZLB to bind in the log-linearized baseline model. We use a version of the shooting algorithm from Fair and Taylor (1983) to solve for the IRF of the non-linear model in the following steps:

1. Initialize the algorithm by fixing the end point Y_T at $T=75$.
2. Use the linear solution as an initial guess for the algorithm.
3. We split the IRF into 10 subsamples of length 11. In each sub-sample we take the first two observations as initial condition, and the last observation as terminal condition.
4. With these initial conditions, we calculate a new path for the endogenous variables $\{Y_t^1\}_{t=0}^T$ for each sub-sample (Type I iteration).
5. Use this new path as future expectations and keep iterating until the solution converges on a path $\{Y_t^K\}_{t=0}^T$ in each sub-sample (Type II iteration).
6. For each sub-sample we update the initial and terminal conditions using iterated values from the preceding and proceeding sub-samples until convergence.
7. After convergence across sub-samples is achieved increase T by 1 period until the change in the IRF from increasing T falls below a critical value (Type III iteration).

We achieved convergence after $T=76$ periods (total runtime was 10 minutes while the run time for the log-linearized model is a tiny fraction of a second). The IRF from the linear and non-linear solution are plotted in the figure below. We calculate the distance between the two IRFs by summing the absolute value of their differences and dividing by the absolute sum of the IRF coefficients of the nonlinear solution. This metric yields errors for output, consumption, and inflation of approximately 1%. Since only the moments of output, consumption and inflation are of interest to us, this result suggests that the linear approximation introduces only small errors, and will not quantitatively affect our results.

