Precautionary Pricing: The Disinflationary Effects of ELB Risk *

Robert Amano[†] Thomas J. Carter[‡] Sylvain Leduc[§]

May 1, 2018

Abstract

We construct and globally solve a model that offers a potential explanation for the weak performance of inflation as a longer-lasting phenomenon in the post-Great Recession era. In our model, risk-averse agents' forward-looking anticipation that the economy may experience episodes during which nominal rates are driven to their effective lower bound (ELB) leaves a significant impact on their behaviour even during "normal" times when the economy is away from the ELB. This impact includes a tendency to set substantially lower prices relative to an otherwise comparable model which abstracts from the ELB. As a result of this "precautionary pricing" behaviour, average inflation during periods when the ELB is lax tends to undershoot the central bank's target by up to sixty basis points, nearly two-thirds of which can be directly attributed to ELB risk. In addition, the model predicts a substantially negative inflation risk premium, in line with recent data.

PRELIMINARY DRAFT.
DO NOT CITE OR CIRCULATE WITHOUT THE AUTHORS' PERMISSION.

*The views expressed are the authors' and do not necessarily reflect those of the Bank of Canada.

†Bank of Canada, Ottawa, ON, Canada. *E-mail*: RAmano@BankOfCanada.ca

‡Bank of Canada, Ottawa, ON, Canada. *E-mail*: TCarter@BankOfCanada.ca

§Bank of Canada, Ottawa, ON, Canada. *Email*: SLeduc@BankOfCanada.ca

1 Introduction

In recent years, inflation in the advanced world has been surprisingly sluggish. Though weak inflation could initially be attributed to the Global Financial Crisis and ensuing recession, its persistence alongside a significant recovery in the real economy led many policymakers and academics to speculate that longer-run forces may be at work. For example, Yellen (2017), referring to weakness in US inflation, notes that "my colleagues and I are not certain that it is transitory...It may be that there is something more endemic or long-lasting that we need to pay attention to."

In addition, estimates of consistently negative inflation risk premia in the post-crisis era suggest that investors may perceive a significant risk that inflation could fall *even lower*, since nominal bonds' tendency to appreciate in this contingency makes them useful as insurance instruments. For example, Chen, Engstrom, and Grishchenko (2016), and Grishchenko and Huang (2013) estimate that US inflation risk premia were positive in the early 2000s but have since turned negative, beginning around the time of the Great Recession. Camba-Mendes and Werner (2017) similarly provide evidence that several measures of US and Eurozone inflation risk premia turned negative post-2008.

In this paper, we reconcile these trends in the context of a framework that offers a potential explanation for weak inflation as a longer-lasting phenomenon in the post-Great Recession era. In our model, risk-averse agents' forward-looking anticipation that the economy may experience episodes during which nominal rates are driven to their effective lower bound (ELB) leaves a significant impact on their behaviour, even during "normal" times when the economy is currently away from the ELB. This impact includes, among other things, a tendency for firms to set substantially lower prices relative to an otherwise comparable model which abstracts from the ELB.

As a result of this "precautionary pricing" behaviour, our baseline model has the property that the average rate of inflation, conditional on the policy rate being above the ELB, is less than 1.4%, despite the fact that the central bank is assumed to pursue a two percent inflation target. In contrast, an otherwise comparable model that abstracts from the ELB predicts average inflation around 1.8% after taking account the various non-linearities captured by our global solution method, which tend to push the average inflation rate slightly below target. Therefore, accounting for the relatively small probability that the ELB may bind in the future nonetheless results in a disinflationary bias roughly 3 times larger (60 versus 20 basis points) compared to an economy where policy is always unconstrained. In addition, the possibility that the ELB may bind in the future translates into negative inflation risk premia. In particular, our baseline calibration implies an average inflation risk premium of about -0.5 percent.

Since the mechanism driving these results is fundamentally precautionary in nature, we take steps to ensure that the model exhibits a reasonable amount of ELB risk and aversion thereto. We specifically do so by augmenting an otherwise typical New Keynesian framework to include three key ingredients, namely (i) recursive preferences, (ii) a source of long-run risk, and (iii) a demand-shock process admitting relatively deep, long-lived ELB episodes. In particular, we parameterize these ingredients to keep risk aversion and the implied equity premium in empirically plausible ranges while also ensuring that the typical ELB episode is roughly in line with recent experience, both in terms of duration and in terms of implied short- and long-run output losses. The ELB also binds with a frequency well within the range of estimates suggested by a post-war sample of developed countries.

In our framework, long-run risk arises from the interaction of the ELB with investment, namely due to the presence of an investment externality that leads to endogenous growth. At the ELB, monetary policy is unable to provide sufficient stimulus to bring output to its potential, resulting in below-target inflation. This also results in long-run risk, as the lower associated level of investment brings about a period of lower growth whose long-run level effects are never fully reversed, even after the policy rate has escaped the ELB. During periods when the ELB is lax, firms' anticipation of these episodes leads them to set lower prices, all else equal. In addition, this effect is reinforced by the fact that households increase their precautionary savings to buffer against the adverse effects that the ELB places on monetary policy, leading to weaker aggregate demand. Importantly, we show that recursive preferences help to amplify this overall disinflationary effect.

Our work complements a few papers in the literature that share our attention to the way that the risk of future ELB episodes influences economic outcomes during times when the ELB is not binding (e.g., Adam and Billi, 2007; Nakov 2008). More closely related is the work of Hills, Nakata, and Schmidt (2016), in which the authors compare the deterministic and risky steady states of an ELB-constrained New Keynesian model. However, they do so assuming standard — i.e., non-recursive — preferences and thus largely abstract from the points raised above regarding the importance of the precautionary forces at play.

Gourio and Ngo (2016) and Nakata and Tanaka (2016) also explore the asset-pricing implications of the ELB in New Keynesian economies but focus on periods when the ELB binds. For example, Gourio and Ngo (2016) highlight a mechanism that causes the the inflation risk premium to fall during these periods, namely due to the central bank's inability to offset demand shocks. In contrast, we focus on a complementary mechanism that operates even during periods when the economy is far removed from the ELB. Camba-Mendes and Werner (2017) provide some empirical support for this mechanism. In particular, they provide evidence suggesting that the fall in the inflation risk premium in the post-2008 period is not due to a decline in inflation uncertainty but, instead, to a shift in the balance of macroeconomic risks towards future ELB episodes.

Our paper also complements the work of Kung (2015) who examines the behavior of the term structure of interest rates in a model of vertical innovation with nominal rigidities. He shows that this framework can account for the negative empirical relationship between growth and inflation and explores the implications for bond pricing. Guerron-Quintana and Jinnai (2015) embed financial frictions and liquidity shocks into a similar endogenous growth framework to examine the impact of the Great Recession on the trend level of output. These authors, however, abstract from the ELB constraint on monetary policy.

Finally, our work also relates to the approach taken by Anzoategui et al. (2017) who examine the effect of the Great Recession on productivity in a model of endogenous technology adoption, finding that lower R&D investment played an important role in lowering the economy's growth trajectory and productivity. As them, we emphasize the importance of the endogenous response of investment and its impact on growth following recessions and ELB episodes (albeit through the presence of a positive investment externality rather than the expanding-variety mechanism that they consider.) However, our focus is on the implications of the ELB when the economy is operating away from the constraint.

The remainder of the paper is organized as follows. Section 2 describes the model. We then calibrate the model in Section 3 and present our main results in Section 4. Section 5 concludes.

2 Model

In this section, we augment an otherwise standard New Keynesian model to include three key ingredients, namely (i) recursive preferences, (ii) a source of long-run risk, and (iii) a demand shock admitting ELB episodes broadly consistent with recent experience.

To maintain continuity with related literature, we have borrowed much of the model's basic structure from Gourio and Ngo (2017), which itself closely follows Rudebusch and Swanson (2012). Subsections 2.1 through 2.3 describe the various agents populating the model economy, while subsections 2.4, 2.5, and 2.6 respectively elaborate on long-run risk, shocks, and asset prices.

2.1 Households

Households have recursive preferences of the form

$$V_t = \frac{c_t^{1-\sigma}}{1-\sigma} - \frac{\chi A_t^{1-\sigma} \ell_t^{1+\nu}}{1+\nu} - \beta[\mathbb{E}_t\{(-V_{t+1})^{1-\gamma}\}]^{\frac{1}{1-\gamma}}, \tag{1}$$

where c_t and ℓ_t respectively denote consumption and labour, while A_t denotes trend productivity and has been included to ensure the existence of a balanced-growth path.

In any given period, households face a budget constraint of the form

$$w_t\ell_t + r_t^k k_{t-1} + \frac{B_{t-1}R_{t-1}}{P_t} + D_t = c_t + i_t + \frac{\xi_t B_t}{P_t} + k_{t-1}\frac{\psi}{2}\left(\frac{i_t}{k_{t-1}} - \frac{\widetilde{i}}{k}\right)^2,$$

where w_t denotes the real wage rate; k_{t-1} denotes capital acquired in the previous period, which can be rented at rate r_t^k; B_{t-1} denotes nominal bonds acquired in the previous period, which pay nominal interest at gross rate R_{t-1}; P_t denotes the price level; D_t denotes dividends paid by the firms described below, which households are assumed to own; i_t denotes investment; and ξ_t is a shock on which we will elaborate momentarily. We also a convex investment-adjustment cost, $\frac{\psi}{2}\left(\frac{i_t}{k_{t-1}} - \frac{\widetilde{i}}{k}\right)^2$, where $\frac{\widetilde{i}}{k}$ denotes the value that the ratio i_t/k_{t-1} takes along the economy's balanced-growth path, while ψ is a parameter.

The law of motion for capital takes the form

$$k_t = (1-\delta)k_{t-1} + \frac{i_t}{\widetilde{Q}}, \tag{2}$$

where $\widetilde{Q}$ denotes the relative price of capital along the economy's balanced-growth path, which we allow to differ from unity for calibration purposes.

Appendix A describes the household's optimality conditions. We focus here on the Euler equation associated with nominal bonds, which is given by

$$1 = \mathbb{E}_t \left(\frac{m_{t+1}}{\xi_t} \cdot \frac{R_t}{\Pi_{t+1}} \right), \tag{3}$$

where $\Pi_{t+1} := P_{t+1}/P_t$ denotes the gross rate of inflation, while m_{t+1} denotes the real stochastic discount factor and can be shown to take the standard form

$$m_{t+1} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \left[\frac{-V_{t+1}}{\left[\mathbb{E}_t \left[(-V_{t+1})^{1-\gamma} \right]\right]^{\frac{1}{1-\gamma}}} \right]^{-\gamma}. \tag{4}$$

From these expressions, we see that low realizations of the shock ξ_t place upward pressure on households' demand for government bonds. ξ_t is thus meant to capture a "flight-to-quality" during which households aim to redirect funds from consumption and investment to the relative safety of bonds which offer a guaranteed nominal return. We will therefore refer to ξ_t as a "demand shock" at many points in the sequel. Similar shocks figure prominently in, e.g., Amano and Shukayev (2012) and Coibion et al. (2012).

2.2 Firms

The final goods used for consumption and investment are produced by aggregating a unit measure of intermediate goods. Specifically,

$$y_t = \left(\int_0^1 y_{it}^{\frac{\theta-1}{\theta}} di \right)^{\frac{\theta}{\theta-1}},$$

where y_t denotes the total quantity of final goods, while y_{it} gives the quantity of intermediate good i in particular. Assuming perfect competition and zero profits, demand for intermediate good i is then given by

$$y_{it} = y_t \left(\frac{P_{it}}{P_t} \right)^{-\theta},$$

where P_{it} denotes the nominal price of the intermediate good in question, with

$$P_t = \left(\int_0^1 P_{it}^{1-\theta} di \right)^{\frac{1}{1-\theta}}.$$

Each intermediate good is supplied on a monopolistically competitive basis using a technology of the form

$$y_{it} = k_{it}^{\alpha_k} (A_t \ell_{it})^{\alpha_\ell}, \;\; \alpha_k \in (0,1), \;\; \alpha_\ell := 1 - \alpha_k. \tag{5}$$

Intermediate good producers are assumed to face Rotemberg-style nominal frictions — more specifically, their real net profits in a given period read as

$$y_t \left(\frac{P_{it}}{P_t} \right)^{-\theta} \left(\frac{P_{it}}{P_t} - MC_t \right) - y_t \cdot \frac{\varphi}{2} \left(\frac{P_{it}}{P_{i,t-1}} - \Pi^* \right)^2,$$

where Π^* denotes the central bank's inflation target, while MC_t denotes real marginal costs — i.e.,

$$MC_t = \frac{1}{Z_t} \left(\frac{r_t^k}{\alpha_k} \right)^{\alpha_k} \left(\frac{w_t/A_t}{\alpha_\ell} \right)^{\alpha_\ell}.$$

Standard arguments then yield a non-linear Phillips curve of the form

$$\varphi \left(\Pi_t - \Pi^* \right) \Pi_t = (1-\theta) + \theta MC_t + \mathbb{E}_t \left[m_{t+1} \cdot \varphi \left(\Pi_{t+1} - \Pi^* \right) \Pi_{t+1} \cdot \frac{y_{t+1}}{y_t} \right],$$

along with the following market-clearing condition:

$$\underbrace{y_t - y_t \cdot \frac{\varphi}{2} \left(\Pi_t - \Pi^*\right)^2 - k_{t-1} \frac{\psi}{2} \left(\frac{i_t}{k_{t-1}} - \frac{\widetilde{i}}{k} \right)^2}_{=:GDP_t} = c_t + i_t.$$

2.3 Central bank

A central bank sets nominal rates according to a truncated Taylor rule of the form

$$R_t = \max \left\{ 1, R^* \left(\frac{\Pi_t}{\Pi^*} \right)^{\phi_\Pi} \left(\frac{GDP_t / A_t}{\widetilde{GDP}} \right)^{\phi_{GDP}} \right\},$$

where R^* denotes the nominal rate consistent with closure of the output and inflation gaps, while $\widetilde{GDP}$ denotes the value that the ratio GDP_t/A_t takes on the economy's balanced-growth path.

2.4 Long-run risk

We follow Guerron-Quintana and Jinnai (2015), Kung (2015), Kung and Schmid (2015), Comin et al. (2017), and several others in introducing long-run risk through an endogenous-growth mechanism. More specifically, we follow the especially tractable precedent set by the first two of these four references in assuming that a learning-by-doing mechanism links investment with positive spillovers — specifically, $A_t = k_{t-1}$. The stationary system which then arises after detrending is given in section B of the appendix.

2.5 Shocks

To enable the model to admit ELB episodes broadly consistent with recent experience, we find it useful to divide the demand shock ξ_t into distinct short- and medium-run components: $\xi_t = \xi_t^{SR} \xi_t^{MR}$. The short-run component follows an AR(1) process similar to that in, e.g., Gourio and Ngo (2017):

$$\log \xi_t^{SR} = \rho \log \xi_{t-1}^{SR} + \epsilon_t, \;\; \epsilon_t \overset{iid}{\sim} \mathcal{N}(0, \sigma^2). \tag{6}$$

In contrast, we eschew an AR(1) specification for the medium-run component, since it is well-known that this would result in unrealistically low duration and/or severity for the typical in-model ELB episode, as argued in Richter and Throckmorton (2015) and Coibion et al. (2016), among others. Instead, we adopt the regime-switching process proposed in Coibion et al. (2016), which itself builds on earlier work by Eggertsson and Woodford (2003), Christiano et al. (2011), Werning (2011), Carlstrom et al. (2015), and several others. We have illustrated this process in Figure 1. In any given period, the medium-run component ξ_t^{MR} can take one of two values, $\xi_t^{MR} = 1$ or $\xi_t^{MR} = \underline{\xi}^{MR} < 1$. These respectively correspond to "normal" times and an extended episode of relatively low aggregate demand. We assume that the economy spends most of its time operating under the former of these two regimes. However, in any period when the economy is currently experiencing normal conditions, there's a small probability η with which the next period will witness the onset of a low-demand episode. In this case, the episode is assumed to last some $\overline{T}$ periods, after which conditions revert back to normal. Parameters η, $\overline{T}$, and $\underline{\xi}^{MR}$ thus grant us independent control over the risk, duration, and severity of low-demand spells and will ultimately enable us to model episodes comparable to the US Great Recession, as shown in our next section.

2.6 Asset prices

We consider several asset prices in our analysis, most notably including the equity premium and inflation risk premium. We compute the former as the difference between between the expected return on capital and the real risk free rate, multiplied by a standard leverage factor of three. That is,

$$3 \times \left(R_t^k - R_t^{real}\right),$$

where

$$R_t^k := \mathbb{E}_t \left[\frac{r_{t+1}^k - \frac{i_{t+1}}{k_t} - \frac{\psi}{2}\left(\frac{i_{t+1}}{k_t} - \widetilde{\frac{i}{k}}\right)^2 + Q_{t+1}\left[(1-\delta) + \frac{1}{\widetilde{Q}}\frac{i_{t+1}}{k_t}\right]}{Q_t} \right],$$

with

$$1 = \mathbb{E}_t \left(\frac{m_{t+1}}{\xi_t} \cdot R_t^{real} \right).$$

As for the the inflation risk premium, we follow the approach in Gourio and Ngo (2017), which involves pricing two geometrically-declining consols. The first is a nominal consol which, if purchased at time t, pays one dollar at $t+1$, λ dollars at $t+2$, λ^2 dollars at $t+3$, etc. Its price P_t^{nom} therefore satisfies

$$P_t^{nom} = \mathbb{E}_t \left[m_{t+1} \left(\frac{1 + \lambda P_{t+1}^{nom}}{\Pi_{t+1}} \right) \right],$$

while an otherwise comparable real consol admits price

$$P_t^{real} = \mathbb{E}_t \left[m_{t+1} \left(1 + \lambda Q_{t+1}^{nom} \right) \right].$$

The implied break-even inflation rate ($BEIR_t$) for these consols is thus given by the difference between their log-yields,

$$BEIR_t = \log\left(\frac{1}{P_t^{nom}} + \lambda \right) - \log\left(\frac{1}{P_t^{real}} + \lambda \right),$$

which can then be decomposed into expected log-inflation (ELI_t) and an inflation risk premium (IRP_t). In particular:[1]

$$ELI_t = (1-\lambda)\mathbb{E}_t[\log(\Pi_{t+1})] + \lambda \mathbb{E}_t\left(ELI_{t+1}\right)$$

$$IRP_t = BEIR_t - ELI_t.$$

3 Calibration and solution method

To maintain continuity with previous literature, we adopt a calibration very similar to those in Gourio and Ngo (2017) and Rudebusch and Swanson (2012), respectively "GN" and "RS" herein. See Table 1 for a summary.

With regard to household preferences, we respectively set the elasticity of substitution and Frisch elasticity of labour to $1/\sigma = 0.5$ and $1/\nu = 2/3$, consistent with microevidence from, e.g., Vissing-Jørgensen (2002) and Pistaferri (2003). As for the Epstein-Zin parameter γ, we discipline it via the in-model equity premium described in our previous section. In particular, we choose γ to place the average equity premium around 3.1%, toward the lower end

[1] Though we focus on the inflation *risk* premium, very similar results would obtain if we instead focused on the inflation *term* premium, which would first involve deriving the risk-neutral prices

$$Q_t^{nom} = \mathbb{E}_t\left(\frac{m_{t+1}}{\Pi_{t+1}}\right)\mathbb{E}_t\left(1+\lambda P_{t+1}^{nom}\right),$$

and

$$Q_t^{real} = \mathbb{E}_t\left(m_{t+1}\right)\mathbb{E}_t\left(1+\lambda Q_{t+1}^{real}\right),$$

then computing the difference between the nominal term premium

$$\log\left(\frac{1}{P_t^{nom}}+\lambda\right) - \log\left(\frac{1}{Q_t^{nom}}+\lambda\right)$$

and real term premium

$$\log\left(\frac{1}{P_t^{real}}+\lambda\right) - \log\left(\frac{1}{Q_t^{real}}+\lambda\right).$$

of the range normally considered in the asset-pricing literature. The implied coefficient of relative risk aversion, computed using the approach in Swanson (2018), is about 113, very close to the best-fit calibration in RS. The only remaining preference parameter is then households' subjective discount factor, which we select to imply a 3% real rate along the economy's balanced-growth path, though precautionary savings in our full model have the upshot of placing the average real rate close to 2%.

Turning to the parameters governing the nominal side of the economy, we set the elasticity of substitution to the relatively standard value $\theta = 7.66$. This implies a markup of 15% along the balanced-growth path, consistent with the range of estimates reported in Basu (1992) and Basu and Fernald (1997). As for the parameter scaling the Rotemberg adjustment cost, we follow GN in setting $\varphi = 238.11$. The oft-cited mapping in, e.g., Ascari and Rossi (2012), associates this choice with a Calvo parameter around 0.85, toward the upper end of the range normally considered in the New Keynesian literature and very close to the estimate recently reported in Del Negro et al. (2015). We also adopt a relatively high value for the Taylor coefficient, $\phi_\Pi = 3.0$, though we note that this value is still low in comparison with the estimate reported in Gust et al. (2017), along with the calibrations often assumed for other globally-solved New Keynesian models (e.g., Hills et al., 2016, Nakata and Tanaka, 2016).

Turning next to the model's remaining technological parameters, we set the capital share and depreciation rate to the standard values $\alpha_k = 1/3$ and $\delta = 0.02$, respectively. We then calibrate the relative price of investment along the balanced-growth path so as to place the economy's balanced-growth rate around 1.5%, consistent with current estimates on potential growth for many developed countries. We also set the investment adjustment cost parameter to $\psi = 5.6$, consistent with Eberly (1997) and Erceg and Lindé (2014).

The only outstanding parameters are those governing the shock processes. With regard to the AR(1) shock, we have assumed a process somewhat less persistent and volatile than that in GN, namely for numeric-stability reasons. As for the medium-run shock, our intention is to calibrate it so as to ensure that the typical flight-to-quality episode is comparable to but less severe than the US Great Recession. For this reason, Figure 2 plots US GDP per capita over the period 2000–2017, along with the level of real GDP per capita that would have obtained had the economy avoided a recession and instead continued growing at the average rate observed the 2000–2007 period. Figure 3 plots the Federal funds rates over the same period. From these figures, we see that the recession was associated with (i) a peak-to-trough reduction in real GDP per capita around 5.5%, (ii) an 8.1% deviation from the pre-crisis trend after three years, and (iii) a nearly seven-year spell during which the Federal funds rate was 25 basis points or lower. In comparison, we have chosen the parameters governing the severity and duration of flight-to-quality episodes (respectively $\underline{\xi}^{MR}$ and $\overline{T}$) such that the the average episode is associated with (i) a peak-to-trough reduction in real GDP around 4.5%, (ii) a 1.1% deviation from the pre-shock path after three years, and (iii) a roughly eight-quarter period during which the ELB binds. Note that this entails $\overline{T} = 12$ in particular. As for the parameter governing the ex-ante risk of low-demand episodes, we have set $\eta = 0.9\%$, implying an unconditional probability of being at the ELB around 6.6%, a figure well within the range of estimates reported in Coibion et al. (2016) on the basis of post-war experience in a sample of developed countries.

For a given set of parameter values, we solve the model using an approach very similar to that in Fernandez-Villaverde et al. (2015). In particular, we first solve for macroeconomic outcomes while ignoring the bond prices described in subsection 2.6. We then solve for those prices taking the macroeconomic solution as given, with the decline rate λ chosen to place duration around five years along the economy's balanced-growth path.

4 Results

In this section, we first provide some intuition for the model's mechanism by examining the response of the economy to shocks. We then turn to the impact of the ELB on the longer-run behavior of the economy, especially with respect to inflation.

4.1 Model dynamics

Figure 4 illustrates the economy's response to a medium-run demand shock — that is, an unexpected shift from $\xi_t^{MR} = 1$ to $\xi_t^{MR} = \underline{\xi}^{MR}$ lasting $\overline{T} = 12$ quarters. We compute this response by taking differences between this scenario and one where no such episode occurs, holding the short-run shock ξ_t^{SR} constant at its steady-state value of one. Looking first at the baseline model's response in solid blue, we see that the medium-run demand shock drives nominal rates to the ELB for a full eight quarters, during which the economy experiences substantial deflation, along with significant declines in (productivity-normalized) GDP and its components. The endogenous-growth mechanism also links these episodes with slower growth and, by extension, permanent level shifts in (unnormalized) GDP and the continuation value for households.

In contrast, the red dashed lines in Figure 4 consider an otherwise comparable model that abstracts from the ELB. In this case, the central bank's ability to provide greater monetary accommodation results in a lower real rate and more muted impacts on economic activity, growth, and household continuation values.

Figure 5 repeats for a negative short-run demand shock, initializing from $\xi_t^{SR} = 1$ with the medium-run shock ξ_t^{MR} now held constant at its "normal" value of one. All of the variables under consideration exhibit standard responses, with relatively modest long-run effects via the endogenous-growth channel. That said, the comparison between the baseline model in solid blue and the no-ELB version of the model in dashed red suggests that the ELB plays a modestly amplifying role despite its remaining lax for the full duration of the exercise.

Broadly speaking, these findings associate the medium-run demand shock with a significant amount of long-run risk to the extent that it exerts a permanent effect on the level of economic activity and household continuation values. Agents' forward-looking efforts to insure themselves against this risk—and the signature that those efforts leave on macroeconomic outcomes during "normal" times when the ELB is lax—are the subject of our next subsection.

4.2 Disinflationary effects of ELB risk

To gauge the macroeconomic effects of ELB risk, Table 2 reports various moments for different versions of the model, all computed based on 100,000-quarter simulations. More specifically, the top panel in the table reports unconditional averages, while the middle and bottom panels restrict attention to periods during which the ELB is lax and binding, respectively. In the left-most column, the table reports the results from the baseline model, while the middle columns display the findings from otherwise comparable models that abstract from the ELB and recursive preferences. The right-most column presents the outcomes in the deterministic steady state.

Focusing first on the baseline model, the table shows the presence of a significant disinflationary bias. Compared to the deterministic steady state, in which the inflation rate is 2 percent, the unconditional inflation rate in the baseline model is only 1.19 percent. This unconditional disinflationary bias is partly driven by strong deflation when monetary policy is at the ELB, in which case the deflation rate reaches 1.56 percent. However, inflation, at 1.39 percent, is also substantially lower than the central bank's 2 percent target even when we restrict attention to periods when the ELB constraint is not currently binding. Given the relatively conservative nature of our calibration, we interpret these results as a lower bound on the disinflationary effects of ELB risk. We also note that the inflation risk premium is negative, reaching -0.47 percent.

The ELB constraint thus has macroeconomic repercussions even when it is not currently binding. This partly stems from firms' anticipation that substantial deflation would occur should the ELB bind in the future, leading them to set lower prices when operating under normal conditions, namely due to nominal rigidity.

Recursive preferences tend to enhance this channel to the extent that they associate a larger stochastic discount factor with future states in which the ELB is expected bind. This can be seen by comparing the baseline model against an otherwise comparable model which abstracts from recursive preferences. In particular, we see that the average degree of inflation undershooting during normal times is about three times as large with recursive preferences. Note also that the inflation risk premium is only -0.03 percent.

Table 2 also indicates that the key features of our model place significant downward pressure on real interest rates. For example, the average real rate observed during periods when the ELB is lax is forty to fifty basis points lower under the baseline version of the model, as compared against versions which abstract from either the ELB or recursive preferences. This is a natural consequence of the tendency for these two model features to enhance households' precautionary-savings incentives and has obvious bearing on the debate surrounding the low r-star estimates recently produced by, e.g., Holston, Laubach, and Williams (2016) — in particular, it suggests that frameworks which abstract from these model features are liable to overstate the level to which real rates should be expected to return in the long run. Stronger precautionary-savings incentives under the baseline version of the model also manifest in the form of a moderately higher level of investment and, by extension, somewhat faster growth.

Finally, we assess the robustness of our results along several important dimensions and report those results in Table 3. In particular, we examine the sensitivity of the real interest rate and the inflation risk premium with respect to changes in the inflation coefficient in the Taylor rule, the inflation target, the severity of ELB episodes, the degree of price rigidity, and the risk of medium-run demand shocks.

Our baseline calibration assumes that monetary policymakers place a weight of 3 on inflation in the Taylor rule. In the column labelled "Higher ϕ_Π", we raise this coefficient to 4. The greater emphasis on reducing the departures of inflation from target leads to a more accommodative monetary policy in response to medium-run demand shocks, which translates into less precautionary savings and a higher real interest rate than in our baseline model. In turn, the inflation risk premium is somewhat less negative.

Raising the inflation target to 3 percent leads to a similar outcome, as can be seen in the column labelled "Higher Π^*". Given that the remainder of the baseline calibration is unchanged, raising the inflation target implies that the average time spent at the ELB declines and the associated drop in output is more muted. As a result, the disinflationary and precautionary-savings effects discussed above are weaker than in our baseline model.

In the column labelled "Higher $\underline{\xi}^{MR}$," we also examine the role played by the magnitude of the medium-run demand shock. In this exercise, we adjust $\underline{\xi}^{MR}$ so that the average ELB episode is associated with a peak-to-trough drop in real GDP around 1%, as opposed to the 4.5% drop implied by our baseline calibration. Despite this more muted decline in economic activity, ELB risk continues to imply a negative inflation risk premium even when the economy is away from the ELB.

In the penultimate column ("Higher φ"), we consider a higher degree of price rigidity and increase the coefficient controlling price adjustments by ten per cent. Since prices are now more rigid, the ELB episode entails a more muted decline in inflation, implying a less negative inflation risk premia. Associated with these effects is a decline in precautionary savings, resulting in a higher real interest rate. Overall, this exercise highlights the "curse of flexibility" at the ELB.

Finally, the last column of the table ("Lower η") reports results for a case where we reduce the probability of medium-run shocks. In particular, we halve this probability from 0.9% under the baseline parameterization to 0.45%. From the table, we see that our results are little changed by this alternative calibration.

5 Conclusion

Motivated by ongoing weakness in inflation in the aftermath of the Global Financial Crisis, we examined the role of ELB risk in the inflation process. We did by augmenting an otherwise standard New Keynesian model to include three key features: recursive preferences, long-run risk (via endogenous growth), and a demand-shock process that is broadly consistent with recent experience at the ELB. We found that the prospect of future ELB episodes triples the degree of disinflationary bias that should be expected during "normal" times when the ELB is lax, suggesting that it may be difficult for central banks to achieve their inflation targets in the post-Great Recession era.

Looking ahead, our results also lead to some significant questions for monetary policy. Should central banks adjust their policy rules to eliminate this systematic undershooting of the inflation objective? Is there some simple way for monetary authorities to divorce their implicit targets from their explicit objectives? What are the implications for optimal inflation? We plan on exploring these and other questions in future research.

References

Amano, R. and Shukayev, M. (2012). Risk premium shocks and the zero bound on nominal interest rates. *Journal of Monetary, Credit, and Bank-*

ing, 44(8):1475–1505.

Anzoategui, D., Comin, D., Gertler, M., and Martinez, J. (2017). Endogenous technology adoption and R&D as sources of business cycle persistence. Mimeo. Available at `https://www.dartmouth.edu/~dcomin/Publications_files/ACGM.pdf`.

Ascari, G. and Rossi, L. (2012). Trend inflation and firms price-setting: Rotemberg vs. Calvo. *Economic Journal*, 122(563):1115–1141.

Basu, K. (1992). A geometry for non-walrasian general equilibrium theory. *Journal of Macroeconomics*, 14(1):87–103.

Basu, S. and Fernald, J. G. (1997). Returns to scale in U.S. production: Estimates and implications. *Journal of Political Economy*, 105(2):249–83.

Carlstrom, C. T., Fuerst, T. S., and Paustian, M. (2015). Inflation and output in New Keynesian models with a transient interest rate peg. *Journal of Monetary Economics*, 76:230–243.

Christiano, L. J., Eichenbaum, M., and Rebelo, S. (2011). When is the government spending multiplier large? *Journal of Political Economy*, 119(1):78–121.

Coibion, O., Dordal-i-Carreras, M., Gorodnichenko, Y., and Wieland, J. (2016). Infrequent but long-lived zero lower bound episodes and the optimal rate of inflation. *Annual Review of Economics*, 8:497–520.

Coibion, O., Gorodnichenko, Y., and Wieland, J. (2012). The optimal inflation rate in New Keynesian models: should central banks raise their inflation targets in light of the zero lower bound? *Reviiew of Economic Studies*, 79(4):1371–1406.

Comin, D., Gertler, M., Ngo, P., and Santacreu, A. (2017). Stock price fluctuations and productivity growth. Mimeo. Available at http://academic.csuohio.edu/ngop/CGPS082216.pdf.

Del Negro, M., Giannoni, M. P., and Schorfheide, F. (2015). Inflation in the Great Recession and New Keynesian models. *American Economic Journal: Macroeconomics*, 7(1):168–96.

Eberly, J. (1997). International evidence on investment and fundamentals. *European Economic Review*, 41:297–308.

Eggertsson, G. B. and Woodford, M. (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 34(1):139–235.

Erceg, C. and Lindé, J. (2014). Is there a fiscal free lunch in a liquidity trap? *Journal of the European Economics Association*, 12:915–44.

Fernandez-Villaverde, J., Gordon, G., Guerron-Quintana, P., and Rubio-Ramirez, J. F. (2015). Nonlinear adventures at the zero lower bound. *Journal of Economic Dynamics & Control*, 57:182–204.

Gourio, F. and Ngo, P. (2017). Risk premia at the ZLB: a macroeconomic interpretation. Mimeo. Available at https://dornsife.usc.edu/assets/sites/524/docs/draftJan2017.pdf.

Guerron-Quintana, P. A. and Jinnai, R. (2015). Liquidity shocks and asset prices. Mimeo. Available at https://sites.google.com/site/pabloaguerronquintana/liquid_price.pdf.

Gust, C., Herbst, E., Lopez-Salido, D., and Smith, M. (2017). The empirical implications of the interest-rate lower bound. *American Economic Journal: Macroeconomics*, 107(7):1971–2006.

Hills, T., Nakata, T., and Schmidt, S. (2016). The risky steady state and the interest rate lower bound. Finance and economics discussion series, divisions of research & statistics and monetary affairs, Federal Reserve Board 2016-009.

Kung, H. (2015). Macroeconomic linkages between monetary policy and the term structure of interest rates. *Journal of Financial Economics*, 115(1):42–57.

Kung, H. and Schmid, L. (2015). Innovation, growth, and asset prices. *Journal of Finance*, 70(3):1001–1037.

Nakata, T. and Tanaka, H. (2016). Equilibrium yield curves and the interest rate lower bound. Finance and economics discussion series, divisions of research & statistics and monetary affairs, Federal Reserve Board 2016-085.

Pistaferri, L. (2003). Anticipated and unanticipated wage changes, wage risk, and intertemporal labor supply. *Journal of Labor Economics*, 21(3):729–54.

Richter, A. W. and Throckmorton, N. A. (2015). The zero lower bound: frequency, duration, and numerical convergence. *The B.E. Journal of Macroeconomics*, 15(1):1–26.

Rudebusch, G. and Swanson, E. (2012). The bond premium in a DSGE model with long-run real and nominal risk. *American Economic Journal: Macroeconomics*, 4(1):105–143.

Vissing-Jørgensen, A. (2002). Limited asset participation and the elasticity of intertemporal substitution. *Journal of Political Economy*, 110(4):825–53.

Werning, I. (2011). Managing a liquidity trap: monetary and fiscal policy. NBER working paper 17344.

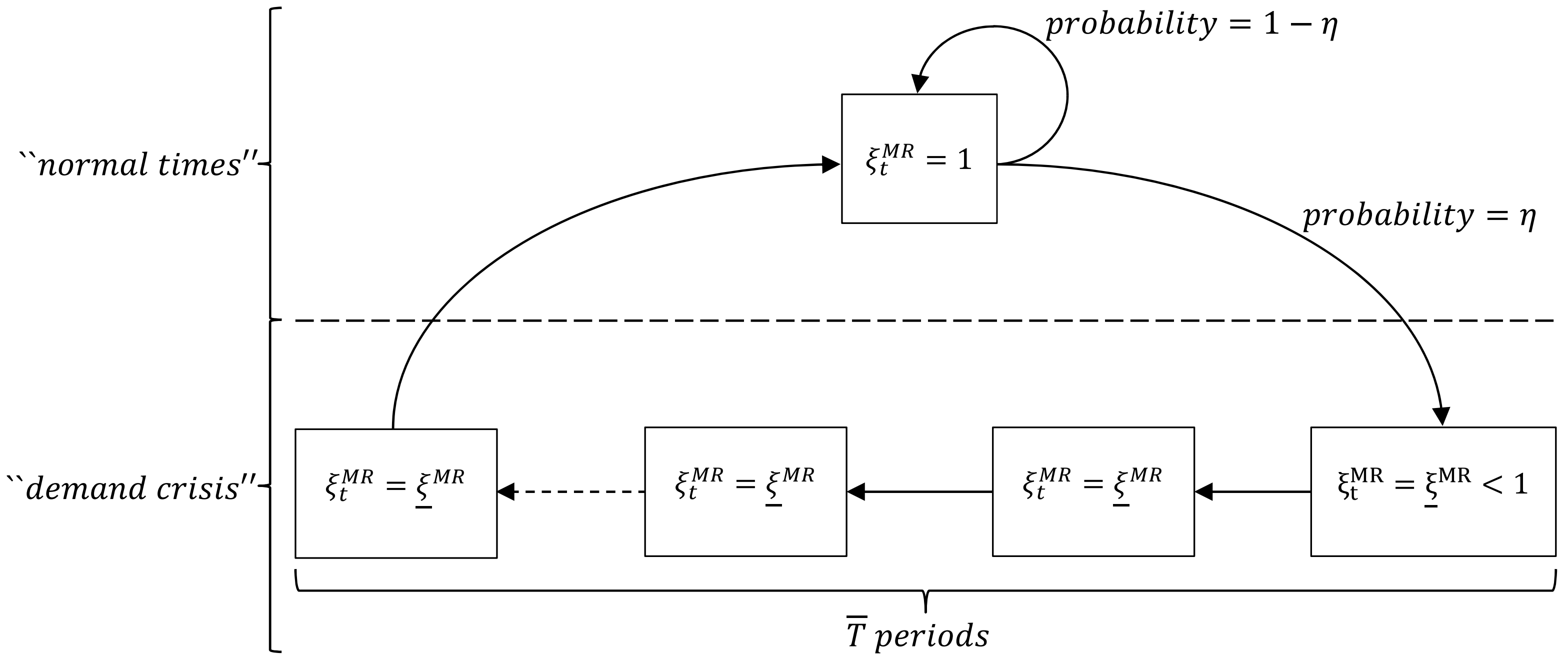

Figure 1: Demand-shock process

Table 1: Baseline calibration

Parameter	Value	Description and notes
σ	2	Inverse IES, set to value in GN and RS; see also Vissing-Jørgensen (2002)
ν	1.5	Inverse Frisch elasticity of labour, set to value in GN and RS; see also Pistaferri (2003)
χ	33.4	Multiplicative constant in the disutility of labour, chosen to set labour supply along BGP = 1/3, as in GN
γ	-158	Epstein-Zin parameter, set to place the average equity premium around 3.1% (implied CRRA = 113, close to RS)
β	0.9999	Discount factor, set to place the average real rate around 2%
θ	7.66	Elasticity of substitution across intermediate goods, set to value in GN and Fernandez-Villaverde et al. (2015); see also Basu (1992) and Basu and Fernald (1997)
φ	238.11	Rotemberg parameter, set to value in GN, implies Calvo parameter ≈ 0.85; see also Del Negro et al. (2015)
ϕ_{Π}	3.00	Taylor coefficient, set to value in GN; c.f. Hills et al. (2016), Nakata and Tanaka (2016), and Gust et al. (2017)
ϕ_{GDP}	0.13	Weight on the output gap in the Taylor rule, set to value in GN
α_k	1/3	Capital share
δ	0.02	Depreciation rate
ψ	5.6	Capital adjustment cost; c.f. Eberly (1997) and Erceg and Lindé (2014)
$\widetilde{Q}$	5.06	Relative price of investment on the BGP, chosen to set growth rate along = 1.5% p.a.
ρ	0.725	Persistence of short-run shocks
σ	0.1%	Standard deviation of productivity shocks
$\underline{\xi}^{MR}$	0.99	Size of the medium-run shock; see main text for details
η	0.9%	Risk of flight-to-quality episodes; see main text for details
$\overline{T}$	12	Duration of flight-to-quality episodes; see main text for details

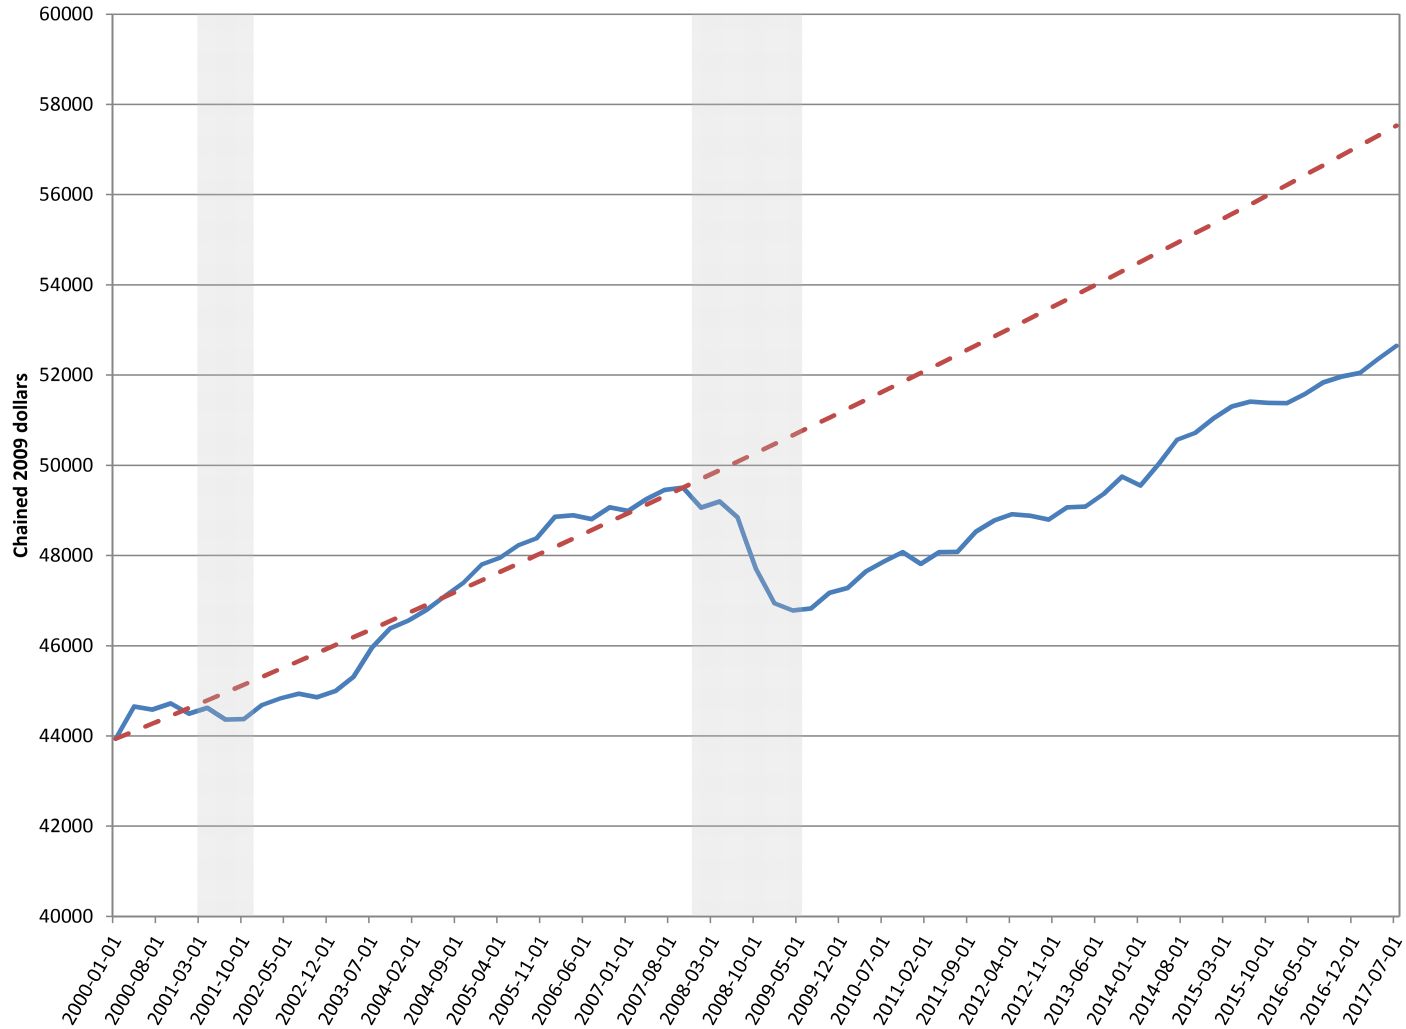

Figure 2: GDP per capita (solid blue) and pre-crisis trend (dashed red), 2000–17

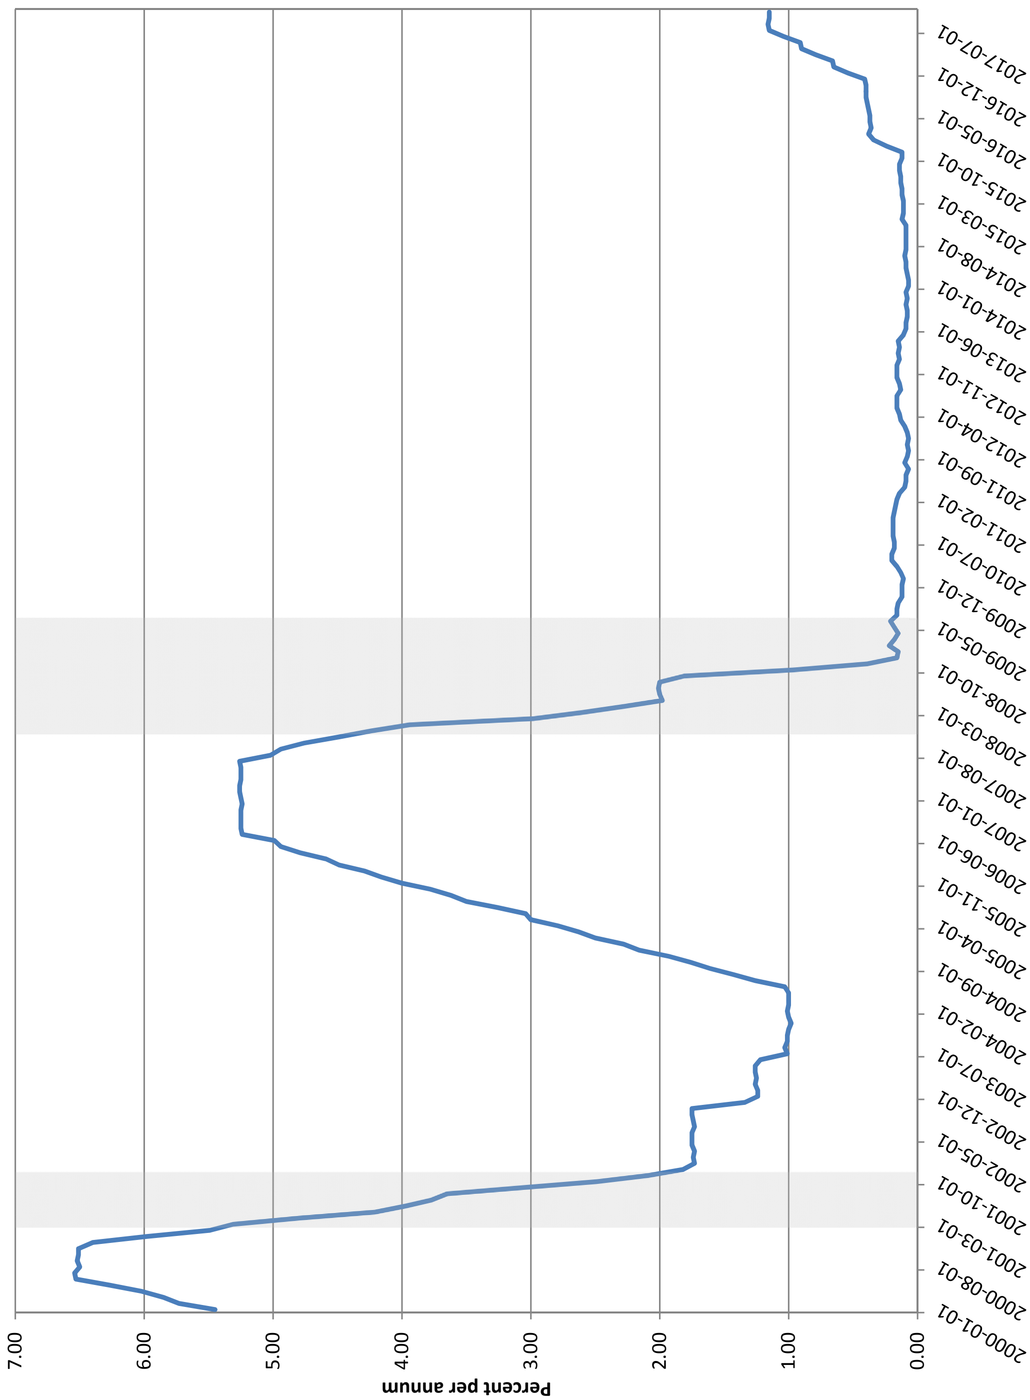

Figure 3: Fed funds rate, 2000–17

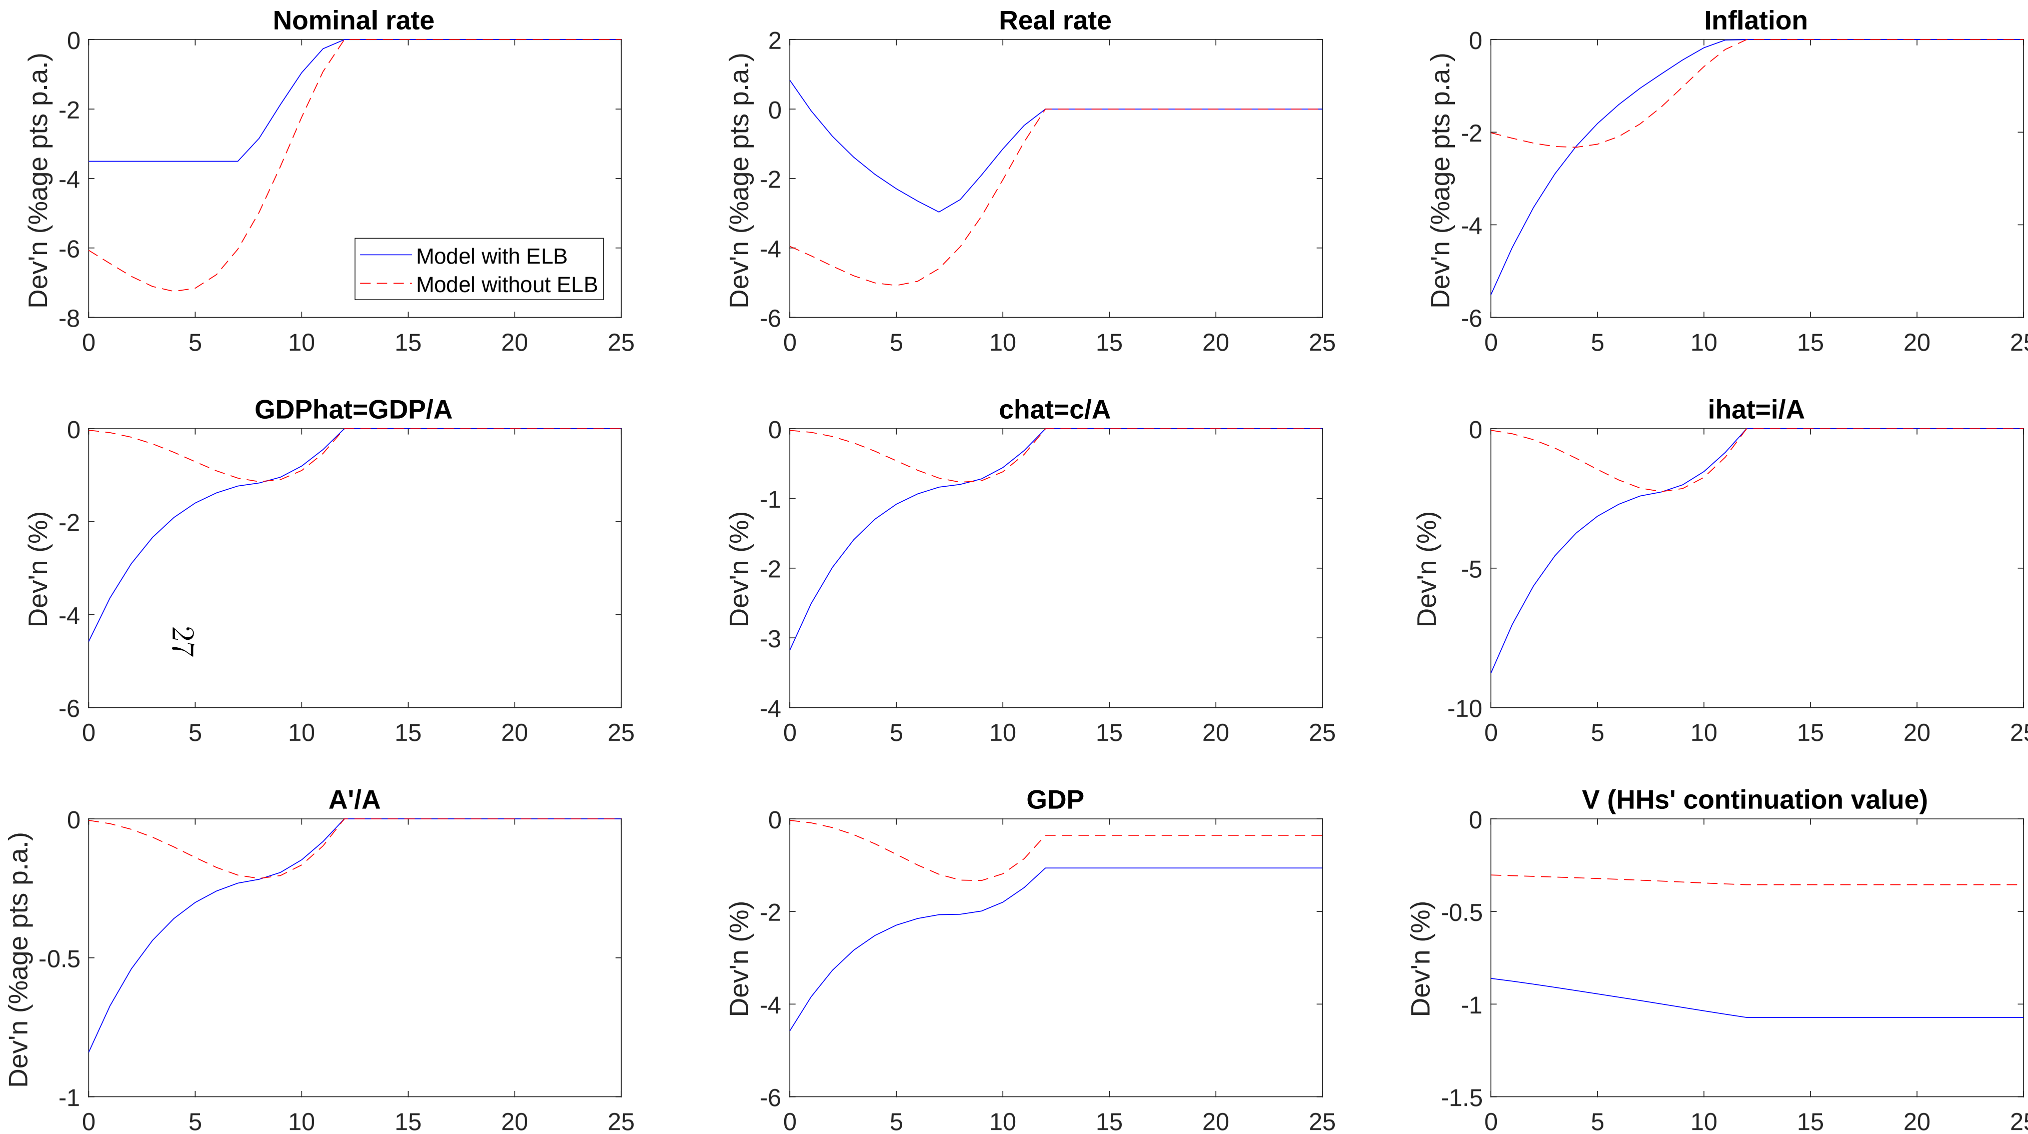

Figure 4: IRFs: medium-run demand shock

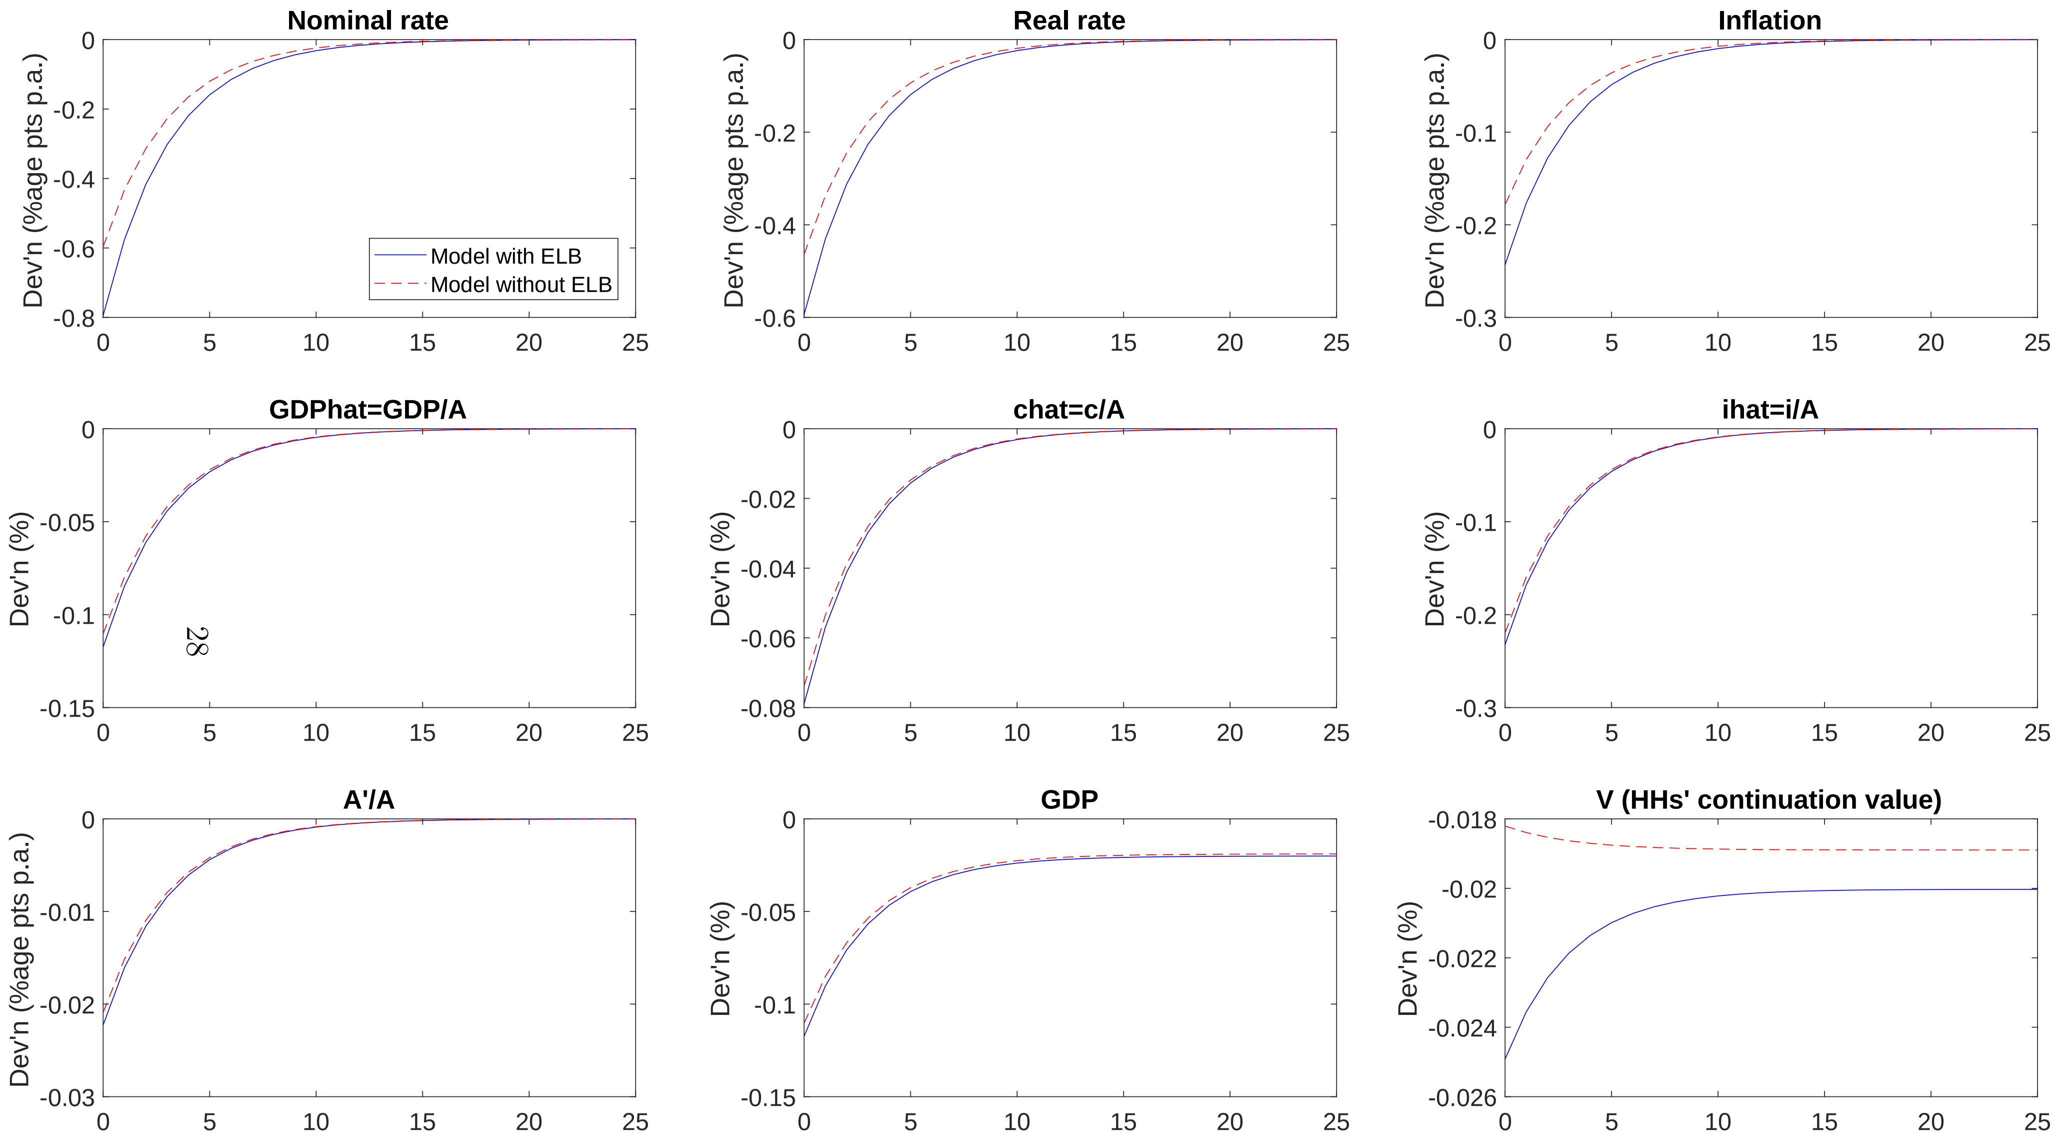

Figure 5: IRFs: negative short-run demand shock

Table 2: First moments

		Baseline	No ELB	Non-rec. pref's	Deterministic steady state
Unconditional	Nominal rate (net percent p.a.)	3.21	4.43	4.15	5.02
	Real rate (net percent p.a.)	2.16	2.64	2.64	3.02
	Inflation (net percent p.a.)	1.19	1.79	1.51	1.99
	Inflation risk premium (net percent p.a.)	-0.47	-0.11	-0.04	0.00
	Normalized investment (ihat=i/A)	0.121	0.120	0.120	0.120
	Productivity growth (net percent p.a.)	1.55	1.50	1.49	1.49
ELB lax	Nominal rate (net percent p.a.)	3.44	4.43	4.47	5.02
	Real rate (net percent p.a.)	2.24	2.64	2.74	3.02
	Inflation (net percent p.a.)	1.39	1.79	1.78	1.96
	Inflation risk premium (net percent p.a.)	-0.47	-0.11	-0.03	0.00
	Normalized investment (ihat=i/A)	0.121	0.120	0.120	0.120
	Productivity growth (net percent p.a.)	1.58	1.50	1.53	1.49
ELB tight	Nominal rate (net percent p.a.)	0	N/A	0	N/A
	Real rate (net percent p.a.)	0.97	N/A	1.35	N/A
	Inflation (net percent p.a.)	-1.56	N/A	-2.04	N/A
	Inflation risk premium (net percent p.a.)	-0.47	N/A	-0.16	N/A
	Normalized investment (ihat=i/A)	0.12	N/A	0.11	N/A
	Productivity growth (net percent p.a.)	1.12	N/A	1.02	N/A

Table 3: Sensitivity analysis

		Baseline	Higher ϕ_Π	Higher Π^*	Higher $\underline{\xi}^{MR}$	Higher φ	Lower η
Unconditional	Real rate (net percent p.a.)	2.16	2.38	2.59	2.71	2.18	2.45
	Inflation risk premium (net percent p.a.)	-0.47	-0.33	-0.23	-0.12	-0.44	-0.44
ELB lax	Real rate (net percent p.a.)	2.24	2.53	2.79	2.85	2.28	2.49
	Inflation risk premium (net percent p.a.)	-0.47	-0.33	-0.21	-0.12	-0.44	-0.44
ELB tight	Real rate (net percent p.a.)	0.97	0.40	-0.74	0.06	0.83	1.17
	Inflation risk premium (net percent p.a.)	-0.47	-0.38	-0.43	-0.24	-0.44	-0.47

APPENDIX

A Households optimality conditions

Households' behaviour can be summarized by (3) and (4) in the main text, along with three additional equations, namely the supply function for labour,

$$\chi A_t^{1-\sigma} \ell_t^{\nu} = w_t c_t^{-\sigma}; \tag{7}$$

the Euler equation for capital,

$$1 = \mathbb{E}_t \left[m_{t+1} \cdot \frac{1}{Q_t} \left[\begin{array}{l} r_{t+1}^k - \frac{i_{t+1}}{k_t} - \frac{\psi}{2} \left(\frac{i_{t+1}}{k_t} - \frac{\widetilde{i}}{k} \right)^2 \\ \qquad\qquad + Q_{t+1} \left[(1-\delta) + \frac{1}{\widetilde{Q}} \frac{i_{t+1}}{k_t} \right] \end{array} \right] \right]; \tag{8}$$

and a pricing condition pinning down the relative price of capital,

$$Q_t = \widetilde{Q} \left[1 + \kappa \left(\frac{i_t}{k_{t-1}} - \frac{\widetilde{i}}{k} \right) \right]. \tag{9}$$

B Stationary system

Letting $\hat{x} := x_t / A_t \; \forall x \in \{c, i, GDP, w, y\}$, $\hat{V}_t := V_t / A_t^{1-\sigma}$, and $G_t^k := k_t / k_{t-1}$, the relevant system reads as follows:

$$\hat{V}_t = a \left(\frac{\hat{c}_t^{1-\sigma}}{1-\sigma} - \frac{\chi \ell_t^{1+\nu}}{1+\nu} \right) - \beta (G_t^k)^{1-\sigma} [\mathbb{E}_t[(-\hat{V}_{t+1})^{1-\gamma}]]^{\frac{1}{1-\gamma}} \tag{10}$$

$$\chi \ell_t^{\nu} = \hat{w}_t \hat{c}_t^{-\sigma} \tag{11}$$

$$m_{t+1} = \beta \left(\frac{\hat{c}_{t+1} G_t^k}{\hat{c}_t} \right)^{-\sigma} \left[\frac{-\hat{V}_{t+1}}{[\mathbb{E}_t[(-\hat{V}_{t+1})^{1-\gamma})]^{\frac{1}{1-\gamma}}} \right]^{-\gamma} \tag{12}$$

$$1 = \mathbb{E}_t \left(\frac{m_{t+1}}{\xi_t} \cdot \frac{R_t}{\Pi_{t+1}} \right) \tag{13}$$

$$1 = \mathbb{E}_t \left[m_{t+1} \cdot \frac{1}{Q_t} \left[\begin{array}{r} r^k_{t+1} - \hat{i}_{t+1} \; - \; \frac{\psi}{2} \left(\hat{i}_{t+1} - \frac{\widetilde{i}}{k} \right)^2 \\ \\ + \; Q_{t+1} \left[(1-\delta) + \frac{\hat{i}_{t+1}}{\widetilde{Q}} \right] \end{array} \right] \right] \tag{14}$$

$$Q_t = \widetilde{Q} \left[1 + \kappa \left(\hat{i}_t - \frac{\widetilde{i}}{k} \right) \right]$$

$$\frac{\hat{w}_t \ell_t}{\alpha_\ell} = \frac{r^k_t}{\alpha_k} \tag{15}$$

$$MC_t = \left(\frac{r^k_t}{\alpha_k} \right)^{\alpha_k} \left(\frac{\hat{w}_t}{\alpha_\ell} \right)^{\alpha_\ell} \tag{16}$$

$$\varphi(\Pi_t - \Pi^*)\Pi_t = (1-\theta) + \theta MC_t + \mathbb{E}_t \left[m_{t+1} \cdot \varphi \left(\Pi_{t+1} - \Pi^* \right) \Pi_{t+1} \cdot \frac{\hat{y}_{t+1} G^k_t}{\hat{y}_t} \right] \tag{17}$$

$$\hat{y}_t = \ell_t^{\alpha_\ell} \tag{18}$$

$$G\hat{D}P_t = \hat{y}_t - \hat{y}_t \cdot \frac{\varphi}{2} \left(\Pi_t - \Pi^* \right)^2 - \frac{\psi}{2} \left(\hat{i}_{t+1} - \frac{\widetilde{i}}{k} \right)^2 \tag{19}$$

$$\hat{c}_t + \hat{i}_t = G\hat{D}P_t \tag{20}$$

$$G^k_t = (1-\delta) + \frac{\hat{i}_t}{\widetilde{Q}} \tag{21}$$

$$R_t = \max\left\{1, R^* \left(\frac{\Pi_t}{\Pi^*}\right)^{\phi_\Pi} \left(\frac{G\hat{D}P_t}{\widetilde{GDP}}\right)^{\phi_{GDP}}\right\} \tag{22}$$