FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Precautionary Pricing: The Disinflationary Effects of ELB Risk

Robert Amano
Bank of Canada

Thomas J. Carter
Bank of Canada

Sylvain Leduc
Federal Reserve Bank of San Francisco

October 2019

Working Paper 2019-26

https://www.frbsf.org/economic-research/publications/working-papers/2019/26/

Suggested citation:

Amano, Robert, Thomas J. Carter, and Sylvain Leduc. 2019. "Precautionary Pricing: The Disinflationary Effects of ELB Risk," Federal Reserve Bank of San Francisco Working Paper 2019-26. https://doi.org/10.24148/wp2019-26

Precautionary Pricing: The Disinflationary Effects of ELB Risk *

Robert Amano† Thomas J. Carter‡ Sylvain Leduc§

October 15, 2019

Abstract

We construct a model to evaluate the role that the risk of future effective lower bound (ELB) episodes plays as a factor behind the persistently weak inflation witnessed in many advanced economies since the Great Recession. In our model, a range of precautionary channels cause ELB risk to affect inflation and other macroeconomic outcomes even during "normal times" when nominal rates are far away from the ELB. This behavior is enhanced through a growth channel that captures possible long-lasting output declines at the ELB. We show that ELB risk substantially weighs on inflation even when the policy rate is above the ELB. Our model also predicts substantially below-target inflation expectations and negative inflation risk premia.

*We thank David Lopez-Salido, Sebastian Schmidt, and seminar and conference participants at the Bank of Canada, Ohio State University, the Riksbank, the Swiss National Bank, the 2018 CEBRA annual meetings, and the 2018 CEF annual meetings. The views expressed are the authors' and do not necessarily reflect those of the Bank of Canada, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

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

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

§Federal Reserve Bank of San Francisco, San Francisco, CA, United States. *Email*: sylvain.leduc@sf.frb.org

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 has led many policymakers and academics to speculate that longer-run forces may be at work. For example, toward the end of her tenure, Federal Reserve Chair Janet Yellen, referring to weakness in U.S. inflation, noted 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."[1]

In this paper, we use a quantitative framework to argue that one of the "long-lasting" forces driving weak inflation in the post-Great Recession era may be the spectre of a return to the effective lower bound (ELB) on nominal interest rates. Our model features occasional "crisis" episodes during which the ELB constrains monetary policy for an extended period of time, along with a strong growth channel linking these episodes with long-lasting output declines such as the one witnessed during and after the Great Recession. We then highlight a range of precautionary channels via which agents' forward-looking anticipation of future crisis episodes has a significant effect on inflation and other macroeconomic outcomes even in "normal", non-crisis states of the world.

In particular, we highlight a novel "precautionary pricing" channel that leads firms to set substantially lower prices in our framework, relative to an otherwise comparable model that abstracts from the ELB. For example, our benchmark model predicts that the average rate of inflation in "normal" non-crisis states should only be about 1.7 percent, despite the central bank pursuing a 2 percent inflation target. Most of this undershooting can specifically be attributed to the effects of ELB risk, since an otherwise comparable

[1] As quoted in the *Wall Street Journal* (Torry, 2017).

model that abstracts from the ELB predicts that average inflation should be very close to target in normal times.

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 risk aversion. We specifically do so by augmenting an otherwise typical New Keynesian framework to include three key ingredients: (i) high risk aversion via 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 place risk aversion within the range considered in the macro-finance literature, while also ensuring that the typical ELB episode is roughly in line with recent experience in terms of duration and severity. The ELB also binds with a frequency 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 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 return output to its potential level, 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 being equal, partly through the above-noted precautionary pricing channel. In addition, this effect is reinforced by the fact that households increase their precautionary savings to buffer against future ELB episodes, leading to weaker aggregate demand.[2] Importantly, we show that high risk aversion helps to amplify the overall disinflationary effects of ELB risk.

[2] Generally speaking, the negative effects of ELB risk share many similarities with those induced by uncertainty shocks in the model of Leduc and Liu (2016).

Our framework is also useful for interpreting the weakness in rates of inflation compensation derived from financial instruments since the Great Recession and determining whether this decline is due to lower inflation expectations or a lower inflation risk premium. Understanding the role of each component is important, particularly since central banks pay close attention to long-run inflation expectations to gauge their credibility and since expectations partly impact current inflation. Inflation compensation rates derived from nominal and inflation-protected securities have declined substantially since the financial crisis. For instance, the inflation compensation rate on 10-year U.S. Treasury securities was 2.4 percent, on average, between March 2004 and November 2007. Between 2012 and August 2019, inflation compensation has declined by roughly 60 basis points.

Consistent with the post-crisis evidence, ELB risk lowers inflation compensation by a similar magnitude. Importantly, this partly reflects expectations of only 1.4 percent inflation in the medium run. In addition, the low inflation compensation rate reflects a -0.25 inflation risk premium. This estimate is broadly in line with Chen et al. (2016) and Grishchenko and Huang (2013), who find that U.S. inflation risk premia were positive in the early 2000s but have since turned negative, beginning around the time of the Great Recession. Similarly, Camba-Mendez and Werner (2017) provide evidence that several measures of U.S. and euro-zone inflation risk premia turned negative after 2008.

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 et al. (2016), which compares the deterministic and risky steady states of an ELB-constrained New Keynesian model. However, they do so assuming standard non-recursive preferences and thus largely abstract from the points raised above regarding the importance of the precautionary forces at play in our analysis. In

related work, Lansing (2018) studies a New Keynesian model in which the ELB occasionally binds as a result of switching between the economy's two steady states. Applying the model to the data, he finds that agents place a significant weight on the possibility of being in the deflationary equilibrium when making their forecasts, thus helping to account for the persistent undershooting of the Federal Reserve's inflation target since mid-2012.

Gourio and Ngo (2017) 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 (2017) highlight a mechanism that causes the inflation risk premium to fall during these periods due to the central bank's inability to offset demand shocks. In contrast, we focus on a complementary mechanism that operates even when the economy is far removed from the ELB. Camba-Mendez 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. Mertens and Williams (2018) use options data to measure the impact of ELB risk on the forecast densities for future U.S. interest rates and inflation. In particular, they find that the forecast density for inflation has shifted to the left since 2012, implying somewhat lower expected inflation.

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 research and development investment played an important role in lowering the economy's growth trajectory and productivity. Similar to their approach, 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. 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: (i) recursive preferences, (ii) a source of long-run risk, and (iii) a demand shock process admitting deep, long-lived 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 elaborate on long-run risk, shocks, and asset prices, respectively.

2.1 Households

A representative household has 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 denote consumption and labor, respectively; σ is the inverse of the intertemporal elasticity of substitution; χ is a parameter scaling the disutility of labor; ν is the inverse of the elasticity of labor supply; A_t denotes trend productivity and has been included to ensure the existence of a balanced-growth path; β is the discount factor; γ is a parameter governing the level of risk aversion; and $\mathbb{E}_t(\cdot)$ denotes an expectation conditional on information available at time t.

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 which we describe in Subsection 2.5. Note that we assume 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 parameter scaling adjustment costs.

In addition, the law of motion for capital takes the form

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

where δ is the depreciation rate and $\widetilde{Q}$ denotes the relative price of capital along the economy's balanced-growth path.

We relegate most of the household's optimality conditions to the appendix and 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} represents 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 "flight-to-quality" effects that induce households to redirect funds from consumption and investment to the relative safety of bonds which offer a guaranteed nominal return. Going forward, we therefore refer to ξ_t as a "demand shock". 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; y_{it} gives the quantity of intermediate good i; and θ is the elasticity of substitution between these

intermediate goods. 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, while

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

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

$$y_{it} = k_{i,t-1}^{\alpha_k} (A_t \ell_{it})^{\alpha_\ell}, \tag{5}$$

where $\alpha_k \in (0,1)$ is the capital share; $\alpha_\ell := 1 - \alpha_k$ is the labor share; and $k_{i,t-1}$ and ℓ_{it} respectively denote firm-specific inputs of capital and labor. 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^* is the central bank's inflation target; the parameter φ scales the cost of adjusting prices; and MC_t denotes real marginal costs — i.e.,

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

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

$$\varphi \left(\Pi_t - \Pi^* \right) \Pi_t = (1-\theta) + \theta MC_t + \mathbb{E}_t \left[m_{t+1} \cdot \varphi \left(\Pi_{t+1} - \Pi^* \right) \Pi_{t+1} \cdot \frac{y_{t+1}}{y_t} \right], \tag{6}$$

along with the following market-clearing condition:

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

2.3 Central bank

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

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

where ϕ_Π governs the central bank's response to deviations from the inflation target Π^*, while ϕ_{GDP} governs its response to deviations of output from the balanced-growth path. The parameter R^* denotes the nominal rate consistent with closure of the output and inflation gaps, while $\widetilde{GDP/A}$ represents 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) and Comin et al. (2017) in introducing long-run risk through an endogenous-growth mechanism. More specifically, we assume that a learning-by-doing mechanism links investment with positive spillovers that individual firms fail to internalize — formally, $A_t = k_{t-1}$.

As shown by Guerron-Quintana and Jinnai (2015), one benefit of this approach to modelling endogenous growth is that it removes the capital stock as a potential state variable once the model has been detrended to account for productivity growth.[3] For instance, after taking account of the externality,

[3] The stationary system which arises after detrending is given in the appendix.

the aggregate production function reads as

$$y_t = k_{t-1}^{\alpha_k}(A_t \ell_t)^{\alpha_\ell} = k_{t-1}\ell_t^{\alpha_\ell} \tag{7}$$

and is thus linear in $k_{t-1} = A_t$. This simplifies our numerical analysis substantially given our use of a global solution method.

2.5 Shocks

To enable the model to admit ELB episodes of significant severity and duration, we divide the demand shock ξ_t (introduced in Subsection 2.1) 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):

$$\log \xi_t^{SR} = \rho \log \xi_{t-1}^{SR} + \epsilon_t, \;\; \epsilon_t \overset{iid}{\sim} \mathcal{N}(0, \sigma^2). \tag{8}$$

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 illustrate 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 "crisis" episode characterized by low aggregate demand. We assume that the economy spends most of its time operating under the normal regime. However, in any normal period, there remains a small probability η of a crisis episode in the following period. 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 allow us to adjust the risk, duration, and

severity of low-demand spells, as shown in our next section.

2.6 Asset prices

We consider several asset prices in our analysis, including the equity and inflation risk premia. We compute the former as the difference 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

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

and

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

As for 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

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

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:

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

$$IRP_t := BEIR_t - ELI_t.$$

Though we focus on the inflation risk premium, we also report results for the inflation term premium, which we compute by first deriving the risk-neutral prices

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

and

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

then taking the difference between the implied nominal and real term premia, respectively, given by

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

and

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

3 Calibration and solution method

In this section, we discuss our calibration and global solution method. Table 1 reports the parameter values used in our benchmark analysis. With

regard to household preferences, we set the elasticity of substitution and Frisch elasticity of labor to $1/\sigma = 0.5$ and $1/\nu = 2/3$ respectively, consistent with microevidence from, e.g., Vissing-Jørgensen (2002) and Pistaferri (2003). We set households' subjective discount factor to imply a 3 percent real interest rate along the economy's balanced-growth path, though precautionary savings lead to a significantly lower average real rate in the full model.

As for the Epstein-Zin parameter γ, we set it such that the implied coefficient of relative risk aversion, computed using the approach in Swanson (2016), is 50 – a figure in line with the range considered in much of the macro-finance literature (e.g., Campbell and Cochrane, 1999; Piassezi and Schneider, 2006; Swanson, 2016; Tallarini, 2000; Van Binsbergen et al., 2012). Moreover, in subsequent sensitivity analysis, we further entertain coefficients of relative risk aversion up to 110 – a level somewhat higher than that considered in Gourio and Ngo (2017) but in line with the estimates reported in Rudebusch and Swanson (2012).

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 percent 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 Gourio and Ngo (2017) in setting $\varphi = 238.11$. The oft-cited mapping in 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 in 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 to place the economy's balanced-growth rate around 1.5 percent, broadly consistent with 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).

As mentioned earlier the parameters $\underline{\xi}^{MR}$, η and $\overline{T}$ influence the severity, probability and duration of the medium-run demand shock, respectively. The first of these parameters has been calibrated so that the typical in-model ELB episode is associated with a peak-to-trough reduction in real GDP of about 5 percent, roughly in line with the 5.5 percent peak-to-trough drop in real GDP per capita that the United States experienced during the ELB episode associated with the Great Recession (see Figures 2 and 3). As for the parameters governing the probability and duration of the medium-run demand shock, these have been jointly calibrated to place the average in-model ELB episode around 8 quarters in length and the unconditional probability of a binding ELB around 7 percent. An 8-quarter average duration is a bit less than the 10-15 quarter range within which Coibion et al. (2016) place the average duration of post-war ELB episodes in a sample of industrialized economies, while an unconditional probability of 7 percent is consistent with the 6-11 percent range that they estimate from the same sample. However, since there is a significant amount of uncertainty around these probabilities, we report, in Subsection 4.5, some sensitivity analysis using alternative probabilities.

Finally, we assume that the short-run demand shock follows a process that is slightly less persistent and volatile than that in Gourio and Ngo (2017), which has the benefit of ensuring greater numerical stability. However, the main results of the paper are not significantly impacted by these assumptions.

As for our solution method, we rely on a global method very similar to that in Fernandez-Villaverde et al. (2015), thus preserving all of the model's non-linearities. More specifically, we first solve for macroeconomic outcomes while ignoring asset prices, then solve for asset prices taking the macroeconomic solution as given.

4 Results

In this section, we first provide some intuition for the model's behavior by examining its response to shocks (subsection 4.1). We then turn to the impact of the ELB on the longer-run behavior of the economy (subsection 4.2) before dissecting the underlying mechanism (subsection 4.3) and exploring its sensitivity to the level of risk aversion (subsection 4.4) and other key parameters (subsection 4.5).

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 $\xi_t^{MR} = \underline{\xi}^{MR} < 1$ lasting $\overline{T} = 12$ quarters. We compute this response by taking the 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, we see that the medium-run demand shock drives nominal rates to the ELB for eight quarters, during which time 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.

Figure 4 also reports responses to the medium-run demand shock in 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 substantially more muted impacts on economic activity, growth, and household continuation values.

Turning to the effects of a negative short-run demand shock, Figure 5 shows that the economy responds in a standard fashion, with output, consumption, and inflation all declining, while monetary policy responds by lowering the nominal interest rate.[4] Growth temporarily falls, in line with the decline in investment, leading to relatively modest long-run effects via the endogenous-growth channel. However, comparing the responses from the benchmark model to those from the no-ELB version of our framework suggests that the ELB plays a modest amplifying role, despite remaining lax for the full duration of the shock.

Overall, these results suggest that the medium-run demand shock generates a significant amount of long-run risk, as it exerts a permanent effect on the level of economic activity and household continuation values, especially in the presence of the ELB. Agents' forward-looking efforts to insure themselves against this risk—and the impact that those efforts leave on outcomes during "normal times" when the medium-run shock remains dormant—are the subject of our next subsection.

4.2 Long-run effects of ELB risk

To gauge the long-run macroeconomic effects of ELB risk, Table 2 reports various moments for different versions of the model, all computed based on 1,000,000-quarter simulations. More specifically, the top panel in the table reports unconditional averages, while the middle and bottom panels respectively restrict attention to "normal times" during which the medium-run shock remains high (i.e., $\xi_t^{MR} = 1$) and "crisis" episodes during which the medium-run shock exerts sustained downward pressure on aggregate demand (i.e., $\xi_t^{MR} = \underline{\xi}^{MR} < 1$). The table reports results for the benchmark model

[4] In this exercise, the medium-run shock ξ_t^{MR} is held constant at its normal value of one.

and for otherwise comparable models under which we abstract from the ELB and/or assume standard non-recursive preferences, thus lowering the coefficient of relative risk aversion close to unity. In addition, we report the steady-state values of all variables of interest. Note that we are particularly interested in the results reported in the middle panel, since they allow us to assess the extent to which ELB risk impacts the economy during normal times.

Focusing first on the benchmark model, the table points to the presence of a significant disinflationary bias. Compared to an inflation rate of 2 percent in deterministic steady state, the unconditional average inflation rate in the benchmark model is only 1.4 percent. This unconditional disinflationary bias is partly driven by the average deflation of roughly 0.9 percent witnessed during crisis episodes.[5] However, even when we restrict attention to normal times, inflation averages roughly 1.7 percent, substantially lower than the central bank's 2 percent target. Agents' anticipation of future crises thus has important macroeconomic repercussions even during normal times. Moreover, the ELB plays a quantitatively important role in generating these results, since average inflation in normal times comes in very close to target when we abstract from this constraint. Risk aversion also plays an important role, since the average margin by which the central bank undershoots its inflation target in normal times shrinks by about half when we shift attention from the benchmark model to an otherwise comparable model under which we assume low risk aversion.

In addition, the magnitude of households' precautionary saving motives can be seen by looking at the real interest rate outcomes reported in Table 2. For instance, the average real rate observed in normal times is about

[5] Deflation of this magnitude did not occur during the Great Recession, in contrast to our model's predictions. As shown by Del Negro et al. (2015), introducing financial frictions into a standard DSGE model can help reconcile this fact with the sharp contraction in economic activity witnessed around the same time. Given the nonlinearities present in our framework, adding financial frictions would substantially complicate the numerical analysis. We leave this issue to future work.

40 basis points lower under the benchmark model, compared with an otherwise comparable economy without an ELB. This occurs because the ELB prevents monetary policy from fully offsetting the impact of the medium-run shock on aggregate demand, thus creating consumption risk against which households aim to protect themselves by saving more. That said, we also note that the precautionary savings motive in our model gives rise to an offsetting mechanism that tends to raise aggregate demand, all else equal. This mechanism specifically arises because of the presence of capital accumulation and endogenous growth: owing to their downward effect on the real interest rate, precautionary-savings incentives also translate into a moderately higher level of investment and, by extension, somewhat faster growth.

In Table 3, we turn our attention to financial variables and the impact of ELB risk on inflation compensation, computed based on the difference between nominal and real yields. Under our benchmark calibration, the breakeven inflation rate is predicted to be low, even in normal times, when the second panel of the table places the average breakeven rate around 1.2 percent. In terms of composition, the model predicts that this weakness partly reflects a low expected rate of inflation of 1.4 percent, coupled with a negative inflation risk premium of 0.25 percent, a figure in line with the empirical evidence since 2008 (e.g., Camba-Mendez and Werner, 2017; Fleckenstein et al., 2017).[6] In addition – and consistent with the precautionary behavior emphasized above – the top panel of the table shows that the presence of ELB risk boosts the equity premium from an unconditional average of 1.1 percent in the absence of the ELB to about 1.9 percent once the constraint on monetary policy is taken into account. As shown in the middle

[6] Campbell et al. (2017) provide evidence consistent with the notion that nominal bonds have shifted from being "inflation bets" to "deflation hedges." Chen et al. (2016) and Grishchenko and Huang (2013) find that the inflation risk premium implied by nominal yields and a TIPS-based measure of inflation compensation switched signs from positive to negative around the Great Recession. Fleckenstein et al. (2017) study deflation risk based on data from inflation swaps and options markets and also find evidence of a change in sign around 2009.

panel of the table, this positive effect also occurs – albeit to a smaller degree – if we shift attention from unconditional averages to averages conditional on being in normal times.

4.3 Inspecting the mechanism: precautionary pricing

To get some sense for the channels underlying the results documented above, it is useful to revisit the model's Phillips curve, as given by equation (6). If we average across all "normal" states of the world, then this equation can be written as

$$\mathbb{E}[f(\Pi_t)|\xi_t^{MR}=1] = (1-\theta) + \theta\mathbb{E}(MC_t|\xi_t^{MR}=1) + \mathbb{E}[\tilde{m}_{t+1}f(\Pi_{t+1})|\xi_t^{MR}=1],$$

where $f(\Pi_t) := \varphi\left(\Pi_t - \Pi^*\right)\Pi_t$ is an increasing function over the range of inflation rates realized in our simulations;[7] $\tilde{m}_{t+1} := m_{t+1}(y_{t+1}/y_t)$ denotes the "effective" discount factor associated with firms' pricing decisions; and $\mathbb{E}(\cdot|\xi_t^{MR}=1)$ denotes an average conditional on currently being in "normal times." Further noting that

$$\mathbb{E}[\tilde{m}_{t+1}f(\Pi_{t+1})|\xi_t^{MR}=1]$$

$$= \mathbb{E}(\tilde{m}_{t+1}|\xi_t^{MR}=1)\mathbb{E}[f(\Pi_{t+1})|\xi_t^{MR}=1] + cov[\tilde{m}_{t+1}, f(\Pi_{t+1})|\xi_t^{MR}=1]$$

and

$$\mathbb{E}[f(\Pi_{t+1})|\xi_t^{MR}=1]$$

$$= (1-\eta)\mathbb{E}[f(\Pi_t)|\xi_t^{MR}=1] + \eta\mathbb{E}[f(\Pi_{t+1})|\xi_t^{MR}=1, \xi_{t+1}^{MR}=\underline{\xi}^{MR}],$$

[7] More specifically, we find that inflation is generally below target in our simulations, but still consistently high enough that $f'(\Pi_t) = \varphi(2\Pi_t - \Pi^*) > 0$.

we then reach the following expression for average inflation in normal times:

$$\mathbb{E}[f(\Pi_t)|\xi_t^{MR}=1]$$

$$= \frac{\left[\begin{array}{l} (1-\theta)\ +\theta\mathbb{E}(MC_t|\xi_t^{MR}=1) \\ \qquad\qquad +\eta\mathbb{E}(\tilde{m}_{t+1}|\xi_t^{MR}=1)\mathbb{E}[f(\Pi_{t+1})|\xi_t^{MR}=1,\xi_{t+1}^{MR}=\underline{\xi}^{MR}] \\ \qquad\qquad +cov[\tilde{m}_{t+1},f(\Pi_{t+1})|\xi_t^{MR}=1] \end{array}\right]}{1-(1-\eta)\mathbb{E}(\tilde{m}_{t+1}|\xi_t^{MR}=1)} \quad (9)$$

This expression identifies four distinct determinants of average inflation outcomes in normal times, namely: (i) average marginal costs in normal times, as measured by the term $\mathbb{E}(MC_t|\xi_t^{MR}=1)$; (ii) expectations about the future behavior of inflation should the economy fall into a crisis episode, as measured by the term $\mathbb{E}[f(\Pi_{t+1})|\xi_t^{MR}=1,\xi_{t+1}^{MR}=\underline{\xi}^{MR}]$; (iii) the covariance between future inflation outcomes and firms' effective stochastic discount factor, as measured by the term $cov[\tilde{m}_{t+1},f(\Pi_{t+1})|\xi_t^{MR}=1]$; and (iv) the average rate at which firms discount the future when making pricing decisions, as measured by the term $\mathbb{E}(\tilde{m}_{t+1}|\xi_t^{MR}=1)$. While (i) is common to all New Keynesian models, and (ii) has been studied by, e.g., Reifschneider and Williams (2000), Coenen et al. (2004), Adam and Billi (2007), Nakov (2008), and Hills et al. (2016), (iii) introduces a novel "precautionary pricing" incentive into firms' behavior – one from which most of the previous literature has abstracted, partly due to reliance on linearized and/or perfect foresight solutions, coupled with the low risk aversion typically implied by more standard, non-recursive preferences. Moreover, to the extent that firms inherit their stochastic discount factor from households, (iv) represents a novel channel via which households' precautionary-savings behavior influences firms' pricing decisions beyond its usual impact on aggregate activity

and marginal costs.

To gauge the relative importance of these channels in driving the strong disinflationary bias from ELB risk, we calculate the change in the average rate of inflation in normal times when a specific channel in equation (9) is set to its value implied by an otherwise comparable model without the ELB.[8] Using this approach, the middle column of Table 4 identifies precautionary pricing as the strongest source of disinflationary pressure in our baseline model. Expectations about future inflation should the economy fall into a crisis also constitute a significant source of disinflationary pressure. On the other hand, the marginal cost channel plays a key offsetting role because ELB risk introduces a sizeable premium into the return on capital, making marginal costs somewhat higher in the baseline model, relative to an otherwise comparable model without an ELB. These results point to risk-related considerations in general – and precautionary pricing in particular – as key determinants of firms' pricing behavior in the baseline model. In contrast, the results reported in the right-hand column of Table 4 indicate that the precautionary pricing channel plays a relatively minor role at the low levels of risk aversion normally considered in the New Keynesian literature.

4.4 The role of risk aversion

Given the above-noted role played by risk-related factors in influencing firms' pricing behavior, a natural question to ask is how different degrees of risk aversion might impact our results. We explore this issue in Figures 6 and 7, which were constructed by varying the coefficient of relative risk aversion from the low value close to unity that would be implied by non-recursive

[8] To express results in units of inflation, we use equation (9) to compute $400\left[f^{-1}\left\{\mathbb{E}\left[f\left(\Pi_t\right)|\xi_t^{MR}=1\right]\right\}-1\right]$. We first calculate this moment using our model. We then repeat this calculation setting the term capturing a specific channel in equation (9) to its value implied by an otherwise comparable model without the ELB. All other terms are kept to their values under our model with the ELB. Table 9 reports the difference between these two numbers. Positive (negative) values thus indicate sources of (dis)inflationary pressure.

preferences all the way up to a maximum of 110 – a figure more than double that assumed under our baseline calibration and in line with the estimates reported in Rudebusch and Swanson (2012). More specifically, Figure 6 reports averages for the real rate, inflation, and various asset-pricing premia, all conditional on the economy being in "normal times". Figure 7 then shifts attention to averages conditional on the economy being in a "crisis" episode. Both figures report results for versions of the model with and without an ELB, while Figure 6 also includes values for the precautionary pricing term derived in our previous subsection – i.e., $cov[\tilde{m}_{t+1}, f(\Pi_{t+1})|\xi_t^{MR} = 1]$.

Figure 6 points to a strong downward effect of the ELB on inflation and the real rate in normal times, even for relatively low levels of risk aversion. For example, with minimal risk aversion, inflation and the real rate both average roughly twenty basis points lower in normal times when we take account of the effects of the ELB, relative to an otherwise comparable version of the model under which abstracts from the ELB. Moreover, the results reported earlier in Table 3 suggest that minimal risk aversion still places the average break-even rate in normal times around 1.5 percent in a version of the model with an ELB. Even though the lower panels of Figure 6 indicate that relatively high degrees of risk aversion are needed to achieve a quantitatively meaningful impact on some of the asset-pricing premia under consideration, these results suggest that the macroeconomic effects of ELB risk can be significant even at relatively low levels of risk aversion.

In addition, the downward pressures that the ELB tends to place on inflation and the real rate both grow significantly as we increase risk aversion and thus make firms and households more concerned about the possibility of future ELB episodes, increasing precautionary-saving and -pricing incentives in turn. For example, with maximal risk aversion, average inflation in normal times is nearly 60 basis points below target, as opposed to the average deviation of 35 basis points predicted by our benchmark calibration. Similarly, the average real rate in normal times undershoots its value in de-

terministic steady state by about twice the margin implied by our benchmark calibration, while the inflation risk premium roughly doubles on average.

It is also worthwhile to note that the model's precautionary-saving mechanism has implications for the behavior of the inflation rate during crisis episodes. In normal times, higher risk aversion generally leads to lower inflation outcomes, largely due to the associated intensification of agents' precautionary-saving and -pricing incentives. In contrast, Figure 7 suggests that higher risk aversion has the opposite effect, leading *higher* inflation during crises. This reflects the fact that the ELB tends to bind during crisis episodes: to the extent that more-risk averse households engage in more precautionary savings, the Fisher equation then tends to translate the implied downward pressure on real rates into higher inflation outcomes. Mertens and Williams (2018) and Lansing (2018) highlight a similar effect associated with exogenous declines in the level of the equilibrium real interest rate.

4.5 Additional robustness

Finally, we assess the robustness of our results along several important dimensions and report those results in Table 5. In particular, we examine the sensitivity of inflation, real interest rates and the inflation risk premium to changes in the inflation coefficient in the Taylor rule, the inflation target, the severity of crisis episodes, the degree of price rigidity, the risk of medium-run demand shocks, and the level of the real interest rate in steady state.

Our benchmark calibration assumes that monetary policymakers place a weight of 3 on inflation in the Taylor rule. In the column labeled "Higher ϕ_{Π}," we raise this coefficient to 4. This greater emphasis on reducing the departures of inflation from target leads to higher average inflation outcomes, along with a more stimulative monetary-policy response to medium-run demand shocks, which in turn leads to less precautionary savings, a higher real interest rate, and an inflation risk premium closer to zero. This is in line with the analysis in Nakata and Schmidt (2016) showing that appointing

a conservative central banker, as in Rogoff (1985), reduces the deflationary bias in economies with occasionally binding constraints.

We next consider raising the inflation target, as is sometimes proposed as a way to reduce the probability of ELB episodes (see, e.g., Blanchard et al., 2010). Increasing the inflation target to 3 percent leads to outcomes similar to those in the previous exercise, as seen in the column labeled "Higher Π^*." Given that the rest of the calibration is unchanged relative to our benchmark model, raising the inflation target also implies that the incidence, duration and severity of ELB episodes all decline, weakening the ELB's disinflationary effects, along with households' precautionary-saving incentives, and the precautionary-pricing incentives facing firms.

In the column denoted "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 medium-run demand shock scenario depicted in Figure 4 is now associated with a peak impact on real GDP of about 3 percent, as opposed to the roughly 5 percent drop implied by our benchmark calibration. Despite this more muted decline in economic activity, we see that ELB risk continues to imply a sizeably negative inflation risk premium while exerting significant downward pressure on inflation and the real rate.

In the next column, "Higher φ," we consider a higher degree of price rigidity by increasing the coefficient scaling the cost of price adjustments by 10 percent. Since prices are now more rigid, ELB episodes entail a more muted decline in inflation, implying a somewhat less negative inflation risk premium. However, we see overall that our results are little changed by this alternative calibration.

In the penultimate column "Lower η," we report results for a case where we reduce the probability of medium-run shocks. In particular, we halve this probability from 0.9 percent under the benchmark parameterization to 0.45 percent. Since this reduction weakens households' precautionary-savings incentives while also weakening the precautionary pricing channel, the reduc-

tion is associated with higher average outcomes for the real rate and inflation, along with a less-negative inflation risk premium.

Finally, the last column of the table examines the effect of lowering the steady-state value of the real interest rate to 2 percent. As a result of the fact that this leaves significantly less room for monetary stimulus in response to medium-run demand shocks, the precautionary-savings and -pricing incentives respectively facing households and firms both increase significantly, placing greater downward pressure on inflation and real rates while making the inflation risk premium more negative.

Overall, the results from this subsection suggest that our qualitative conclusions regarding the importance of ELB is generally robust to perturbations of key model parameters.

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. To do so, we augmented an otherwise standard New Keynesian model to include three key features: high risk aversion (via recursive preferences), long-run risk (via endogenous growth), and a demand-shock process admitting ELB episodes broadly consistent with recent experience. Recursive preferences and long-run risk are attractive features for our purposes since they allow us to understand the role that ELB risk can play even in relatively normal times when the nominal interest rate is away from its lower bound.

In the presence of high but plausible levels of risk aversion and long-run risk, we found that the prospect of future ELB episodes leads to a substantial disinflationary bias that should be expected even during normal times. Even though the central bank has a formal inflation target, inflation expectations in the presence of ELB risk fall significantly below target. In addition, the inflation risk premium is notably negative, as has been observed during the

years following the Great Recession, and real rates come under significant downward pressure.

We emphasized several important channels contributing to these effects. One channel involves an increase in precautionary saving as agents anticipate the negative effects of future ELB episodes, depressing aggregate demand, inflation and real rates even in normal periods. At the same time, firms face precautionary-pricing incentives associated with the strongly negative covariance between the stochastic discount factor and future inflation outcomes. More specifically, firms anticipate that ELB episodes are characterized by declines in inflation and large increases in the stochastic discount factor, leaving them with strong incentives to pursue relatively conservative pricing strategies during normal times.

Despite its simplicity, our model also embeds a mechanism that mitigates the decline in aggregate demand in response to ELB risk. This effect arises because the decline in the real interest rate triggered by precautionary saving also raises investment and thus boosts productivity growth.

Looking ahead, our results also lead to important questions for monetary policy. Should central banks adjust their policy rules to account for the effects of ELB risk on inflation? Is there some simple way for monetary authorities to divorce their implicit targets from their formal objectives? What are the implications for optimal inflation? We plan to explore these and other questions in future research.

References

Adam, K. and Billi, R. (2007). Discretionary monetary policy and the zero lower bound on nominal interest rates. *Journal of Monetary Economics*, 3:728–752.

Amano, R. and Shukayev, M. (2012). Risk premium shocks and the zero

bound on nominal interest rates. *Journal of Monetary, Credit, and Banking*, 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` as of September 27, 2019.

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.

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

Camba-Mendez, G. and Werner, T. (2017). The inflation risk premium in the post-Lehman period. ECB working paper no. 2033.

Campbell, J. and Cochrane, J. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Campbell, J. Y., Sunderam, A., and Viceira, L. M. (2017). Inflation bets or deflation hedges? the changing risks of nominal bonds. *Critical Finance Review*, 6:263–301.

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.

Chen, A., Engstrom, E., and Grishchenko, O. (2016). Has the inflation risk premium fallen? Is it now negative? *FEDS Notes*, April 4, Board of Governors of the Federal Reserve System.

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

Coenen, G., Orphanides, A., and Wieland, W. (2004). Price stability and monetary policy eectiveness when nominal interest rates are bounded at zero. *The B.E. Journal of Macroeconomics: Advances in Macroeconomics*, 1:1–25.

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? *Review 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` as of September 27, 2019.

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.

Fleckenstein, M., Longstaff, F. A., and Hanno, L. (2017). Deflation risk. *The Review of Financial Studies*, 30:2719–2760.

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` as of September 27, 2019.

Grishchenko, O. and Huang, J. (2013). The inflation risk premium: Evidence from the TIPS market. *Journal of Fixed Income*, 22:5–30.

Guerron-Quintana, P. A. and Jinnai, R. (2015). Liquidity shocks and asset prices. Mimeo. Available at `https://sites.google.com/site/pabloaguerronquintana/home/research` as of Septebmer 27, 2019.

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. FEDS working paper no. 2016-009, Board of Governors of the Federal Reserve System.

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.

Lansing, K. J. (2018). Endogenous regime switching near the zero lower bound. Federal Reserve Bank of San Francisco working paper no. 2017-24.

Leduc, S. and Liu, Z. (2016). Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, 82:20–35.

Mertens, T. and Williams, J. C. (2018). What to expect from the lower bound on interest rates: Evidence from derivatives prices. Federal Reserve Bank of San Francisco working paper no. 2018-03.

Nakata, T. and Schmidt, S. (2016). The risk-adjsuted monetary policy rule. FEDS working paper no. 2016-061, Board of Governors of the Federal Reserve System.

Nakata, T. and Tanaka, H. (2016). Equilibrium yield curves and the interest rate lower bound. FEDS working paper no. 2016-085, Board of Governors of the Federal Reserve System.

Nakov, A. (2008). Optimal and simple monetary policy rules with zero floor on the nominal interest rate. *International Journal of Central Banking*, 4:73–127.

Piassezi, M. and Schneider, M. (2006). Equilibrium yield curves. *NBER Macroeconomics Annual*, 21:389–442.

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

Reifschneider, D. and Williams, J. (2000). Three lessons for monetary policy in a low-inflation era. *Journal of Money, Credit and Banking*, 32:936–66.

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.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *The Quarterly Journal of Economics*, 100:1169–1189.

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.

Swanson, E. (2016). A macroeconomic model of equities and real, nominal, and defaultable debt. Mimeo. Available at `http://www.socsci.uci.edu/~swanson2/papers/ezap.pdf` as of September 27, 2019.

Tallarini, T. (2000). Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45.

Torry, H. (2017). Yellen says she's uncertain weak inflation is transitory. *Wall Street Journal*, October 15, available at `https://www.wsj.com/articles/yellen-says-gradual-rate-increases-should-help-sustain-economys-growth-15080724` as of September 22, 2019.

Van Binsbergen, J. H., Fernández Villaverde, J., Koijen, R. S. J., and Rubio Ramírez, J. (2012). The term structure of interest rates in a dsge model with recursive preferences. *Journal of Monetary Economics*, 59(7):634–648.

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.

Table 1: Baseline calibration

Parameter	Value	Description and notes
σ	2	Inverse IES, set to value in Gourio and Ngo (2017) and Rudebusch and Swanson (2012); see also Vissing-Jørgensen (2002)
ν	1.5	Inverse Frisch elasticity of labor, set to value in Gourio and Ngo (2017) and Rudebusch and Swanson (2012); see also Pistaferri (2003)
χ	33.4	Multiplicative constant in the disutility of labor, chosen to set labor supply along BGP = 1/3, as in Gourio and Ngo (2017)
γ	-69	Epstein-Zin parameter, set to imply a CRRA=50
β	0.9999	Discount factor; see main text for details
θ	7.66	Elasticity of substitution across intermediate goods, set to value in Gourio and Ngo (2017) and Fernandez-Villaverde et al. (2015); see also Basu (1992) and Basu and Fernald (1997)
φ	238.11	Rotemberg parameter, set to value in Gourio and Ngo (2017), implies Calvo parameter ≈ 0.85; see also Del Negro et al. (2015)
ϕ_{Π}	3.00	Taylor coefficient, set to value in Gourio and Ngo (2017); 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 Gourio and Ngo (2017)
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 balanced growth rate = 1.5%
ρ	0.725	Persistence of short-run shocks; see main text for details
σ	0.1%	Standard deviation of productivity shocks; see main text for details
$\underline{\xi}^{MR}$	0.99	Size of the medium-run shock; see main text for details
η	0.9%	Risk of crisis episodes; see main text for details
$\overline{T}$	12	Duration of crisis episodes; see main text for details

Table 2: Average macroeconomic outcomes

	Baseline model		Low risk aversion		Deterministic steady state
	Without ELB	With ELB	Without ELB	With ELB	
Unconditional averages					
Inflation	1.8	1.4	1.8	1.52	2
Real rate	2.66	2.47	2.65	2.65	3.06
Nominal rate	4.47	3.81	4.47	4.17	5.12
Normalized investment	0.12	0.12	0.12	0.12	0.12
Expected return on capital	3.03	3.1	3.02	3.04	3.06
Rental rate on capital	0.14	0.14	0.14	0.14	0.14
Growth rate	1.49	1.52	1.49	1.49	1.5
Averages, conditional on "normal times"					
Inflation	1.97	1.65	1.98	1.8	2
Real rate	3.05	2.62	3.04	2.81	3.06
Nominal rate	5.01	4.16	5.02	4.55	5.12
Normalized investment	0.12	0.12	0.12	0.12	0.12
Expected return on capital	3.05	2.91	3.05	2.83	3.06
Rental rate on capital	0.14	0.14	0.14	0.14	0.14
Growth rate	1.51	1.56	1.5	1.53	1.5
Averages, conditional on "crises"					
Inflation	0.26	-0.88	0.26	-1.05	2
Real rate	-0.88	1.05	-0.89	1.16	3.06
Nominal rate	-0.46	0.69	-0.46	0.69	5.12
Normalized investment	0.12	0.12	0.12	0.12	0.12
Expected return on capital	2.83	4.8	2.82	4.89	3.06
Rental rate on capital	0.14	0.13	0.14	0.13	0.14
Growth rate	1.39	1.17	1.39	1.13	1.5

Table 3: Average financial outcomes

	Baseline model		Low risk aversion	
	Without ELB	With ELB	Without ELB	With ELB
Unconditional averages				
Inflation compensation	1.74	1.15	1.78	1.47
Expected inflation	1.8	1.39	1.8	1.51
Inflation risk premium	-0.05	-0.25	-0.02	-0.04
Nominal term premium	-0.12	-0.31	0	-0.03
Real term premium	-0.09	-0.14	0	-0.01
Inflation term premium	-0.03	-0.18	0	-0.02
Equity premium	1.12	1.9	1.11	1.17
Averages, conditional on "normal times"				
Inflation compensation	1.76	1.16	1.8	1.49
Expected inflation	1.8	1.4	1.81	1.52
Inflation risk premium	-0.04	-0.25	-0.01	-0.03
Nominal term premium	-0.12	-0.32	0	-0.03
Real term premium	-0.09	-0.14	0	-0.01
Inflation term premium	-0.04	-0.18	0	-0.02
Equity premium	0.01	0.86	0.01	0.06
Averages, conditional on "crises"				
Inflation compensation	1.58	1.05	1.61	1.33
Expected inflation	1.74	1.34	1.74	1.45
Inflation risk premium	-0.16	-0.29	-0.14	-0.13
Nominal term premium	-0.1	-0.27	0	-0.03
Real term premium	-0.08	-0.12	0	-0.01
Inflation term premium	-0.03	-0.15	0	-0.02
Equity premium	11.13	11.27	11.12	11.19

Table 4: Disinflationary bias: contributions from different channels

Channel (see equation 9)	Baseline model	Low risk aversion
Marginal costs – i.e., $\mathbb{E}(MC_t \vert \xi_t^{MR} = 1)$	8.11	3.87
Expected inflation conditional on a crisis occurring – i.e., $\mathbb{E}[f(\Pi_{t+1}) \vert \xi_t^{MR} = 1, \xi_{t+1}^{MR} = \underline{\xi}^{MR}]$	-3.46	-3.66
Precautionary pricing – i.e., $cov[\tilde{m}_{t+1}, f(\Pi_{t+1}) \vert \xi_t^{MR} = 1]$	-4.71	-0.31
Effective discount factor – i.e., $\mathbb{E}(\tilde{m}_{t+1} \vert \xi_t^{MR} = 1)$	-0.02	-0.01

Table 5: Sensitivity analysis

	Baseline	Higher ϕ_{Pi}	Higher Π^*	Higher $\underline{\xi}^{MR}$	Higher φ	Lower η	Lower R_{SS}^{real}
Unconditional averages							
Inflation	1.4	1.62	2.61	1.58	1.44	1.64	1.09
Real rate	2.47	2.56	2.6	2.61	2.48	2.7	1.2
Inflation risk premium	-0.25	-0.16	-0.13	-0.15	-0.23	-0.19	-0.37
Averages, conditional on "normal times"							
Inflation	1.65	1.84	2.83	1.79	1.67	1.78	1.38
Real rate	2.62	2.78	2.89	2.81	2.65	2.79	1.2
Inflation risk premium	-0.25	-0.15	-0.12	-0.15	-0.23	-0.19	-0.37
Averages, conditional on "crises"							
Inflation	-0.88	-0.34	0.64	-0.38	-0.7	-1.06	-1.46
Real rate	1.05	0.54	0.03	0.76	0.94	1.18	1.17
Inflation risk premium	-0.29	-0.22	-0.25	-0.21	-0.27	-0.27	-0.37

Figure 1: Medium-run demand-shock process

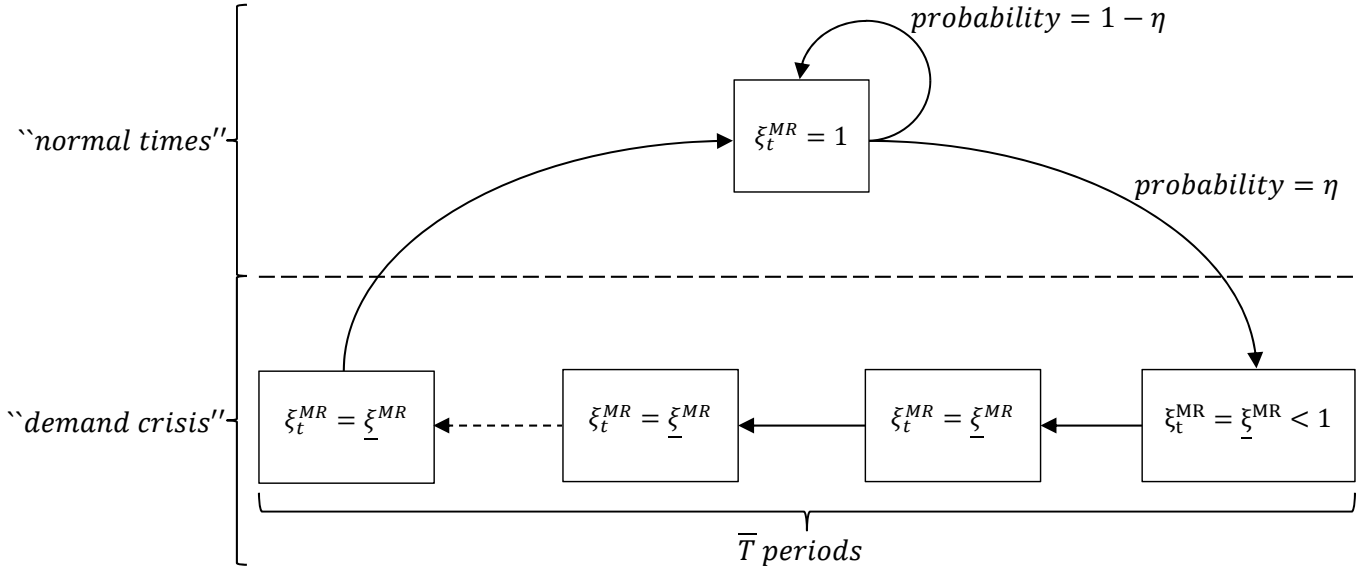

Figure 2: GDP per capita (solid blue) and pre-crisis trend (dashed red), 2000–17

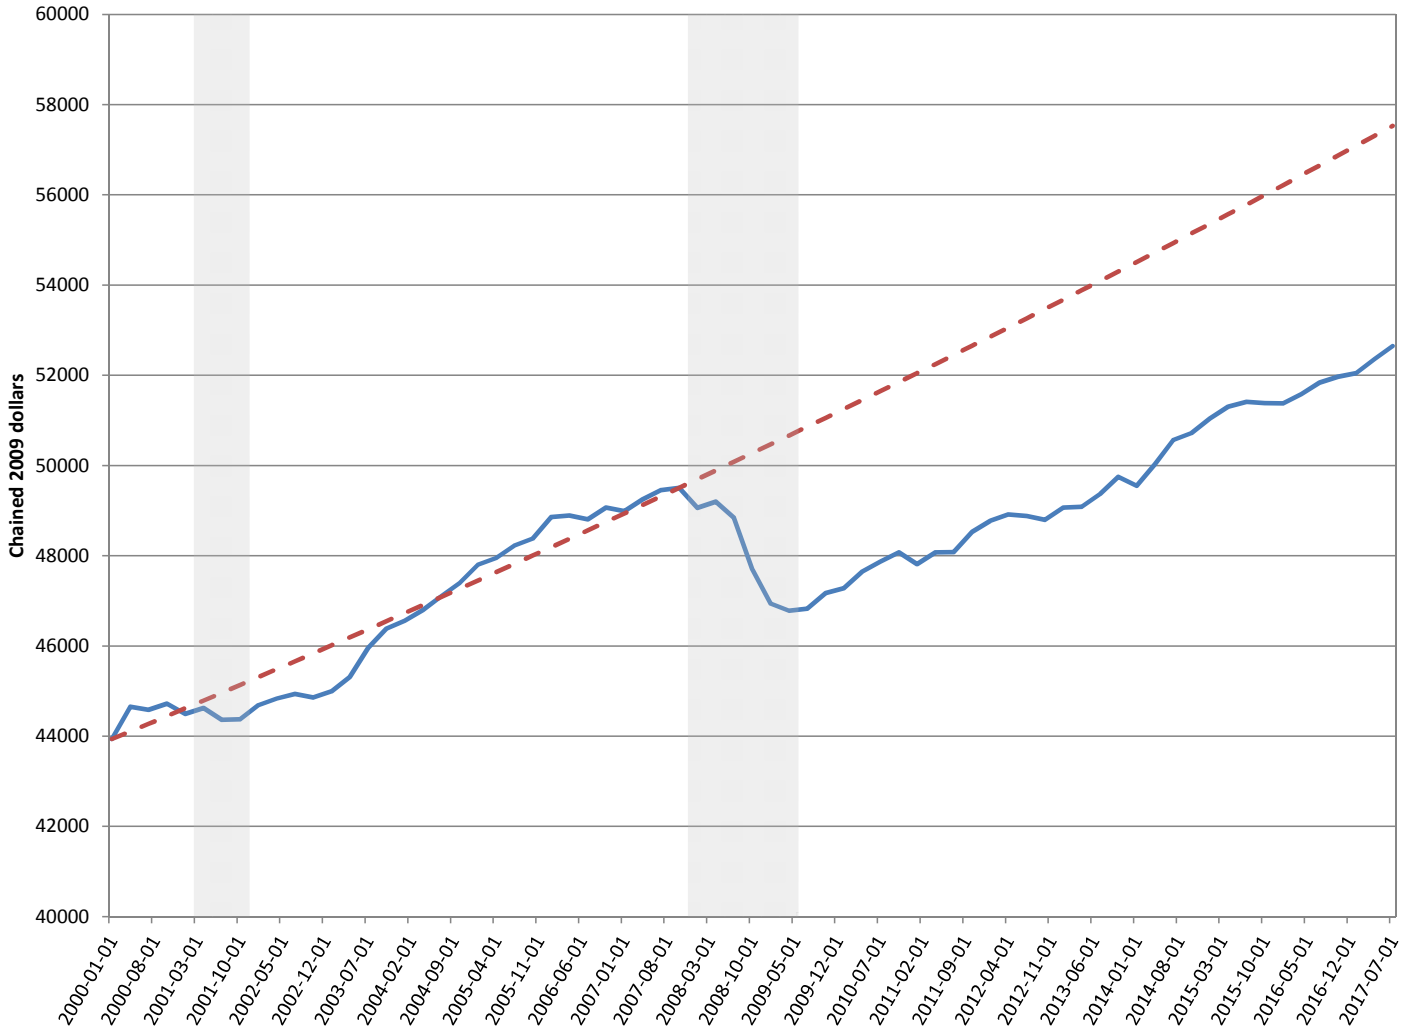

Figure 3: Fed funds rate, 2000–17

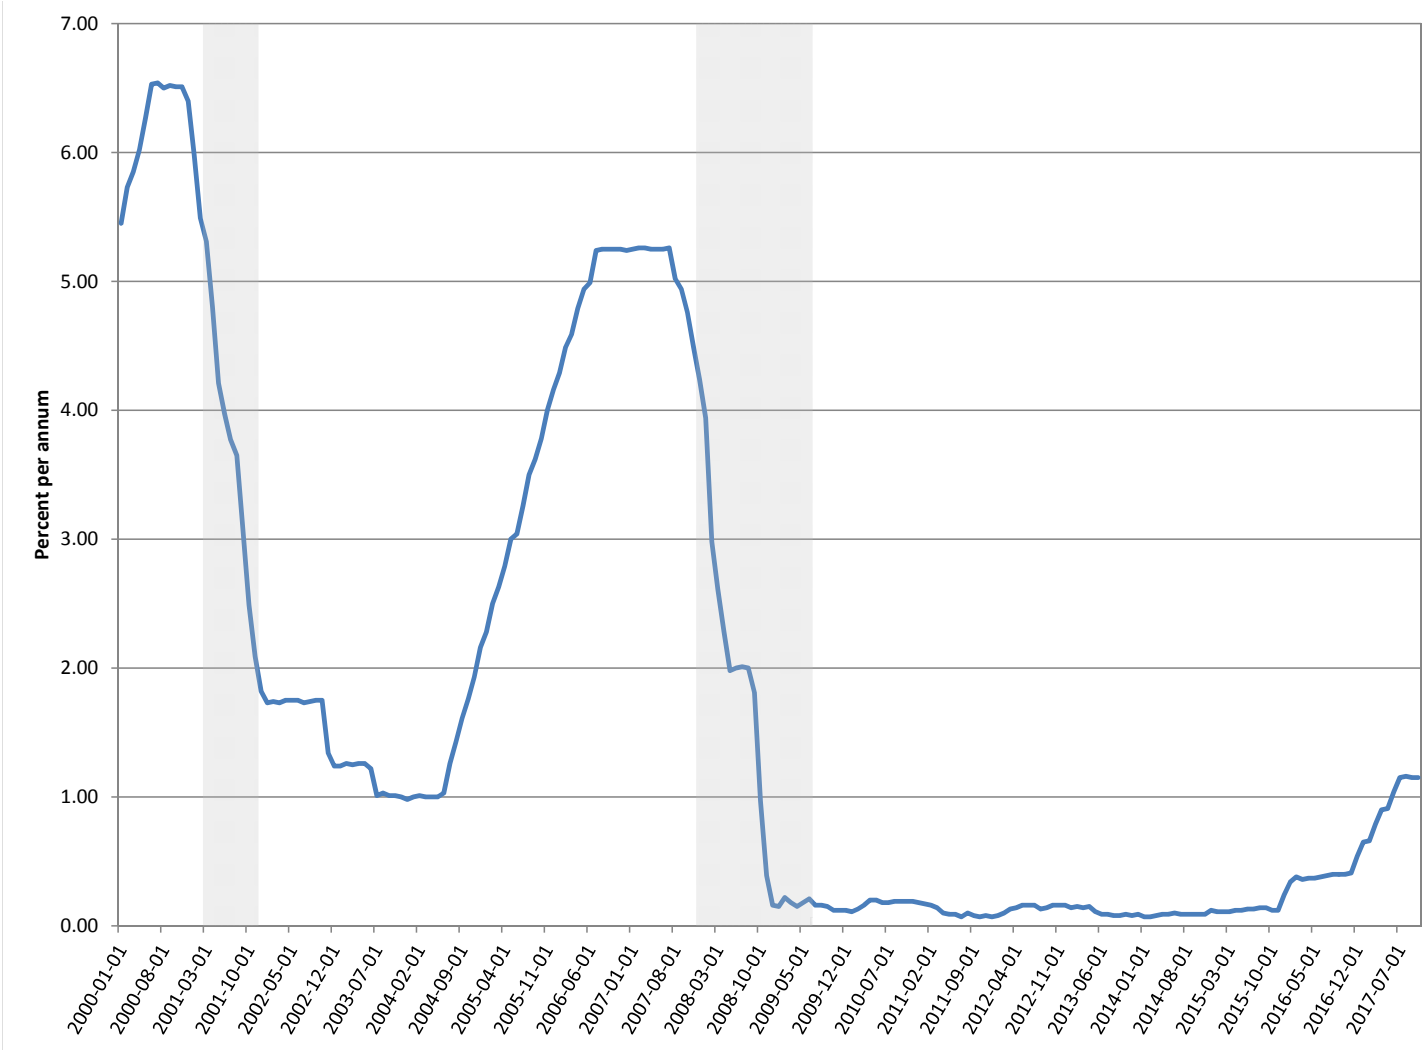

Figure 4: Responses to the medium-run demand shock

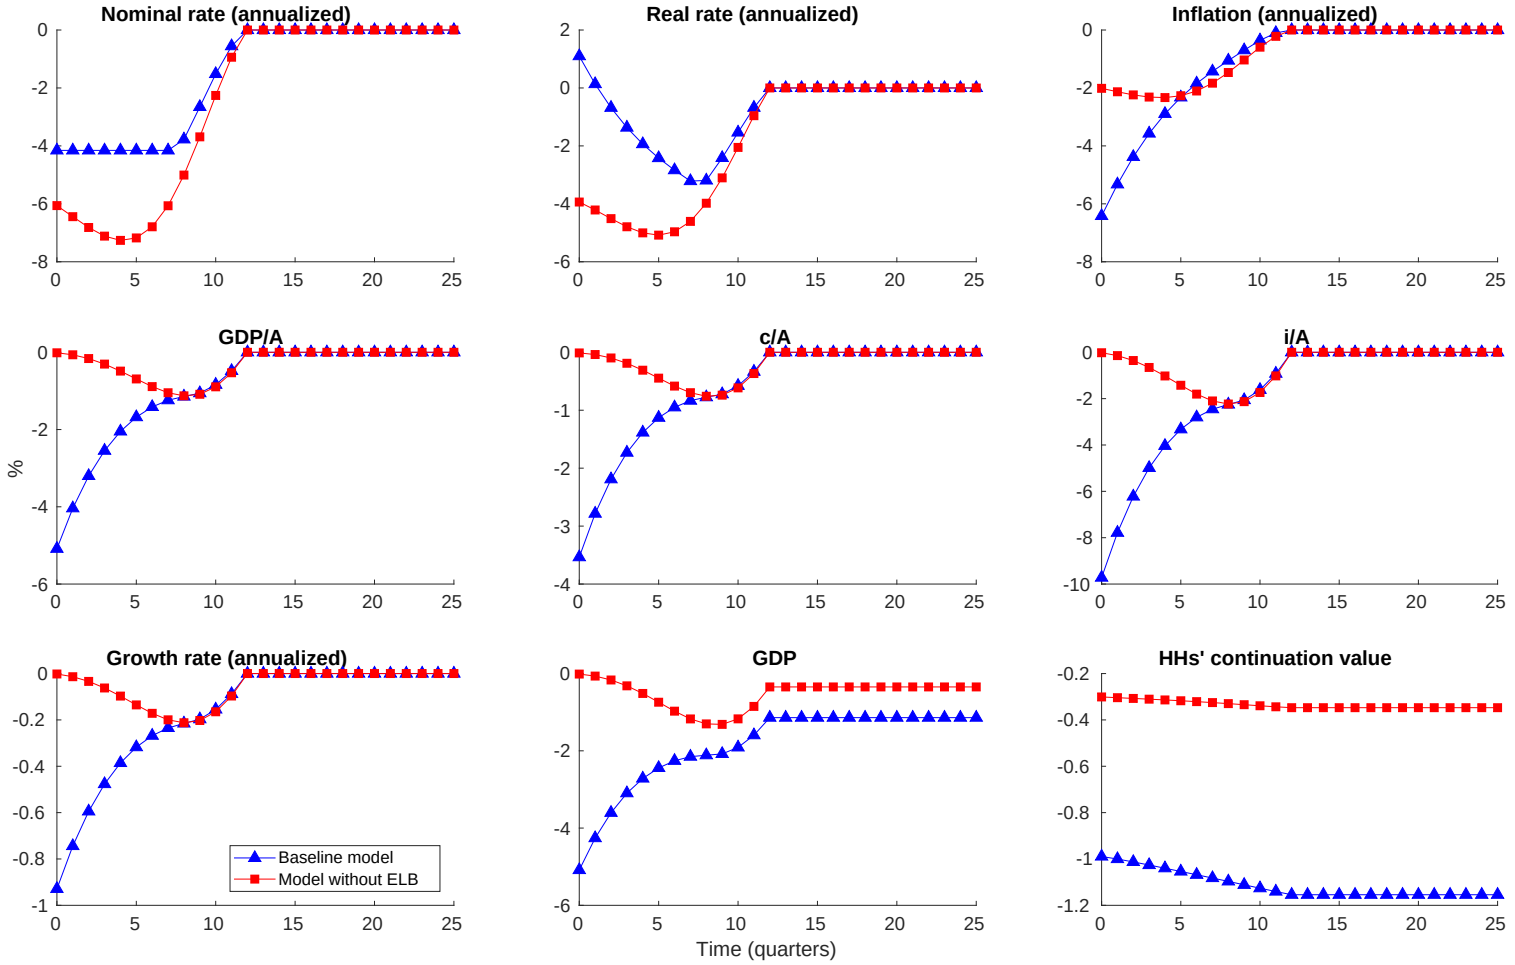

Figure 5: Responses to a negative short-run demand shock

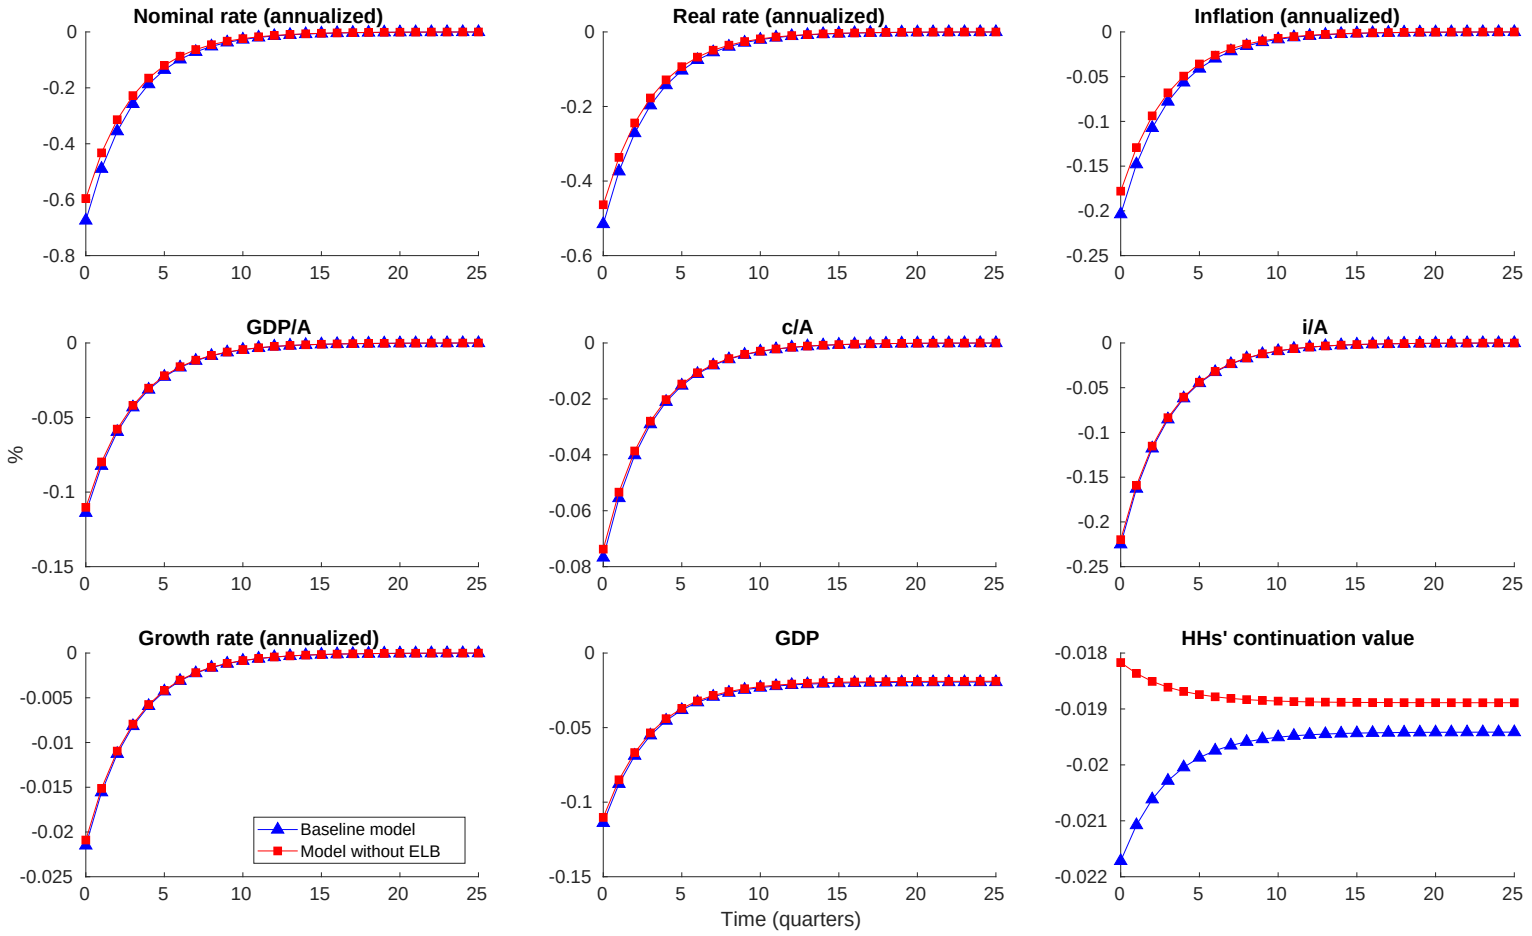

Figure 6: Average outcomes in "normal times" under various degrees of risk aversion

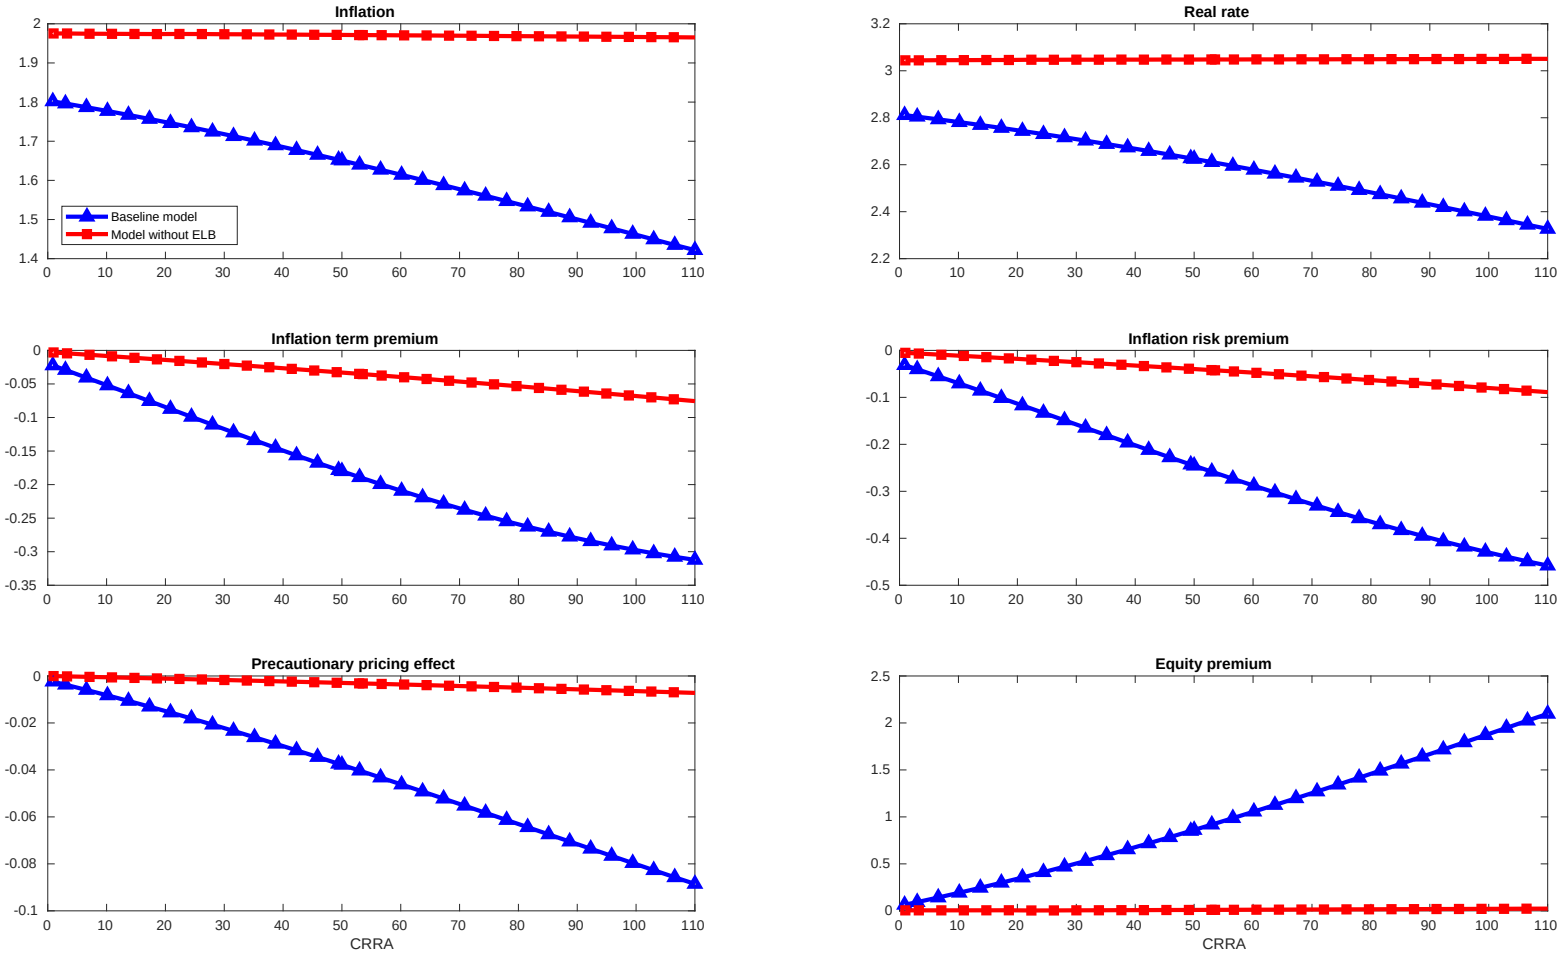

Figure 7: Average outcomes during "crises" under various degrees of risk aversion

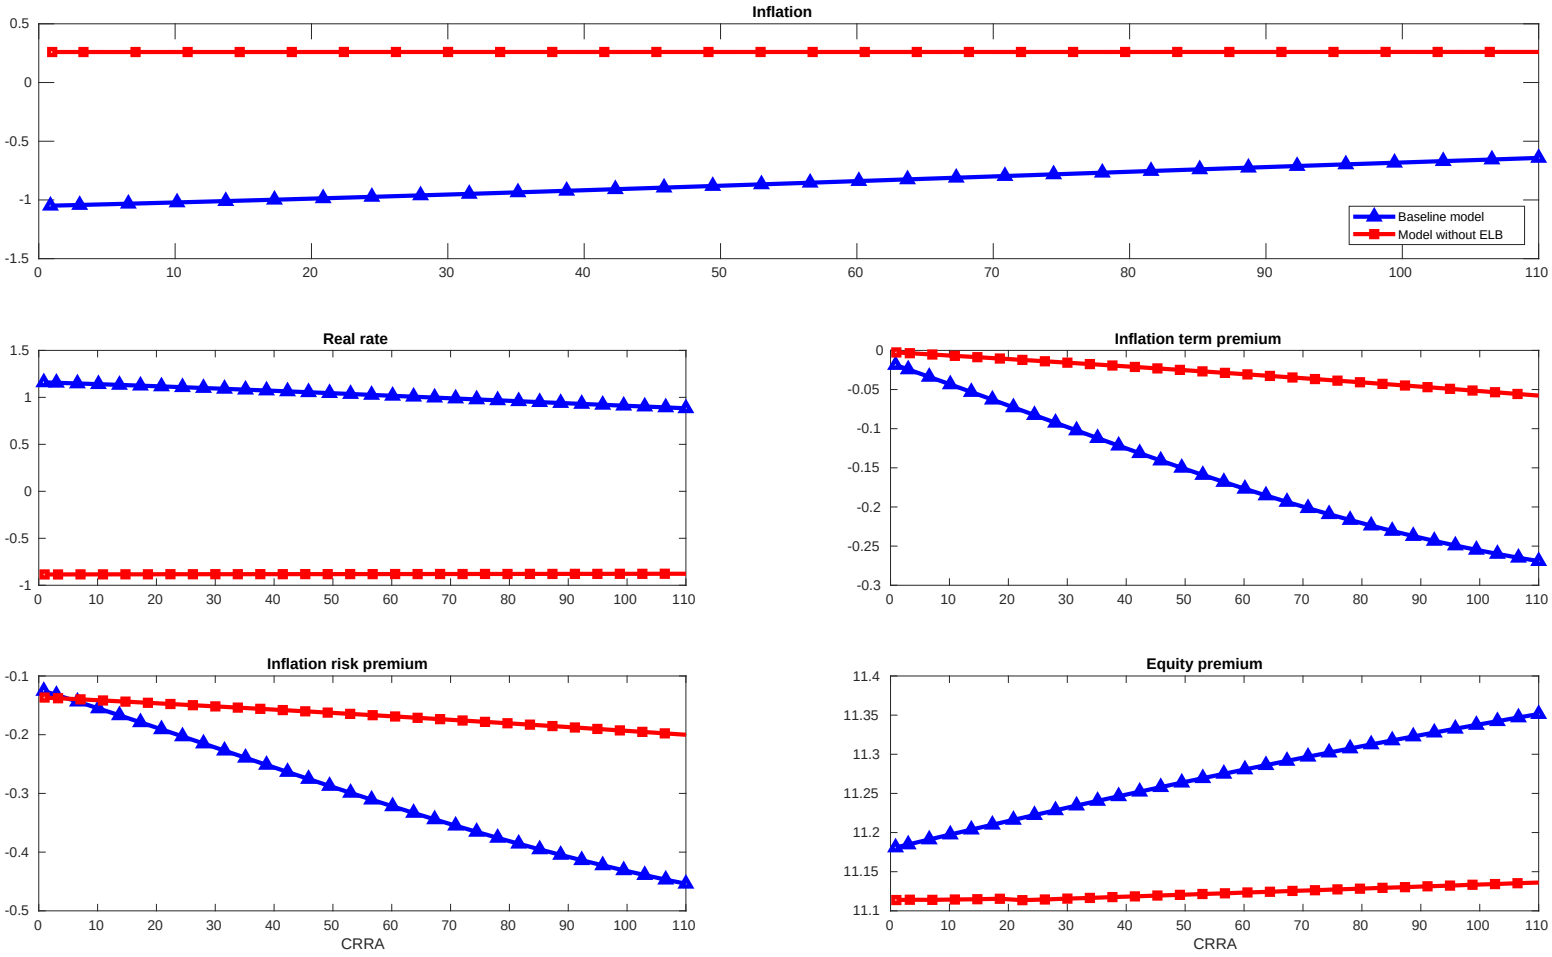

APPENDIX

A Stationary system

Letting $\hat{x} := x_t/A_t \ \forall x \in \{c, i, GDP, w, y\}$, $\hat{V}_t := V_t/A_t^{1-\sigma}$, $G_t^k := k_t/k_{t-1}$, and $IAC(\hat{i}_t) := \frac{\psi}{2}\left(\hat{i}_t - \frac{\widetilde{i}}{k}\right)^2$, the relevant system reads as follows:

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

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

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

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

$$1 = \mathbb{E}_t \left[m_{t+1} \cdot \frac{1}{Q_t} \left[\begin{array}{l} r_{t+1}^k - \hat{i}_{t+1} \ - \ IAC(\hat{i}_{t+1}) \\ \\ \qquad\qquad + \ Q_{t+1}\left[(1-\delta) + \frac{\hat{i}_{t+1}}{Q_{t+1}}\right] \end{array} \right] \right] \tag{14}$$

$$Q_t = \widetilde{Q}\left[1 + \psi\left(\hat{i}_t - \frac{\widetilde{i}}{k}\right)\right] \tag{15}$$

$$\frac{\hat{w}_t \ell_t}{\alpha_\ell} = \frac{r_t^k}{\alpha_k} \tag{16}$$

$$MC_t = \left(\frac{r_t^k}{\alpha_k}\right)^{\alpha_k} \left(\frac{\hat{w}_t}{\alpha_\ell}\right)^{\alpha_\ell} \tag{17}$$

$$\varphi(\Pi_t - \Pi^*)\Pi_t = (1-\theta) + \theta MC_t + \mathbb{E}_t\left[m_{t+1} \cdot \varphi\left(\Pi_{t+1} - \Pi^*\right)\Pi_{t+1} \cdot \frac{\hat{y}_{t+1}G_t^k}{\hat{y}_t}\right] \quad (18)$$

$$\hat{y}_t = \ell_t^{\alpha_\ell} \quad (19)$$

$$G\hat{D}P_t = \hat{y}_t - \hat{y}_t \cdot \frac{\varphi}{2}\left(\Pi_t - \Pi^*\right)^2 - IAC(\hat{i}_t) \quad (20)$$

$$\hat{c}_t + \hat{\hat{i}}_t = G\hat{D}P_t \quad (21)$$

$$G_t^k = (1-\delta) + \frac{\hat{\hat{i}}_t}{\widetilde{Q}} \quad (22)$$

$$R_t = \max\left\{1, R^*\left(\frac{\Pi_t}{\Pi^*}\right)^{\phi_\Pi}\left(\frac{G\hat{D}P_t}{\widetilde{GDP/A}}\right)^{\phi_{GDP}}\right\} \quad (23)$$