BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel 2016-19

Should Monetary Policy Lean Against Housing Market Booms?

by Sami Alpanda and Alexander Ueberfeldt

www.bank-banque-canada.ca

Bank of Canada Staff Working Paper 2016-19

April 2016

Should Monetary Policy Lean Against Housing Market Booms?

by

Sami Alpanda[1] and Alexander Ueberfeldt[2]

[1]Department of Economics
University of Central Florida
Orlando, Florida 32816
Sami.Alpanda@ucf.edu

[2]Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
uebe@bankofcanada.ca

ISSN 1701-9397

Acknowledgements

We thank Rhys Mendes for fruitful discussions. We also thank Greg Bauer, Oleksiy Kryvtsov, Gino Cateau, Sharon Kozicki, Césaire Meh, Malik Shukayev, Joel Wagner and seminar participants at the Bank of Canada, and at the Norges Bank–Bank of Canada–HEC Research Meeting in Dynamic Macroeconomics, for suggestions and comments. All remaining errors are our own.

Abstract

Should monetary policy lean against housing market booms? We approach this question using a small-scale, regime-switching New Keynesian model, where housing market crashes arrive with a logit probability that depends on the level of household debt. This crisis regime is characterized by an elevated risk premium on mortgage lending rates, and, occasionally, a binding zero lower bound on the policy rate, imposing large costs on the economy. Using our set-up, we examine the optimal level of monetary leaning, introduced as a Taylor rule response coefficient on the household debt gap. We find that the costs of leaning in regular times outweigh the benefits of a lower crisis probability. Although the decline in the crisis probability reduces volatility in the economy, this is achieved by lowering the average level of debt, which severely hurts borrowers and leads to a decline in overall welfare.

JEL classification: E44, E52, G01
Bank classification: Monetary policy framework; Financial stability; Economic models; Housing

Résumé

La politique monétaire devrait-elle prendre une approche préventive à l'égard des booms immobiliers? Nous étudions cette question à l'aide d'un petit modèle néokeynésien à changement de régime, dans lequel la probabilité d'effondrement du marché du logement est régie par une fonction logit et dépend du niveau d'endettement des ménages. Le régime de crise est caractérisé par une prime de risque élevée pour les taux hypothécaires et, à l'occasion, une contrainte de non-négativité du taux directeur, ce qui entraîne des coûts considérables pour l'économie. Sur cette base, nous examinons le degré de prévention optimal de la politique monétaire, en intégrant à la règle de Taylor un coefficient de réaction à l'écart de l'endettement des ménages par rapport à sa valeur tendancielle. Nous constatons qu'en période normale, les coûts d'une politique monétaire préventive surpassent les avantages associés à une plus faible probabilité de crise. Cependant, bien que cette moindre probabilité fasse diminuer la volatilité au sein de l'économie, ce résultat nécessite une réduction du niveau moyen d'endettement, ce qui nuit aux emprunteurs et se traduit par un recul du bien-être global.

Classification JEL : E44, E52, G01
Classification de la Banque : Cadre de la politique monétaire; Stabilité financière; Modèles économiques; Logement

Non-Technical Summary

We evaluate the scope for monetary policy to lean against the probability of a housing market crash that triggers a severe recessionary episode.

The analysis is conducted in the context of a small-scale, regime-switching New Keynesian model that endogenizes the likelihood of a large housing market correction. Following the recent empirical literature, the likelihood of a correction is a function of the credit gap, i.e. the deviation of the level of real credit from its trend level.

In the model, the risk premium on household borrowing rates is countercyclical, rising sharply in the crisis regime, which increases the cost of borrowing. Also, the central bank hits the zero lower bound on the policy interest rate, resulting in an increase in inflation and output volatility during recessionary episodes. Thus, the resulting recessionary episodes are severe, with a cumulative aggregate consumption loss of 14 per cent over 10 quarters.

We assess the effectiveness of leaning using an expanded Taylor rule, in which monetary policy responds to the upside deviations of real credit from trend, and we compare the resulting welfare to that from a baseline case with a standard Taylor rule.

We find that leaning successfully reduces the tail risks inherent in the debt cycle dynamics. However, aggregate welfare is lower with leaning. Intuitively, the marginal benefit of leaning comes at the higher marginal cost of reducing the average consumption of the borrowers. This is a first-order penalty to solve a second-order problem: namely, the reduction of economic volatility. The insurance cost of reducing the likelihood of a tail event is simply too high, suggesting that, in general, central banks should not lean against housing booms.

1 Introduction

Household indebtedness in Canada is near an all-time high. The household debt-to-disposable income ratio in the third quarter of 2015 was 163.7 per cent, close to the peak values observed in the United States prior to the recent financial crisis. Even after adjusting the official Canadian measure to better conform with the debt and income definitions used in the United States, the ratio is still around 154 per cent; lower than the U.S. peak, but still significantly higher than the current value in the United States following several years of post-crisis deleveraging (see Figure 1).

The recent increase in Canadian household debt was accompanied by a sharp rise in house prices, with regular mortgages and home equity loans as the main drivers of new household borrowing. It is difficult to assess exactly what portion of these house price movements reflects changes in fundamentals (such as population dynamics), and what portion reflects pure speculation. Using an empirical model of house price determination, Bauer (2014) estimates that house prices in Canada were close to 20 per cent overvalued relative to fundamentals at the end of 2014, a magnitude that mirrors the household debt gap, calculated as the per cent deviation of real household debt from its long-run log-linear trend (see Figure 2).[1] Other indicators, such as the house price-to-income and the house price-to-rent ratios, reveal even larger "overvaluation" estimates. According to the *Economist* magazine, in mid-2015 the house price-to-income ratio in Canada was about 34 per cent and the house price-to-rent ratio was nearly 89 per cent above their long-term averages.[2]

The fact that household debt is near its historic high, and house prices are overvalued relative to benchmarks, raises macroeconomic and financial stability concerns. Financial imbalances, especially those that involve high debt and leverage, may significantly increase the probability and the impact of tail events (i.e., crises). Housing booms in many advanced and emerging economies were followed by "busts," imposing significant costs on the economy (Jorda et al., 2015). Demirgüç-Kunt and Detragiache (1998) document that banking crises in developed and developing countries are typically preceded by a sharp increase in bank lending to the private sector. Büyükkarabacak and Valev (2010) show that the rise in bank lending to households, rather than to corporations, was the primary culprit in most banking-crisis episodes. More recently, Schularick and Taylor (2012) utilize a logit probability model with panel data, and find that a rapid increase in bank loans to households and businesses significantly increases the probability of a financial crisis within the next five years. Bauer (2014) uses a similar methodology to find that countries with more overvaluation in their housing markets face a significantly higher probability of a sharp correction following a house price boom. He estimates that, for Canada, this probability was around 10 per cent on an annualized basis at

[1] Our stationary model in Section 2 does not capture the trend in the household debt-to-income ratio. As such, we are attributing this to fundamental factors (such as financial innovation), and assessing financial risk based on the household debt "gap" (i.e., short-term deviations of household debt from its long-run linear trend). Of course, the long-run trend in the debt-to-income ratio may itself be indicative of financial imbalances. This would not alter our main conclusions in this paper, however. Since monetary policy generates only temporary effects on the level of household debt, it would likely not be the policy of choice when dealing with long-lasting imbalances captured in the trend.

[2] The long-run average refers to 1970Q1-2015Q2. See http://www.economist.com/blogs/dailychart/2011/11/global-house-prices for more on the global house price index published by the *Economist*.

the end of 2014, which reflects a significant increase from a level of around 2 per cent in 2002.[3]

Crises are costly events, which countries would rather avoid. In many bust episodes observed around the world, asset prices fell sharply, credit availability became more limited, and the economy went into protracted recessions as households and businesses, and the financial institutions that lent to them, went into deleveraging mode. There is ample evidence in the literature showing that recessions following financial crises, especially those that are accompanied by high leverage, are far costlier than the average recession, and last longer as agents try to repair their balance sheets following the crisis, which dampens the recovery (Koo, 2008).

For central banks, the question remains as to whether monetary policy should "lean" against emerging financial imbalances, especially those related to the household sector and housing. On the one hand, leaning could reduce the frequency and severity of financial crises, allowing the economy to largely avoid deep and persistent recessions (such as the recent Great Recession) that impose a substantial welfare loss on agents. These crises and post-crisis recoveries are typically accompanied by large and persistent deviations of inflation from its target and large negative output gaps, which conventional monetary policy cannot easily reverse due to the zero lower bound (ZLB) constraint. Thus, within an inflation-targeting (IT) framework, leaning may ensure that inflation is less likely to miss its target by a large margin in the future due to a financial crisis. Another potential benefit of leaning is to reduce the amplification (i.e., the financial accelerator) effects of high leverage on output and inflation, which increases volatility in the system (Kiyotaki and Moore, 1997; Bernanke et al., 1999).[4] On the other hand, leaning leads to reduced levels of debt, hurting borrowers, who partly rely on leverage to finance their consumption and housing expenditures. Furthermore, leaning may actually lead to greater volatility of inflation relative to its target during normal times, especially if the financial accelerator effects of leverage are muted. Note that in a *flexible* IT framework, such as the one in Canada or the United States, monetary policy "leaning" can be implemented through altering the horizon within which inflation is expected to return to target. For example, in an environment where inflation is below target but household debt developments pose financial stability risks, the path of the policy rate may remain accommodative and yet follow a slightly steeper trajectory with leaning than in its absence. As a result, inflation would be expected to come back to target slightly later than the usual 6- to 8-quarter horizon. From the perspective of a policy-maker, who is trying to minimize a standard loss function that depends on overall inflation and output volatility, leaning can end up leading to higher losses if these short-run inflation and output deviations are large relative to the longer-term benefits of the reduced frequency and severity of crises.[5]

In this paper, we assess the relative benefits and costs of leaning against housing market booms

[3] The literature linking credit developments to subsequent financial crises is vast. See Reinhart and Rogoff (2010), Jorda et al. (2015), and Emanuelsson et al. (2015) for a more comprehensive list of relevant papers.

[4] There may also be a case for leaning, if monetary policy itself is the main source of financial imbalances through the risk-taking channel of the monetary policy transmission mechanism. We abstract from this issue in our paper.

[5] Leaning could also reduce the credibility of central banks, since agents start to view large and persistent deviations of inflation from its target as a weakening of the central bank's commitment to the target. We abstract from this issue in our paper.

within the context of a small-scale, regime-switching New Keynesian model. The core of the model is a simplified version of Iacoviello (2005), where borrowing and lending occur between two types of households, and the borrowers are subject to a borrowing constraint. To this set-up we add the possibility of the economy switching to a crisis regime, which is associated with a significant increase in the risk premium on mortgage lending, a large credit contraction, as well as a steep decline in economic activity and inflation. Crises can be especially costly in our set-up, because the ZLB constraint on the policy rate may become binding, rendering monetary policy ineffective. The probability of switching from the normal to the crisis regime is time-varying, and endogenously determined based on the aggregate household debt level (Woodford, 2012; Ajello et al., 2015). In normal times, housing market booms with a lot of household debt creation can occasionally arise in the model economy due to favourable and persistent housing demand shocks.

We calibrate the model parameters to match certain features of the Canadian economy, and compute the solution of the model using projection methods to better capture the non-linearities inherent in our model: namely, the ZLB constraint on the policy rate and the "asymmetric leaning" of monetary policy (i.e., responding only to positive debt gaps, but not to negative debt gaps).[6] Our solution technique combines the envelope condition method (ECM) of Maliar and Maliar (2013), which iterates on the value function derivatives to find the policy functions, and the sparse grid method as implemented by Klimke and Wohlmuth (2005), in order to keep the problem manageable.

In our benchmark experiment, we find that leaning is not optimal based on a utility-based welfare criterion. In particular, leaning is able to reduce the volatility of economic variables through the reduction of crisis probabilities, but this comes about as the result of the reduction in the average level of household debt, which significantly hurts borrowing households, who partly rely on leverage to finance their consumption and housing expenditures. In particular, leaning implies that interest rates are relatively high during housing booms. This limits the extent of leveraging undertaken by borrowing households, which in turn reduces the volatility of inflation, output and debt mainly through the reduction of crisis probabilities, but also through the reduction in the strength of the financial accelerator mechanism. Through the lens of a utility-based welfare criterion, the benefits of reducing second moments are surpassed by the costs to borrowers implied by the decline in the average debt levels (i.e., first-moment effects).

In a follow-up experiment, we consider the possibility of "symmetric leaning" with respect to debt, whereby negative debt gaps lead to a relatively lower interest rate path than what the standard Taylor rule would imply, along with positive debt gaps leading to a relatively higher interest rate path, as before. Since leaning in this case shifts both the upper and the lower tails of the debt distribution toward the mean, it barely affects the average crisis risk. Thus, the effects of leaning on the crisis probability are muted, making other aspects of leaning more visible. Specifically, the volatility costs of leaning can be large in this case because household debt is significantly more persistent relative to inflation and output, and follows a lower-frequency cycle. Abstracting from the

[6] We also allow for the borrowing constraint on impatient households to be occasionally binding in our computational procedure, but this constraint turns out to be always binding in our simulations.

probability effect, leaning increases the volatility of inflation and output considerably as the interest rate now adjusts vis-à-vis household debt, which follows a different pattern over the business cycle from that of inflation and output. In our set-up, a favourable household demand shock leads to a substantial increase in household debt, but also to a (very small) decline in non-housing consumption and inflation, since agents now value non-housing goods relatively less. This worsens the trade-off for leaning, since the shock in the normal regime drives debt and inflation (and output) in opposite directions. In principle, demand-type shocks moving inflation, output and debt in the same direction should not pose a trade-off for the policy-maker, and leaning should only strengthen the monetary policy response to inflation and output. However, this is not the case in our set-up, because household debt and inflation (and output) are off-phase over the business cycle. In particular, household debt is far more persistent, and peaks with a sizable lag, when the economy is hit with housing demand shocks. Thus, symmetric leaning leads to a significant increase in inflation and output volatility due to this lack of synchronization between household debt and the other macroeconomic aggregates.[7] However, as a complement to our benchmark result, we find that symmetric leaning is welfare improving despite not being effective at reducing the average crisis probability. The key to this result is that the benefits of symmetric leaning are first order (i.e., higher average debt levels and consumption for borrowers), while the costs are second order.

We also show that our results regarding the benefits of leaning are unchanged when we abstract from the ZLB constraint on the policy rate. Recently, several central banks have lowered their deposit rates into negative territory, suggesting that the lower bound on policy rates is not strictly at zero.[8] Furthermore, central banks possess a variety of "unconventional" tools, such as quantitative easing, that could potentially serve a similar purpose as lowering the short-term interest rate. Thus, the ZLB may not be as important a constraint as assumed in our benchmark case. Moreover, there is always the possibility of using macroprudential policies, such as countercyclical loan-to-value (LTV) ratios, instead of monetary leaning to address financial imbalances. As is well known, macroprudential policies can be more targeted toward the source of the imbalance, thereby reducing the adverse side effects of the policy on inflation and output (Alpanda et al., 2014). These considerations may reduce the case for leaning even further.

There is also room to be skeptical regarding the effectiveness of monetary policy leaning in reducing household debt in the first place, especially when one differentiates between the stock and the flow of household debt and considers fixed-rate mortgages (Svensson, 2013; Alpanda and Zubairy, 2014; Gelain et al., 2015). In particular, while monetary tightening would reduce new household loans (i.e., the flow of debt), the real value of the existing stock of debt may actually increase as a result of unexpected disinflation, akin to the debt deflation spiral envisaged in Fisher (1933). In our set-up, leaning is quite effective in reducing household debt. While this is largely consistent with

[7] As Gelain et al. (2015) show, leaning against the household debt "gap" when debt is persistent can introduce indeterminacy into the system even for small values of the leaning parameter, likely for this lack of synchronization.

[8] Theoretically, the ZLB on the policy rate exists because, at negative deposit rates, banks and depositors would switch their reserves and deposits into cash. In practice, this may entail significant storage and transaction costs, so agents may be willing to accept some negative interest on their liquid holdings.

the findings in the cross-country study of Bauer and Granziera (2016), the benefits of leaning in our model may be somewhat upward biased, yet not enough to tip the scale in favour of leaning.

An additional insight can be gained by contemplating a higher inflation target of 2.5 per cent instead of the 2 per cent used in our benchmark case. Interestingly, the higher target is associated with higher welfare in our set-up relative to the benchmark case. This is partly due to the slight reduction in the probability of a crisis, but, more importantly, it is due to the sizable reduction in the frequency of hitting the ZLB. It should be noted, however, that the higher inflation target does not change the implications of leaning for welfare; in other words, leaning against household debt is still not optimal with the higher inflation target.

1.1 Related literature

Monetary policy leaning against household imbalances has received considerable attention in the literature using extensions of the Iacoviello (2005) set-up.[9] These papers do not incorporate the possibility of a crisis regime, capturing the need for monetary leaning through the financial accelerator effects of household debt. The justification for an active policy against financial imbalances, and the reason why policy in the form of leaning could potentially raise welfare, is purely due to the pecuniary externality arising from the borrowing constraint and the financial accelerator mechanism (Lorenzoni, 2008; Korinek, 2011; Bianchi and Mendoza, 2010). In particular, a change in asset prices affects the borrowing constraints of all borrowers, but this side effect is not internalized by a single agent who is deciding whether to purchase more housing through additional borrowing. In contrast, our model features an additional, and potentially more important, type of externality that arises due to the effect of aggregate household debt on the probability of a crisis. In particular, each agent's debt level is small relative to the aggregate; hence, agents, although they are aware of the link between aggregate debt and crisis probabilities, do not internalize their own debt's contribution to the overall crisis probability.

Our paper is closest to Ajello et al. (2015), who also consider optimal monetary leaning within the context of a simple New Keynesian model with an endogenous probability of crises tied to the level of credit. We differ from, and to some degree complement, their paper, in important ways. First, we use a standard infinitely-lived agent set-up in our model, while Ajello et al. (2015) consider only a two-period economy. A two-period set-up may potentially bias the results against leaning. As they also acknowledge in a footnote, leaning today would have benefits in terms of reducing the crisis probability for an extended period of time, since household debt levels are very persistent. Second, Ajello et al. (2015) do not include any shocks in their model except for the crisis shock itself. This could also potentially bias the results against leaning, because shocks (such as housing demand shocks, as we have in our model) introduce an asymmetry into the model due to the convex functional form of the crisis probability in the relevant region of debt. In particular, favourable shocks that raise the household debt level also increase the probability of a crisis, but more so than

[9] A very partial list includes Basant Roi and Mendes (2007), Christensen and Meh (2011), Rubio (2011), Gelain et al. (2013), Lambertini et al. (2013), Alpanda and Zubairy (2014), Gelain et al. (2015), and papers cited therein.

the decline in crisis probability one would observe with adverse shocks. If these non-crisis shocks are normally distributed, optimal policy would feature more leaning in absolute value with respect to positive shocks than with negative shocks. Thus, the optimal level of leaning is likely to be stronger than the 3 basis points (bps) found in the benchmark case of Ajello et al. (2015), when there are other shocks present in the economy apart from the crisis shock itself.[10] Third, Ajello et al. (2015) link the macro variables in the model to the level of credit in reduced form, similar to Woodford (2012), while we use the standard borrowing constraint framework in Iacoviello (2005) to capture these links. Thus, in our set-up, leaning has the additional benefit of reducing the financial accelerator effects of leverage, apart from the decline in crisis probability.[11] Finally, Ajello et al. (2015) assume that agents do not have rational expectations in terms of understanding how changes in aggregate debt affect the probability of switching to the crisis regime, and assume that agents view the crisis probability as a constant, while agents in our set-up are fully rational. Thus, although agents cannot by themselves change the probability of a crisis and treat the crisis probability as an externality in our set-up, they know that once a positive housing demand shock hits, the crisis probability would increase in the medium term.

The next section introduces the model. Section 3 describes the calibration and the computation of the model. Section 4 presents the results, and Section 5 concludes.

2 Model

The model is a closed-economy, regime-switching dynamic stochastic general-equilibrium (DSGE) model with housing and household debt. Similar to Iacoviello (2005), there are two types of households in the economy, patient and impatient households (i.e., savers and borrowers), and the borrowing of impatient households is constrained by the collateral value of their housing. The aggregate borrowing level affects the probability of switching to the crisis regime, similar to Woodford (2012) and Ajello et al. (2015). The rest of the model is standard. On the production side, goods producers rent labour services from each household to produce an output good that can be used for consumption. Goods prices are sticky due to the presence of Rotemberg (1982) price adjustment costs. Monetary policy is conducted via a Taylor rule, with the policy rate being subject to the zero lower bound.

[10] Note, however, that the ZLB constraint may also introduce an additional asymmetry into the model in the presence of non-crisis shocks, which could move optimal leaning in the other direction. In particular, adverse demand shocks (which also reduce debt) would get the economy closer to the ZLB, which the policy-maker may want to avoid as much as possible, leading to a stronger policy response.

[11] In their appendix, Ajello et al. (2015) also consider a feedback effect from debt to output in an ad hoc fashion, and show that the extent of optimal leaning in this case would be far larger than in their baseline.

2.1 Households

2.1.1 Patient households (savers)

The economy is populated by a unit measure of infinitely-lived savers, whose intertemporal preferences over consumption, $c_{P,t}$, housing, $h_{P,t}$, and labour supply, $n_{P,t}$, are described by the following expected utility function:[12]

$$E_t \sum_{\tau=t}^{\infty} \beta_P^{\tau-t} \left(\log c_{P,t} + \xi_t \log h_{P,t} - \frac{n_{P,t}^{1+\vartheta}}{1+\vartheta} \right), \tag{1}$$

where t indexes time, $\beta_P < 1$ is the time-discount parameter, and ϑ is the inverse of the Frisch-elasticity of labour supply. ξ_t is a housing preference shock, and follows an exogenous AR(1) process:

$$\log \xi_t = (1 - \rho_\xi (r_t)) \log \bar{\xi} + \rho_\xi (r_t) \log \xi_{t-1} + \varepsilon_{\xi,t}, \text{ with } \varepsilon_{\xi,t} \sim N(0, \sigma_\xi), \tag{2}$$

where the persistence parameter of the shock switches based on the economic regime, $r_t \in \{0, 1\}$, and 0 denotes the normal regime, while 1 denotes the crisis regime. We assume that the shock's persistence is lower in the crisis regime (i.e., $\rho_\xi(1) < \rho_\xi(0) = \rho_\xi$), ensuring that housing demand returns faster to its long-run value of $\bar{\xi}$ following a crisis relative to the normal regime.[13]

The patient households' period budget constraint is given by

$$c_{P,t} + q_t (h_{P,t} - h_{P,t-1}) + \frac{B_t}{(1 + \chi(r_t)) P_t} + \frac{D_t}{P_t} \leq w_{P,t} n_{P,t} + \frac{R_{t-1} B_{t-1}}{P_t} + \frac{R_{t-1}^m D_{t-1}}{P_t} + \frac{\Pi_t}{P_t} + T_t, \tag{3}$$

where q_t denotes the relative price of housing, P_t is the aggregate price level, $w_{P,t}$ is the real wage rate for patient households, and Π_t denotes the nominal profits of goods producers, which patient households receive in a lump-sum fashion, just as they receive the transfers T_t. Patient households lend to impatient households and the government in nominal amounts of D_t and B_t, respectively, and receive predetermined nominal interest rates of R_t^m and R_t in return.

The $\chi(r_t)$ term in the budget constraint above is a portfolio preference term, which drives a wedge between the expected returns from government bonds and mortgages. We assume that $\chi(r_t)$ takes on only two values: 0 in the normal regime and $\bar{\chi} > 0$ in the crisis regime. Thus, in the crisis regime, savers are incentivized to increase their holdings of government debt (i.e., flight-to-safety).[14]

The transition between the normal and crisis regimes is governed by a Markov chain with tran-

[12] Following Iacoviello (2005), we normalize the size of each type of household (patient and impatient) to a unit measure, and determine the economic importance of each type by their respective shares in labour income.

[13] This assumption is not crucial for any of the results, but lowers the incidence of consecutive crises in our simulations. In particular, with high persistence, the debt overhang triggered by the first crisis significantly outlasts the duration of the crisis regime, keeping the risk of a second crisis elevated even after the economy switches back to the normal regime. With lower persistence, imbalances are significantly reduced over the crisis duration.

[14] See Smets and Wouters (2007), Alpanda (2013), and Alpanda et al. (2014) for more on this portfolio preference term.

sition probabilities given by

	Normal ($r_t = 0$)	Crisis ($r_t = 1$)
Normal ($r_{t-1} = 0$)	$1 - \gamma_t$	γ_t
Crisis ($r_{t-1} = 1$)	δ	$1 - \delta$

, (4)

where the probability of switching from the crisis to the normal regime, δ, is assumed to be a constant, as in Woodford (2012). γ_t is the time-varying probability of having a crisis in period t conditional on being in the normal regime in $t-1$. This transition probability is endogenously determined based on the debt position of households with a logit specification:

$$\gamma_t = \frac{\exp\left(\omega_1 \frac{d_{t-1}}{\bar{d}} - \omega_2\right)}{1 + \exp\left(\omega_1 \frac{d_{t-1}}{\bar{d}} - \omega_2\right)}, \tag{5}$$

where ω_1 and ω_2 are parameters, $d_{t-1} = D_{t-1}/P_{t-1}$ denotes real debt brought from the previous period, and $\bar{d}$ is the steady-state value of debt in the normal regime.[15]

The patient households' objective is to maximize utility subject to the budget constraint and appropriate No-Ponzi conditions. The first-order condition with respect to consumption equates the marginal utility gain from consumption to the marginal cost of spending out of the budget captured by the Lagrange multiplier, $\lambda_{P,t}$. Similarly, the optimality condition for labour equates the marginal rate of substitution between labour and consumption to the real wage rate. The optimality condition for housing equates the marginal cost of acquiring a unit of housing to the marginal utility gain from housing services and the discounted value of expected capital gains, which can be written as

$$q_t = \xi_t \frac{c_{P,t}}{h_{P,t}} + E_t \left[\beta_P \frac{\lambda_{P,t+1}}{\lambda_{P,t}} q_{t+1} \right]. \tag{6}$$

The first-order condition for government bonds equates the marginal utility cost of forgone consumption from saving to the expected discounted utility gain from the resulting interest income:

$$1 = E_t \left[\beta_P \frac{\lambda_{P,t+1}}{\lambda_{P,t}} \frac{R_t \left(1 + \chi\left(r_t\right)\right)}{\pi_{t+1}} \right]. \tag{7}$$

Note that an increase in the risk-premium term would lead patient households to reduce their consumption expenditures, and increase their savings in the form of risk-free government bonds. Since the latter are in zero supply in our model, patient households will instead increase their housing demand while reducing consumption expenditures. Due to arbitrage between government bonds and lending to impatient households, the equilibrium mortgage rate is given by $R_t^m = R_t \left(1 + \chi\left(r_t\right)\right)$.

[15] We could instead make the crisis probability depend on the house price gap (Bauer, 2014), rather than the credit gap. This would increase the number of states in the model, and therefore raise the computational burden. We thus use the credit gap in the logit specification, which broadly conforms with the specifications in Schularick and Taylor (2012) and Ajello et al. (2015).

Thus, patient households would also be happy to increase their savings in the form of mortgages as their returns increase with a positive risk premium, but the demand of impatient households for mortgages declines substantially, driving equilibrium lending levels down.

2.1.2 Impatient households

The economy is also populated by a unit measure of infinitely-lived impatient households. Their utility function is identical to that of patient households, except that their time-discount factor is assumed to be lower in order to facilitate borrowing and lending across these two types of consumers; hence, $\beta_I < \beta_P$. The impatient households' period budget constraint is given by

$$c_{I,t} + q_t \left(h_{I,t} - h_{I,t-1} \right) + \frac{R^m_{t-1} D_{t-1}}{P_t} \leq w_{I,t} n_{I,t} + \frac{D_t}{P_t}, \tag{8}$$

where $w_{I,t}$ denotes the wage rate of impatient households.

Impatient households face a borrowing constraint in the form of

$$d_t \leq \rho_d d_{t-1} + (1 - \rho_d) \phi q_t h_{I,t}, \tag{9}$$

where ϕ is the fraction of assets that can be collateralized for borrowing (i.e., the LTV ratio), and ρ_d determines the persistence of debt as in Iacoviello (2015).

The first-order conditions of the impatient households with respect to consumption and labour are similar to those of patient households. For housing, the optimality condition equates the marginal cost of acquiring a unit of housing with its marginal utility and expected capital gains, but the marginal cost is dampened by the shadow gain from the relaxation of the borrowing constraint given the increase in the level of housing. This condition can be written as

$$\left[1 - \mu_t \left(1 - \rho_d \right) \phi \right] q_t = \xi_t \frac{c_{I,t}}{h_{I,t}} + E_t \left[\beta_I \frac{\lambda_{I,t+1}}{\lambda_{I,t}} q_{t+1} \right], \tag{10}$$

where $\mu_t \geq 0$ is the Lagrange multiplier on the borrowing constraint, which is strictly positive when the borrowing constraint is binding and is equal to 0 when it is not. Similarly, the optimality condition for borrowing is given by

$$1 - \mu_t = E_t \left[\beta_I \frac{\lambda_{I,t+1}}{\lambda_{I,t}} \left(\frac{R^m_t}{\pi_{t+1}} - \mu_{t+1} \rho_d \right) \right], \tag{11}$$

which equates the marginal gain from borrowing minus the shadow price of tightening the borrowing constraint with the expected interest costs. Note that borrowing today relaxes the borrowing constraint in the future as well due to the persistence term; hence, this benefit is subtracted from the expected marginal cost term on the right-hand side.

2.2 Goods production

There is a unit measure of monopolistically competitive goods producers indexed by j. Their technology is described by the following production function:

$$y_t(j) = z n_{P,t}(j)^{\psi} n_{I,t}(j)^{1-\psi} - f_t, \tag{12}$$

where $y_t(j)$ denotes output of firm j, ψ is the share of patient households in the labour input, z is the level of aggregate total factor productivity, and f_t is a time-varying fixed cost of production.

Goods are heterogeneous across firms, and are aggregated into a homogeneous good by perfectly-competitive final-goods producers using a standard Dixit-Stiglitz aggregator. The demand curve facing each firm is given by

$$y_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\eta} y_t, \tag{13}$$

where y_t is aggregate output, and η is the elasticity of substitution between the differentiated goods. Thus, the gross markup of firms at the normal regime steady-state is given by $\theta = \eta/(\eta-1)$.[16]

Firm j's profits in period t are given by

$$\frac{\Pi_t(j)}{P_t} = \frac{P_t(j)}{P_t} y_t(j) - w_{P,t} n_{P,t}(j) - w_{I,t} n_{I,t}(j) - \frac{\kappa}{2}\left(\frac{P_t(j)}{\pi^* P_{t-1}(j)} - 1\right)^2 y_t, \tag{14}$$

where price stickiness is introduced through quadratic adjustment costs with κ as the level parameter, and π^* is the inflation target.

A firm's objective is to choose the quantity of inputs, output and its own output price each period to maximize the present value of profits (using the patient households' stochastic discount factor), subject to the demand function they are facing for their own output from the goods aggregators. The first-order condition for prices yields the following New Keynesian Phillips curve for domestic goods inflation:

$$\left(\frac{\pi_t}{\pi^*} - 1\right)\frac{\pi_t}{\pi^*} = E_t\left[\beta_P \frac{\lambda_{P,t+1}}{\lambda_{P,t}}\left(\frac{\pi_{t+1}}{\pi^*} - 1\right)\frac{\pi_{t+1}}{\pi^*}\frac{y_{t+1}}{y_t}\right] - \frac{\eta - 1}{\kappa}\left(1 - \frac{w_{P,t}}{\theta\psi y_t / n_{P,t}}\right). \tag{15}$$

Note that, at the optimum, the marginal product of each input is equated to its respective marginal cost; hence, the relative demand for the two types of households' labour are related to the two wage rates as follows:

$$\frac{n_{P,t}}{n_{I,t}} = \frac{\psi}{1-\psi}\frac{w_{I,t}}{w_{P,t}}. \tag{16}$$

[16] The fixed cost, f_t, is set to $\theta - 1$ times, aggregate output to ensure that pure profits are equal to zero at all times.

2.3 Monetary policy

Monetary policy is conducted using a Taylor rule on the nominal interest rate, which is subject to the zero lower bound; hence,

$$R_t = \max\left\{1,\ R_{t-1}^{\rho}\left[\overline{R}\left(\frac{\pi_t}{\pi^*}\right)^{a_\pi}\left(\frac{y_t}{\overline{y}}\right)^{a_y}\left(\max\left\{\frac{d_t}{\overline{d}},1\right\}\cdot\mathbf{1}_{\{r_t=1\}}\right)^{a_d}\right]^{1-\rho}\right\}, \tag{17}$$

where ρ is the smoothing term in the Taylor rule, and a_π, a_y, and a_d are the long-run response coefficients for inflation, the output gap, and the household debt gap, respectively. $\overline{R}$, $\overline{y}$, and $\overline{d}$ denote the steady-state values of the nominal interest rate, output, and household debt in the normal regime.

Note that the second *max* operator on the right-hand side of the equation ensures that leaning is implemented in an asymmetric fashion; i.e., it is active only when the household debt gap is positive, and not when the debt gap is negative. Thus, interest rates follow the standard Taylor rule when the debt gap is negative, but they are slightly higher than what the standard Taylor rule would prescribe when the debt gap is positive.[17]

Finally, we introduce an indicator function $\mathbf{1}_{\{r_t=0\}}$, which is equal to 1 in normal times and 0 during crisis episodes. As a result, we lean against household debt only during normal times, and not during crisis periods. While the consequences of leaning during a crisis are small, this would have a negative welfare impact, since leaning hurts borrowers, especially when they are already suffering from a contracting economy. Thus, introducing this additional asymmetry improves the chances for leaning to be beneficial.

2.4 Market clearing conditions and timing of events

The goods market clearing condition is given by

$$c_{P,t} + c_{I,t} = y_t - \frac{I_\pi \kappa}{2}\left(\frac{\pi_t}{\pi^*} - 1\right)^2 y_t, \tag{18}$$

where $I_\pi \in [0,1]$ determines the extent to which the price adjustment costs reduce the real resources in the economy, while the rest are treated as lump-sum transfers to patient households.[18]

We assume that government bonds are in zero supply; hence, $B_t = 0$ for all t. The stock of housing is assumed to be in fixed supply as in Iacoviello (2005); hence,

$$h_{P,t} + h_{I,t} = \overline{h}. \tag{19}$$

The timing of events is as follows. The economy enters period t with an aggregate state vector

[17] In Section 4, we also consider the implications of symmetric leaning, whereby monetary policy responds to negative, as well as positive, household debt gaps.

[18] The choice of I_π does not qualitatively affect the main results regarding the optimality of leaning, but it does have a quantitative effect on the volatility of the economy.

of $d_{t-1}, h_{P,t-1}, R_{t-1}, \xi_{t-1}$, and r_{t-1}. Note that the past mortgage rate, R^m_{t-1}, is known as well, since $R^m_{t-1} = (1 + \chi(r_{t-1})) R_{t-1}$. Furthermore, the crisis probability in period t, γ_t, is also known, since this depends on the lagged value of aggregate household debt, d_{t-1}. At the beginning of period t, the innovations for the $AR(1)$ housing demand shock, $\varepsilon_{\xi,t}$, as well as the crisis regime, r_t, are realized. Next, agents choose consumption, housing, labour supply, etc., and markets clear. The state vector passed over to period $t+1$ is given by $d_t, h_{P,t}, R_t, \xi_t$, and r_t. The model's equilibrium is defined as prices and allocations, such that households maximize the expected discounted present value of utility and firms maximize expected profits subject to their constraints, and all markets clear.

3 Calibration and Computation

3.1 Calibration

We calibrate the parameters using the steady-state relationships of the model and values typically used in the related literature. Table 1 summarizes the list of parameter values.

The trend inflation factor, π^*, is set to 1.005, corresponding to a 2 per cent annual inflation target. The time-discount factor of patient households, β_P, is set to 0.9925 to match an annualized 3 per cent real risk-free interest rate. The discount factor of impatient households, β_I, is set to 0.97 following Iacoviello and Neri (2010). The parameter ϑ is calibrated to ensure that the Frisch elasticity of labour supply is 0.5, while the level parameter for housing in the utility function, $\bar{\xi}$, is set to 0.12 following Iacoviello and Neri (2010).

The price markup parameter, θ, is set to 1.1 to reflect a 10 per cent net markup in prices. The price adjustment cost parameter, κ, is set to 100, which generates a Phillips curve slope that is largely consistent with estimates in the related DSGE literature. We set I_π equal to 0.5, implying that only half of the price adjustment costs pose a direct burden on real resources, while the rest is transferred back to patient households in a lump-sum fashion. The wage share of patient households, ψ, is set to 0.748, largely in line with Iacoviello (2005) after adjusting the patient households' income share for capital income.

We calibrate the LTV ratio, ϕ, to 0.75, which is close to the average LTV ratio on outstanding mortgages in Canadian data. The persistence parameter in the borrowing constraint, ρ_d, is set to 0.96, based on the deleveraging speed observed in the United States following the financial crisis. Note that this parameter also implies that the average duration of household loans is 25 quarters.

The housing demand shock follows an AR(1) process, for which we have to specify the persistence parameters in each regime, $\rho_\xi(r_t)$, and the standard deviation of its shock innovations, σ_ξ. We set the latter to 0.05, while the persistence parameter in the normal regime, ρ_ξ, is set to 0.985. These two parameters are set to generate housing gaps of a similar magnitude and persistence as those observed in the data (Figure 2). The persistence parameter of the housing shock in the crisis regime is assumed to be half of its value in the normal regime (i.e., $\rho_\xi(1) = \rho_\xi/2$), to ensure that a significant amount of deleveraging occurs during the crisis regime.

For the Taylor rule coefficients, we use the mean values of the prior distributions used in Smets and Wouters (2007). In particular, the response coefficients for inflation and the output gap, a_π and a_y, are set to 1.5 and 0.125, respectively, and the smoothing parameter, ρ, to 0.75. We set the leaning parameter, a_d, to 0 in the base case.

Crisis-related parameters We next turn to the parameters that determine the probability and the severity of crises. The probability of entering a crisis is governed by a logit function, which in turn is characterized by two parameters, ω_1 and ω_2. To obtain these parameters in our context, we make use of the cross-country results in Bauer (2014). That study relates monetary policy shocks and the deviations of house prices from fundamentals to the probability of a crisis characterized by a large house price correction (greater than 10 per cent in real terms). We first compare the relationship between the house price gap in Bauer (2014) and the household debt gap relevant in our model. We construct the latter using a linear trend of real household credit, as in the top panel of Figure 2. The two gap measures are highly correlated, and of comparable magnitudes for the most recent period (see Figure 2 bottom panel). In particular, when we run a simple linear regression of the house price gap on the debt gap (with no constant term), we obtain a regression coefficient of 0.998 with a highly significant p-value.[19] We then estimate the mapping from the Bauer model's implied crisis probability to the debt gap. For this, we run the regression

$$\log\left(\frac{p_t}{1-p_t}\right) = -\omega_1 + \omega_2\left(\frac{d_t}{\bar{d}_t}\right) + \varepsilon_t, \tag{20}$$

where p is the probability of a large real house price correction in Bauer (2014). This results in parameter estimates of $\omega_1 = 9.68$ and $\omega_2 = 5.07$. Figure 3 shows the relationship. The unexplained variation visible in the figure is due to other factors that are present in the Bauer (2014) model but not in ours. However, most of the systematic variation in the probability is explained by the credit gap.[20]

In the model, the severity of a crisis can be measured by the cumulative loss in output during the crisis, which in turn is equal to the product of the average output fall per period and the duration of the crisis regime.[21] Both of these aspects are difficult to measure in the data, since the size of the output loss depends on the underlying trend assumed for real output. If, for example, the crisis has a permanent negative impact on the level or the growth rate of trend output, the cumulative output loss might be very large, potentially infinite. On the other hand, if we focus very narrowly on the acute crisis period and define recovery as the return of output to its previous peak level, then

[19] We run a regression $(q_t/\bar{q}_t) = \beta\left(d_t/\bar{d}_t\right) + \varepsilon_t$, where $\bar{q}_t$ indicates the house price trend found in Bauer (2014) and $\bar{d}_t$ refers to the linear trend in household debt.

[20] Alternatively, we could have picked the parameters in Ajello et al. (2015), which are in line with the data presented in Schularik and Taylor (2012). Given our focus on Canada and the current housing market risk, we preferred the estimate in Bauer (2014). Note also that the logit probability function implied by Ajello et al. (2015) is substantially flatter in the relevant range of debt; thus, using their specification would have reduced the case for leaning even further.

[21] In our set-up, a longer average crisis duration also leads to a larger drop in the output gap per period, everything else equal, because labour supply falls more when households anticipate a longer crisis, leading to a larger output loss.

the cumulative output loss would be relatively small. Our model abstracts from the possibility that crises may have permanent effects on the level or the growth rate of output. We try to compensate for this by conducting an experiment where we increase the cumulative output loss, bringing it closer to the upper bound of estimates in the literature.

Our target range for the cumulative output loss from a crisis is between 7.9 and 27.7 per cent. The lower bound of this is based on Schularick and Taylor (2012), who find that the cumulative real output loss over the five years following a financial crisis in the post-WWII period is 7.9 per cent. The upper bound is based on the recent United States experience. In particular, the cumulative difference between the pre-crisis linear trend of output and actual output is around 27.7 per cent for the period 2008Q2 - 2015Q2.[22]

To get a reasonable mid-level number in our target range, we set δ to 0.1, implying an average crisis duration of 10 quarters. Moreover, we set $\bar{\chi}$ to 0.0125, which implies that the spread between mortgages and the policy rate is 5 per cent on an annualized basis for the duration of a crisis. For our benchmark model, these parameter choices generate a cumulative output loss of 14.1 per cent. It should be noted that this number represents a reasonable mid-level, since the boundaries focus on GDP and therefore include investment, which is absent from our model.[23] If we consider the cumulative real aggregate consumption loss following the Great Depression for the United States (using the narrow definition of consumption declining and returning to its previous peak level), we find a consumption downturn that lasts 11 quarters, with a peak decline of 2.7 per cent and a cumulative decline of 14.9 per cent. Moreover, a comprehensive cross-country study by the International Monetary Fund (IMF 2009) confirms that private consumption losses are, on average, much smaller than output losses following a financial crisis (see their Figure 8).

3.2 Computation

We compute the solution of the model using projection methods to better capture the non-linearities inherent in our model: namely, the ZLB constraint on the policy rate and the asymmetric leaning of monetary policy with respect to the household debt gap.[24] Our solution technique combines the envelope condition method (ECM) of Maliar and Maliar (2013), which iterates on the value function derivatives to find the policy functions, and the sparse grid method as implemented by Klimke and Wohlmuth (2005), in order to keep the size of the problem manageable. Details regarding the computational strategy are provided in the appendix.

[22] This may also be an underestimate, since it excludes any further output losses from the crisis post-2015Q2.

[23] Note that investment appears to be the biggest contributor to the decline in real GDP, since Schularick and Taylor (2012) also find a cumulative real investment loss over the five years following a crisis of 25.7 per cent.

[24] As noted earlier, even though we allow the borrowing constraint to be occasionally binding, it always binds in our simulations.

4 Results

In this section, we first analyze the dynamics of the model economy following a switch to the crisis regime and a housing demand shock using impulse responses. We then investigate the optimal degree of leaning in the Taylor rule, which maximizes welfare based on a weighted average of households' utility, or minimizes a standard loss function based on the variance of inflation and output. We conclude this section by conducting robustness checks on our welfare results; in particular, we investigate how the benchmark results would change if (i) the ZLB did not pose a constraint on monetary policy, (ii) leaning was conducted symmetrically with respect to both positive and negative debt gaps, (iii) the central bank had a 2.5 per cent inflation target, instead of 2 per cent, and (iv) the cumulative loss was much larger than in the benchmark case.

4.1 Model dynamics

In Figure 4, we present the impulse responses of model variables following a switch to the crisis regime, which is induced by a sharp rise in the risk premium on mortgage debt for 10 periods. Note that even though the crisis lasts for 10 periods in our example, agents in the model place a $\delta = 10$ per cent probability on returning to the normal regime in each period. Hence, this is not a "perfect foresight" exercise, whereby agents know exactly that the crisis is going to last 10 periods. The increase in the risk premium during a crisis leads to a persistent decline in the borrowing of impatient households, who, as a result, reduce their demand for housing and consumption goods. Patient households reduce their consumption as well, while they use up savings to purchase the housing offered by impatient households. Nevertheless, this is not enough to reverse the fall in the overall demand for housing, leading to a fall in house prices. The decline in overall demand for consumption goods leads to a significant fall in (non-housing) output and inflation as well, which forces the central bank to cut the policy rate, until it almost reaches the ZLB.[25] Given the sizable smoothing parameter in our Taylor rule, it takes several quarters before the policy rate reaches close to the ZLB. After the impact period, the fall in the policy rate is accompanied by a fall in the mortgage rate, R^m, but the latter still hovers around its normal-regime steady-state value due to the risk premium in the crisis regime.

Figure 5 presents the impulse responses of model variables to a positive housing demand shock, conditional on staying in the normal regime throughout the impulse horizon.[26] We first consider the case when monetary policy does not lean against household debt (i.e., $a_d = 0$). The increase in housing demand leads to a rise in house prices, which relaxes the borrowing constraint of impatient households, allowing them to increase their borrowing persistently and purchase more housing. Since housing now provides relatively more utility, impatient households also substitute away from

[25] In general, hitting the ZLB depends on the interest rate level before the crisis and the severity of the crisis. In Figure 13, we consider a more severe crisis scenario where the policy rate is quickly reduced to the ZLB and stays there through the duration of the crisis.

[26] Again, agents expect that the economy could switch to the crisis regime with γ_t probability, but this is never realized over the impulse horizon.

consumption in order to increase their housing purchases. As a result, the demand for non-housing consumption goods falls, which leads to a decline in aggregate non-housing output and inflation. The patient households increase their consumption expenditures slightly, due to a declining relative price of consumption goods, while they sell part of their housing stock to the borrowers at higher prices. Note that the housing demand shock has a significant effect on the dynamics of debt, while its impact on overall output and inflation is quantitatively small. Furthermore, the peak response of debt is significantly later than the peak responses of inflation and output, due to the persistence term in the borrowing constraint of impatient households.

The main cost of monetary policy leaning arises from the successful reduction of household debt during normal times, as shown in Figure 5, where the red dashed line depicts the case with the leaning parameter a_d set equal to 0.024.[27] This reduction of household debt leads to first-order effects on the consumption of impatient households, and therefore their welfare. Moreover, strong leaning as shown here increases the volatility of output and inflation during normal times, since the business cycle is out of phase with the debt cycle. This disconnect between the peak responses of output and inflation compared to household debt also matters, though less than the debt reduction itself. Ultimately, as debt builds up, the policy rate increases and stays above its benchmark case for an extended period. As a result, leaning is relatively successful in reducing the debt cycle, with the peak magnitude of the household debt gap being significantly reduced. But, this comes at a possible cost to macroeconomic stability; in particular, inflation stays below its target for a prolonged period of time and by a large margin. Similarly, the output gap closes relatively more slowly than in the benchmark case. The impulse responses of macroeconomic variables are also visibly more volatile, since the interest rate now tries to respond to credit, along with inflation and the output gap, even though the former is clearly off-phase over the business cycle relative to the latter two variables.

These two impulse responses in combination give us a sense of what drives the main results of the paper. First, the size of the two main forces in the model are clearly disproportionate, since crisis events are huge in comparison to regular business cycle movements. Thus, a policy that successfully reduces the risk of a crisis has major implications for the volatility of inflation and output. Second, leaning can clearly influence household debt volatility; however, that will hurt borrowers during housing market boom periods. This leaves us with two quantitative questions: (a) Is leaning successful enough, when reducing the probability of a crisis, to overcome the potentially higher volatility in normal times? (b) Do the benefits of lower volatility from leaning outweigh the welfare costs to borrowers? Our answers to these questions are yes to (a) and no to (b).

4.2 Optimal degree of leaning

In this subsection, we analyze the optimal degree of leaning within the context of the Taylor rule in (17) and variations thereof. In order to do this, we randomly generate housing demand shocks, and

[27] As we explain later, we consider a_d over the interval $[0, 0.024]$ in our simulations. Our upper bound for a_d implies a near 200 bps increase in the policy rate in annualized terms given a positive 23 per cent debt gap, which is the biggest gap we found in the recent data.

run 125 simulations of our model, each for a length of 4,000 periods.[28] We start each simulation from the normal-regime steady state, and burn the initial 300 periods.

To find the optimal degree of leaning, we search for the Taylor rule parameter, a_d, that maximizes social welfare, W, which is defined as the weighted average of the utility-based welfare measures for the patient and impatient households, and is given by

$$\max W = (1 - \beta_P)\, V_P + (1 - \beta_I)\, V_I. \tag{21}$$

Note that the weights are picked so that the same constant consumption stream would result in equal welfare across the two types of agents, following Lambertini et al. (2013). The welfare of each household type is given by

$$V_i = \frac{1}{1-\beta_i} \frac{1}{N \times T} \sum_{j=1}^{N} \sum_{t=0}^{T} u\left(c_{i,t}^j, h_{i,t}^j, n_{i,t}^j\right) \text{ for } i \in \{P, I\}, \tag{22}$$

where N denotes the number of simulations and T is the number of periods in each simulation. To make it easier to comprehend this measure, we convert it into lifetime consumption equivalence numbers (LTCE) relative to the no-leaning benchmark case. So, the reported welfare numbers provide the relative gain (or loss) in consumption needed to compensate the households for changing from a policy regime without leaning to one with leaning. Across experiments, we keep as our utility reference point the case of no leaning under the base-case calibration.

To contemplate volatility benefits, we also consider a standard loss function that depends on the variance of inflation and the output gap with equal weights:

$$L = var\,(\pi) + var\,(y)\,. \tag{23}$$

In order to discuss the success of leaning, we also determine the probability of entering into a crisis regime conditional on being in normal times under each policy. To obtain a good approximation of this probability, we integrate the logit function using the estimated Epanechnikov kernel density of household debt conditional on being in normal times.

In all the experiments below, we consider variations in a_d over the interval $[0, 0.024]$. Our range for a_d might appear small but that is misleading. To see this, consider the case of a_d equal to 0.024 and a positive 23 per cent debt gap, which is the biggest gap we found in the recent data. These would imply an additional 200 bps increase in the policy rate on an annualized basis. We think that this is a reasonable upper bound for the degree of leaning a central bank might implement.

4.2.1 Benchmark case

In our benchmark case, leaning is found to be undesirable when we use the welfare criterion (see Figure 6). Specifically, as we increase the leaning parameter, a_d, welfare goes down monotonically

[28] To generate the switching between regimes, we also use random numbers picked from a uniform distribution.

by up to 0.23 percentage points in terms of lifetime consumption equivalents. In contrast to this welfare result, we find that leaning is successful when it comes to reducing volatility in the economy, thereby reducing the value of the loss function. Specifically, the unconditional standard deviations of output, inflation and debt, as well as the probability of a crisis, decrease significantly. The key driver behind all of this is the success of leaning in dealing with household debt.

As we already saw in the previous section, leaning dampens the fluctuations of household debt by discouraging impatient households from borrowing. This effect is successful enough to shift the right tail of the distribution to the left, as shown in Figure 7, leading to reductions in the debt mean and volatility. In turn, the probability of a crisis is reduced from 1.04 per cent to 0.93 per cent. On the other hand, the lower average debt levels hurt borrowers by reducing their consumption and housing expenditures. This affects expectations and creates disincentives to work, while reducing inflation pressures, all of which contribute to a lower mean level of inflation with a negligible impact on output.

It is instructive to take a close look at the lower volatility of inflation and output, since these are clearly visible improvements with the potential to outweigh the cost of leaning. To better understand these reductions, we conditioned volatilities by the crisis regime (see Figure 8). From this, we clearly see that the effect of leaning on the crisis probability is much stronger than its effect on the crisis-specific standard deviations of output and inflation. In particular, conditional on each regime, leaning creates a very slightly U-shaped response for the standard deviation of output and inflation, since small amounts of leaning are beneficial, but leaning with $a_d > 0.01$ creates additional volatility. What increases volatility most, however, is alternating between the two regimes, dominating the aforementioned U-shapes. Therefore, reduction of the crisis probability is the key driver behind the lower overall second moments.

Returning to the welfare consequences of leaning, the first-order losses hurting the borrowers clearly dominate the second-order benefits of reducing the economy's volatility. In effect, what makes leaning a success by the loss-function criterion (i.e., lower debt level, which reduces the crisis probability) also makes it a failure in the welfare realm, since it inflicts first-order pain on the borrowers.[29]

4.2.2 Other experiments

To highlight the importance of various aspects of the model, we conduct additional experiments and show the results in comparison to the benchmark case.

What is the importance of the ZLB?

We start by assessing the importance of the ZLB constraint on the policy rate. Figure 9 compares the results with the ZLB constraint (blue solid line) and without (red dashed line). As would be expected, welfare is higher without the zero lower bound. Relative to the benchmark economy, the

[29] We also considered a steeper logit function, within the realm of reasonable estimates based on the Bauer (2014) model, since this would increase the volatility benefits while keeping the costs of leaning about the same. However, qualitatively all the results were similar to our benchmark case.

benefits are 6 bps in LTCE.[30] However, with regards to leaning, all the qualitative results stay intact and the effects are found to be small quantitatively.

Note that, without the ZLB, the mean level of debt and the probability of a crisis are slightly higher relative to the benchmark case with the ZLB. They also decrease slightly less than in the benchmark case as the leaning parameter increases. These results are related to the fact that, without the ZLB constraint, monetary policy is better able to support borrowers during a crisis, as shown by the lower volatility of debt relative to the benchmark economy. As a result, crises are not as costly as in the benchmark case, and need not necessarily be avoided as much. Hence, with no ZLB, more debt is accumulated than otherwise, and the average crisis risk goes up regardless of the degree of leaning.

In summary, the costs of leaning are slightly lower without the ZLB constraint; however, leaning still leads to significant debt-level reductions. Thus, leaning is still found to be non-optimal. A contributing factor to this result is that "cleaning" through expansionary monetary policy after a crisis is easier when the ZLB constraint is not a concern. These results suggest that taking negative interest rates and unconventional monetary policy options into account is unlikely to change our benchmark assessment regarding leaning prior to the crisis (vs. cleaning after the crisis).

Is leaning symmetrically a good alternative?

We next consider the case of symmetric leaning. In our experiments so far, we allowed leaning for only positive debt gaps, whereas we now allow monetary policy to systematically respond to both positive and negative debt gaps. The results of this exercise are summarized in Figure 10. Symmetric leaning has two counteracting forces on the probability of a crisis. In particular, it increases the mean level of debt in the economy, while reducing its variance. The overall effect is a nearly constant crisis probability despite leaning. Hence, our model suggests that, in order to reduce the probability of crises, monetary policy needs to lean asymmetrically, responding only to positive debt gaps, as in our benchmark model.

The results from this experiment are also insightful in understanding the costs of leaning asymmetrically: namely, by raising the average debt level and lowering the standard deviation of debt, symmetric leaning helps the borrowers. In particular, for low levels of leaning, the corresponding welfare gains are actually higher than the losses associated with the increases in the volatility of inflation and output, as well as with the declines in their levels. In particular, for $a_d = 0.024$, leaning results in an overall 15 bps increase in LTCE relative to no leaning. Thus, the costs of asymmetric leaning stem from the fact that they are born by the borrowers and are first order, while the benefits are second order and mostly go to the lenders.[31]

[30] In the welfare plot, we show both the benchmark and the alternative welfare measures relative to the deterministic steady-state welfare in the normal regime.

[31] In a model similar to ours, but without regime switching, Gelain et al. (2015) find that symmetric leaning leads to higher volatility of output and inflation. Based on that insight, they suggest that the debt gap part of the Taylor rule be replaced with credit growth in order to reduce the normal time volatility increase associated with leaning. Doing so for our benchmark calibration does not lead to qualitatively different results, since the debt-probability effect dominates the in-regime gains.

Could targeting a higher rate of inflation reduce the cost of leaning?

Eliminating the ZLB constraint may not be a policy choice, but changing the inflation target certainly is. In this subsection, we consider the implications of raising the inflation target from 2 per cent to 2.5 per cent on the costs of leaning. As we can see in Figure 11, a higher inflation target results in slight welfare increases in our set-up, and is mildly helpful when it comes to reducing volatility in the economy, as measured by the loss-function criterion.[32] What is particularly interesting is that a higher inflation target implies a higher average debt level but a lower level of volatility of debt, which leaves the probability of a crisis only marginally lower.

It is instructive to realize that a higher inflation target leads to a higher standard deviation of inflation, accompanied by a lower one for output, as would be expected. What is not expected is that these two effects cancel each other out, at least for low levels of leaning. Raising the inflation target also considerably lowers the probability of hitting the ZLB (see Figure 12). In the benchmark model, without leaning, the ZLB is hit 1.81 per cent of the time. Raising the inflation target by 50 bps reduces that frequency to 0.12 per cent. Moreover, under the benchmark policy, leaning raises the frequency of hitting the ZLB by 8 bps, while, with a higher inflation target, leaning actually reduces the frequency. Focusing on the benchmark economy first, with more leaning the average inflation goes down, implying that the average interest rate in the economy also declines. This, in turn, leads to a higher chance of entering the crisis period with a low rate that triggers the ZLB. In contrast, with a higher inflation target, the average policy interest rate is higher and further away from a rate that is associated with hitting the ZLB during a crisis. Moreover, leaning concentrates the interest rate closer to that mean, further decreasing the likelihood of a ZLB situation arising.

How important is the size of the crisis?

Finally, one concern might be that the size of our crisis shock is too small to make the probability reduction matter. To address this, we reset key parameters in our model to increase the aggregate consumption loss during a crisis. However, since households are consumption smoothing in our model, a pure and very large increase in the risk-premium shock would be insufficient: households would increase their labour supply to compensate. We therefore add an adverse consumption shock in the crisis regime. Concretely, we change preferences to

$$E_t \sum_{\tau=t}^{\infty} \beta_j^{\tau-t} \left(\nu_t \log c_{j,t} + \xi_t \log h_{j,t} - \frac{n_{j,t}^{1+\vartheta}}{1+\vartheta} \right) \text{ for } j \in \{P, I\},$$

$$\text{where } \nu_t = \nu(r_t) = \left\{ \begin{array}{c} 1 \\ 0.985 \end{array} \text{ for } r_t = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}. \tag{24}$$

Thus, during a crisis episode, households cut their consumption relative to non-crisis times.[33] More-

[32] In this experiment, we reset π^* both in the pricing problem and the resource constraint. This biases the welfare toward a higher inflation target, since the resource cost might be relative to the previous target of 2 per cent, which would add to the cost of a policy change.

[33] When measuring welfare for this case in LTCE, we now have to divide the regular expression by a term $E\left(\sum_{t=0}^{\infty} \beta^t \nu_t\right)$.

over, we increase the spread shock size slightly to $\bar{\chi} = 1.3$ per cent, from $\bar{\chi} = 1.2$ per cent before. Both these changes in the consumption and spread shocks would make downturns more dramatic, but as a side effect they would also generate a much wider debt gap distribution relative to what is observed in the data. To deal with this issue, we decrease the persistence and the standard deviation of the housing demand shock to $\rho_\xi = 0.9725$ and $\sigma_\xi = 0.04$. The end result of these adjustments is primarily on the cumulative output loss associated with an average crisis, which increases from 14.1 per cent in our benchmark case to 26.4 per cent now, bringing us close to the high end of the output loss estimates during crises.

The implications of this exercise are summarized in the impulse responses presented in Figure 13, and simulation moments presented in Figure 14. These essentially indicate a magnified and more volatile version of our benchmark results, along with a larger incidence of a binding ZLB for the policy rate. We find that, despite a lower mean debt level, the higher volatility of debt leads to a higher crisis probability on average. This allows leaning to become more successful than before in reducing the probability of a crisis, thereby reducing the volatility of key variables in the economy. However, as in the benchmark case, the pain inflicted through reducing average consumption and housing of borrowers still outweighs the gains from the reductions in volatility. In sum, leaning is overall welfare-reducing, despite more severe crises.

5 Conclusion

In this paper, we assess the relative benefits and costs of leaning against household imbalances within the context of a small-scale, regime-switching New Keynesian model. We find that leaning is generally welfare-reducing. In particular, leaning is able to reduce the volatility of the main economic variables through the reduction of crisis probabilities, but this comes at the expense of lower average debt levels, hurting borrowing households by reducing their consumption and housing expenditures. These first-moment effects on welfare turn out to be stronger than the benefits in terms of the reduction in second moments when we use utility-based welfare as our optimality criterion. Nevertheless, using an ad hoc loss function that depends on the volatility of inflation and the output gap, leaning can be shown to be beneficial, since this criterion favours the second-moment effects of leaning.

An interesting additional insight is gained from contemplating a higher inflation target of 2.5 per cent. In our set-up, this is beneficial from a welfare perspective, given the slight reduction in the probability of a crisis and the sizable effect of reducing the frequency of hitting the ZLB.

In future work, we would like to extend the model to incorporate features that have been shown in the literature to significantly improve the fit of DSGE models in capturing business cycle dynamics: namely, capital and housing accumulation, capital utilization, adjustment costs in investment, habit formation in consumption, indexation in inflation, and a variety of other shocks. This, however, would likely require the use of perturbation methods instead of computationally costly projection methods.[34]

[34] For example, the Matlab-based RISE (rationality in switching environments) toolbox utilizes perturbation methods to compute solutions for regime-switching models with an endogenous probability of switching, and can be applied to large-scale DSGE models (Maih, 2015).

References

[1] Ajello, A., T. Laubach, D. Lopez-Salido, and T. Nakata (2015). "Financial Stability and Optimal Interest-Rate Policy," mimeo, Federal Reserve Board.

[2] Alpanda, S. (2013). "Identifying the Role of Risk Shocks in the Business Cycle Using Stock Price Data," *Economic Inquiry*, 51, 304-335.

[3] Alpanda, S., G. Cateau, and C. Meh (2014). "A Policy Model to Analyze Macroprudential Regulations and Monetary Policy," Bank of Canada Working Paper 2014-6.

[4] Alpanda, S., and S. Zubairy (2014). "Addressing Household Indebtedness: Monetary, Fiscal or Macroprudential Policy?," Bank of Canada Working Paper No. 2014-58.

[5] Basant Roi, M., and R. R. Mendes (2007). "Should Central Banks Adjust Their Target Horizons in Response to House-Price Bubbles?," Bank of Canada Discussion Paper No. 2007-4.

[6] Bauer, G. H. (2014). "International House Price Cycles, Monetary Policy and Risk Premiums," Bank of Canada Working Paper No. 2014-54.

[7] Bauer, G.H. and E. Granziera (2016). "Monetary Policy, Debt and House Price Corrections," Bank of Canada Working Paper, forthcoming.

[8] Bernanke, B. S., M. Gertler, and S. Gilchrist (1999). "The Financial Accelerator in a Quantitative Business Cycle Framework." in *Handbook of Macroeconomics Volume 1C*, ed. by J. B. Taylor and M. Woodford, 1341-93. Amsterdam: Elsevier Science, North-Holland.

[9] Bianchi, J., and E. G. Mendoza (2010). "Overborrowing, Financial Crises and 'Macroprudential' Taxes," NBER Working Paper No. 16091.

[10] Büyükkarabacak, B. and N. T. Valev (2010). "The role of household and business credit in banking crises," *Journal of Banking and Finance*, 34, 1247-1256.

[11] Christensen, I., and C. Meh (2011). "Countercyclical loan-to-value ratios and monetary policy," manuscript, Bank of Canada.

[12] Demirgüç-Kunt, A., and E. Detragiache (1998). "The Determinants of Banking Crises in Developing and Developed Countries," *IMF Staff Papers*, 45, 81-109.

[13] Dorich, J., M. Johnston, R. Mendes, S. Murchison and Y. Zhang. 2013. "ToTEM II: An Updated Version of the Bank of Canada's Quarterly Projection Model," Bank of Canada Technical Report No. 100.

[14] Emanuelsson, R., O. Melander, and J. Molin (2015). "Financial risks in the household sector," *Economic Commentaries* (Riksbank), 6, 1-35.

[15] Fisher, I. (1933). "The Debt-Deflation Theory of Great Depressions," *Econometrica*, 1(4), 337-357.

[16] Gelain, P., K. J. Lansing, and C. Mendicino (2013). "House Prices, Credit Growth, and Excess Volatility: Implications for Monetary and Macroprudential Policy," *International Journal of Central Banking*, 9, 219-276.

[17] Gelain, P., K. J. Lansing, and G. J. Natvik (2015). "Leaning against the credit cycle," Norges Bank Working Paper No. 04-2015.

[18] Iacoviello, M. (2005). "House Prices, Borrowing Constraints, and Monetary Policy in Business Cycles," *American Economic Review*, 95, 739-764.

[19] Iacoviello, M. (2015). "Financial Business Cycles," *Review of Economic Dynamics*, 18, 140–163.

[20] Iacoviello, M., and S. Neri (2010). "Housing Market Spillovers: Evidence from an Estimated DSGE Model," *American Economic Journal: Macroeconomics*, 2, 125-164.

[21] International Monetary Fund (IMF) (2009). "From Recession to Recovery: How Soon and How Strong?" in World Economic Outlook (WEO) Crisis and Recovery, April 2009, Chapter 3, 103-138.

[22] Jorda, O., M. Schularick, and A. M. Taylor (2015). "Leveraged Bubbles," NBER Working Paper No. 21486.

[23] Kiyotaki, N., and J. Moore (1997). "Credit Cycles," *Journal of Political Economy*, 105, 211-48.

[24] Klimke, A. (2007). "Sparse Grid Interpolation Toolbox - User's Guide," IANS Report No. 2007/17, University of Stuttgart.

[25] Klimke, A., and B. Wohlmuth (2005). "Algorithm 847: spinterp: Piecewise Multilinear Hierarchical Sparse Grid Interpolation in MATLAB," *ACM Transactions on Mathematical Software*, 31(4), 561–579.

[26] Koo, R. C. (2008). *The Holy Grail of Macroeconomics – Lessons from Japan's Great Recession*, John Wiley & Sons, New York.

[27] Korinek, A. (2011). "Systemic Risk-Taking: Amplification Effects, Externalities, and Regulatory Responses," ECB Working Paper No. 1345.

[28] Lambertini, L., C. Mendicino, and M. T. Punzi (2013). "Leaning against boom-bust cycles in credit and housing prices," *Journal of Economic Dynamics and Control*, 37, 1500-1522.

[29] Lorenzoni, G. (2008). "Inefficient Credit Booms," *Review of Economic Studies*, 75, 809-833.

[30] Maih, J. (2015). "Efficient Perturbation Methods for Solving Regime-Switching DSGE Models," Norges Bank Working Paper No. 01-2015.

[31] Maliar, L., and S. Maliar (2013). "Envelope condition method versus endogenous grid method for solving dynamic programming problems," *Economics Letters*, 120, 262–266.

[32] Reinhart, C. M. and K. S. Rogoff (2010). *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press.

[33] Rotemberg, J. J. (1982). "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies*, 49, 517–531.

[34] Rubio, M. (2011). "Fixed and Variable-Rate Mortgages, Business Cycles, and Monetary Policy," *Journal of Money, Credit and Banking*, 43, 657-688.

[35] Schularick, M., and A. M. Taylor (2012). "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008," *American Economic Review*, 102, 1029-1061.

[36] Smets, F., and R. Wouters (2007). "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586-606.

[37] Stokey, N. L., and R. E. Lucas, with E. C. Prescott (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

[38] Svensson, L. E. O. (2013). ""Leaning Against the Wind" Leads to Higher (Not Lower) Household Debt-to-GDP ratio," manuscript, Stockholm University.

[39] Woodford, M. (2012). "Inflation Targeting and Financial Stability," NBER Working Paper No. 17967.

A Computational Appendix

In this section, we outline an envelope condition method (ECM) approach to solving our model, along with a sparse grid approximation technique. The ECM approach was originally proposed in Maliar and Maliar (2013), and is based on iterating on the *derivatives* of value functions, as opposed to the standard value function iteration approach or the collocation approach on the Euler equations. Its main advantage is the replacement of complex root finding problems with simple algebra; thus, it is faster than value function approximation and more robust in terms of convergence relative to the collocation method. We adapt the method to our needs by, first, writing the optimization problems of the patient and impatient households as functional equations, and deriving the corresponding first-order and envelope conditions that characterize equilibrium (see Stokey et al., 1989). We also write the Phillips curve and the labour demand expressions of the firm, along with the Taylor rule of the central bank and the feasibility conditions, consistent with the recursive notation used in the rest of the model. As outlined below, the core of the algorithm works through initializing the value function derivatives (i.e., initializing the approximating Chebyshev polynomials' parameters, θ), and updating them until sufficient convergence is achieved.

To implement the sparse grid approximation, we use a toolbox developed by Andreas Klimke, as outlined in Klimke (2007) with background material in Klimke and Wohlmuth (2005). The specific functions we employ from the package are *spset* to determine the sparse grid options, *spval* to obtain the gridpoints, and *spinterp* to evaluate the approximation to the value function derivatives. Specifically, we use a Chebyshev polynomial approximation, with a dimension-adaptive grid of at least 100 points but not more than 10,000 points. The minimum approximation tolerance we require is 10^{-2}. Initially, we use a standard sparse grid with up to 6^4 points. Once we have sufficiently converged with our ECM approach, we allow for up to 10,000 points and switch on the option to adjust the dimension of the gridpoints. Using a lot of points slows the process down considerably, while not helping very much with the convergence. It should be noted that the dimension-adaptation option of the grid is detrimental to convergence, since it creates additional fluctuations purely due to changes of the grid. Our approach represents a balancing act between accuracy, speed and obtaining a solution, i.e. functional convergence.

The general outline of the algorithm is as follows:

1. We find the deterministic steady state (conditional on the normal regime) for all the variables and the value function derivatives.

2. We use the steady-state values in step 1 as initial guesses for the value function derivatives, and construct a sparse grid over a compact subset of the state space, $S = (d, h_P, R, \xi, r)$. Note that the regime state, r, takes on a value of 0 or 1, and for the housing demand shock. We also create a quadrature over the innovations of the housing demand shock with nodes ε_j and weights ω_j for $j = 1, ..., J$.

3. We use the value function derivatives given in step 2 to solve for all the current decisions

of households, including the state variables affecting the next period. The ECM approach simplifies this step, and allows us to obtain the optimal choices fairly easily.

4. Using the solution from step 3 and the previous value function derivatives, we update the value function derivatives using the equilibrium conditions below:

$$V_d^P(S) = \frac{R^m}{\pi}\beta_P E V_d^P(S'), \tag{25}$$

$$V_h^P(S) = u_h(c_P, h'_P, n_P) + \beta_P E V_h^P(S'), \tag{26}$$

$$V_b^P(S) = \frac{R(1+\chi(r))}{\pi}\beta_P E V_b^P(S'), \tag{27}$$

$$V_d^I(S) = \left(\rho_d - \frac{R^m}{\pi}\right)\mu + \frac{R^m}{\pi}\beta_I E V_d^I(S'), \tag{28}$$

$$V_h^I(S) = -u_h(c_I, h'_P, n_I) - (1-\rho_d)\phi q \mu + \beta_I E V_h^I(S'), \tag{29}$$

$$\Gamma(S) = \beta_P E\left[\lambda'_P\left(\frac{\pi'}{\pi^*} - 1\right)\frac{\pi'}{\pi^*}y'\right]. \tag{30}$$

The first five of these expressions are obtained by combining the first-order and envelope conditions of the functional equations of patient and impatient households, while the last one is a function we define to be able to iterate on the Phillips curve equation. To evaluate the expectations on the right-hand side of these expressions, we use the quadrature over the stochastic shock and the Markov chain probabilities. For example,

$$\begin{aligned} V_d^P(d, h_P, R, \xi, r) &= \frac{R^m}{\pi}\beta_P E V_d^P(S') \\ &\approx \frac{R^m}{\pi}\beta_P \left\{ \begin{array}{ccc} \sum_{r'=0}^{1}\gamma(d,r')\sum_{j=1}^{J}\omega_j V_d^P\left(d', h'_P, R', \xi^{\rho_\xi(r)}\exp(\varepsilon_j), r'\right) & \text{for} & r=0 \\ \sum_{r'=0}^{1}\delta(r')\sum_{j=1}^{J}\omega_j V_d^P\left(d', h'_P, R', \xi^{\rho_\xi(r)}\exp(\varepsilon_j), r'\right) & & r=1 \end{array} \right\} \end{aligned}$$

where

$$\gamma(d, r') = \left\{ \begin{array}{c} 1-\gamma(d) \\ \gamma(d) \end{array} \text{ for } r' = \begin{array}{c} 0 \\ 1 \end{array} \right\},$$

$$\delta(r') = \left\{ \begin{array}{c} \delta \\ 1-\delta \end{array} \text{ for } r' = \begin{array}{c} 0 \\ 1 \end{array} \right\}.$$

5. We iterate on step 4 above until we reach convergence on the value function derivatives. At each step, we also check the convergence of the policy functions that we need to update the states, in particular for d' and h'_P.

Table 1. Benchmark Calibration

	Symbol	Value
Inflation target (gross, qtr.)	π^*	1.005
Discount factor	$(\beta_P,\ \beta_I)$	$(0.9925,\ 0.97)$
Inverse labour supply elasticity	ϑ	2
Level parameter for housing in utility	ξ	0.12
Gross markup in price	θ	1.1
Price adjustment parameters	κ	100
Share of patient HH in labour income	ψ	0.748
LTV ratio on mortgage debt	ϕ	0.75
Persistence in the borrowing constraint	ρ_d	0.96
Switch prob. from crisis to normal regime	δ	0.10
Parameters in the logit specification	$(\omega_1,\ \omega_2)$	$(9.68,\ 5.07)$
Risk premium in the crisis regime, quarterly	$\overline{\chi}$	0.0125
Taylor rule - persistence	ρ	0.75
- inflation	a_π	1.5
- output gap	a_y	0.125
- debt gap (no leaning)	a_d	0
Persistence of housing demand process	$(\rho_\xi(0), \rho_\xi(1))$	$(0.985, 0.985/2)$
St.-dev. of housing demand shock innovation	σ_ξ	0.05
Fraction of price adj. costs affecting resources	I_π	0.5

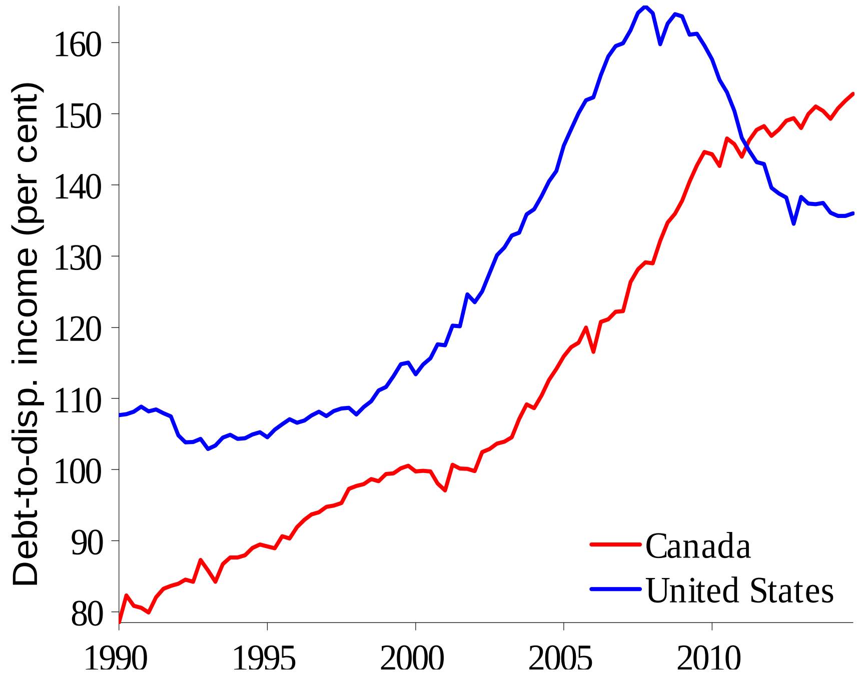

Figure 1: Debt-to-disposable income in Canada and the United States. The Canadian disposable income is adjusted according to the U.S. definition in order to make the comparison more accurate *Source*: Statistics Canada, Bank of Canada, Federal Reserve Board, U.S. Bureau of Economic Analysis

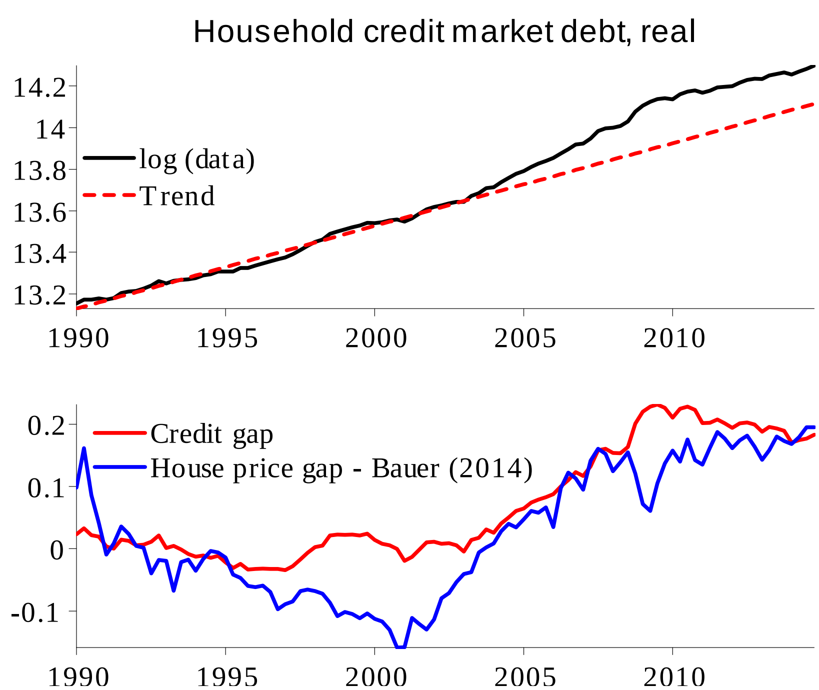

Figure 2: Credit gap in Canada with comparison to house price gap based on Bauer (2014)

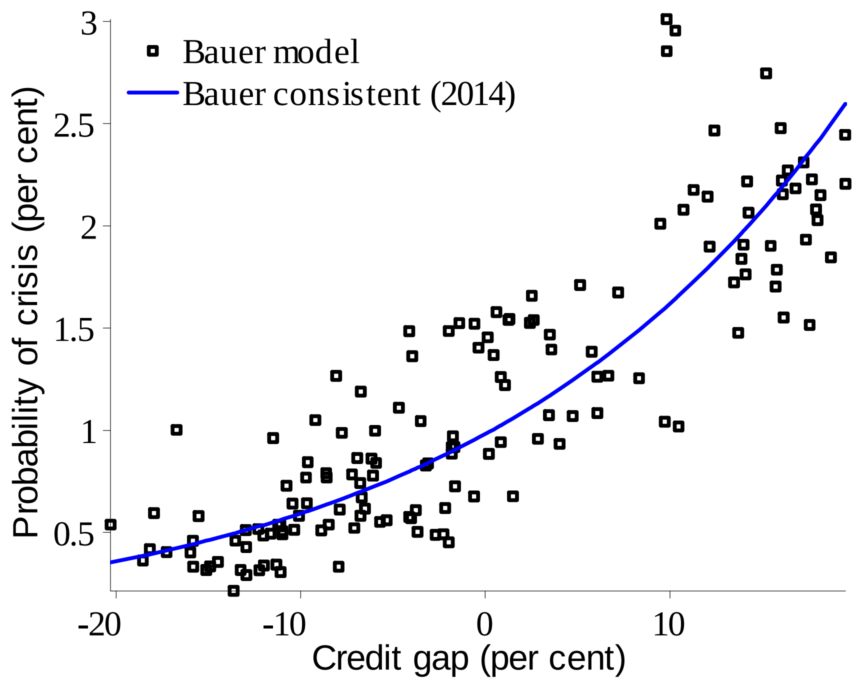

Figure 3: Probability of a house price correction in Canada within the next quarter (in %)
Source: Bauer (2014)

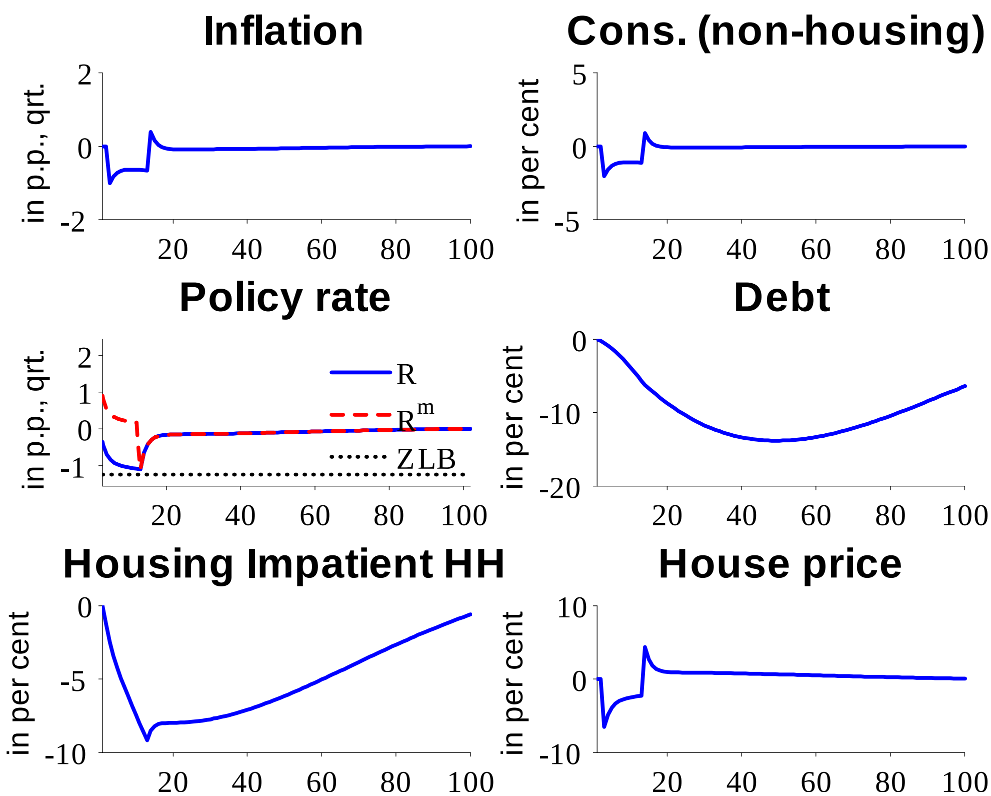

Figure 4: Impulse responses of key model variables following a crisis shock induced by a sharp rise in the risk premium on mortgage debt for 10 periods

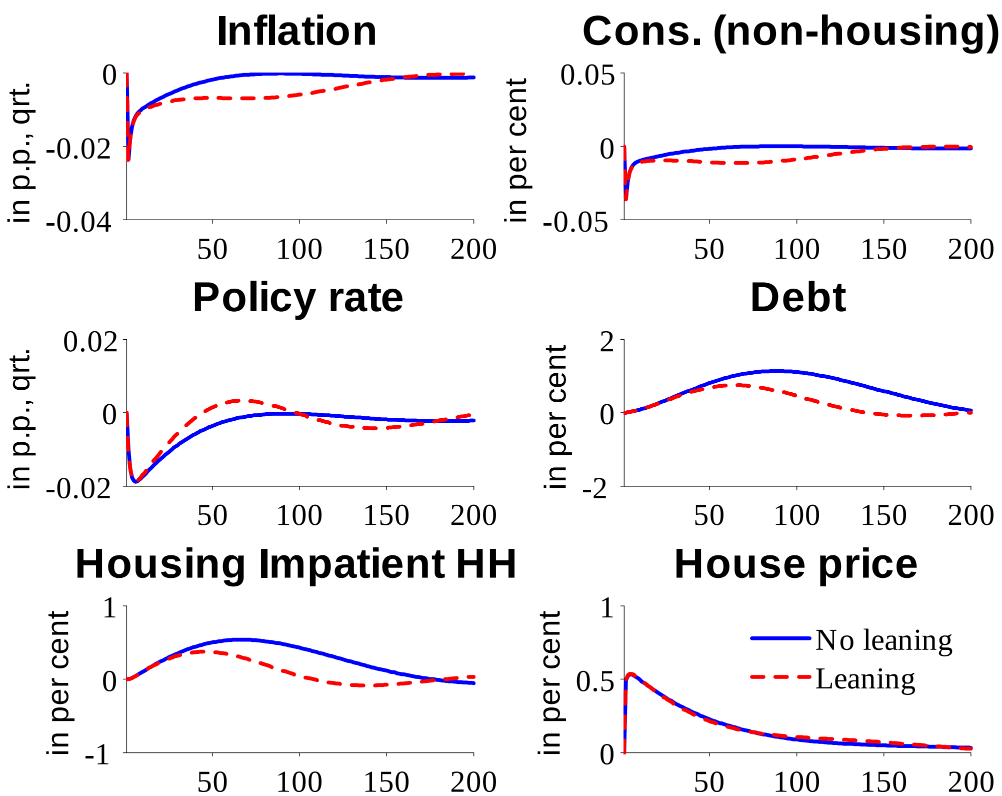

Figure 5: Impulse responses of key model variables following a housing demand shock in the normal regime. Blue lines denote the case with no leaning (i.e., $a_d = 0$) and dashed red lines denote the case with monetary policy leaning (with $a_d = 0.032$)

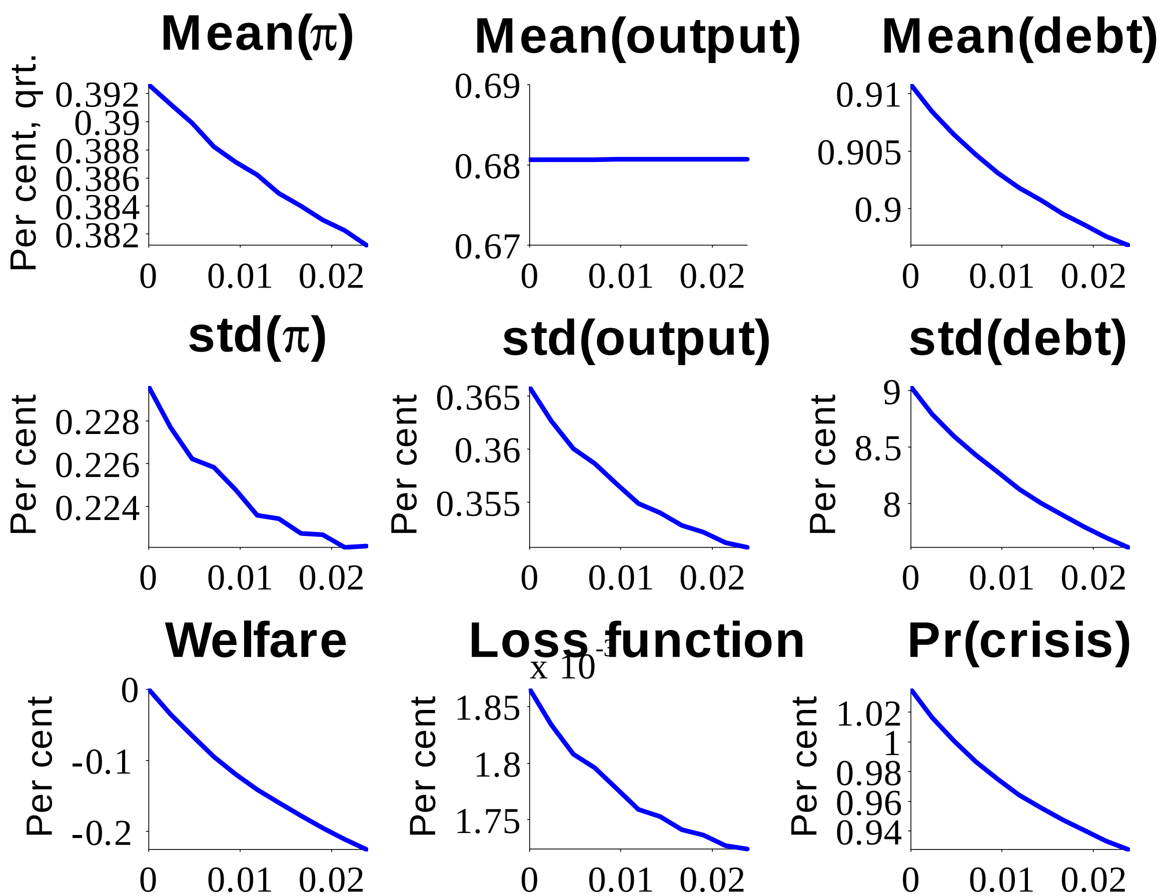

Figure 6: First and second moments as well as social welfare, loss function and crisis probability depending on the degree of leaning, a_d

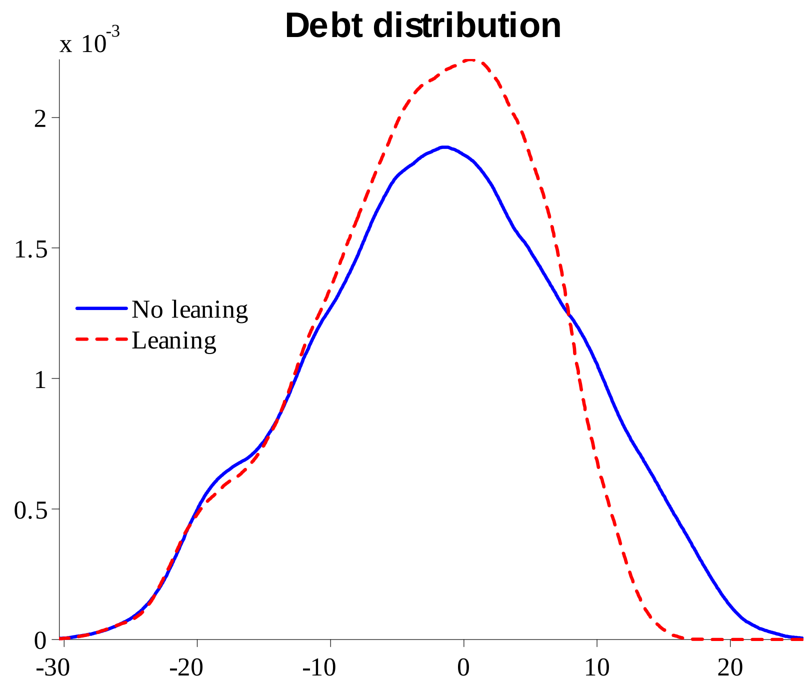

Figure 7: Debt distribution (over time) for the benchmark calibration and the case of no leaning (blue line, $a_d = 0$) and leaning (red dashed line, $a_d = 0.032$)

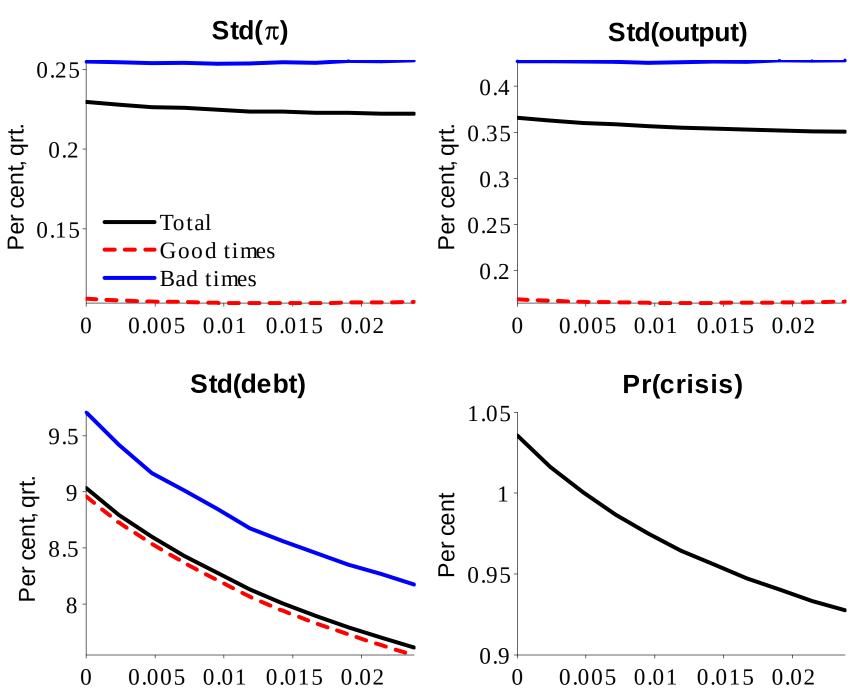

Figure 8: Standard deviations of output, inflation and debt for the benchmark economy for the two regimes

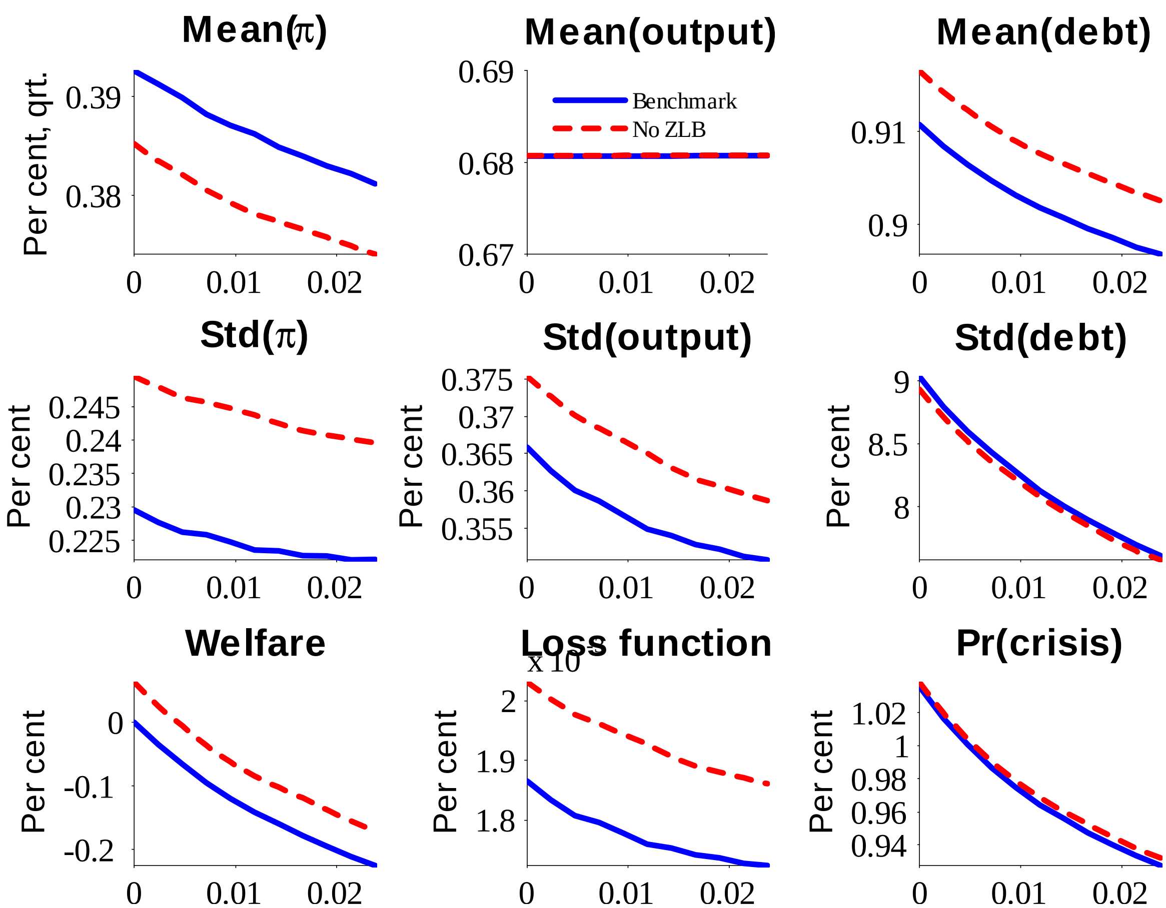

Figure 9: Comparison of leaning implications for model with lower zero lower bound (ZLB) (blue solid line) to one without binding ZLB (red dashed line)

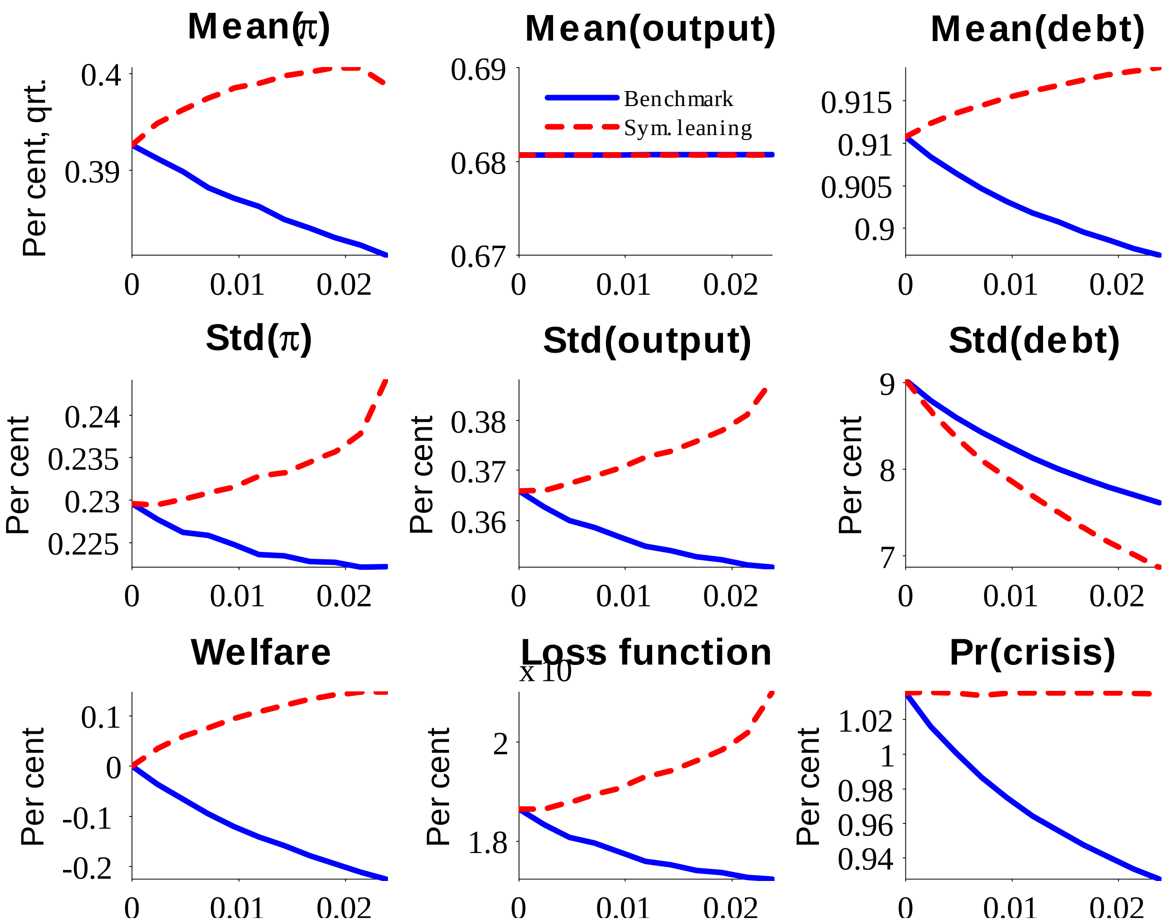

Figure 10: Comparison of leaning implications for model with asymmetric (blue solid line) to one with symmetric (red dashed line) leaning

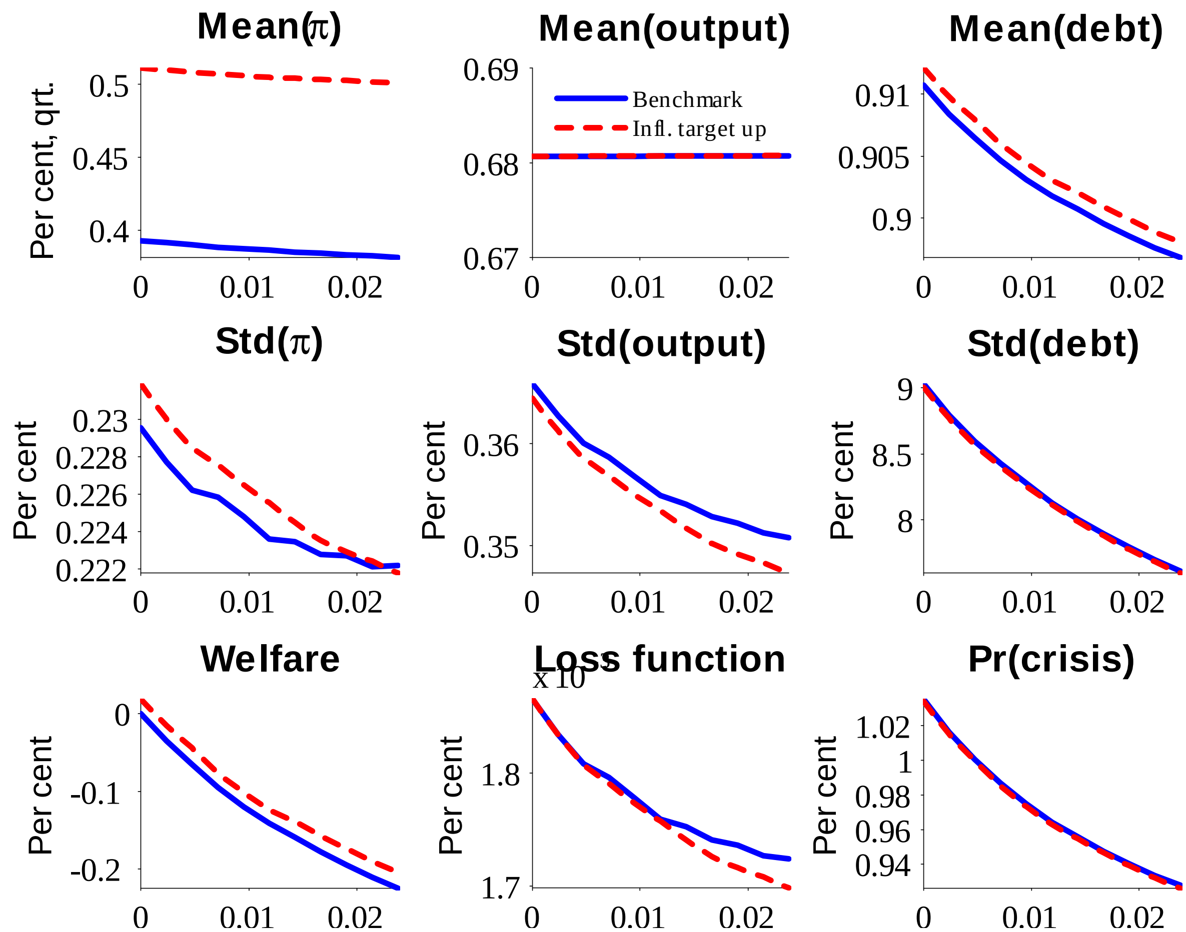

Figure 11: Comparison of leaning implications for model with a 2% inflation target per annum (blue solid line) to one with a 2.5% target (red dashed line)

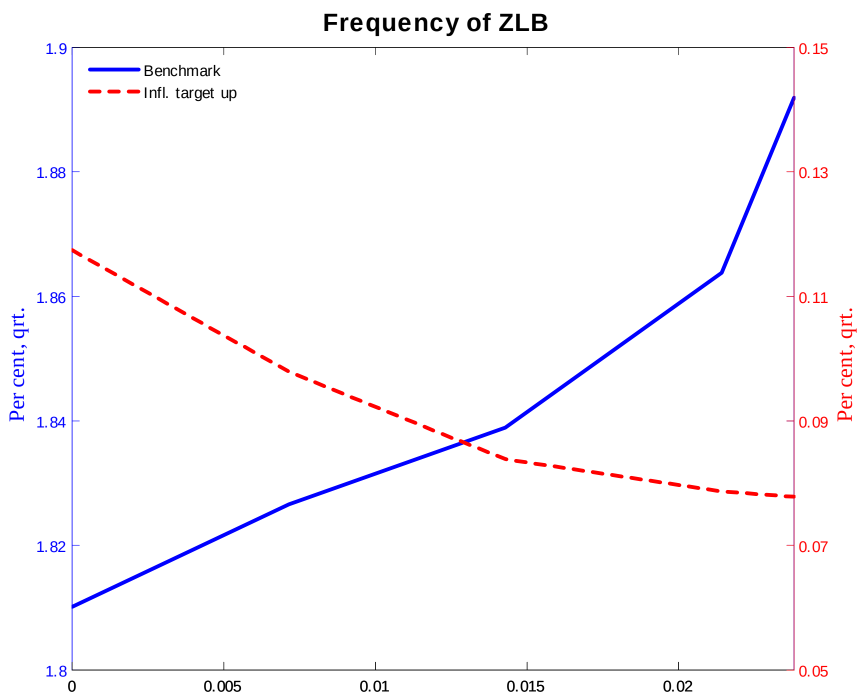

Figure 12: Comparison of hitting the ZLB with a higher inflation target of 2.5% per annum to hitting it at the 2% benchmark rate

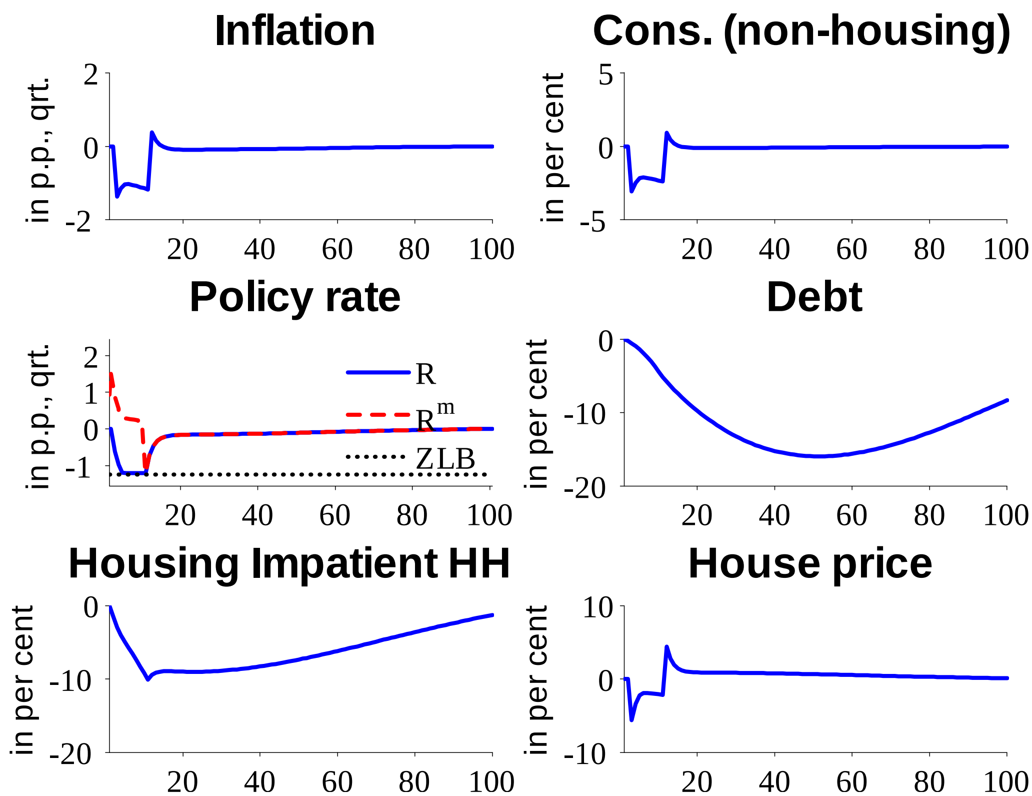

Figure 13: For an extreme crisis scenario, impulse responses are shown of key model variables following a crisis shock induced by a sharp rise in the risk premium on mortgage debt for 10 periods

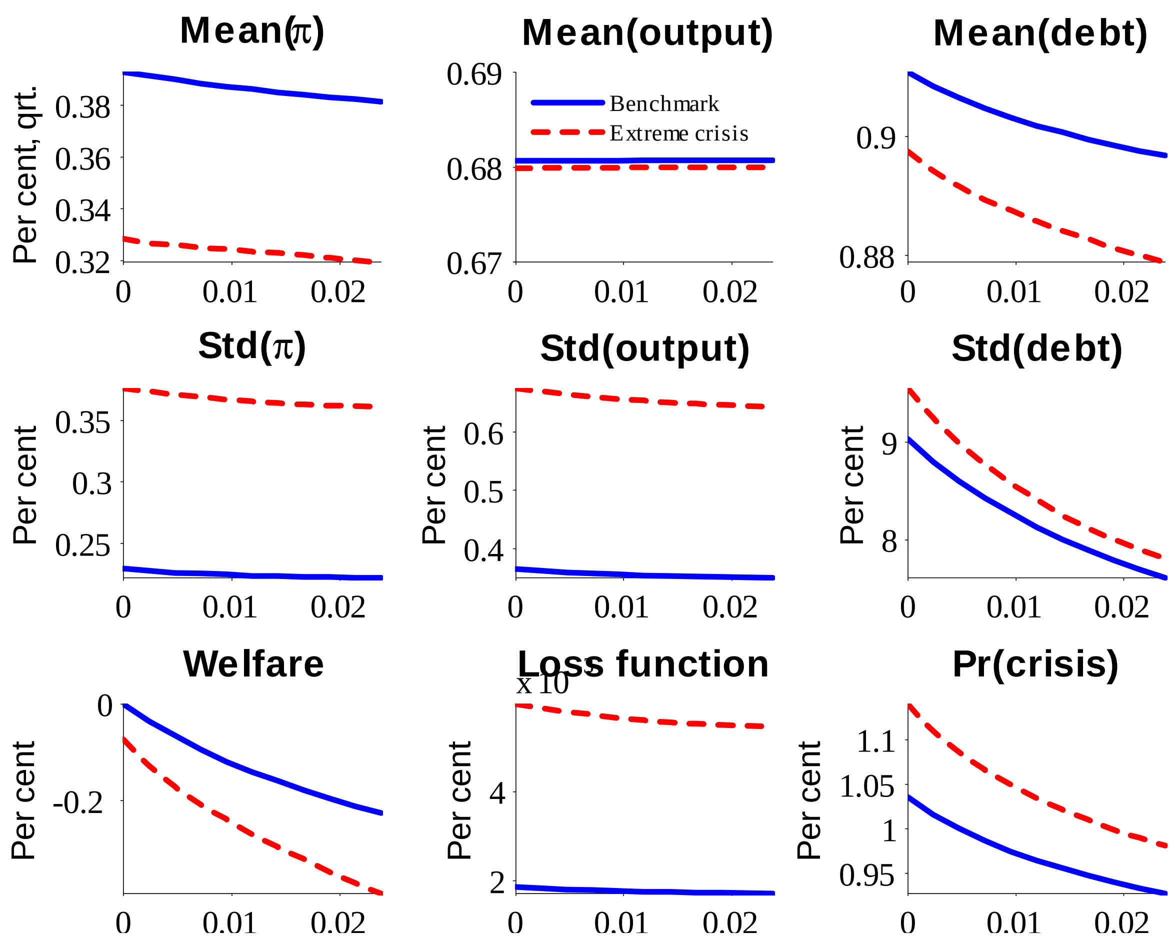

Figure 14: Comparison of leaning implications for model with benchmark-sized crisis (blue solid line) to one with very large-sized crises (red dashed line)