Household Credit Uncertainty and the Real Interest Rate

Pierre Mabille*

December 2016

[Preliminary version]

Abstract

I introduce aggregate uncertainty about the debt limits faced by heterogeneous households in a Bewley-Huggett economy with a financial sector and government bonds in exogenous supply. Uncertainty about households' debt contracts results from aggregate shocks to the balance sheets of financial intermediaries supplying credit to households, which sometimes give rise to credit crises. In the rational expectations equilibrium, agents internalize the existence of (endogenous) stochastic credit regimes, which creates an additional precautionary savings motive. When the economy transits from a regime of easy credit to a regime of tight credit, the interest rate drops because of the combined effects of i) deleveraging and ii) the precautionary savings motive created by the uncertainty about future debt limits. A low and persistent interest rate environment arises, with households holding more assets and less debt. Consumption is smoother over the cycle than in a model in which the transition to a tight credit regime is unexpected. Thus an extended period of low interest rates is not necessarily associated with a massive drop in aggregate consumption, which is a feature of the post-financial crisis US data that classical models of the liquidity trap fail to generate.

*Department of Economics, NYU Stern. Email: `pmabille@stern.nyu.edu`. I thank Felipe Alves, Zhen Huo, Virgiliu Midrigan and Stijn Van Nieuwerburgh for comments.

1 Introduction

Real interest have been persistently low since the end of the Great Recession in June 2009. While fluctuating between 1 and 4 percent before and during the financial crisis, Constant Maturity rates on US Treasury Inflation-Indexed Securities (CMTIIS) of 5 to 20 years rarely exceeded 1 percent in the post-financial crisis period, and often ventured into lower, if not negative values (figure 2). The enduring low interest rate environment in the US has called for a number of explanations, most of which focus on the real causes of the long-term decline in real interest rates for over three decades. They include ”secular stagnation”, a higher propensity to save and a lower propensity to invest mainly due to demographic and technological factors, and an increased external demand for safe assets (Bean et al. (2015)). However, while the secular decline in real rates is associated with slower aggregate consumption growth in the long run, this relationship may fail to hold in the short-run. In fact, almost the opposite is true about the post-financial crisis period, where the further ratchet down in real CMTIIS rates in mid-2009 happened at the time when real Personal Consumption Expenditures, a measure of US households aggregate consumption, started growing again after reaching the Recession trough (figure 3). In addition, while models of the liquidity trap (such as Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2016)) account well for credit crises' boom-bust short-run transitions, they often fail to provide an explanation for prolonged periods with low rates and moderate consumption fluctuations, as for instance following the recent crisis.

The goal of this paper is to reconcile the ”liquidity trap view” of credit crises with the moderate post-crises consumption fluctuations observed in the data. It explores the idea that financial crises generate uncertainty about households debt contracts, and therefore contribute to drive down real interest rates in a persistent way that does not necessarily entail a collapse in aggregate consumption. It analyzes how short-term financial frictions and ”expectations of instability in the near future” shape macroeconomic variables (households consumption, savings, labor supply) and real interest rates in an incomplete markets model with heterogeneous agents, a financial sector and a government. In the model, short-term financial frictions come from imposing a simplified realistic financing structure upon households-entrepreneurs balance sheets. Households have limited access to financial markets, and can only purchase government bonds and accumulate debt through financial intermediaries, in order to insure themselves against idiosyncratic income risk. In particular, households borrow by issuing risk free bonds that can only be bought by financial intermediaries, which set a debt limit on every such transaction. They save by buying shares of financial intermediaries – which can hereby be seen as mutual funds – that entitle them to the latter's profits. Financial intermediaries (henceforth referred to as "banks" for simplicity) are endowed with exogenous resources, which fluctuate over time in a stochastic fashion, and can thus be interpreted as capital or any external source of funding for the financial sector. They make short-term loans to households and invest in risk free government bonds. Bank capital is an aggregate shock which hits intermediaries' balance sheets, and translates into fluctuations in households credit conditions. To study how uncertainty about households debt limits contribute to create a low interest rate environment, I restrict households debt contracts to take the form of line of credits. In the latter, banks choose the credit limit on households debt, but not the actual amount of debt taken up by households within that limit. Depending on the shocks to bank capital – which can be thought of as "financial crises" or "expansions" – households' credit lines can be tighter or looser. The stochastic nature of bank capital thus contributes to making households debt limits stochastic. In the rational expectations equilibrium, agents internalize aggregate credit uncertainty, which creates an additional precautionary savings motive. When the economy transits from a regime of easy credit to a regime of tight credit, the interest rate drops because of the combined effects of deleveraging – as in models of the liquidity trap – and the precautionary savings motive created by the uncertainty about future debt limits. A low interest rate environment arises, with households holding more bonds and less debt, and moderate consumption fluctuations.

The main contribution of this paper is to build a quantitative model of credit crises in which an extended period of low real interest rates is not necessarily associated with a massive drop in aggregate consumption – a feature of the post-financial crisis US data. Another contribution is to provide a stylized framework to study the effect of both fiscal and macroprudential policy, in an environment where households internalize credit uncertainty and seek to insure themselves against the latter by increasing their demand for safe and liquid assets. The last contribution is methodological. First, I solve a reduced-form model without a financial sector, in which the economy enters and exits stochastic credit regimes over time, which are governed by a Markov process, and associated with tighter or looser borrowing constraints. Agents' decisions are functions of their expectations about future debt limits and the distribution of agents across states. I therefore solve the model numerically using a variant of the Krusell and Smith (1998) algorithm, in which agents forecast prices as functions of current and past aggregate state variables. The equilibrium is the outcome of a fixed point problem in which agents' expectations coincide with actual aggregates. I underscore the quantitative differences between models of the liquidity trap in which credit crises occur unexpectedly and then unfold in a deterministic fashion, and models in which aggregate uncertainty about access to credit is known and internalized by households. Then, I solve the full model with a financial sector. Methodologically, this solution method is novel, and combines the endogenous gridpoints method for problems with multiple occasionally binding constraints on the households side, with the resolution of a discrete and continuous optimization problem for banks, and the treatment of aggregate risk in a version of the Krusell-Smith algorithm.

Results I present the results for the reduced-form version of the model, without a financial sector. In this economy, households insure themselves against idiosyncratic income risk by buying and selling government bonds subject to stochastic debt limits, which represent regimes of ”easy” or ”tight” credit. In equilibrium, households internalize the aggregate risk created by stochastic credit regimes. They demand more bonds to insure themselves against credit regime fluctuations, and possibly avoid deleveraging when the economy enters a tight credit regime. As a result, an environment with low real interest rates arises, with less households debt and smoother consumption.
In comparison to an economy in which debt limit fluctuations are unexpected, the real interest rate is lower in both regimes of tight and easy credit. It drops more sharply when the economy is hit by a credit shock, because of the combined effects of households deleveraging and the precautionary savings motive created by uncertainty about future debt limits. The conditional likelihood of a tight credit shock when the economy is in a regime of easy credit, and the persistence of tight credit regimes govern the ”average level” and the ”slope” of the interest rate's trajectory. For instance, in an economy in which the probability to transit from an easy to a tight credit regime is 0.02, and the probability to remain in a tight credit regime is 0.05 (credit shocks are unlikely to happen in a regime of easy credit, but are slightly more likely to persist once the economy is in a tight credit regime), the net interest rate in a high credit regime is 10% lower than its corresponding value when shocks to debt limits are unexpected. In addition, it decreases by three times the amount it does in the latter economy when it enters a low credit regime[1].
The long-run mean of aggregate consumption is lower than in an economy with unexpected shocks, within a given credit regime and aggregating across credit regimes. The debt distribution is less skewed to the left, because poorer households seek to insure themselves against credit uncertainty by accumulating bonds. In equilibrium, richer households hold less bonds, which contributes to the

[1]The interest rate's average level in easy credit regimes is 2.3% annually, against 2.5% when credit shocks are unexpected. In the period after the credit shock, it falls to -8.3% against 0.5%. By ”slope” I loosely refer to the difference of those, here -6 and -2 percentage points.

decline in aggregate consumption, and also makes the bond distribution less skewed to the right. The counterpart is that consumption is smoother over time, because poorer households are less indebted, and therefore the deleveraging effect of economies with unexpected shocks is dampened. This is consistent with the view that a low interest rate environment is associated with lower aggregate consumption, but that consumption fluctuations can be moderate following a tightening in debt limits. Finally, across credit regimes, a loosening of debt limits is associated on average with a lower cross-sectional variance of consumption, but a higher variance of the bond distribution. This is consistent with results in the literature (see below) that "financial liberalizations" increase wealth inequality and decrease consumption inequality, while the reverse holds in tight credit regimes.

Discussion On the theoretical side, this paper quantifies the main example taken by Farhi and Werning (2015) to study optimal macroprudential policies in a model of credit crunch[2]. They consider a three-period economy where the realization of a random variable in the second period determines if the borrowing constraint is tight or loose. A recession and a liquidity trap arise when the economy enters a credit crisis. Security markets are complete, and agents form expectations about the likelihood of a credit crisis in the first period, then make consumption and savings decisions. This paper considers the same economy in an infinite horizon setting with incomplete markets, and in which agents differ along their idiosyncratic productivity levels but not their patience. They form expectations about the transitions between easy and tight credit regimes. The amounts of debt they hold in equilibrium reflects their expectation about both the likelihood and persistence of tight credit regimes over time. Boz and Mendoza (2014) study an infinite-horizon economy in which households learn about the collateral constraint they face when borrowing. In their model, learning about the risk of a new financial environment interacts with Fisherian amplification and produces a boom-bust cycle in debt, asset prices and consumption. Perceived probabilities of credit regime changes are endogenous, and agents are more or less optimistic about credit regimes depending on the history of the aggregate state. The fact that agents' perceived transition probabilities differ from the empirical, true probabilities lead them to make optimization mistakes in the short-run. As a result, when calibrated to the US data, the model predicts large increases in household debt during 1998-2006, followed by a collapse in both households debt and consumption. Unlike theirs, this paper maintains the assumption of rational expectations, with agents knowing the true probabilities of credit regime changes, but not when exactly these changes occur. In equilibrium, prices adjust

[2]Section 5.1., "Liquidity Trap and Deleveraging" (p. 22). Their example builds itself on Eggertsson and Krugman (2012). They also have nominal frictions, which this paper does not have.

more than quantities, and a persistently low real interest rate following a tight credit regimes is not necessarily associated with a massive drop in consumption. As stressed in the results description, this paper also stands in contrast with Guerrieri and Lorenzoni (2016), who study the transitional dynamics of an economy that has been hit by an exogenous, unexpected tightening of households' borrowing limits. While such models of the liquidity trap account well for credit crises' boom-bust transition, where both real rates and consumption collapse following a tightening of debt limits[3], they cannot generate post-crises periods in which real interest rates stay low with a massive drop in consumption. By making households internalize the uncertainty about debt contracts, this paper is able to generate both a low interest rate environment and moderate consumption fluctuations.

On the quantitative side, this paper shifts the focus of the question in Nakajima and Rios-Rull (2014), by asking how access to credit at the intensive margin – instead of the extensive margin – affects the nature of the business cycle. It does not model the production sector with aggregate uncertainty about Total factor Productivity (TFP), but rather focuses on how shocks to the financial sector's balance sheet translates into shocks to households credit conditions. Households debt limits become the main source of aggregate uncertainty, insofar as a significant fraction of households are on their borrowing limits for realistic calibrations of the model. Therefore, while aggregate uncertainty about TFP does not induce significant wealth redistribution across agents (as Krusell and Smith (1998) already pointed out), credit uncertainty is able to generate significant fluctuations in asset prices and quantities. Finally, two other quantitative papers are also interested in the effect of shifting households' borrowing constraints. Kaplan et al. (2015) engineer aggregate shocks to credit market conditions in a life-cycle model with housing and a liquid asset. These shocks stand for changes in credit regimes (such as an increase in maximum LTV constraints). In contrast, this paper abstracts from housing and retirement accounts[4] to focus on the financial sector as a source of short-run aggregate shocks to households credit conditions. Favilukis et al. (2016) also study the general equilibrium impact on the real interest rate and macroeconomic variables of shifts in households' borrowing constraints. They study the impact of exogenous "financial market liberalization" on house prices, interest rates and households' ability to insure against risk. As opposed to

[3] Other papers include for instance Eggertsson and Krugman (2012).

[4] Those admittedly represent an important fraction of illiquid assets held by US households. I abstract from them for two reasons. First, it allows to keep the model tractable when adding a financial sector to a Bewley-Huggett economy. Second, the housing market and retirement accounts are arguably less likely to impact households' financing conditions in the (very) short run than the immediate health of the financial sector. Of course this is less true in the long run, insofar as households' borrowing constraints are largely determined by the value of their collateral.

comparing stochastic steady states of distinct economies with tight and loose borrowing constraints across long simulations, this paper considers a single economy in which borrowing constraints can be tight or loose, and the possibility of fluctuations is internalized by agents. It also models the source of debt limit fluctuations by adding a financial sector. As in their model, a tightening of borrowing constraints leads to a drop in the equilibrium interest rate. In addition, endogenous debt limit fluctuations make low real rates persistent. I also find that a tightening of borrowing limits is associated on average with lower wealth inequality, and slightly higher consumption inequality.
Finally, in a New-Keynesian setting with rational expectations, Fernández-Villaverde et al. (2015) underscore the trade-off between generating long spells of low interest rates (in particular at the zero lower bound) and moderate drops in consumption, as in the post-2009 US data. They suggest that this trade-off can be alleviated by introducing a wedge in households' Euler equations, so that a collapse in the real rate does not entail a collapse in consumption. This paper can be read as an illustration of their insight, which focuses on financial frictions and aggregate uncertainty about debt limits in an incomplete markets model.

Outline The paper is organized as follows. Section 2 describes the model of the economy with heterogeneous households, a financial sector and a government, in which aggregate shocks hit banks' balance sheets and translate into shocks to households' debt limits. Section 3 presents the numerical solution of the reduced-form model, in which debt limits follow a two-state exogenous Markov process. Two variants of the Krusell-Smith algorithm to solve the model with credit uncertainty are described. The first one solves directly for prices as functions of the aggregate states, while the second one solves for agents' forecast rule about moments of the bond distribution, and their equilibrium relation with prices[5]. Section 4 presents the numerical results and compares them to those of an identical economy in which shocks to debt limits are unexpected.

[5] Julia code available upon request.

2 Model

2.1 Households

Households-entrepreneurs There is a continuum of infinitely-lived households $i \in [0, 1]$ ordering consumption and labor supply flows according to the utility function

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t U(c_{it}, n_{it})\right],$$

where c_{it} is consumption and n_{it} is labor effort of household i. Each household produces consumption goods using the linear technology

$$y_{it} = \theta_{it} n_{it},$$

where θ_{it} is an idiosyncratic shock to the labor productivity of household i, which follows a Markov chain on the space $\Theta = \left\{\theta^1, ..., \theta^S\right\}$, with transition matrix Π_θ. Assume $\theta^1 = 0$ and interpret this realization of the shock as unemployment.

By assumption households can only access financial markets through financial intermediaries (access cost to financial markets is not modeled). Households save by buying shares of the financial intermediaries ("bank capital"), which earn a dividend π^a ("bank profit"), but trade is restricted so that they cannot sell those shares. The number of shares available for sale is fixed at $\overline{A} = 1$. This captures the pattern of households saving in the US, with a large fraction of of non-housing households positive net wealth being stocks and bonds bought through various financial funds. They can borrow from the financial intermediaries, which set the credit limit on those loans. This captures the fact that most households borrow through financial intermediaries under conditions set by the latter, not directly on competitive financial markets.

Households' expenditures consist of the principal and interests on loans taken the previous period, savings taking the form of shares in bank capital, consumption and taxes. Resources are the dividend and price from selling the shares of bank capital inherited from the previous period, new loans taken from the banks, and labor income. The household budget constraint and borrowing/lending constraints are

$$c_{it} + p_t^a a_{it+1} + p_t^b b_{it+1} \leq y_{it} + (p_t^a + \pi_t^a) a_{it} + b_{it},$$

$$0 \geq b_{it+1} \geq -\phi_t,$$

$$1 \geq a_{it+1} \geq 0.$$

a_{it} are shares of bank capital at the beginning of period t (chosen in $t-1$), with ex-dividend price p_t^B; they bear a risky net rate of return from t to $t+1$ defined by $r_t^a = \frac{p_{t+1}^a + \pi_{t+1}}{p_t^a} - 1$ (here ex-post). The total amount of bank shares is fixed and normalized to 1. b_{it} are loans outstanding at the beginning of period t (chosen in $t-1$), taking the form of zero coupon risk free bonds sold by households and bought by the lender (here the banking sector); they bear a net interest rate $r_t^b = \frac{1}{p_t^b} - 1$.

Three assumptions are made about financial frictions, which restrict trade between agents:

- Households can only borrow from the banks, and they cannot trade these loans on secondary markets (i.e. they cannot hold $b_{it} > 0$).
- Household debt is bounded below by the credit limit ϕ set endogenously (in t for $t+1$ loans) by financial intermediaries. This corresponds e.g. to households credit card debt.
- Households have limited access to financial markets: they can only save by buying bank shares, and they cannot sell these shares on secondary markets (i.e. they cannot hold $a_{it} < 0$).

The tax schedule is the following: $\tau_{it} = \tau_t$ if $\theta_{it} > 0$ and $\tau_{it} = \tau_t - \nu_t$ (unemployment benefits) if $\theta_{it} = 0$.

This is a model with two assets (REFERENCES), thus households face a portfolio choice. They choose their consumption, labor supply and net savings. If the latter are negative, they borrow (by taking loans b from the banks) more than they save (by buying shares of the bank a). If they are positive, the reverse holds. In theory nothing prevents the households from only borrowing, only saving, or both borrowing and saving. In equilibrium there is a region of optimal portfolios in the (b, a) space that is determined by the credit limit and the optimality conditions of the household' problem. Let $\lambda(\theta, a, b)$ denote the distribution of households over productivity levels, bank share holdings and loans.

2.2 Government

The government chooses the aggregate supply of bonds B_t, the unemployment benefit ν_t and the lump-sum tax τ_t, so as to satisfy the budget constraint

$$B_t + \nu_t u_t = p_t^g B_{t+1} + \tau_t,$$

where $u_t = \Pr(\theta_{it} = 0)$ is the fraction of unemployed agents in the population[6]. Assume that the supply of government bonds and the unemployment benefits are kept constant at B and ν, while the tax τ_t adjusts to ensure government budget balance. The governments' budget constraint rewrites

$$r_t^g B = (1 + r_t^g)(\tau_t - \nu u_t),$$

where r_t^g is the real interest rate on government bonds from t to $t+1$.

2.3 Financial intermediaries

The financial sector of the economy consists of a continuum of measure 1 of perfectly competitive financial intermediaries, which grant loans to households. They can be thought of as credit card companies or "banks". Banks' expenses consist of loans to households and government bonds. Loans to households take the forms of lines of credit, in which banks choose the credit limit ϕ – an upper bound on the amount of loans granted to households – up to which households can borrow at a net lending rate r^b.

By assumption, macroprudential regulation restricts the banks' credit limit choice to take two values, $\phi \in \left\{\underline{\phi}, \overline{\phi}\right\}$ ($0 < \underline{\phi} < \overline{\phi}$). In some sense it is for the bank a capacity problem, since it chooses the maximum amount of loans that it will be able to grant to households. Depending on households, this constraint will be occasionally binding, but it will not bind in the aggregate (since in equilibrium there must be households saving by holding bank equity, and it is unlikely that all of these are indebted up to the credit limit). The financial sector is competitive, so banks do not internalize the effect of choosing ϕ on aggregate laws of motion.

On the asset side of their balance sheets, banks also invest in risk free government bonds B^g earning a net rate of return r^g. Banks' resources consist of equity (or capital, or net worth) held by households as a saving vehicle, the credit line fraction unused by households (the difference between the credit limit and the amount of loan actually taken by households), and of exogenous capital Ω.

Aggregate risk Exogenous, stochastic bank capital is the source of aggregate risk in the economy. Ω can be either "high" or "low", $\Omega \in \left\{\underline{\Omega}, \overline{\Omega}\right\}$ ($\underline{\Omega} < \overline{\Omega}$. The transition between values of bank capital

[6] The continuum of agents has measure 1, so by the LLN and without renormalization, the fraction of unemployed is equal to the probability of being unemployed. u is constant after realizations of θ_{it} are drawn from the ergodic distribution, if the stochastic process for θ_{it} is independent from the aggregate shock.

is governed by a discrete Markov process whose transition probability matrix is

$$\Pi_\Omega = \begin{pmatrix} \pi_{ll} & \pi_{lh} \\ \pi_{hl} & \pi_{hh} \end{pmatrix}.$$

In the calibration, I assume that aggregate risk due to bank capital fluctuations and households idiosyncratic income risks are independent.

2.4 Recursive formulation of agents' problems

At the beginning of period t, banks choose ϕ and B. Then households choose c, a, b, n, taking ϕ as given. The aggregate states are bank capital Ω and the distribution λ of households across productivity (income) levels, asset holdings and loans.

Households Household's i individual state at the beginning of period t is $(\theta_{it}, a_{it}, b_{it}; \Omega_t, \lambda_t)$. Households forecast the future price of loans p^b_{t+1} and the future price of banks' shares p^a_{t+1}, based on the current and next period values of the aggregate state Ω_t and Ω_{t+1}. Since p^b_{t+1} and p^a_{t+1} depend on the distribution of households in $t+1$, they also forecast the distribution λ_{t+1}. In recursive form, the households' problem is (omitting from i subscripts and aggregate states for ease of notation):

$$V(b_t, a_t, \theta_t) = \max_{c_t, n_t, b_{t+1}, a_{t+1}} U(c_t, n_t) + \beta \mathbb{E}_t V(b_{t+1}, a_{t+1}, \theta_{t+1}),$$

subject to the constraints

$$c_t + p^a_t a_{t+1} + p^b_t b_{t+1} \leq \theta_t n_t + (p^a_t + \pi^a_t) a_t + b_t,$$

$$0 \geq b_{t+1} \geq -\phi_t,$$

$$a_{t+1} \geq 0.$$

Substituting for c_t using the budget constraint, and denoting μ^{b+}_t, μ^{b-}_t and μ^a_t the multipliers associated with the financing constraints, optimality conditions with respect to n_t, b_{t+1} and a_{t+1} are the following:

$$-U_n(c_t, n_t) = U_c(c_t, n_t)\theta_t \text{ and } n > 0, \text{ or } -U_n(c_t, n_t) \geq U_c(c_t, n_t)\theta_t \text{ and } n = 0$$

$$-p^b_t U_c(c_t, n_t) + \beta \mathbb{E}_t V_b(b_{t+1}, a_{t+1}, \theta_{t+1}) - \mu^{b+}_t + \mu^{b-}_t = 0$$

$$-p^b_a U_c(c_t, n_t) + \beta \mathbb{E}_t V_a(b_{t+1}, a_{t+1}, \theta_{t+1}) + \mu^a_t = 0$$

Making use of the envelope condition, the last two Euler equations can be rewritten as

$$p_t^b U_c(c_t, n_t) = \beta \mathbb{E}_t \left[U_c(c_{t+1}, n_{t+1}) \right] + \mu_t^{b-} - \mu_t^{b+}$$

$$p_t^a U_c(c_t, n_t) = \beta \mathbb{E}_t \left[U_c(c_{t+1}, n_{t+1})(p_{t+1}^a + \pi_{t+1}^a) \right] + \mu_t^{a-}$$

Because of the financing constraints, their interpretation differs from a model in which households are free to trade all available assets (both bonds and stocks). For loans, at the optimum the current marginal benefit from taking an additional unit of loans (for households, selling a zero coupon bond at price p^b) is equalized with the discounted marginal costs of doing so, namely repaying one unit of loan (the bond's face value) the next period (and thus forego one unit of consumption), adjusting for the shadow values of relaxing and tightening the borrowing constraint. For assets, the current marginal loss from buying an additional unit of bank shares is equated to the discounted marginal benefit it, namely earning banks' profit next period and the ex-dividend price, as well as the shadow value of relaxing the no-short selling constraint.

To solve the model numerically, I adapt Krusell-Smith algorithm by restricting agents to forecast a vector $\mathbf{m}$ of moments of the distribution (mean, variance), and to use linear laws of motion. So for computations, let household i's state in period t be $\left(\theta_{it}, a_{it}, b_{it}; \Omega_t, \mathbf{m_t^A}; p_t^a, p_t^b\right)$. Let A_t denote aggregate savings – equal to $\int a_{it} d\lambda_t = \int a_{it} d\lambda_t$ – which must equal $\overline{A} = 1$ in equilibrium, the aggregate amount of banks' shares.

Usually in Krusell-Smith algorithms, the last two states $\left(p_t^a, p_t^b\right)$ allow to ensure that the asset market (bank loans and shares) clears during the simulation step. It is still the case for p_t^a (bank shares are in exogenous supply of 1), but not for p_t^b. Indeed, banks don't choose the actual aggregate amount of loans B_{t+1}, it results from the households' decisions. Banks only choose the maximum possible amount of loans ϕ_t they are willing to grant. The aggregate amount of loans must satisfy $B_{t+1} \leq \phi$, which holds by construction because $b_{it+1} \geq -\phi_t$ for all i and t. However, p_t^b is still needed because it is in the period t information set, based on which households forecast future prices and moments of the distribution.

Households' forecasts are ($B_{\Omega,\Omega'}^1$ is a matrix of coefficients):

$$\mathbf{m_{t+1}} = b_{\Omega,\Omega'}^0 + B_{\Omega,\Omega'}^1 \mathbf{m_t} + b_{\Omega,\Omega'}^2 p_t^a + b_{\Omega,\Omega'}^3 p_t^b,$$

$$p_{t+1}^a = c_{\Omega,\Omega'}^0 + C_{\Omega,\Omega'}^1 \mathbf{m_t} + c_{\Omega,\Omega'}^2 p_t^a + c_{\Omega,\Omega'}^3 p_t^B$$

$$p_{t+1}^b = d_{\Omega,\Omega'}^0 + D_{\Omega,\Omega'}^1 \mathbf{m_t} + d_{\Omega,\Omega'}^2 p_t^a + d_{\Omega,\Omega'}^3 p_t^b$$

Given the aggregate laws of motion, I solve the agent's problem using the endogenous gridpoints method for problems with occasionally binding constraints among endogenous variables (as in Hintermaier and Koeniger (2010)).

Financial intermediaries I solve the banks' problem by value function iteration with linear interpolation. As it is clear from their optimality conditions, banks also forecast future prices and the distribution of agents. For consistency, they use the same laws of motion forecasts as households. Banks maximize the expected discounted flow of current and future dividends π_t^B, weighted by the sequence stochastic discount factor $\Lambda_{t,t+j}$. Favilukis et al. (2016) discuss the choice of the SDF to evaluate the value of a mutual fund sector held by households, which is analogous to the financial intermediaries sector. They give evidence that the choice of the SDF is close to innocuous for the results. The same tests can be run in the numerical solution of the full model. Banks choose government bond and credit limit sequences to solve the following problem:

$$V_t^a = \max_{\{\phi_{t+j}, B_{t+1+j}\}_{j\geq 0}} \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \Lambda_{t,t+j} \pi_{t+j}^a = \max_{\{\phi_{t+1}, B_{t+1}\}} \pi_t^a + \beta \mathbb{E}_t \Lambda_{t,t+1} V_{t+1}^a,$$

subject to the one-period budget constraints and domain restrictions:

$$\pi_t^a + p_t^b \phi_t + p_t^g B_{t+1}^g \leq \Omega_t + B_t^g + B_t + (\phi_{t-1} - B_t),$$

$$\phi_t \in \left\{\underline{\phi}, \overline{\phi}\right\},$$

$$\pi_t^a \geq 0.$$

$0 \leq B_t = -\int b_{it} d\lambda_t \leq \phi_t$ denotes the aggregate amount of loans taken by households.
The budget constraint binds at the optimum, hence

$$\pi_t^a(\phi_t, B_{t+1}^g) = \Omega_t + B_t^g + B_t + (\phi_{t-1} - B_t) - p_t^b \phi_t - p_t^g B_{t+1}^g.$$

The bank's problem can be decomposed by first solving for the optimal amount of government bonds given a credit limit, and then choosing the optimal one of the two credit limits:

$$V_t^a = \max\left(\underline{v_t^a}, \overline{v_t^a}\right),$$

where

$$\underline{v_t^a} = \max_{B_{t+1}^g} \pi_t^a(\underline{\phi}, B_{t+1}^g) + \beta \mathbb{E}_t \left[\Lambda_{t,t+1} V_{t+1}^a | \phi_t = \underline{\phi}\right]$$

and

$$\overline{v_t^a} = \max_{B_{t+1}^g} \pi_t^a(\overline{\phi}, B_{t+1}^g) + \beta \mathbb{E}_t \left[\Lambda_{t,t+1} V_{t+1}^a | \phi_t = \overline{\phi}\right]$$

2.5 Equilibrium

I look at the recursive competitive equilibrium of the economy in which shocks to households' debt limits result from aggregate shocks to banks' balance sheets.

Definition 1 *Given government taxes τ_t, the* RCE of an economy with stochastic debt limits *consists of* i) *a collection of household's decision rules for consumption, next period bank shares bonds and labor supply* $(c(\theta,a,b;\Omega,\lambda), a'(\theta,a,b;\Omega,\lambda), b'(\theta,a,b;\Omega,\lambda), n(\theta,a,b;\Omega,\lambda))$*;* ii) *a collection of value functions for households and financial intermediaries,* $V(\theta,a,b;\Omega,\lambda)$ *and* $V^a(\theta,a,b;\Omega,\lambda)$*;* iii) *pricing functions* $\left(p^a(\Omega,\lambda), p^b(\Omega,\lambda), p^g(\Omega,\lambda), w(\Omega,\lambda)\right)$*;* iv) *a law of motion* Ψ *for the distribution* λ*, such that:*

- *given prices and banks' choices, the household's decision rules solve the household's problem, and V is the associated value function;*
- *given prices, the banks' decision rules solve the banks' problem, and V^a is the associated value function;*
- *efficiency units of labor are paid at their marginal product, $w = 1$;*
- *the labor market clears, $\int n(\theta,a,b;\Omega,\lambda)d\lambda = 1$;*
- *the goods market clears, $\int c(\theta,a,b;\Omega,\lambda)d\lambda = \int y(\theta,a,b;\phi,\lambda)d\lambda = \int \theta n(\theta,a,b;\Omega,\lambda)d\lambda$;*
- *the bond market clears, $\int b'(\theta,a,b;\Omega,\lambda)d\lambda = \bar{B}$;*
- *the equity market clears, $\int a'(\theta,a,b;\Omega,\lambda)d\lambda = \bar{A}$;*
- *the tax satisfies the government budget constraint;*
- *for every pair (Ω,Ω'), the aggregate law of motion Ψ is generated by the exogenous Markov chain Π_ϕ and the decision rule a', b' as follows*[7]

$$\lambda'\left(A \times B \times \theta'\right) = \Psi_{(A\times B\times\theta')}(\Omega,\Omega',\lambda) = \int_{\mathcal{A}\times\mathcal{B}\times\Theta} Q_{\Omega,\Omega'}\left((\theta,a,b),(\theta',A,B)\right)d\lambda,$$

where $Q_{\Omega,\Omega'}$ is the transition function between two periods where the aggregate shock transits from Ω to Ω', defined by

$$Q_{\Omega,\Omega'}\left((\theta,a,b),(\Theta,A,B)\right)d\lambda = \mathbb{1}\left\{a'(\theta,a,b;\Omega,\lambda)\right\}\mathbb{1}\left\{b'(\theta,a,b;\Omega,\lambda)\right\}\sum_{\theta'\in\Theta}\Pi_\theta(\theta'|\theta).$$

[7] Formally, the state space is $\mathcal{S}, \Sigma_{\mathcal{S}}$, with $\mathcal{S} = \mathcal{A}\times\mathcal{B}\times\Theta$, where Θ is the discrete set of productivity levels, and $\mathcal{A},\mathcal{B}$ the sets of possible asset and bond holdings (restricted to be compact). $\Sigma_{\mathcal{S}}$ is the sigma-algebra associated with $\mathcal{S}$.

3 Numerical solutions

This section presents the model calibration, and proposes two algorithms for solving the reduced-form model in which debt limit fluctuations are exogenous and follow a two-state Markov chain. In one algorithm, agents forecast bond prices directly as functions of current and lagged debt constraints. This solution is close to the method implemented in (Khan and Thomas, 2013) (direct price forecasts). In the other, agents forecast moments of the bond distribution and their relations to bond prices (indirect price forecasts). In equilibrium, agents' forecasts for the laws of motion of aggregate variables coincide with their actual paths.

3.1 Calibration

The reduced-form model is solved and simulated at quarterly frequency. I calibrate it to match in a steady state without stochastic debt limits the empirical targets summarized in Table 2. Borrowing limits in the easy and in the tight credit regimes are respectively chosen to be $\phi = 1.75$ and $\phi' = 1.3$. The size of the shock is chosen so that the debt-GDP ratio drops by 5 percentage points in the new steady state, from 18% to 13%[8]. The calibration of the pre-crisis household debt-GDP ratio around 20% is the same as in Nakajima and Rios-Rull (2014) and Guerrieri and Lorenzoni (2016). I assume a relatively low probability of entering a bad credit regime, and a slightly higher probability to remain in the latter once the economy has entered it[9].

The productivity process is discretized using Tauchen method. The asset grid has 1200 equally-spaced points from -3 to 30.

Post-crisis real interest rates The model is calibrated to match the mean of 5- to 20-year TIIS yields from October 2008 to August 2016. Time-series averages of 5, 7, 10 and 20-year TIIS yield are respectively (in percent) 0.0740, 0.4199, 0.70 and 1.2273. Their standard deviations are 0.9120, 0.8766, 0.8084, 0.7485 percent.

3.2 Direct price forecasts

Since bonds are in exogenous fixed supply, there is no forecasting equation for bonds, as opposed to capital in Krusell and Smith (1998). Households forecast bond prices directly, as a function of

[8] The Krusell-Smith solution of the model turns out to be unstable for debt-to-GDP ratio drops larger than 10%.

[9] Later on this value should be calibrated to match the average frequency of credit/banking crises in the data.

current and past credit regimes. Since there are two credit regimes, the number of prices to guess is 2^n, where n is the number of regimes the agent keeps track of. The longer the expected duration of the tight credit regime, the larger the number of past credit regimes needed to accurately model the dynamics of the economy in that regime.

Algorithm

1. Guess 2^n prices as a function of current and past credit limits. I assume that the agents keep track of three debt limits to forecast bond prices[10]

$$q_t = \frac{1}{1+\tilde{r}_t} = q(\phi_t, \phi_{t-1}, \phi_{t-2}).$$

   There are $2^3 = 8$ possible sequences $(\phi_t, \phi_{t-1}, \phi_{t-2}) \in \{\phi_l, \phi_h\}^3$, so we guess a $(2 \times 2 \times 2)$ matrix of prices. The state space is $(b, \theta; \phi, \phi_{-1})$.

2. Solve the household's problem and obtain the decision rule $b_{t+1}(b_t, \theta_t; \phi_t, \phi_{t-1}, \phi_{t-2})$, which is an approximation to the original decision rule $b_{t+1}(b_t, \theta_t; \phi_t, \lambda_t)$. I use the endogenous grid method from Carroll (2006).

   We define a grid over the additional aggregate state variable. At every point on the grid $(b, \theta; \phi, \phi_{-1}, \phi_{t-2})$, solve for the choice $b' = b^*$ that satisfies the following Euler equation:

$$\begin{aligned}
& U_c\left((1+\tilde{r}(\phi, \phi_{-1}, \phi_{t-2}))b + y(\theta) - b^*\right) \\
\geq \;& \beta \sum_{\phi', \theta'} \left[(1+\tilde{r}(\phi', \phi, \phi_{-1}))U_c\left((1+\tilde{r}(\phi', \phi, \phi_{-1}))b^* + y(\theta') - b'(b^*, \theta'; \phi', \phi, \phi_{-1})\right)\right] \\
\times \;& \pi(\phi', \theta' | \phi, \theta),
\end{aligned}$$

   where $\tilde{r}$ depends on current and past credit limits. (In practice I compute the RHS of the Euler equation for a given initial guess of the consumption policy, and compute the implied current optimal consumption. I then use the budget constraint to solve for the endogenous amount of bonds leading to these policies, then solve the problem again and iterate until the consumption policy converges at the 10^{-8} criterion under the sup norm.)

3. In the following steps, simulate the economy for N individuals and T periods. Do it one time for a given draw of random numbers, and check if (i) prices converge under the sup norm and

[10] Khan and Thomas (2013) also pick the current value and two lagged values of the aggregate credit shock. Note that this kind of numerical strategy is not costly in Julia, where loops are as fast as vectorized operations.

(ii) the bond market clears. In each period t compute the (simulated) aggregate demand for bonds B_t from the cross-sectional distribution.

Specifically, I simulate a time series of credit constraint shocks ϕ_t and a panel of idiosyncratic productivity values for agents θ_{it}, respectively initialized at their stationary distributions, where the latter are computed from the transition matrices. I draw the initial distribution of bonds b_{it} across agents from the stationary distribution of bonds in a steady state of the economy with a permanently loose credit regime ($\phi_t = \phi_h \; \forall t \geq 0$)[11].

4. Construct the time series for a 2^n-valued categorical variable characterizing the credit regime at date $t \geq 1$, defined as combination $(\phi_t, \phi_{t-1}, \phi_{t-2}) \in \{\phi_l, \phi_h\}^3$.

5. Check market clearing every period by comparing the demand for bonds $B_t = \int b_{t+1} d\lambda_t$ to the bond supply $\bar{B}$, for each of the credit regimes for which we have guessed bond prices. If in a given credit regime $(\phi_t, \phi_{t-1}, \phi_{t-2}) = (\phi', \phi, \phi_{-1})$, $B_t > \bar{B}$ (households save more than there are bonds to be purchased), then decrease the interest rate $r(\phi', \phi, \phi_{-1})$ guessed for that credit regime. Otherwise, increase that interest rate.

6. Finally, assess whether the solution is accurate enough by checking the market clearing condition for bonds every period in the simulated economy. If fit at the solution is not satisfactory, add additional past credit limits to the state space. Repeat until the economy is close enough to market clearing every period.

Unlike in transitional dynamics algorithms, in the stochastic simulation we do not necessarily want to put more weight on the first periods compared to later ones. One reason being that we draw initial values form the stationary distribution of bonds and productivity. More importantly, it is not possible to exactly clear the bond market with only $2^n < T$ prices: in this sense the laws of motion from state to state computed here are only an approximate laws of motion. This creates the issue of the measure of bond market clearing that one chooses in this setting, where we only have a limited number of prices to clear the bond market. One approach taken here is to compute the average distance on the bond market for all periods of a given credit regime, for each credit regime. Then depending on its value, increase or decrease the interest rate. This allows the bond market to clear ”on average”. More specifically, the updating rule I consider minimizes the average bond market distance in each regime by adjusting the interest rate in this regime, proportionally to

[11]This corresponds to the pre-crisis steady state of the transitional dynamics model of Guerrieri and Lorenzoni (2016).

the frequency of that regime over the entire simulation. The algorithm converges to bond market clearing according that distance, at less than the 10^{-4} criterion, and so do prices over successive iterations.

3.3 Indirect price forecasts

Aggregate shocks to households' credit constraints induce the second and higher (centered) moment of the bond distribution to move over time. For instance, its variance decreases when the credit constraint tightens (less agents are indebted, and rich agents hold less bond to make the latter possible), and increases when it loosens. I therefore modify the baseline KS algorithm to have agents forecasting the law of motion of the variance and higher of the bonds distribution over time. Intuitively, this alleviates the concern raised in Guerrieri and Lorenzoni (2016) that the single aggregate quantity of bonds outstanding might be a poor forecaster of the trajectory of households at the edges of the state space, in particular credit-constrained households. I assume that agents use a linear forecast rule for the variance of the bonds distribution:

$$\text{Var}B_{t+1} = b^0_{\phi,\phi'} + b^1_{\phi,\phi'}\text{Var}B_t + b^2_{\phi,\phi'}r_t,$$

where the b coefficients depend on the current and next period borrowing constraints $\phi, \phi' \in \{\phi_l, \phi_h\}$.

Another possible specification to replicate more closely the U-shaped response of variables in the KS economy would include an additional term for previous values of aggregate bond holdings.

Similarly, agents forecast next period interest rate (from $t+1$ to $t+2$) according to the following linear law of motion:

$$r_{t+1} = a^0_{\phi,\phi'} + a^1_{\phi,\phi'}\text{Var}B_t + a^2_{\phi,\phi'}r_t.$$

The state space used for the numerical solution of the model is $(\theta_{it}, b_{it}; \phi_{t-1}, \phi_t, \text{Var}B_t, r_t)$ at time t for agent i. Note the differences with the "theoretical" state space from the model above. First, it includes the last period borrowing constraint ϕ_{t-1} (such that current bonds $b_t \geq -\phi_{t-1}$. This allows to speed up computations, by directly computing analytically the consumption of an agent who was constrained and is still constrained in the current period from the budget constraint, for various productivity levels. Second, it includes the interest rate r_t faced by the agent in the current period, on his bond holdings choice b_{t+1} from t to $t+1$. The reason is that, since bonds are in exogenous supply, r_t depends on the aggregate demand for bonds B_{t+1}, hence on the distribution of b_{it+1} across agents, which itself depends on r_t. Thus this is a fixed point problem in r_t. In the KS algorithm, since agents forecast next period's state space, we need to ensure that the bond market clears every period. Having r_t in the state space allows to implement this additional step by having agents' current demand for bonds be menus that depend on r_t, and then find the interest rate r_t^* that clears the bond market (this step is implemented using a nonlinear solver). The flip side of the coin is that it increases the dimension of the state space.

Algorithm

1. Specify a functional form for the law of motion for the variance of the bond distributio and the interest rate (see above).

2. Guess the matrices of coefficients $\left\{a^0_{\phi,\phi'}, a^1_{\phi,\phi'}, a^2_{\phi,\phi'}, b^0_{\phi,\phi'}, b^1_{\phi,\phi'}, b^2_{\phi,\phi'},\right\}$.

3. The state for household i at date t is $(\theta_{it}, b_{it}; \phi_{t-1}, \phi_t, \text{Var}B_t, r_t)$. We include the current aggregate states used for the agent's forecasts, and the current interest rate. This additional dimension is required to find the market-clearing interest rate in the simulation step. Define additional grids over those.
   Solve the household's problem and obtain the decision rule $b_{t+1}\left(\theta_{it}, b_{it}; \phi_{t-1}, \phi_t, \text{Var}B_t, r_t\right)$, which is an approximation to the original decision rule $b_{t+1}(\theta_{it}, b_{it}; \phi_t, \lambda_t)$.
   I use the endogenous grid method from Carroll (2006).
   I also compute analytically the consumption of agents constrained in the current and previous periods from the budget constraint, as $c_{it} = \theta_{it} n_{it} - \phi_{t-1} + \frac{\phi_t}{1+r_t}$, where the labor supply n_{it} satisfies the intra-temporal optimality condition between labor and consumption (I use a nonlinear solver to jointly solve for c_{it} and n_{it} satisfying those two conditions).

4. Simulate the economy for N individuals and T periods. In each period t compute (the simulated) B_t and its moments from the cross-section distribution (stochastic simulation method). For instance $N = 10,000$ and $T = 1,500$ (with 500 period burnt).

   (a) Fix the random number generator seeds, and draw once and for all a panel of exogenous productivity levels $\{\theta_{it}\}_{it}$, and a time series of exogenous borrowing constraints $\{\phi_t\}_t$, starting from the loose borrowing constraint (to be consistent with the initial bond distribution). Also draw the initial bond distribution (the cross-section $\{b_{i0}\}_i$ from the stationary distribution in the deterministic steady state where the loose borrowing constraint prevails forever. Ensure that the bond market clears given the exogenous bond supply, by redistributing bonds equally across agents if it doesn't. Compute the variance of the initial bond distribution. Thus we have all the components of the initial state space.

(b) Then use the policy functions for which we have solved above to compute simulated values for consumption, bond holdings and labor supply. Policy functions are linearly interpolated on a fine grid (the model was solved on a coarse grid).

(c) Every period $t \geq 0$, use a non-linear solver to solve for the bond market-clearing interest rate r_t^*. Specifically, set up a function that computes every agent's demand for bonds and aggregate them for a given interest rate, and then solve for the zero of that function minus the exogenous bond supply. This step is somewhat computationally costly.

(d) At every period, re-compute agents' consumption, bond demand and labor supply using the (linearly interpolated) policy functions with the market-clearing clearing r_t^*, and the time series of variances (and higher moments) of the bond distribution implied by those interest rates.

5. Since the law of motion is time-invariant, we can separate the dates $t, t+1$ in the sample where the current and next period borrowing constraint is ϕ_l, ϕ_h. Split the sample into four sub-samples, depending on the current and next period values of the borrowing constraints, $\phi, \phi' \in \{\phi_l, \phi_h\}$. In each sub-sample, run an OLS regression of the simulated (compatible with bond market clearing) $\{\text{Var}B_{t+1}\}_t$ and $\{r_{t+1}\}_t$ on $\{\text{Var}B_t\}_t$ and $\{r_t\}_t$, to estimate the coefficients for each sub-sample.

6. If estimated coefficients differ from their preceding values, update them and go back to step 2.

7. Continue until convergence: fixed point algorithm in the coefficients of the laws of motion.

8. Assess whether the solution is accurate enough. One possibility is to compute R^2s for each of the regressions. If the fit at the solution is not satisfactory, add moments of the distribution (such as skewness, kurtosis) or try a different functional form for the law of motion.

4 Numerical results

4.1 Benchmark model with unexpected debt limit fluctuations

I first calibrate the model parameters to match the same empirical targets as above. This involves computing two steady states of the economy, the first one with debt limit ϕ_h and the second one, after the credit crisis, with debt limit ϕ_l.

Figure 4 presents the results for the transitional dynamics experiment. I implement the algorithm described in appendix A.1. As the borrowing limit gradually declines, constrained households decrease their debt, and unconstrained households increase their asset position. As a result, aggregate debt to GDP decreases from 18% to 13%. Because agents accumulate more assets and hold less debt, they consume less, with consumption decreasing by 1% compared to its pre-crisis level. As a result of agents' deleveraging, the real interest rate drops to clear the bond market, and reaches negative values. Figure 5 shows how the decrease in the bond demand schedule generates a decrease in the interest rate when government bonds are in fixed supply.

4.2 Reduced-form model with stochastic debt limits

I present results from simulations of $N = 20,000$ households for $T = 1,500$ periods. The Julia code runs for about 20 minutes on an Intel Core i7-4770S desktop with 3.10 GHz and 8.00 Gb of RAM.

The initial distribution of productivity across agents at $t = 0$ is drawn from the stationary distribution. For the first two periods $t = 0, 1$, the bond distribution is drawn from the stationary distribution when the economy has been forever in the high credit regime ϕ_h[12]. I redistribute assets across agents so that the bond market clears by construction in the first two periods.

4.2.1 The real interest rate

In equilibrium, the pricing function relates the interest rate to a sequence of credit regimes (in the direct price forecasting solution of the model),

$$(\phi_t, \phi_{t-1}, \phi_{t-2}) \in \{\phi_l, \phi_h\}^3 .$$

Table 1 gives the values of the interest rate in the $2^3 = 8$ credit regimes.

[12]This corresponds to the steady state before the crisis of the model with transitional dynamics.

(l,l,l)	(l,h,l)	(h,l,l)	(h,h,l)	(l,l,h)	(l,h,h)	(h,l,h)	(h,h,h)
< -8.96	-0.36	0.18	2.04	-8.96	-3.51	-0.92	2.3

Table 1: Price function: equilibrium values of the interest rates in the various credit regimes $(\phi_t, \phi_{t-1}, \phi_{t-2})$. The interest rate is in annual percentages. The regime (ϕ_l, ϕ_l, ϕ_l) is never reached because of its low probability in transition matrix, the corresponding value is not informative because the initial guess is never updated. Values are in annual percentage terms.

The possibility of a crisis regime creates a additional need for precautionary savings, so the real interest rate needs to be low for the bond market to clear. Therefore it is lower on average over the whole simulation, and the collapse when reaching a credit regime is stronger than in transitional dynamics. This pattern hinges on the credit shocks being expected as opposed to unexpected. Overall, the distribution of interest rates over the business cycles is shifted downwards.

The interest rate reaches its lowest value when the economy has been in a tight credit regime for two periods after exiting a loose credit regime. Two effects jointly explain the sharp drop in the interest rate. During the credit crisis poor indebted agents strive to decrease their debt when they are at the constraint, as in the benchmark economy. Richer unconstrained agents also want to accumulate assets to hedge against future credit regimes changes. Thus agents' expectations about credit regime changes contribute to maintain an extremely low interest rate environment, similar to the one characterizing the recent US data. This effect is stronger if i) credit crises (regimes of tight credit) are persistent, and ii) credit crises are more likely in general. Figure 6 takes a subset of 500 periods in the simulation and displays the dynamics of the interest rate as the economy alternates the two credit regimes.

I also compare long-run simulation averages of the real interest in the indirect price forecast solution of the model. In the latter, the interest rate schedule is continuous and can take more than 8 values. Therefore, long-run averages display less sharp variations than the discrete interest rate schedule of the direct price forecast solution. Here, the interest drops on average by 10% following a tightening of households' debt limit compared to its value in a regime of easy credit (its annualized value becomes 1.69%, against 1.89%). Thus the real interest rate is on average very low in all credit regimes. In addition, periods of low interest rates are persistent. Figure 8 gives evidence of it in an example subsample in the simulations. In this example, the real interest rate fluctuates at low values, between 1% and 2% (annualized). It decreases in credit crises and stays low for time spans ranging from 2 (0.5 year) to 10 periods (2.5 years).

4.2.2 Consumption

I compute long-run statistics for consumption in regimes of easy and tight credit. While consumption drops sharply in the benchmark model, here its long-run average is almost the same in both credit regimes (around 0.42 in the calibration considered), and it is even higher in low credit regimes in some of the simulations (though not by much). The cross-sectional variance of consumption slightly increases following tightening of households' debt limits (from 0.00487 to 0.00488), while the variance of the bond distribution decreases (from 13.52 to 13.48). Thus in the model financial crises increase consumption inequality and decrease wealth inequality. The former increases because poor households still have to deleverage while rich households sell bonds and consume. The latter decreases because poor households are less indebted and hold more bonds, while rich households hold less bonds because they have a lower return. Overall, the model with stochastic debt limits generates qualitatively similar dynamics as the liquidity trap model for household debt to GDP and the interest rate and employment. Quantitatively, however, the former economy has more bonds (less debt), hence lower interest rates and smoother consumption. These patterns match the recent period of durably low interest rates associated with moderate consumption fluctuations observed in the US since 2009.

4.2.3 Decision rules

Figure 7 plots households' decision rules in the various credit regimes. Orders of magnitude are similar to transitional dynamics, but decision rules display different patterns.

Results are consistent with simulations: intuitively, the most striking difference across regimes is for indebted agents, when the economy transits from a regime of loose credit to a regime of tight credit. When entering a tight credit regime, indebted agents with low productivity decrease their consumption compared to the loose credit regime, supply more labor and accumulate less their debt. Agents with higher productivity consume more, supply less labor and decrease less their debt position.

Overall, agents' consumption is low in a regime of persistently easy credit, and they accumulate large amounts of bonds as precautionary savings. When the credit shock hits in (l, h, h), they cut their debt, but may also consume more. The interest rate adjusts to clear the bond market. The worst case for consumption (for a given productivity level) is h, l, l: after having been in a tight credit regime for two periods, to be again in a high credit regime. Households need to deleverage during credit crisis, *and* build up precautionary savings in prevision of future crises. These results differ for example from the conclusions of Lorenzoni (2008) about "inefficient credit booms". In the economy with stochastic debt limits, there is over-accumulation of assets, not excessive borrowing before a credit crisis.

Everything being equal (in particular for given bond holdings), agents with relatively high productivity consume less in (h, h, h) than in (l, h, h) because they accumulate more assets. This is not the case of low income (low productivity) agents, who don't have enough resources to accumulate assets, and who consume because otherwise their utility would be too low. Often those are indebted agents, who cannot accumulate assets, and mechanically consume more in the good credit regime than in the bad one because they can get more debt to consume in the good regime. This is confirmed by bond accumulation policy. Thus it is a feature of this model that whatever their asset position, more productive agents accumulate more assets (they earn more income), even when they are in a position to be credit-constrained. Another feature is that even though agents increase their bond accumulation in tight credit regimes to deleverage (because they are not allowed to hold as much debt as before the shock), they still want to accumulate even more bonds in good times, as a hedge against bad times (they are overall richer in good times, while they are closer to consuming their income in bad times).

Finally, two remarks can be made about households' labor supply policy. First, agents work less when having been in a persistent tight credit regime. Second, it takes larger values of bond holdings for them stop working and become unemployed than in the liquidity trap economy. When they are very productive they even for high bond values, an effect absent from the latter economy. Thus internalizing the possibility of a credit supply shock has the effect of making them work more on average. Higher incentive to work are due to the extra need to accumulate more bonds as precautionary savings. However, because agents working more in credit crises are not necessarily the most productive ones, an increase or decrease in aggregate labor supply does not imply corresponding variations in output and consumption.

4.3 Full model with shocks to banks' balance sheets

[In progress]

5 Conclusion

This paper describes a model of credit crises in which stochastic credit regimes create aggregate uncertainty about the debt limits faced by households. The latter results from shocks to the financial sector's balance sheet, which translate into fluctuating credit conditions over which households form expectations. The model's contribution is twofold. First, it bridges the gap between dynamic quantitative models with incomplete markets and heterogeneous households, and the financial intermediation literature, by studying the implications of adding a stylized realistic financial sector to the former. Second, it analyzes a setting in which fluctuations in households' debt limits are objects over which agents form expectations, rather than being unexpected, and are stochastic rather than having a deterministic trajectory, thus departing from classical models of the liquidity trap. The model can generate extended period of very low interest rates and moderate consumption fluctuations following credit crises, which are a feature of the recent US data.

References

Yann Algan, Olivier Allais, Wouter J. Den Haan, and Pontus Rendahl. Chapter 6 - solving and simulating models with heterogeneous agents and aggregate uncertainty. In Karl Schmedders and Kenneth L. Judd, editors, *Handbook of Computational Economics Vol. 3*, volume 3 of *Handbook of Computational Economics*, pages 277 – 324. Elsevier, 2014.

Charles Bean, Christian Broda, Takatoshi Ito, and Randall Kroszner. Low for long? causes and consequences of persistently low interest rates. Technical Report Geneva Reports on the World Economy 17, Centre for Economics Policy Research, October 2015.

Emine Boz and Enrique Mendoza. Financial innovation, the discovery of risk, and the u.s. credit crisis. *Journal of Monetary Economics*, 62:1–22, 2014.

Christopher D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91:312–320, June 2006.

Gauti B. Eggertsson and Paul Krugman. Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach. *The Quarterly Journal of Economics*, 127(3):1469–1513, 2012.

Emmanuel Farhi and Ivan Werning. A theory of macroprudential policies in the presence of nominal rigidities. June 2015.

Jack Favilukis, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. Theoryhe macroeconomic e ects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. *Journal of Political Economy*, 2016.

Jesús Fernández-Villaverde, Grey Gordon, Pablo Guerrón-Quintana, and Juan F. Rubio-Ramírez. Nonlinear adventures at the zero lower bound. *Journal of Economic Dynamics & Control*, 57: 182–204, June 2015.

Martin Floden and HJespere Linde. Idiosyncratic risk in the u.s. and sweden: Is there a role for government insurance? *Review of Economic Dynamics*, 4(2):406–437, 2001.

Veronica Guerrieri and Guido Lorenzoni. Credit crises, precautionary savings and the liquidity trap. April 2016.

Thomas Hintermaier and Winfried Koeniger. The method of endogenous gridpoints with occasionally binding constraints among endogenous variables. *Journal of Economic Dynamics and Control*, 34(10):2074 – 2088, 2010. ISSN 0165-1889.

Greg Kaplan, Kurt Mitman, and Gianluca Violante. Consumption and house prices in the great recession: Model meets evidence. 2015 Meeting Papers 275, Society for Economic Dynamics, 2015.

Aubhik Khan and Julia K. Thomas. Credit shocks and aggregate fluctuations in an economy with production heterogeneity. *Journal of Political Economy*, 121(6):1055–1107, December 2013.

Per Krusell and Anthony Smith. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, October 1998.

Guido Lorenzoni. Inefficient credit booms. *Review of Economic Studies*, 75:809–833, 2008.

Makoto Nakajima and Jose-Victor Rios-Rull. Credit, bankruptcy and aggregate fluctuations. Working Paper 20617, National Bureau of Economic Research, October 2014.

Christopher J. Nekarda and Valerie A. Ramey. The cyclical behavior of the price-cost markup. Working Paper 19099, National Bureau of Economic Research, June 2013.

Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49, 2005.

A Appendix

A.1 Transitional dynamics algorithm for the economy with unexpected debt limit fluctuations

Algorithm t denotes calendar time, not iterations in the algorithm. The use of bold font denotes vectors, matrices or higher dimensional arrays.

1. Compute the stationary recursive competitive equilibrium (RCE) in the two steady states of the economy $t = T_1$ and $t = T_2$, before and after the credit tightening (i.e. for the initial and the final value of ϕ_t). Obtain policy functions (indexed by time)

$$\{c_t(\theta, b), b_t'(\theta, n), n_t(\theta, b)\}_{t=T_1,T_2},$$

   prices $\{r_t\}_{t=T_1,T_2}$, and stationary distributions $\{\Lambda_t(\theta, b)\}_{t=T_1,T_2}$ (c.d.f – use λ to denote the p.d.f).

2. Guess a sequence of interest rates $\{r_t\}_{t=T_1,...,T_2}$. A first guess can be r_{T_2}.

3. Set up the (deterministic) sequence of borrowing constraints $\{\phi_t\}_{t=T_1,...,T_2}$. Set up a maximum number of tentative iterations to solve the transitional dynamics problem. From this step on solve the problem backward, starting from steady state $t = T_2$, up to T_1.

4. Solve backward for endogenous variables for $t = T_2, ..., T_1$.

   (a) Assign values for the interest rate r_t, the borrowing constraint ϕ_t, taxes τ_t and employment benefits z_t.

   (b) Compute (in closed-form, from the budget constraint) the consumption function $c_{cons,t}(\theta)$ of an agent with $b_t = -\phi_t$ and $b_t = -\phi_{t+1}$ (constrained today and tomorrow),

$$c_{cons,t}(\theta) = -\phi_t + \frac{\phi_{t+1}}{1 + r_t} + \theta n_t + z_t - \tau_t.$$

   (c) Compute the RHS of the Euler equation $\mathbf{RHS} = \mathbf{P}\mathbf{c}^{-\gamma}$ as a $\dim\Theta \times \dim\mathcal{B}$ 2-dimensional array, where $\mathbf{P}$ is the transition matrix for productivity θ and $\mathbf{c} = [c(\theta, b)]_{(\theta,b)\in\Theta\times\mathcal{B}}$ is the 2-dimensional array representing the consumption decision in the state space $\Theta\times\mathcal{B}$. Restrict RHS to consider unconstrained agents, i.e. take $\mathbf{RHS_t} = \mathbf{RHS_t}[:, b \geq -\phi_{t+1}]$.

(d) Loop over $\theta \in \Theta$. For a given θ, back out the $\dim \mathcal{B}$ 1-dimensional array for consumption implied by the RHS of the Euler equation as

$$\mathbf{c_t}(\theta) = \left[(1+r_t)\beta \mathbf{RHS_t}(\theta)\right]^{-1/\gamma}.$$

Deduce the corresponding $\dim \mathcal{B}$ 1-dimensional array for labor supply from the f.o.c as $\mathbf{n_t}(\theta) = \max\left\{0, 1 - \left(\frac{\psi}{\theta}\right)^{1/\eta} \mathbf{c_t}(\theta)^{-\gamma}\right\}$.

(e) Deduce the corresponding $\dim \mathcal{B}$ 1-dimensional array for current bond holdings (a function of tomorrow's bond holdings) from the budget constraint, given the agent is unconstrained tomorrow, as

$$\mathbf{b_t}(\theta) = \mathbf{b_{t+1}}[\mathbf{b_{t+1}} \geq -\phi_{t+1}]/(1+r_t) + \mathbf{c_t}(\theta) - \theta \mathbf{n_t}(\theta) - z_t + \tau_t.$$

This is the inverse of the bond accumulation policy (store it in a 3-dimensional array indexed by t). These points are (usually) not on the grid. Current bond holdings leading the agent to be constrained tomorrow are taken to be $\mathbf{b_{cons,t}}(\theta) = \min\{\mathbf{b_t}(\theta), \mathcal{B}\}$.

(f) Interpolate the fine bond grid $\mathcal{B}_f$ on the function for the inverse of the bond accumulation policy, which maps tomorrow's bond holdings to today's given θ. Thus we get back on the (finest) bond grid.

(g) Consider the case of constrained agents. If the value of current bond holdings leading an agent to choose the lowest bond holdings on the grid $\mathcal{B}$ greater than $-\phi_{t+1}$ is larger than $-\phi_t$, [and the consumption of an agent constrained today and tomorrow is lower than the lowest consumption of an agent unconstrained tomorrow][13], then add the following step for precision. Create a grid $\mathcal{C}_{cons,fine}$ of consumption values

$$\mathbf{c_{cons,t}} \leq c_{cons,fine} \leq \min \mathbf{c_t}.$$

Deduce the corresponding labor and inverse of the bond accumulation policy.

Why is this step needed? Because the inverse of the bond accumulation policy is a monotonic function, if the conditions above hold, then it must be that all values of current bond holdings between $\mathbf{b_t}(\theta, \min_{b \in \mathcal{B}} b)$ and $-\phi_t$ must lead the agent to choose future bonds on the grid that are between the lowest bond holdings without being constrained $\mathbf{b_{t+1}}[\mathbf{b_{t+1}} \geq -\phi_{t+1}]$ and the borrowing constraint tomorrow $-\phi_{t+1}$. By creating these new grids we allow policy functions to be defined over that range of current

[13] I needed to add this condition to make this step stable. Sometimes one condition was satisfied but not the other, and the creation of new grids didn't make sense.

bond holdings which are between $-\phi_t \leq \mathbf{b_t}(\theta, \min_{b\in\mathcal{B}} b)$. This allows to solve more accurately for decision rules, which have more curvature in this region the state space closer to the borrowing constraint.

(h) Define the (augmented) consumption function on the bond grid $\mathcal{B}$ by interpolating the latter on the function mapping current (endogenous) bond holdings (the inverse of the bond accumulation policy) to consumption. Ensure that consumption is positive.

5. Compute the distribution Λ over $\Theta \times \mathcal{B}$ for all $t = T_1, ..., T_2$.

   (a) Start from the stationary distribution at the initial steady state Λ_{T_1}.

   (b) Given θ, interpolate the inverse of the bond accumulation policy on the marginal c.d.f. $\Lambda_t[\theta, :]$. Make sure that the interpolated values are between 0 and 1.

   (c) Update the distribution using the exogenous transition matrix $\Lambda_{t+1} = P\Lambda_t$.

   (d) Compute the implied households' aggregate bond demand, and the difference with the (exogenous) bond supply (which is zero at market clearing).

6. Update the interest rate sequence $\{r_t\}_{t=T_1,...,T_2}$ until the bond market clears in the last step.

A.2 Robustness test

This section discusses the robustness of the Krusell-Smith algorithms, considering first the direct price forecasts solution. It shows that the model with stochastic debt limits converges to the steady state of the liquidity trap economy before the credit crisis, when the transition matrix governing credit shocks is parametrized so that the economy is forever in the good credit regime, $\phi = 1.04$. The simulation is for $N = 10,000$ agents and $T = 1,500$ periods. When starting from an initial guess for $r = 2\%$ annually, it converges to bond market clearing and $r = 2.5\%$, which is the steady value of the interest rate in the steady state with easy credit of the liquidity trap economy. This is an important validity check that the algorithm passes.

Figure 1 gives an example of convergence of the algorithm to the steady state of the economy with a permanent regime of easy credit, starting from an initial guess of 2%, which is lower than what we know (from steady state computations) is the equilibrium interest rate in this economy, 2.5%. The interest rate converges to the steady state interest rate of 2.5% annually. In 30 iterations, the price converges to the 10^{-8} digit, the "average bond market clearing distance" is 10^{-7}.

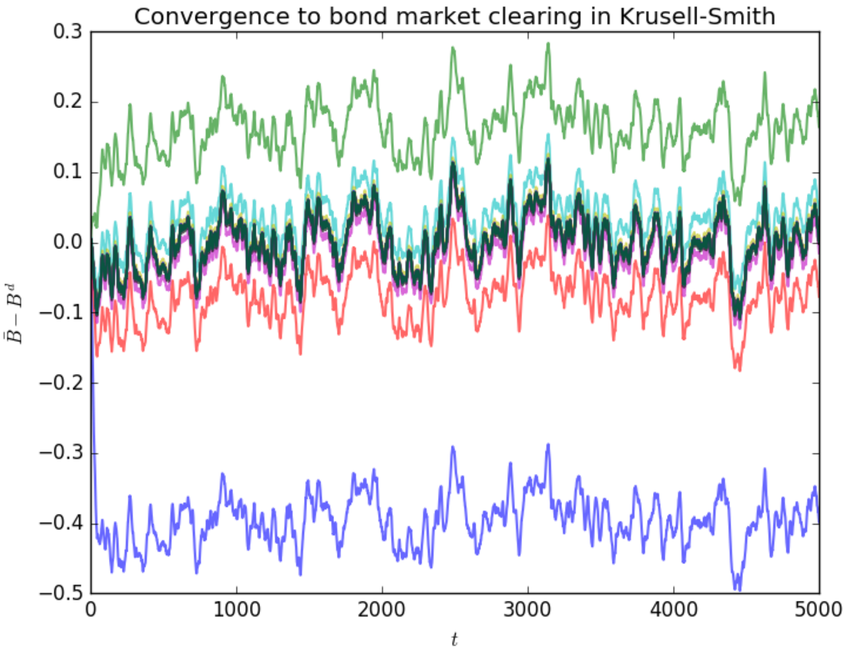

Figure 1: Example of successive iterations to illustrate convergence to bond market clearing. After 30 iterations the interest rate converges to 2.52% annually.

We can make the following observation. The longer the stochastic simulation, the smaller the market clearing distance, because cross-sectional sampling noise decreases with the simulation size. This is a classical feature of Bewley-Huggett-Aiyagari models with aggregate uncertainty, documented by Algan et al. (2014)). This paper tries to strike a balance between precision and manageable computation times, and therefore simulates below an economy $N = 20,000$ agents and $T = 5,000$ periods. At this stage the algorithm does not deliver exact market clearing at every date, but comes close. Increasing the number of period also matters, because the price updating rule depends on the weight of each credit regime in the simulation. The longer the number of periods, the more representative of the aggregate risk the sample is, and the more precisely prices can be updated.

In the solution with indirect price forecasts, a classical measure of accuracy of computations is to look at R^2 in the forecasting linear regressions. The R^2 of the forecasting equation for the variance of the bond distribution is greater than 99% in all subsamples of the simulation, whatever the credit regime in effect. The R^2 of the forecasting equation for the interest rate is somewhat lower if the variance is the only moment of the bond distribution included in the regressions. In this case it ranges from 65% in tight credit regimes to 57% in tight credit regimes followed by easy credit. Once the skewness of the bond distribution is included it becomes close to 1 in all subsamples of the simulations.

A.3 Tables and Figures

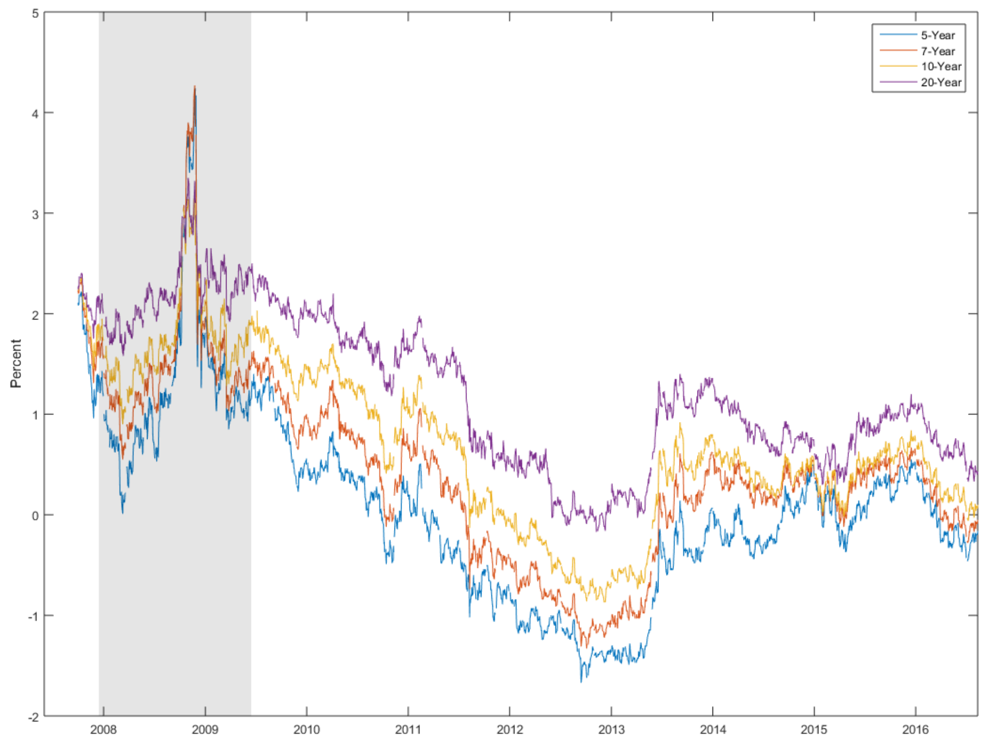

Figure 2: 5- to 20-Year Treasury Inflation-Indexed Security yields, Constant Maturity. Daily frequency. (Source: Board of Governors of the Federal Reserve System)

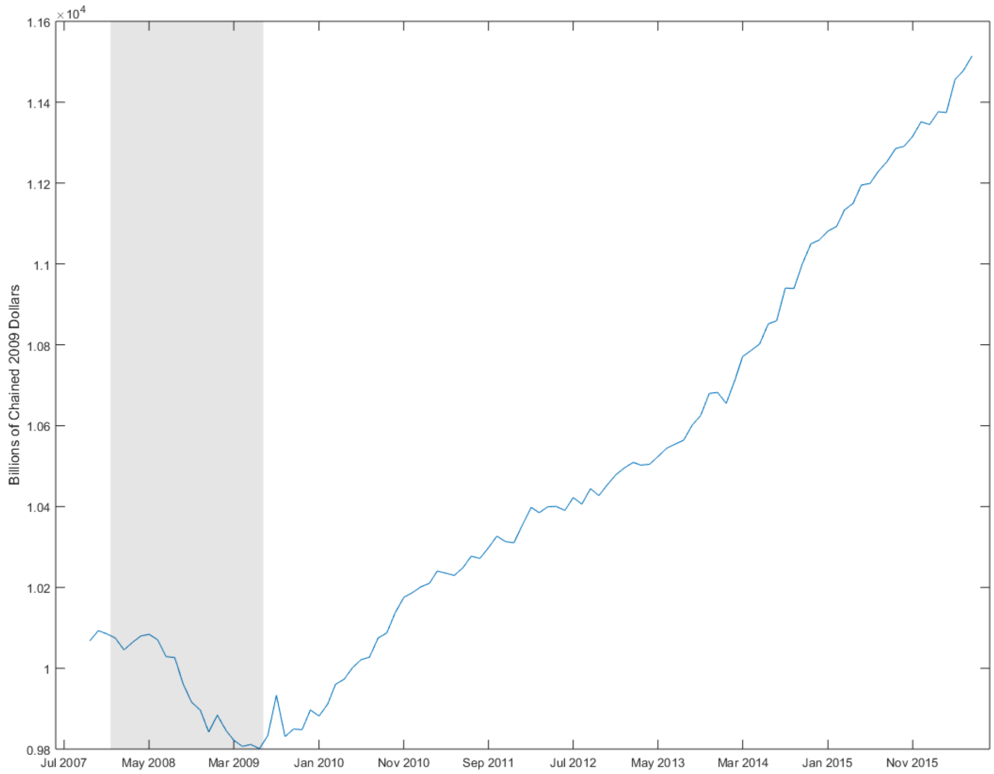

Figure 3: Real Personal Consumption Expenditures, in billions of chained 2009 dollars. Monthly Frequency, seasonally adjusted annual rate. (Source: US Bureau of Economic Analysis)

Parameter	Explanation	Value	Target/source
β	Discount factor	0.9713	Interest rate $r = 2.5\%$
γ	Coefficient of relative risk aversion	4	
η	Curvature of utility from leisure	1.88	Average Frisch elasticity $= 1$
ψ	Coefficient on leisure in utility	12.48	Average hours worked $= 0.4$ (Nekarda and Ramey (2013))
ρ	Persistence of productivity shock	0.967	Persistence of wage process in Floden and Linde (2001)
σ_ϵ	Variance of productivity shock	0.017	Variance of wage process in Floden and Linde (2001)
$\pi_{e,u}$	Transition to unemployment	0.057	Shimer (2005)
$\pi_{u,e}$	Transition to employment	0.882	Shimer (2005)
ν	Unemployment benefit	0.10	40% of average labor income
ϕ_h	High borrowing limit	1.75	Debt-GDP ratio of 0.18
ϕ_l	Low borrowing limit	1.3	5% decrease debt-GDP ratio
$\Pi_\phi(h,l)$	Transition to low credit regime	0.02	
$\Pi_\phi(l,l)$	Persistence low credit regime	0.05	
B	Bond supply	1.60	Liquid assets-GDP ratio of 1.78

Table 2: Parameters, quarterly frequency. The quantities ν, ϕ, B are expressed in terms of yearly aggregate output.

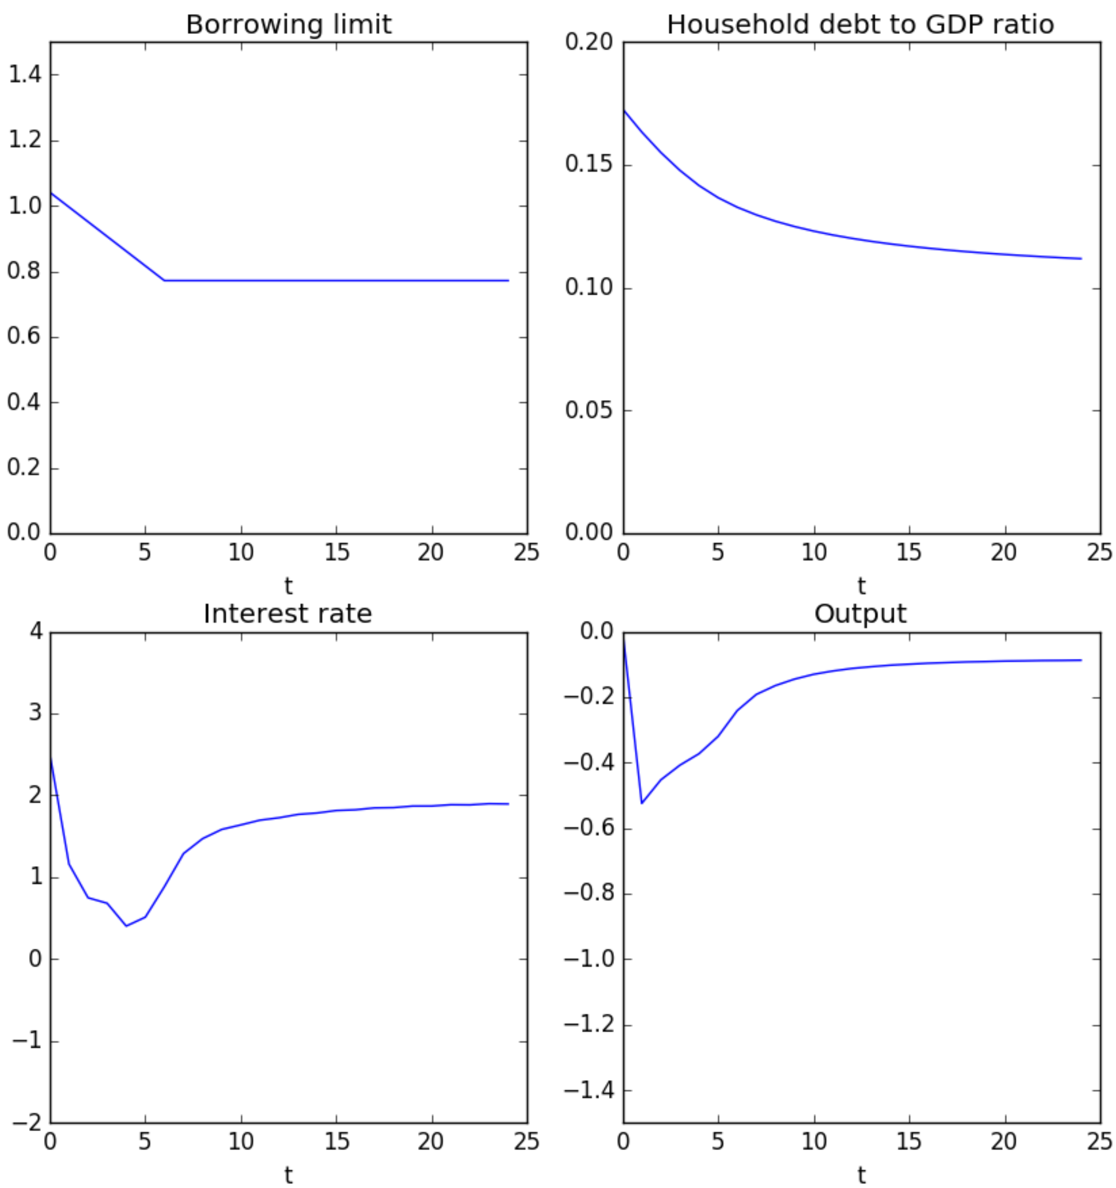

Figure 4: Trajectory of the economy's main variables under TD and same debt limit ϕ_l as in the KS economy. Output is in percentage deviation from its pre-crisis level.

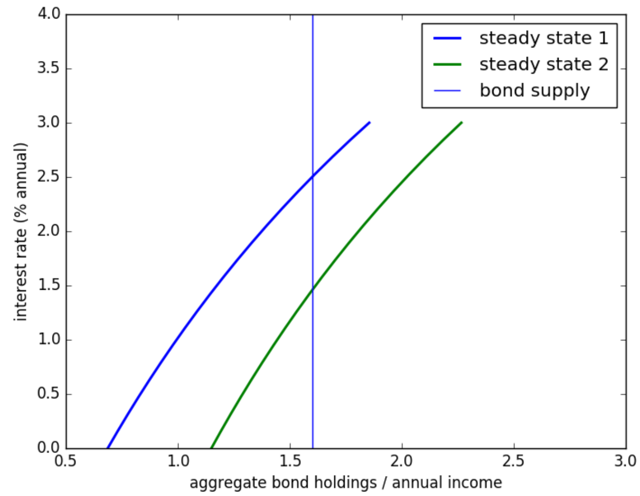

Figure 5: Bond market equilibrium in the steady states with high and low borrowing limits.

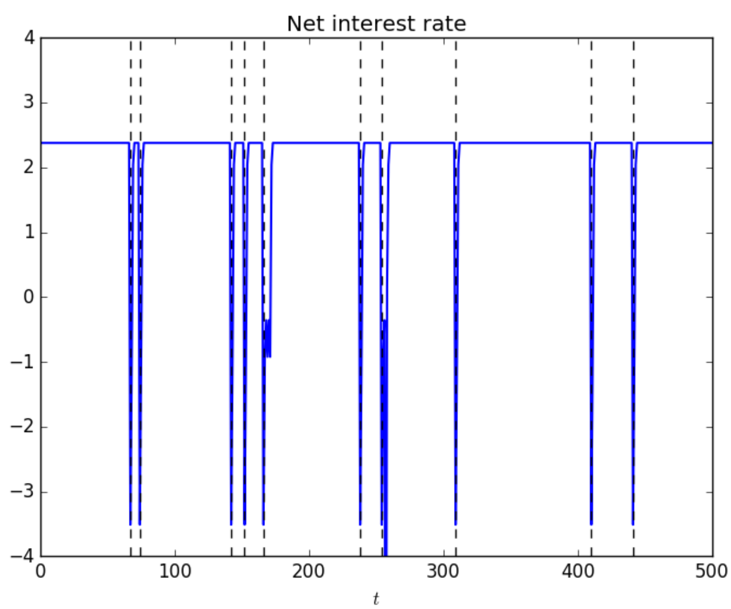

Figure 6: Trajectory of interest rate for 500 periods in the simulation. The decrease beyond graph axis limits between $t = 200$ and $t = 300$ corresponds to two successive tight credit, where $r = -8.96$.

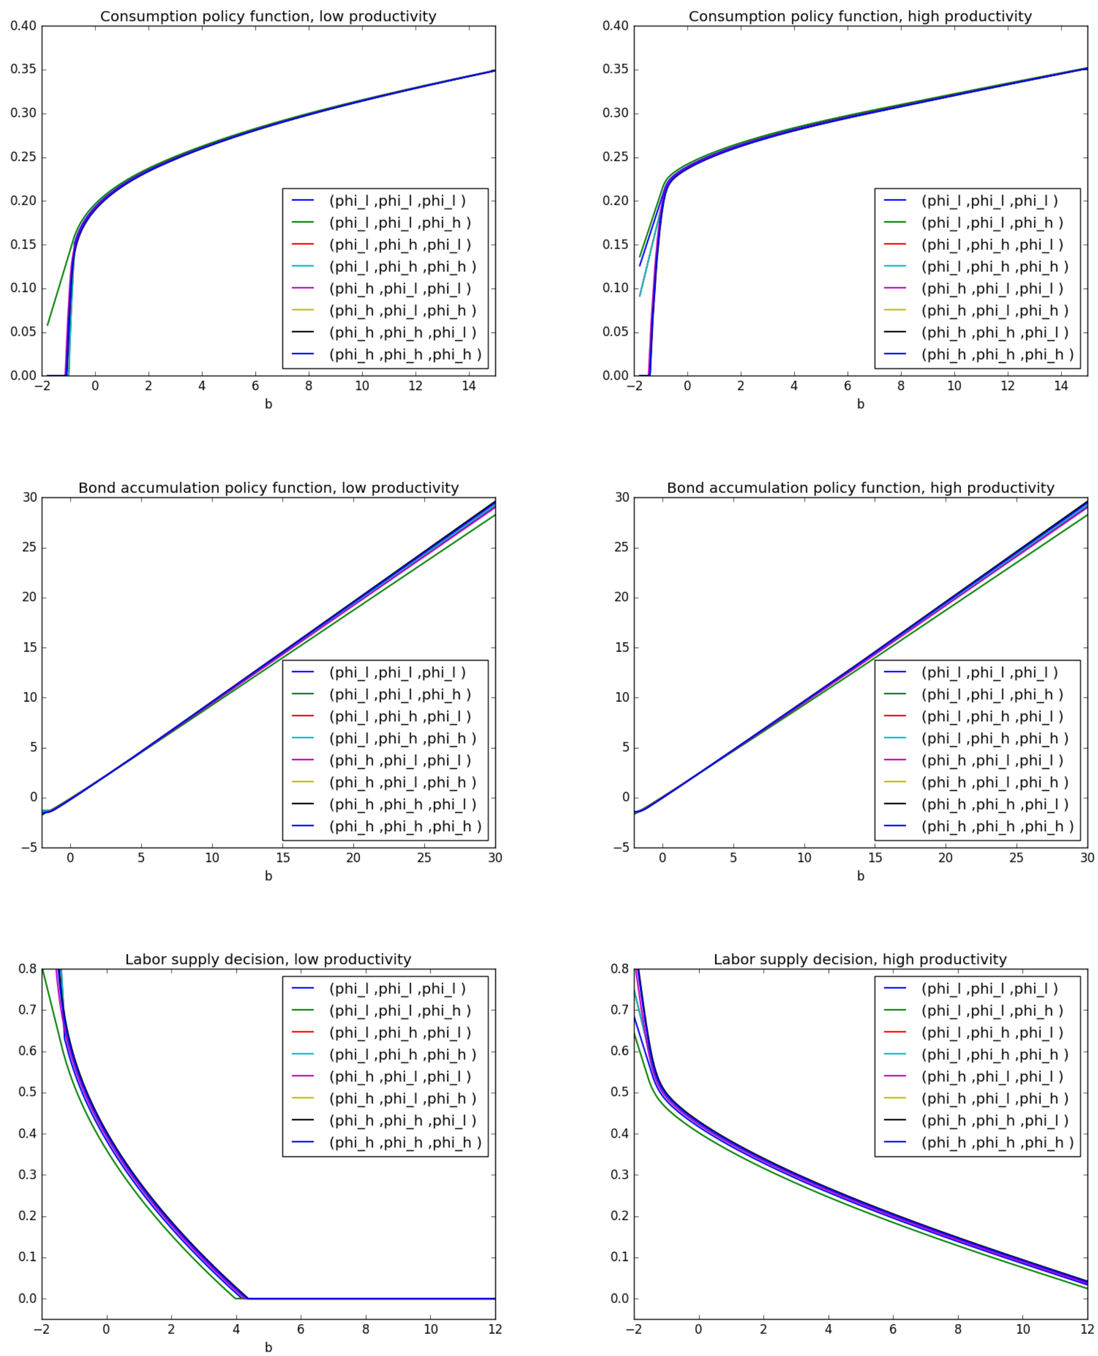

Figure 7: Decision rules for consumption, bond holdings and labor supply as function of households' asset position and credit regimes. Consumption is normalized by annual GDP in good times.

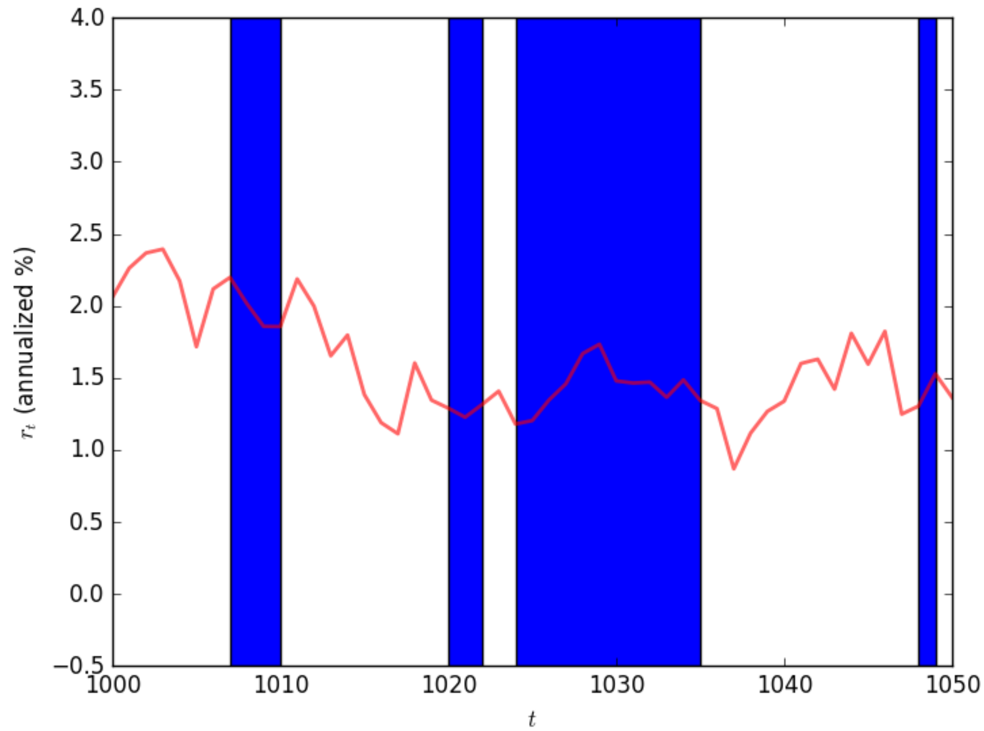

Figure 8: Example of a simulation path for the annualized real interest (in percentage terms). Episodes of tightening of households debt limit are in blue.