

Credit Crises, Precautionary Savings and the Liquidity Trap*

Veronica Guerrieri
University of Chicago and NBER

Guido Lorenzoni
MIT and NBER

First draft: July 2010
This draft: March 2011
Preliminary

Abstract

We use a model *à la* Bewly-Huggett-Ayagari to explore the effects of a credit crunch on consumer spending. Households borrow and lend to smooth idiosyncratic income shocks facing an exogenous borrowing constraint. We look at the economy response after an unexpected permanent tightening of this constraint. The interest rate drops sharply in the short run and then adjusts to a lower steady state level. This is due to the fact that after the shock a large fraction of agents is far below their target holdings of precautionary savings and this generates a large temporary positive shock to net lending. We then look at the effects on output. Here two opposing forces are present, as households can deleverage in two ways: by consuming less and by working more. We show that under a reasonable parametrization the effect on consumer spending dominates and precautionary behavior generates a recession. If we add nominal rigidities two things happen: (i) supply-side responses are muted, and (ii) there is a lower bound on the interest rate adjustment. These two elements tend to amplify the recession caused by the credit tightening.

*Email addresses: vguerrie@chicagobooth.edu; glorenzo@mit.edu. We thank Adrien Auclert and Amir Kermani for excellent research assistance. We also thank for helpful comments Andy Atkeson, VV Chari, John Leahy, Juan Pablo Nicolini, Thomas Philippon, Valery Ramey, Richard Rogerson, Gianluca Violante, Pierre Yared, Mike Woodford and seminar participants at the Minneapolis Fed, the Milton Friedman Institute Conference, the SED annual conference in Montreal, Columbia, Stern Finance, UCLA, UCSD, the Richmond Fed, the AEA meetings, ASU, MSU, Rice, Princeton. Guerrieri thanks the Federal Reserve Bank of Minneapolis for its generous hospitality during part of this research.

1 Introduction

How does an economy adjust from a regime of easy credit to one of tight credit? Suppose it is relatively easy for consumers and firms to borrow and the economy is in some stationary state with a given distribution of net lending positions. An unexpected shock hits—say a shock to the financial system—and borrowing gets harder, say in terms of tighter borrowing limits and/or in terms of higher credit spreads. Now the consumers and firms with the largest debtor positions need to readjust towards lower levels of debt (deleveraging). Since the debtor position of one agent is the creditor position of another, this also means that lenders will have to reduce their holdings of financial claims. How are the spending and production decisions of borrowers and lenders affected by this economy-wide financial adjustment? What happens to aggregate activity? How long does the adjustment last?

In this paper, we address these questions in the context of a workhorse Bewley model. Households borrow and lend to smooth transitory income fluctuations. The model captures two channels in the agents' response: a direct channel, by which constrained agents are forced to reduce their indebtedness, and a precautionary channel, by which unconstrained agents increase their savings as a buffer against future shocks, once they perceive a reduction in their potential borrowing capacity. Both channels increase the net supply of lending in the economy, so the equilibrium interest rate has to fall in equilibrium.

Our analysis leads to two sets of results. First, we look at the short run dynamics of the interest rate and show that they are characterized by a sharp initial fall followed by a gradual adjustment to a new, lower steady state. The reasons for the interest rate overshooting is that, at the initial asset distribution, the agents at the lower end of the distribution try to adjust faster towards their higher target level of net savings. So the initial increase in net lending is stronger. To maintain the asset market in equilibrium, interest rates have to fall sharply. As the asset distribution converges to the new steady state the net lending pressure subsides and the interest rate moves gradually up.

Second, we look at the response of aggregate activity. Overly indebted agents can deleverage in two ways: by spending less or by earning more. In the context of our model, this means that the shock leads both to a reduction in consumer spending and to an increase in labor supply. Whether a recession follows depends on the relative strength of these two forces. In particular, if the consumer's precautionary motive is strong enough, the reduction

in consumer spending dominates, and output declines. As for the case of interest rates, the contraction is stronger in the short run, when the distribution of asset holdings is far from its new steady state and some agents are far below their new savings target.

We then enrich the model introducing durable consumption, which was the most responsive component of consumption in the recent crisis. In this case, households have an interesting portfolio choice: they can invest in liquid bonds or in durable goods which are costly to resell. In this model, a larger fraction of the population is typically liquidity constrained, because households who start from a low level of financial assets and receive positive income shocks respond by increasing their indebtedness, to purchase durable goods, and hence borrow more. After a credit crunch, the households at the constraint are forced to deleverage and have to reduce consumption of both durable and non-durable goods. However, the precautionary motive that affects all the unconstrained households generates a trade-off. The larger desire for precautionary savings induce households to save more, but they tend to allocate their increased savings away from durable consumption goods and towards more liquid assets. Overall, with our calibration, this leads to a drop in durable consumption and a to a larger drop in aggregate consumer spending.

Our results on interest rate dynamics link our analysis to the idea of the liquidity trap. A liquidity trap is a situation where the economy is in a recession and the nominal interest rate is zero. In this situation, the central bank cannot lower the nominal interest rate to boost private spending as it would in normal times. The monetary policy literature—recently Krugman (1998) and Woodford and Eggertsson (2001)—has pointed out that the basic problem in a liquidity trap is that the real interest rate required to achieve full employment, the “natural” interest rate, is unusually low and possibly negative. If inflation is low, in line with the central bank target, or, even worse, if deflation has taken hold, the real interest rate corresponding to a zero nominal rate is higher than the natural rate and private spending is stuck at an inefficiently low level. In the context of a simple representative agent models it is not easy to identify shocks that push the economy in a liquidity trap and the literature has mostly resorted to introducing ad-hoc shocks to intertemporal preferences, which mechanically increase the consumer’s willingness to save. Our analysis shows that shocks to the agents’ borrowing capacity are precisely the type of shocks that can push down the “natural” rate by increasing the net demand for savings in the short run, and thus trigger a liquidity trap. Historically, liquidity traps have typically

arisen following disruptions in the banking system, the most notable examples being the Great Depression, Japan in the 90s, and the current crisis. Our paper shows a natural connection between credit market shocks and the emergence of a liquidity trap.

Our paper is related to different strands of literature. First, there is the vast literature on savings in incomplete-markets economies with idiosyncratic income uncertainty, following the seminal work of Bewley (1977), Huggett (1993), and Aiyagari (1994).¹ Our paper is particularly related to some recent contributions that focus on transitional dynamics after different types of shocks. For example, Mendoza, Rios Rull and Quadrini (2010) look at the response of an economy opening up to international asset trade. Our treatment of durables is related to Carroll and Dunn (1997), which is the first paper to incorporate durables in a model of precautionary savings.

Two papers that explore the effects of precautionary behavior on business cycle fluctuations are Guerrieri and Lorenzoni (2009) and Ragot and Challe (2010). Both papers, derive analytical results under simplifying assumptions that essentially eliminate the wealth distribution from the state variables of the problem. In this paper we take a computational approach, to get a sense of how the adjustment mechanism works when the wealth distribution evolves endogenously. Another related paper is Chamley (2010), a theoretical paper which explores the role of precautionary motive in a monetary environment and focuses on the possibility of multiple equilibria.

The paper is also related to the growing literature that analyzes the real effects of a credit crunch in dynamic general equilibrium models, including Curdia and Woodford (2009), Hall (2009), Jermann and Quadrini (2009), Brunnermeier and Sannikov (2010), Gertler and Karadi (2010), Gertler and Kiyotaki (2010), Del Negro, Eggertsson, Ferrero, and Kiyotaki (2010). Mostly these papers focus on the effects of a credit tightening on firms' investment, rather than on household consumption as we do here. Notable exceptions are Hall (2009), who looks at both consumption and investment, and Midrigan and Philippon (2011) who focus on cross-sectional implications after a drop in home equity in a cash-in-advance model. In independent work, Eggertsson and Krugman (2011) also look at a shock to the borrowing limit as a source of a recession cum liquidity trap. The main difference with our paper, is that they strive for analytic tractability focusing on a model with one borrower and one lender, where borrowing and lending are driven by a binary

¹Heathcote, Storesletten, and Violante (2009) provide an exhaustive review.

preference shocks, while in our model borrowing and lending are driven by idiosyncratic income risk and the purchase of durables. A distinctive feature of our paper is the focus on the dynamics of the distribution of net lending positions. Most of the literature, for reasons of tractability, departs minimally from a representative agent environment, assuming that there are only one borrower and one lender and making assumptions that avoid dealing with the wealth distribution as a state variable. Here instead we are interested in tracking the distribution of net lending positions over time. The slow adjustment of this distribution is behind the long-lasting effects of a credit shock in our model. Finally, our analysis brings to attention the role of labor supply and of durables in financial adjustment.

Finally, there is a growing number of papers that focus on the dynamics of entrepreneurial wealth (Cagetti and De Nardi, 2006, Buera and Shin, 2007). Two recent papers that look at the response of the entrepreneurial sector to a credit shock are Goldberg (2010) and Khan and Thomas (2010). In particular, Goldberg (2010) shares with our paper the emphasis on precautionary behavior and on the scarcity of liquid assets, but focusing on its effects on entrepreneurs' decisions.

The rest of the paper is organized as follows. In Section 2, we introduce the environment and define an equilibrium. We also describe our main calibration exercise and characterize the steady state. In Section 3, we perform our main exercise, that is, we analyze the equilibrium transitional dynamics after an unexpected permanent tightening of the borrowing limit, or credit crunch. Section 4 studies the effects of simple fiscal policies. Section 5 explores a variant of the model with nominal rigidities where the Central Bank sets the interest rate path and studies the effects of a credit crunch under alternative monetary policies. In Section 6, we introduce durables consumption and characterize the steady state and dynamics in that case. Section 7 concludes. The appendix explains the computational strategy.

2 Model

Consider an economy populated by households hit by idiosyncratic income shocks, who smooth income risk by borrowing and lending. The model is a version of a standard Bewley model with endogenous labor supply and no capital. The only asset traded is a one-period risk-free bond. Households are allowed to hold negative amounts of bonds—i.e.,

to borrow—up to an exogenous limit. We will first analyze the steady state equilibrium of this economy for a given borrowing limit. Then we will study the economy transitional dynamics following an unexpected shock that permanently reduces this limit.

There is a continuum of infinitely lived households with preferences represented by the utility function

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

where c_t is consumption and n_t is labor effort. Each household produces consumption goods using the linear technology

$$y_t = \theta_t n_t,$$

where θ_t is an idiosyncratic productivity level which follows a Markov chain on the space $\{\theta^1, \dots, \theta^S\}$. Let $\theta^1 = 0$, so that agents hit by the worst shock are unemployed. For the moment, there are no aggregate shocks.

The household budget constraint is

$$q_t b_{t+1} + c_t + \tilde{\tau}_t \leq b_t + y_t,$$

where b_t are the household bond holdings, q_t is the bond price and $\tilde{\tau}_t$ is a tax. The budget constraint requires that the households's current resources—bonds plus current income—cover consumption, tax payments, and the purchase of new bonds. To allow for unemployment benefits, we assume that $\tilde{\tau}_t = \tau_t - z_t$ if $\theta_{j,t} = 0$ and $\tilde{\tau}_t = \tau_t$ otherwise. That is, all households pay the lump sum tax τ_t and the unemployed receive the unemployment benefit z_t .² The household's debt position is bounded below by the exogenous limit $\phi \geq 0$, that is, bond holdings have to satisfy

$$b_{t+1} \geq -\phi. \tag{1}$$

The interest rate implicit in the bond price is $r_t = 1/q_t - 1$.

The government chooses the aggregate supply of real bonds B_t and the unemployment benefit z_t for all t and sets the lump sum tax τ_t to satisfy the budget constraint:

$$B_t + uz_t = q_t B_{t+1} + \tau_t,$$

²The presence of the unemployment benefit implies that the natural borrowing limit is strictly positive.

where $u = \Pr(\theta_{j,t} = 0)$ is the fraction of unemployed agents in the population. For the moment, we assume that the supply of government bonds and the unemployment benefit are constant at some levels \bar{B} and \bar{z} , while the tax τ_t adjusts to ensure government budget balance. In Section 4, we will consider the effects of other fiscal policies.

The main deviation from Aiyagari (1994) and much of the following literature is that our model does not feature capital accumulation. The standard assumption in models with capital is that claims to physical capital are a perfect substitute for government bonds and other safe and liquid stores of value. Clearly, this would not be a satisfactory assumption here, given that we are trying to capture episodes of financial turmoil and flight to liquidity. To introduce capital in our model requires introducing imperfect substitutability between different assets, possibly due to different risk profiles and/or different resaleability. In section 6, we extend the model by allowing consumers to accumulate durable consumption. In that context, we discuss further the role of the assumptions made on the resaleability of different assets.

Finally, a stark simplification is that there is a single interest rate on bonds r_t , which applies both to positive and negative bond holdings. In other words, household' liabilities and government bonds are perfect substitutes, or, equivalently, there is a perfect intermediation sector which transforms the liabilities of the households in debt into the deposits of the households in credit.

2.1 Equilibrium

Given a sequence of interest rates $\{r_t\}$, let $C_t(b, \theta)$ and $N_t(b, \theta)$ denote the optimal decisions for consumption and labor supply at time t for a household with bond holdings $b_t = b$ and current productivity $\theta_t = \theta$. Notice that, given consumption and labor supply, next period bond holdings are given by the budget constraint. Therefore, the transition for bond holdings is fully determined by the functions $C_t(b, \theta)$ and $N_t(b, \theta)$.

Let $\Psi_t(b, \theta)$ denote the joint distribution of bond holdings and current productivity levels in the population. The household's optimal transition for bond holdings together with the Markov process for productivity yield a transition probability for the individual states (b, θ) . This transition probability can be used to compute the distribution $\Psi_{t+1}(b, \theta)$ given the distribution $\Psi_t(b, \theta)$. We are now ready to define an equilibrium.

Definition 1 *An equilibrium is a sequence of interest rates $\{r_t\}$, a sequence of decision rules for consumption and labor supply $\{C_t(b, \theta), N_t(b, \theta)\}$, a sequence of tax rates $\{\tau_t\}$, and a sequence of joint distributions for bond holdings and productivity $\{\Psi_t(b, \theta)\}$ such that, given the initial distribution $\Psi_0(b, \theta)$:*

(i) $C_t(b, \theta)$ and $N_t(b, \theta)$ are optimal given $\{r_t\}$ for all t ;

(ii) $\Psi_t(b, \theta)$ is consistent with consumption and labor supply policies for all t ;

(iii) the tax τ_t satisfies

$$\tau_t = u\bar{z} + \frac{r_t}{1+r_t}\bar{B};$$

(iv) the bonds market clears:

$$\int \int bd\Psi_t(b, \theta) = \bar{B}.$$

The optimal policies for consumption and labor supply are characterized by two optimality conditions. First, the Euler equation

$$U_c(c_t, n_t) \geq \beta(1+r_t)E_t[U_c(c_{t+1}, n_{t+1})], \quad (2)$$

which must hold with equality if the borrowing constraint $b_{t+1} \geq -\phi$ is not binding. Second, the optimality condition for labor supply

$$\theta_t U_c(c_t, n_t) + U_n(c_t, n_t) = 0, \quad (3)$$

for all households with $\theta_t > 0$.

A key observation is that, when we lower the borrowing limit, agents face more uncertainty in future consumption, as consumption becomes more responsive to income shocks. With prudence in preferences, this implies that for a given average level of consumption tomorrow, the expected marginal utility on the right hand side of the Euler equation will be higher, by Jensen's inequality. This means that for a given level of interest rates, consumption today will fall, as if there was a negative preference shock reducing the marginal utility of consumption today. This will be the core mechanism reducing consumption demand after a credit shock. In this sense, our model with precautionary savings provides a microfoundation for models with preference shocks.

2.2 Calibration

We analyze the model by numerical simulations, therefore we need to specify preferences and calibrate the model parameters. We calibrate the model to capture households' accumulation and decumulation of liquid assets in response to income and employment shocks. Notice that we are abstracting from life-cycle considerations and from many important drivers of individual wealth dynamics, like purchases of durable goods and housing (which we will bring back into the picture in Section 6), health expenses, educational expenses and so on. However, our computational exercise will allow us to identify some general qualitative features of the response of an economy in which borrowers and lenders gradually adjust to a shock to credit access. We believe these qualitative implications survive in environments with richer motives for borrowing and lending.

The utility function is:

$$U(c, n) = \frac{c^{1-\gamma}}{1-\gamma} + \psi \frac{(1-n)^{1-\eta}}{1-\eta}.$$

We will discuss shortly the advantages of a non isoelastic specification for the disutility of labor. The parameters are reported in Table 1. The discount factor β corresponds to a discount factor of 0.98 at a yearly frequency. This value, higher than in most calibrations, is chosen to replicate the low-interest rate environment of the mid 2000. In our baseline, we choose a relatively high coefficient of risk aversion $\gamma = 4$. Clearly, this coefficient is crucial in determining the consumers' precautionary behavior, so we will experiment with different values as well. The parameters η and ψ are chosen so that the hours worked of employed workers are on average 40% of their time endowment, normalized to 1 (following Nekarda and Ramey, 2010), and so that the average Frisch elasticity of labor supply is 2. This high elasticity will be crucial in generating sizeable output responses in the baseline model without nominal rigidities.

β	0.995	ρ	0.974	\bar{B}	1.4
γ	4	σ_ε^2	0.025	ϕ	5.1
η	0.75	z	0.39		
ψ	1.26				

Table 1. Baseline calibration

For the wage process we assume that the log of θ follows a AR1 process with autocorrelation $\rho = 0.974$ and variance $\sigma_\varepsilon^2 = 0.025$. The wage process is approximated by a 5-state Markov chain, following the approach in Tauchen (1991).

For the transitions between employment and unemployment we follow Shimer (2005), who estimates the finding rate and the separation rate from CPS data. At a quarterly frequency, we set the transition from employment to unemployment at 0.1, corresponding to jobs that lasts on average 2.5 years, and the transition from unemployment to employment at 0.833. We assume that when re-employed workers start at the lowest positive value of θ . For the unemployment benefit z , we also follow Shimer (2005) and calibrate it so that it corresponds to 40% of average labor income.

Finally, we need to choose values for the supply of government bonds \bar{B} and for the borrowing limit ϕ . We choose these values to capture US households' balance sheets in 2006, just prior to the recent financial crisis, looking at the Federal Reserve Board flow of funds data. First, we look at the households' holdings of liquid assets, namely the sum of their holdings of all deposits plus treasury securities (line 9 plus line 16), which was 7.1 trillion dollars, or 53% of GDP. We choose \bar{B} to match this ratio. Second, we interpret debt in our model as consumer credit (line 34 which corresponds, essentially, to total household liabilities minus mortgage debt), which was 2.4 trillion dollars, or 18% of GDP.

2.3 Steady State

We first compute the initial steady state decision rules and the initial bond distribution. Figure 1 shows the optimal steady state values of consumption and labor supply for each level of bond holdings. For ease of reading, we plot the policies for only two values of θ , θ^2 and θ^6 .

Different responses at different levels of bond holdings are apparent. At high levels of bond holdings, consumers behavior is close to the permanent-income hypothesis and the consumption function is increasing almost linearly in b . For lower levels of bond holdings, the consumption function becomes concave. Carroll and Kimball (1996) show that this is a typical feature of the consumption function in this class of models. The optimality condition for labor supply implies that the behavior of labor supply mirrors that of consumption, capturing an income effect. In particular, a steeply increasing consumption function at low

levels of b translates into a steeply decreasing labor supply function. As a consequence, the labor supply function is convex in b . Notice that for high levels of b , the substitution effect plays a larger role and labor supply is very responsive to the wage rate (the ratio θ^6/θ^2 is 1.3). For lower levels, labor supply is much less responsive to the wage rate.³

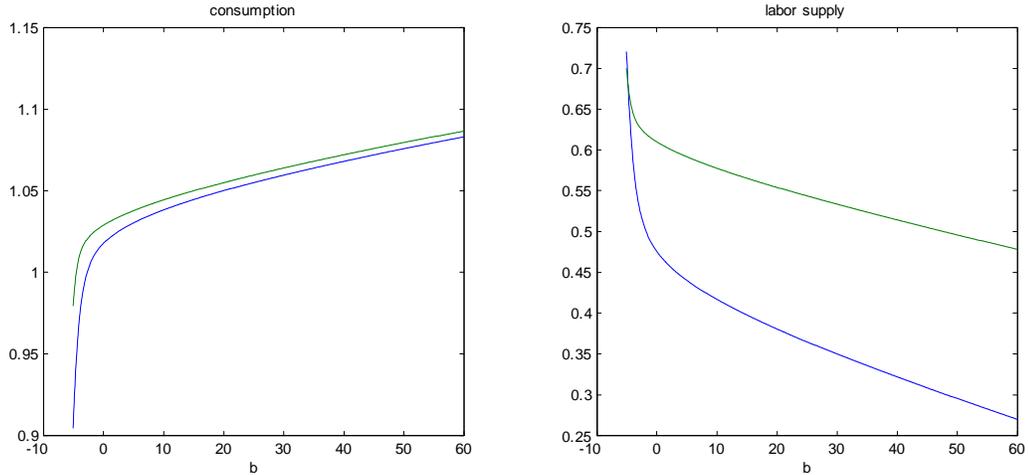


Figure 1: Optimal consumption and labor supply at the initial steady state (for $\theta = \theta^2$ and $\theta = \theta^6$).

3 Credit Crunch

We now explore the response of our economy to a credit crunch. We consider an economy that at time -1 is in steady state, with a borrowing limit equal to ϕ' and a stationary wealth distribution Ψ' . At time 0 , the economy is hit by an unexpected shock leading to a decrease in the borrowing limit to ϕ'' . In particular, we look at the effects of a shock that halves the debt limit from $\phi' = 5.1$ to $\phi'' = 2.55$. As the initial wealth distribution is $\Psi_0 = \Psi'$, which is different from the new steady state distribution Ψ'' , the economy goes through a gradual transition towards the new steady state.

Before looking at the transitional dynamics, let us briefly compare the two steady states. Figure 2 shows how the interest rate is determined in the two steady states. The

³For very low levels of b , labor supply is decreasing in the wage rate. This reflects a strong income effect, associated with a highly persistent wage process.

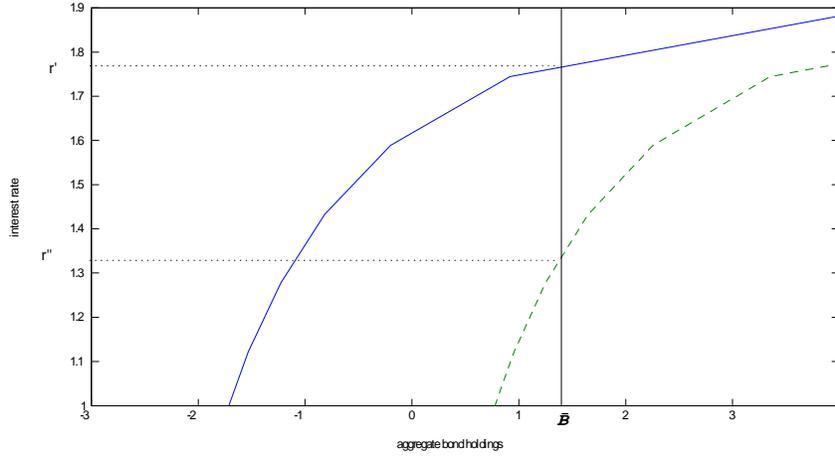


Figure 2: Steady state interest rate determination

solid line shows the average demand for bond holdings in the initial steady state, which is an increasing function of the interest rate, as it is common in Beweley models. The dashed line shows the average demand for bond holdings in the new steady state. The new demand curve is to the right of the old one due to two effects. First, there is a mechanical effect, all households with bond holdings below $-\phi'$ now need to hold at least $-\phi'' > -\phi'$. Second, there is a precautionary effect: for a given interest rate, households accumulate more wealth to stay away from the borrowing limit. Given that the supply of bonds is fixed at \bar{B} , it follows that the new interest rate r'' is lower than r' to convince the households to demand the same quantity of bonds.

3.1 Interest rate dynamics

To study the transitional dynamics, we assume that the borrowing limit ϕ_t adjusts gradually towards its new level along the linear adjustment path

$$\phi_t = \max \{ \phi'', \phi' - \Delta t \}.$$

The reason for this assumption is to ensure that agents at all initial levels of debt can adjust without being forced into default. Since all debt in the model has a one-quarter maturity, a sudden adjustment in the debt limit would make it impossible for many borrowers to roll over their debt. An assumption of gradual adjustment of the debt limit is a simple way

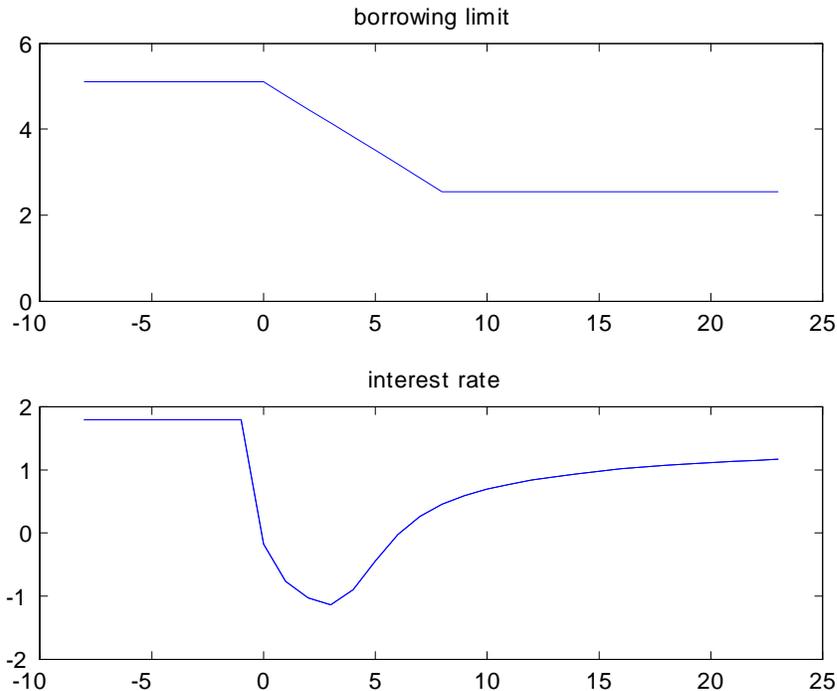


Figure 3: Interest rate dynamics

of capturing the fact that with longer debt maturities agents have some time to adjust to the new regime. In particular, we choose Δ so that the unemployment benefit is sufficient to cover the minimum debt repayment $-b_t - q_t\phi_{t+1}$ for an agent starting at $b_t = -\phi_t$. Given the model parameters and the size of the shock this gives us an adjustment lasting 8 quarters. Default and bankruptcy are clearly an important element of the adjustment to a tighter credit regime, but we abstract from them.

In the top panel of Figure 3, we plot the exogenous adjustment path for ϕ_t . In the bottom panel we plot the interest rate path. The interest rate drops dramatically after the shock, going negative for more than a year. This is our first main result and we now investigate the mechanism behind it to argue that it is fairly general result and not just the outcome of our choice of parameters.⁴

The first observation to explain the interest rate overshooting is that the bond distribu-

⁴We have also explored the robustness of this result numerically and it holds for all the parameter configurations we have tried.

tion converges gradually to its new steady state and that the new steady state distribution is more concentrated than the initial one. Let $F_t(b)$ denote the CDF of the marginal bond distribution, that is, $F_t(b) = \int \Psi_t(b, \theta) d\theta$. Let $F'(b)$ and $F''(b)$ denote the distributions, respectively, at the initial and at the final steady state. The middle panel of Figure 4 reports the densities f' (solid line) and f'' (dashed line), associated, respectively, to F' and F'' . The panel suggests that indeed the distribution in the new steady state is more concentrated. Since the bond supply is fixed at \bar{B} we know that the two distributions have the same mean. To check formally that F' is a mean-preserving spread of F'' , in the bottom panel of Figure 4 we plot the integral $\int_{-\infty}^b (1 - F(\tilde{b})) d\tilde{b}$ for the two distributions. The fact that integral for F'' is always above the one for F' confirms the visual impression from the middle panel. Why is the distribution in the new steady state more concentrated? Two forces are at work here. At low levels of bond holdings, the precautionary behavior induces agents in the new steady state to accumulate bonds faster. At high levels of bond holdings, the low equilibrium interest rate induces agents to decumulate bonds faster. This makes bond holdings to mean revert faster and makes the stationary distribution more concentrated.

Consider now the top panel of Figure 4. This panel plots the average bond accumulation $b_{j,t+1} - b_{j,t}$ (averaged over θ) as a function of the initial bond holdings $b_{j,t}$. The decision rules used for this plot are those that would arise at date 0 *if* the interest rate were to adjust immediately to its new steady state level and stay there at all following dates $t = 0, 1, 2, \dots$. Note that this function is not exactly convex, but almost so. The reason for this convexity is the same reason behind the concavity of the consumption function and the convexity of the labor supply functions in Figure 1.⁵

We are now ready to put the pieces together. Let us make a mental experiment and suppose the interest rate jumps immediately to its new steady state at date 0. If the wealth distribution was already at its new steady state level, the average bond accumulation would average to zero, as we would just be in the new steady state. In other words, if we integrate the function in the top panel weighted by the density f'' in the second panel we get zero. If instead we integrate the same function weighted by the density f' we get a positive

⁵The non-convexity at very low levels of b is due to the fact that at the new steady state, the labor supply for very low levels of b is very high for the low shocks and in that region it is less elastic (given our preferences).

number, because the function in the top panel is (approximately) convex and f' is a mean-preserving spread of f'' . This means that at the conjectured interest rate path households want to accumulate bonds on average. Since the bond supply is fixed this means that at the conjectured interest rate path there is excess demand of bonds. To equilibrate the bonds market we need a lower interest rate in the initial periods.

3.2 Output response

Next, we want to understand what happens to output. Figure 5 shows that output converges to a lower level in the new steady state and overshoots in the short run. The economy goes through a recession and then converges to a permanently lower level of output. The scale for output is percentage deviations from the initial steady state, so, in terms of magnitude, the recession generated in our baseline model is small, with less than a 1/2 of a percent reduction in output. We will consider below mechanisms that can magnify this response. But first let us understand the mechanism behind the recession.

The output response depends both on consumption and on labor supply decisions. Let us focus on the transitional dynamics and try to understand the overshooting in Figure 5. Again, let us make a mental experiment and suppose the interest rate jumped directly to r'' . As argued in the discussion of Figure 1, the consumption and the labor supply policies are, respectively, concave and convex functions of the household's bond holdings. Then, given that the initial distribution is more dispersed than the the new steady state distribution (in the sense of second-order stochastic dominance), average consumption demand is lower than at the new steady state and average output supply is higher. Therefore, at the price r'' there is excess supply in the goods market, which corresponds to the excess demand in the bonds market discussed above.

To clear the goods markets (and the bonds market) the interest rate must be lower on the transition path. As we lower the interest rate towards its equilibrium value, the goods market adjusts on both sides: consumption increases and labor supply falls, due to intertemporal substitution. Therefore, the market clearing output level can, in general, be either above or below its steady state level. Two sets of considerations determine which side of the market dominates the adjustment path: (i) how large are the negative shift in consumption demand and the positive shift in labor supply due to the larger dispersion

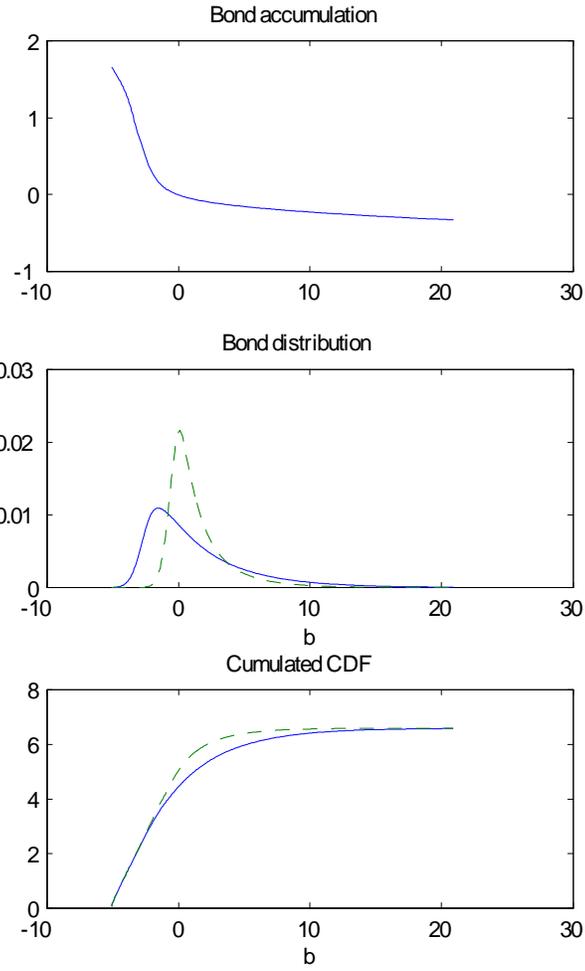


Figure 4: Explaining the overshooting: bond accumulation and bond distribution at the two steady states

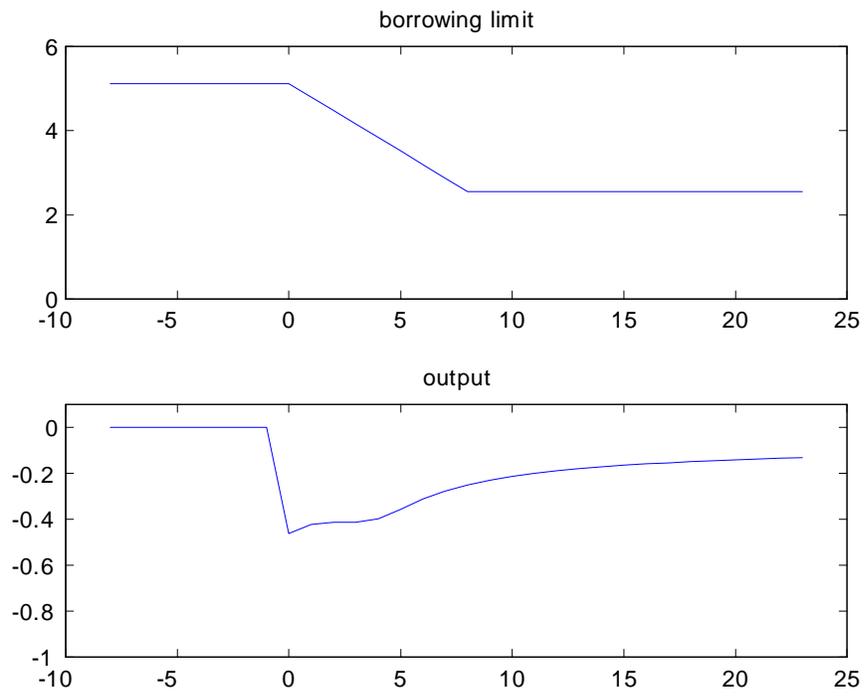


Figure 5: Output dynamics

of bond holdings at the beginning of the transition; and (ii) how elastic are consumption demand and labor supply to a reduction of the interest rate from r'' to the equilibrium level r_0 ? Our parameters imply that the fall in consumption demand is the dominating factor, and output falls below its new steady state value. Building on this discussion, we can now better understand the role of our parameters.

On the demand side, the effect of a decrease in consumption demand is higher when γ is higher. Notice that γ is both the coefficient of relative risk aversion and the inverse elasticity of intertemporal substitution. On the one hand, when households are more risk averse, the precautionary motive is stronger, making their consumption policy more concave. Therefore, the initial shift in consumption demand is stronger. On the other hand, when the elasticity of intertemporal substitution is lower, consumption responds less elastically to the interest rate. Both effects tend to make the recession larger. Figure 6 shows the behavior of interest rate and output for the same economy with $\gamma = 2$ (red lines) instead of $\gamma = 4$ (blue lines). According to the intuition, the precautionary motive is less strong, making the interest rate decrease less in the short run and the recession milder. However, the recession is longer because the agents are less prepared to a credit crunch and hence take longer to adjust their wealth accumulation.

On the supply side, instead, the elasticity of labor supply to the interest rate and its reaction to a shock in ϕ are determined by the parameter η , but in different directions. When η is lower, the labor supply is more elastic to the interest rate, weakening the increase in labor supply. However, at the same time, when η is lower the labor supply function becomes more convex, which implies that the reaction of average labor supply to a shock in ϕ is stronger. However, the choice of a not isoelastic preferences ensure that poor households are less sensitive to a decrease in wealth, making the labor supply policy less convex for any value of η . This ensures that on net, with our parameterization, output tend to overshoots in the short run. Figure 7 shows that when η increases (red lines) so that the Frish elasticity goes from 2 to 1, the short run decline in output is smaller. This confirms that the effect of the elasticity of labor supply to the interest rate dominates the decrease in convexity.

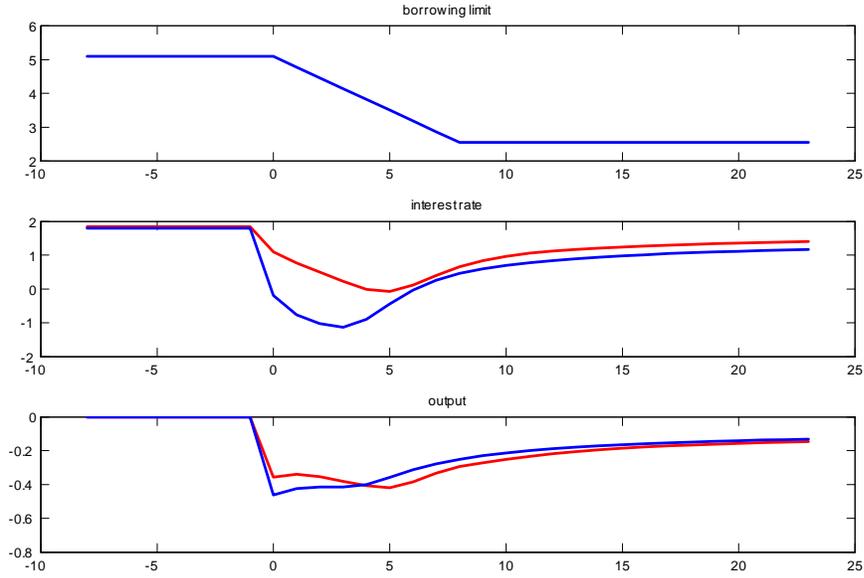


Figure 6: Changing the coefficient of relative risk aversion γ

4 Fiscal policy

We now explore the role of different government policies in mitigating the recession. In particular, let us consider changes in the supply of government bonds. Increasing the supply of bonds can be beneficial for two reasons. First, there is a direct increase in the supply of liquid assets that reduces the downward pressure on the real interest rate. Second, as the government increases bond supply, the associated deficit can be used to reduce taxes or increase transfers in the short run. Since Ricardian equivalence fails in our economy, this has a positive effect on spending.

Our model has the feature, common to many models with government supplied liquidity and lump sum taxation, that an increase in the supply of government bonds B can exactly offset a change in the borrowing limit ϕ . In particular, the only thing that matters for the equilibrium is the sum $B + \phi$. Here, however, we analyze the effects of policies that only partially offset the long run change in ϕ , possibly because of concerns with the distortionary effects of higher taxation in the long run.

Consider, in particular, a policy of increasing gradually the supply of real bonds to a

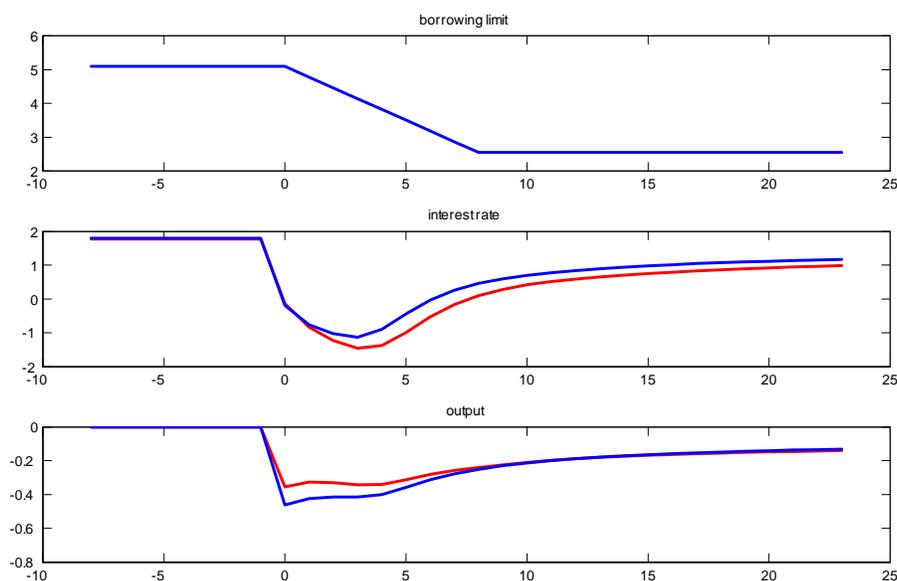


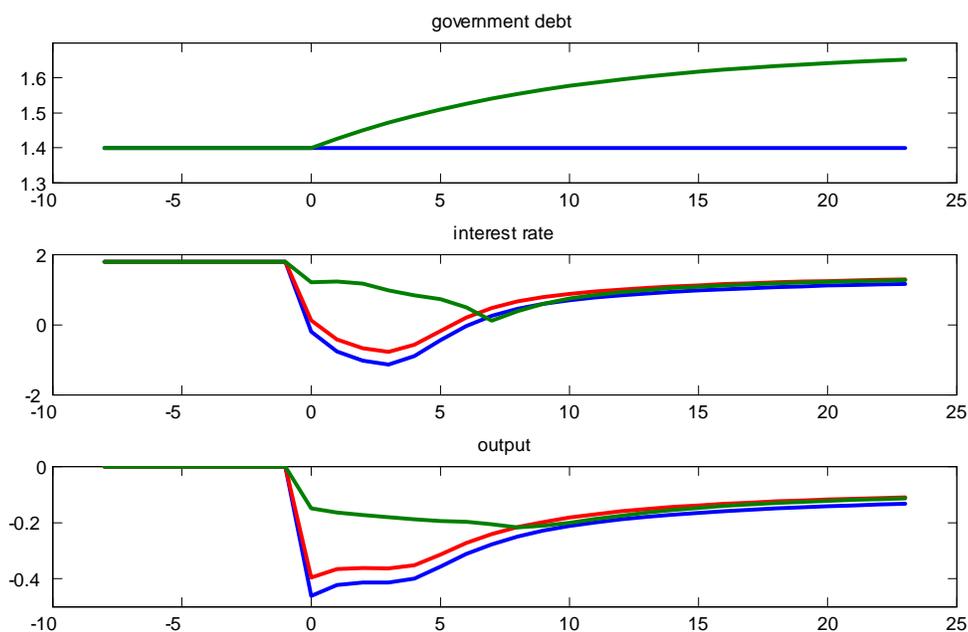
Figure 7: Changing the labor supply elasticity (η)

level that is 20% higher in the new steady state. Namely, assume that B_t follows the path

$$B_t = \rho_b^t B' + (1 - \rho_b^t) B'',$$

for some $\rho_b^t \in (0, 1)$. We then consider two different ways of spending the deficit associated to this increase in bond supply. First, we look at a policy where taxes adjust to balance the government budget in every period. Second, we look at a policy where the government deficit is used to finance a temporary increase in the unemployment benefit. In particular, we let the unemployment benefit to be 50% higher for the first two years after the shock. Figure ?? shows what happens to interest rate and output under these two policies. The red lines represent the policy where the increase in B finances a temporary reduction in the tax τ_t ; the green lines represent the policy where the deficit goes partly to finance an increase in unemployment insurance.

The figure shows that increasing the supply of government bonds help the economy to reduce the overshooting both in interest rate and in output. Moreover, what is particular effective in this economy is to combine this deficit increase with an increase in unemployment insurance. Increasing the unemployment benefit in the short run is more beneficial



than reducing the lump-sum tax because it is targeted to the fraction of the population who is more likely to be credit constrained.

5 Nominal Rigidities

So far we assumed fully flexible prices and that the real interest rate adjusts to its equilibrium path to equilibrate the demand and supply of bonds. In this section we explore what happens in a variant of the model with nominal rigidities. Adding nominal rigidities is a simple way to focus on the response of consumer spending, assuming that output in the short run is demand-determined. Once we move in this direction, we can then let the central bank choose the interest rate path and see how different choices lead to different adjustment paths for the distribution of bond holdings and for real output.

The households' side of the model is as before, but output is now produced by a continuum of monopolistically competitive firms. Each firm produces a good $j \in [0, 1]$ and consumption is a Dixit-Stiglitz aggregate of these goods. Namely, consumption of household

i is given by

$$c_{i,t} = \left(\int_0^1 c_{i,t}(j)^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}}$$

where $c_{i,t}(j)$ is consumption of good j . Each firm produces with a linear technology which produces one unit of good with one efficiency unit of labor. We interpret the shock $\theta_{i,t}$ as a shock to the efficiency of household j labor, so we have a common wage rate w_t per efficiency unit, while the wage rate per hour is given by $w_t\theta_{j,t}$.

The firms are owned by the consumers, so letting Π_t denote total profits, the budget constraint is now

$$q_t b_{j,t+1} + P_t c_{j,t} = b_{j,t} + W_t \theta_{j,t} n_{j,t} - \tau_{j,t} + \Pi_t.$$

Monopolist j faces the demand

$$y_{j,t} = \left(\frac{p_{j,t}}{P_t} \right)^{-\varepsilon} C_t,$$

where C_t is aggregate consumption in the economy.

If prices are flexible, the equilibrium is very similar to that of the perfectly competitive economy of the previous section. The only difference is that the real wage is

$$w_t = \frac{W_t}{P_t} = \frac{\varepsilon - 1}{\varepsilon}, \quad (4)$$

and that households receive profit income on top of labor income. Therefore, the response of the economy to the credit tightening are similar to the ones of the baseline model.

To analyze the case of nominal rigidities, we consider an extreme form of rigidity, in which prices are fully rigid, that is, $P_t = P_{t-1} = 1$. In combination with this extreme assumption, we assume that the central bank chooses a path for the nominal (and real) interest rate r_t which converges to the new steady state level r'' . This ensures that the private benefit from adjusting prices goes to zero in the long run.

To find an equilibrium we choose a path $\{r_t\}$ and look for a sequence of real wage rates $\{w_t\}$ and profits $\{\Pi_t\}$ such that given the optimal consumption and the labor supply decision rules, the bond market clears in each period.

Assume now that the central bank tries to replicate the flexible price path for the interest rate, with the only added constraint that the interest rate cannot go negative. The last panel of Figure 8 shows the output response in this case. The dashed line corresponds to the response in the flexible price case. With nominal rigidities the economy is in a liquidity

trap and the output response is larger. As long as the economy is in the liquidity trap, output dynamics are fully dominated by the demand side.

6 Durable Goods

We now extend the model to consider both durable and non-durable consumption. The introduction of durables is interesting for several reasons. First, durable consumption is the component of consumption that experienced the most noticeable decline in the recent US recession. Therefore, introducing durable goods allows us to get a better sense of the quantitative significance of our mechanism. Second, a large fraction of US household debt is in the form of secured credit used to finance the purchase of durables, especially if we interpret durables broadly to include housing. Therefore, including durables allows us to better calibrate our model to include both unsecured and secured credit (including mortgage debt) and to capture their overall dynamics.

From a more theoretical perspective introducing durables is interesting because it is a way of introducing a form of investment in our model. As argued in section 2, our main deviation from Aiyagari (1994) is that we abstract from capital accumulation. This implies that when consumers' precautionary demand for assets increases it cannot be directed towards increased investment in physical capital, but it all goes towards bonds' accumulation. In a model with durables, consumers face a more interesting portfolio problem, as they can choose to allocate their increased savings either towards bonds or towards investment in durables. Therefore, in principle, a shock that increases the precautionary demand for assets could translate into an increase in durable demand. In this section, we show that in our baseline experiment other forces are at work which make the consumers' demand for durables fall, not increase. Namely, durable goods are less liquid than bonds and are subject to decreasing returns. As households build up precautionary reserves, they shift their portfolio towards more liquid assets, favoring bonds over durable goods. Therefore, the interesting theoretical finding is that we can have a model of precautionary savings where a shock leads at the same time to an increased demand for bonds and to a reduction in the demand for less liquid investment, a true flight to liquidity.

It is important to point out that in a model with both secured and unsecured debt, there are two ways of capturing a credit crunch: a contraction in the unsecured credit limit

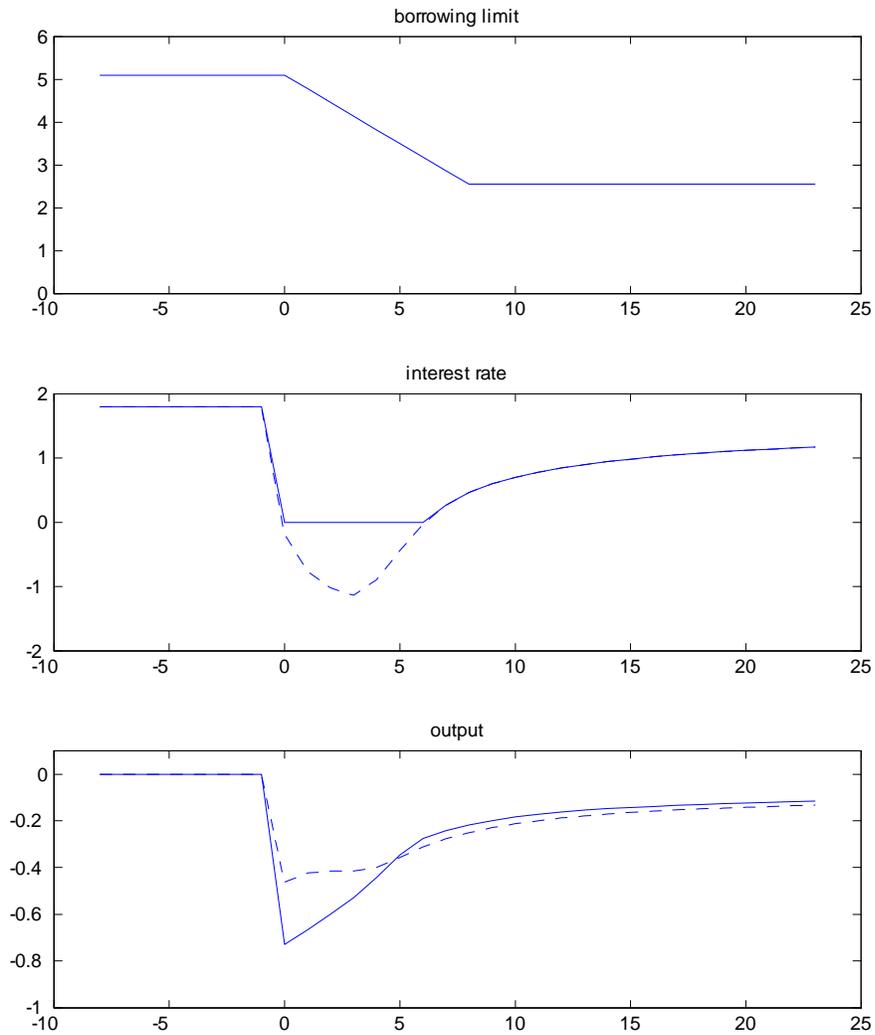


Figure 8: Interest rate and output responses with a lower bound on the interest rate (dashed lines: flexible price benchmark).

and a contraction in the pledgeability of collateral. We focus for now on the second type of credit crunch.

Households preferences are now represented by the utility function

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t U(c_t, k_t, n_t) \right],$$

where c_t is consumption of non-durable goods, k_t represent services from the stock of durables, and n_t is labor effort.

A consumer holding a stock k_t of durable goods has to pay δk_t to cover for depreciation. Then he can choose whether to accumulate or decumulate durables for next period. If he chooses to accumulate durable goods ($k_{t+1} > k_t$), he needs to pay $k_{t+1} - k_t$. If he chooses to decumulate durables ($k_{t+1} < k_t$), he faces real costs of reselling which are proportional to the capital sold and equal to $\zeta (k_t - k_{t+1})$. We capture these assumptions compactly by defining the adjustment cost function

$$g(k_{t+1}, k_t) = \begin{cases} \delta k_t + k_{t+1} - k_t & \text{if } k_{t+1} \geq k_t \\ \delta k_t + (1 - \zeta)(k_{t+1} - k_t) & \text{if } k_{t+1} < k_t \end{cases}.$$

We assume that $1 - \zeta > \delta$, so the consumer can always choose to liquidate part of the durable to cover for depreciation.⁶ The household budget constraint is then

$$q_t b_{t+1} + g(k_{t+1}, k_t) + c_t + \tilde{\tau}_t \leq b_t + y_t,$$

where the tax $\tilde{\tau}_t$ depends on the household's productivity θ_t as in the baseline model.

The production side of the model is as in the benchmark model, with the linear production function, $y_t = \theta_t n_t$, and an exogenous Markov process for θ_t . Notice that durable and non-durable goods are produced with the same technology.

The household borrowing constraint is

$$b_{t+1} \geq -\phi - \phi_k k_{t+1}. \tag{5}$$

⁶Grossman and Laroque (1990) and Gruber and Martin (2003) make a different assumption for durable goods: when agents choose $k_{t+1} \neq k_t$ they have to sell k_t at price $(1 - \zeta) k_t$ and buy k_{t+1} at full price. So they assume

$$\begin{aligned} g(k_{t+1}, k_t) &= \delta k_t + k_{t+1} - (1 - \zeta) k_t \text{ if } k_{t+1} \neq k_t \\ &= \delta k_t \text{ if } k_{t+1} = k_t \end{aligned}$$

The household now has access to both uncollateralized and collateralized debt. The maximum level of uncollateralized debt is given by ϕ . The parameter ϕ_k denotes the fraction of the value of the durable that can be used as collateral. The lower is ϕ_k , the less the household can use borrowed funds to finance durable purchases.

The government budget constraint is unchanged:

$$B_t + uz_t = q_t B_{t+1} + \tau_t,$$

where $u = \Pr(\theta_{j,t} = 0)$ is the fraction of unemployed agents. Again, we fix the supply of government bonds and the unemployment benefit at the levels \bar{B} and \bar{z} , and we let the tax τ_t adjust to ensure the government's budget balance.

6.1 Equilibrium

The main difference with the baseline model is that durable goods are now an additional state variable. The optimal decisions of the agents, for given sequence of interest rates $\{r_t\}$, are now functions of a three-dimensional state variable (b, k, θ) : the initial stock of bonds, the initial stock of durables, and the current productivity level. Given these three states, the household chooses non-durable consumption, $C_t(b, k, \theta)$, labor supply, $N_t(b, k, \theta)$, and next period stocks of durables and bonds, $K_{t+1}(b, k, \theta)$ and $B_{t+1}(b, k, \theta)$.

Let $\Psi_t(b, k, \theta)$ denote now the joint distribution of b, k and θ in the population. Combining the household's optimal transition for bond holdings and durable goods with the exogenous Markov process for productivity, we obtain a transition probability for the individual states (b, k, θ) , and, aggregating, a transition for the distribution Ψ_t .

The definition of equilibrium is the natural generalization of definition 1, where the bonds market clearing condition is now

$$\int \int \int b d\Psi_t(b, k, \theta) = \bar{B}.$$

Market clearing in the goods market takes the form

$$\int \int \int [C_t(b, k, \theta) + g(K_{t+1}(b, k, \theta), k) - \theta N_t(b, k, \theta)] d\Psi_t(b, k, \theta) \leq 0,$$

given that durable and non-durable goods are produced with the same technology.

The optimality conditions for non-durable consumption, labor supply, and bond holdings give the analogs of equations (2) and (3). However, now there is also an optimality

condition for durable goods. Letting $V(b, k, \theta)$ denote the household's value function, we can write this optimality condition as

$$(1 - \zeta) E [V_{t,b}(b_{t+1}, k_{t+1}, \theta_{t+1}) | \theta_t] \leq qE [V_{t,k}(b_{t+1}, k_{t+1}, \theta_{t+1}) | \theta_t] \leq E [V_{t,b}(b_{t+1}, k_{t+1}, \theta_{t+1}) | \theta_t]. \quad (6)$$

The first inequality must hold as an equality if $k_{t+1} < k_t$ and the second must hold as an equality if $k_{t+1} > k_t$. This condition shows the presence of an inaction region, in which households do not adjust their durable holdings.

To illustrate the household's optimal investment behavior it is useful to split the household's problem at each time t in two stages. First the household chooses his consumption and labor effort, which determine its total savings

$$x_t = b_t + y_t - c_t - \tilde{\tau}_t.$$

Then these savings are allocated to the purchase of bonds and durables, to satisfy

$$q_t b_{t+1} + g(k_{t+1}, k_t) = x_t.$$

Let us focus on this second stage. Suppose we vary the current level of savings x_t , for a given initial level of capital k_t and a given level of current productivity θ_t . In figure 9 we plot the optimal values of b_{t+1} and k_{t+1} for different possible values of x_t . When x_t is small the household readjusts downward its durable holdings, selling durables. In particular, in the case depicted, the household needs to sell durables to meet its collateral constraint, which corresponds to the dotted line labeled $b_{t+1} = -\phi - \phi_k k_{t+1}$. In this region, an increase in current savings leads to an increase in durable holdings, as fewer durables are sold, and to a *decrease* in bond holdings, as the household takes full advantage of its increased debt capacity. When the savings x_t are in some intermediate region, it is optimal to keep k_{t+1} constant at its initial level k_t . In this region all changes in current savings translate into an adjustment in next period bond holdings. Finally, when b_t is high enough the household adjusts upward its durable holdings. In this region, an increase in current savings is split between an increase in bond holdings and durables.

The dashed lines in the figure correspond to the adjustment bands, and are independent of the initial level of capital k_t . This implies that we can use the figure to trace the locus of optimal portfolios (b_{t+1}, k_{t+1}) for all possible levels of x_t (for given k_t): follow the upper

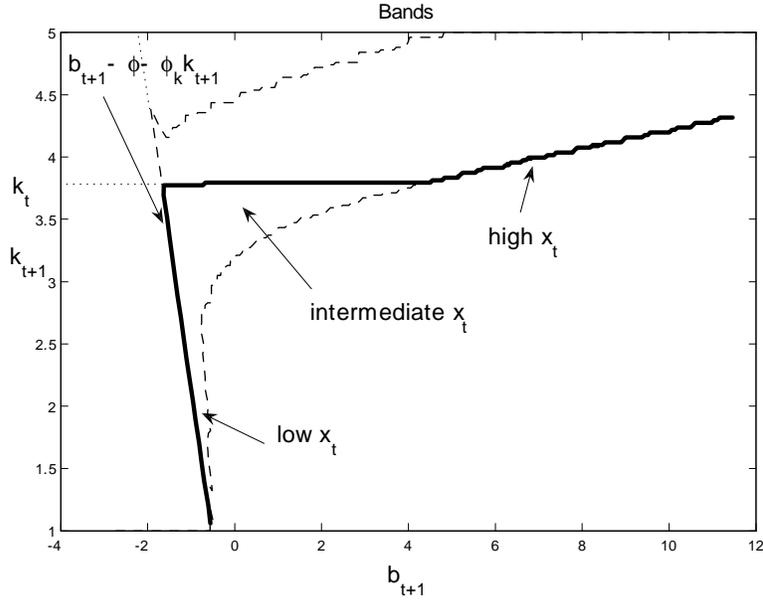


Figure 9: Optimal investment in bonds and durables

band until it hits k_t , then move horizontally until you hit the lower band, then move along the lower band.⁷

6.2 Calibration

Having now a richer model of households' consumption makes our quantitative results more interesting. We keep the calibration as close as possible to our benchmark calibration. Clearly, this may translate into different parameter values given that the model has now

⁷Using (6) and the concavity of the value function it is possible to prove that optimal portfolio choice always takes this form. The lower band corresponds to locus of pairs (b_{t+1}, k_{t+1}) that satisfy $b_{t+1} + \phi + \phi_k k_{t+1} > 0$ and solve the following equations for some h

$$\begin{aligned} q_t E[V_{t,k}(b_{t+1}, k_{t+1}, \theta_{t+1}) | \theta_t] &= E[V_b(b_{t+1}, k_{t+1}, \theta_{t+1}) | \theta_t] \\ q_t b_{t+1} + k_{t+1} &= h. \end{aligned}$$

The upper band corresponds to locus of pairs (b_{t+1}, k_{t+1}) that satisfy $b_{t+1} + \phi + \phi_k k_{t+1} > 0$ and solve the following equations for some x

$$\begin{aligned} q_t E[V_{t,k}(b_{t+1}, k_{t+1}, \theta_{t+1}) | \theta_t] &= (1 - \zeta) E[V_b(b_{t+1}, k_{t+1}, \theta_{t+1}) | \theta_t] \\ q_t b_{t+1} + (1 - \zeta) k_{t+1} &= h. \end{aligned}$$

more ingredients.

The utility function is now:

$$U(c, k, n) = \frac{(c^\alpha k^{1-\alpha})^{1-\gamma}}{1-\gamma} + \psi \frac{(1-n)^{1-\eta}}{1-\eta}.$$

We choose a simple Cobb-Douglas function in aggregating durable and non-durable consumption. Ogaki and Reinhart (1998) provide some support for an elasticity of substitution between durables and non-durables close to 1. This implies that α represents the share of non-durable consumption in total consumption expenditures. To compute the ratio of non-durables over total consumption we consider durables as the sum of durable consumption and imputed consumption of housing services from NIPA. We consider the rest (basically non-durables and non-housing services) as nondurables. We then set $\alpha = 0.7$, looking at the average value between 2000 and 2010. We set the coefficient of risk aversion $\gamma = 4$ as in the benchmark calibration. The specification of the disutility of labor is the same as in the benchmark model. The parameters η and ψ are chosen so that the hours worked of employed workers are on average 40% of their time endowment (normalized to 1) and the average Frisch elasticity of labor supply is 2.

For the accumulation process of durable goods, we have to choose δ and ζ . We set $\delta = 5\%$ to match the depreciation rate from NIPA Fixed Assets Tables. The parameter ζ represents the cost of selling durable goods and it is an important parameter determining the illiquidity of durables. We set $\zeta = 15\%$ and then examine how our results depend on this parameter.

For the wage process and the transitions between employment and unemployment we proceed as in the benchmark calibration. However, for the moment, we computed a yearly version of the model, and so we adjust the transition probabilities to reflect a yearly calibration of the processes. Moreover, for a yearly calibration we set $\beta = 0.96$, as standard in the literature.

Finally, we need to choose values for the parameters ϕ and ϕ_k , which determine the borrowing limits on uncollateralized and collateralized debt, and for the net bond supply \bar{B} . We calibrate the levels of ϕ and ϕ_k to match the ratio of household debt to GDP in 2006 which was approximately 93%. To measure household debt we sum mortgages (9,865 bln) and all other consumer credit (2,416 bln) from the flow of funds (so total debt is 12,281 bln). To calibrate the relative values of ϕ and ϕ_k we assume that when a household is in

debt it uses secured and unsecured debt in proportion to its debt limits ϕ and $\phi_k k_t$. Then we look at the average ratio of unsecured debt to total debt in 2006, which was equal to about 7%. Namely, we identify unsecured debt with revolving credit from the FRB table G.19 (844 bln). Given these parameters, we finally choose the net bond supply \bar{B} to get an interest rate level equal to 2% in the initial steady state, as it was in 2006 in the United States.

β	0.96	ρ	0.974	ϕ	0.13
γ	4	σ_ε^2	0.025	ϕ_k	0.51
η	0.75	z	0.39	δ	0.05
ψ	1.15	α	0.7	ζ	0.15

Table 2. Calibration with durables

6.3 Credit Crunch

We are now ready to show the results of our main exercise. We explore the response of the economy to a unexpected shock to the borrowing constraint. In particular, we consider a contraction to the pledgeability of collateral, that is, a drop in ϕ_k . We start from the economy in steady state with $\phi_k = \phi_k^I$, and we call Ψ^I the associated joint distribution of durable goods and bonds holdings. We look at the effects of a shock that reduces the parameter ϕ_k from $\phi_k^I = 0.51$ to $\phi_k^{II} = 0.41$. With a yearly calibration, we can assume that the adjustment takes place in a single period.

Figure 10 shows the transitional dynamics for the real interest rate, output, and investment in durable goods after the shock. The first thing to notice is that the patterns for the real interest rate is qualitatively the same as in the benchmark model. The interest rate converges to a lower level in the new steady state, but in the short run it drops much further. The main mechanism behind this overshooting is similar as in the benchmark model. Long run output dynamics are different than in the baseline. In particular, there is essentially no long run effect on output and actually there is a slight positive long run effect. We believe this is due to the precautionary motive increasing the stock of durables in the long run. However, output still drops in the short run, and actually the drop is considerably larger than in the baseline model. In particular, we obtain more than a one

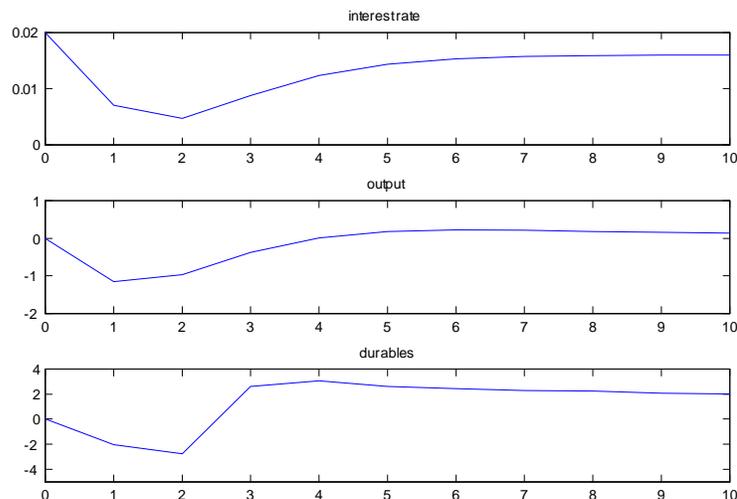


Figure 10: Transitional dynamics of the real interest rate, output, and investment in durable goods.

percentage drop in output, still a relatively small recession, but more than twice as large as the recession we obtained in the baseline.

Why the presence of durable goods amplify the reaction of real activity? The third panel of Figure 10 shows the response of investment in durables, which drop much more than overall output in the short run (by about three percentage points one year after the shock) and then recovers to reach a new, higher steady state. As mentioned above, there is a tension in the model. On the one hand, after a credit crunch, households want to save more for precautionary reasons and this tends to increase investment in durables. On the other hand, households start to rebalance their portfolio in favor of the more liquid assets, hence reducing investment in durables. The picture shows that, with the calibration we propose, the second effect clearly dominates in the short run, but the first effect seems to kick in in the long run. When the shock hits, households reduce their durable consumption, driving the demand for goods further down and making the recession deeper.

7 Concluding Remarks

We have proposed a model with uninsurable idiosyncratic risk to show how a credit crunch can generate a recession due to precautionary motive. This helps to explain why recessions driven by financial market trouble are more likely to drive the economy into a liquidity trap.

In the current version of the paper, we interpret a credit crunch as a tightening of the borrowing limit. More generally, it would be interesting to explore versions of the model with a simple intermediation sector, introducing a spread between the interest faced by borrowers and by savers, as in Curdia and Woodford (2009) and Hall (2009).

Another simplifying assumption in our model is that the unemployment risk is exogenous and not affected by the credit crunch. It would be interesting to develop a version of the model with endogenous job creation. If firms need liquid assets also to invest in new job openings, a credit crunch can generate a drop in vacancy creation and hence an increase in the unemployment risk. This would generate a potentially interesting feedback effect between unemployment risk and the precautionary motive.

Finally, a promising direction for future research could be to explore a variant of the model with risky capital. In this case, households' portfolio choice would be between bonds and claims to capital. It would be interesting to see if portfolio reallocation in favor of safe assets can help to explain the observed drop in asset prices and, through that channel, in investment.

Appendix

Here we describe the algorithm used to compute steady states and transitional dynamics.

To compute the steady state, given a candidate interest rate r , we iterate on the Euler equation and the optimality condition for labor supply to compute the policy functions $C(b, \theta)$ and $N(b, \theta)$ on a discrete grid for the state variable b . In particular, to iterate on the policy functions, we use the endogenous gridpoints approach of Carroll (2006). To compute the invariant distribution $\Psi(b, \theta)$ we derive the inverse of the bond accumulation policy, denoted by $g(b, \theta)$, from the policy functions, and update the conditional bond distribution using the formula $\Psi_{(k)}(b|\theta) = \sum_{\tilde{\theta}} \Psi_{(k-1)}(g(b, \tilde{\theta})|\tilde{\theta})T(\tilde{\theta}|\theta)$ for all $b \geq -\phi$,

where k stands for the k -th iteration and $T(\tilde{\theta}|\theta)$ is the probability of $\theta_{t-1} = \tilde{\theta}$ conditional on $\theta_t = \theta$. Due to the borrowing constraint, the inverse $g(b, \theta)$ is not well defined for $b = -\phi$, but the formula above is still correct if we define $g(-\phi, \theta)$ as the largest value of b such that $b' = -\phi$ is optimal. Finally, we search for the interest rate r that clears the bond market.

To compute transitional dynamics, we get the initial bond distribution $\Psi_0(b, \theta)$ from the initial steady state. We then compute the final steady at $\phi = \phi''$. We choose T large enough that the economy is approximately at the new steady state at $t = T$ (we use $T = 200$ in the simulations reported). Next, we guess a path of interest rates $\{r_t\}$ with $r_T = r''$. We take the consumption policy to be at the final steady state level at $t = T$, setting $C_T(b, \theta) = C''(b, \theta)$, and we compute the sequence of policies $\{C_t(b, \theta), N_t(b, \theta)\}$ using the Euler equation and the optimality condition for labor supply, going backward from $t = T - 1$ to $t = 0$ (also using endogenous gridpoints). Next, we compute the sequence of distributions $\Psi_t(b, \theta)$ going forward from $t = 0$ to $t = T$, starting at $\Psi_0(b, \theta)$, using the optimal policies $\{C_t(b, \theta), N_t(b, \theta)\}$ to derive the bond accumulation policy (using the same updating formula as in the steady state). We then compute the aggregate bond demand B_t for $t = 0, \dots, T$ and update the interest rate path using the simple linear updating rule $r_t^{(k)} = r_t^{(k-1)} - \epsilon(B_t^{(k)} - \bar{B})$. Choosing the parameter $\epsilon > 0$ small enough the algorithm converges to bond market clearing for all $t = 0, \dots, T$. To check that T is large enough, we compare check that $\Psi_T(b, \theta)$ is close enough to $\Psi''(b, \theta)$.

References

- [1] Aiyagari, Rao, 1994, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109 (3), 659-84.
- [2] Bewley, T., 1977, "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16(2), 252-92.
- [3] Brunnermeier, Markus and Yuliy Sannikov, 2010, "A Macroeconomic Model of the Financial Sector," mimeo, Princeton University.
- [4] Carroll, Christopher D., 2006, "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters* 91 (3), 312-320.

- [5] Carroll, Christopher D. and Miles S. Kimball, 1996, "On the Concavity of the Consumption Function," *Econometrica*, 64(4), pp. 981-92.
- [6] Carroll, Christopher D. and Wendy E. Dunn, 1997, "Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets," *NBER Macroeconomics Annual*, 12, 165-217
- [7] Chamley, Christopher, 2010, "Self-fulfilling Liquidity Traps with Money as a Medium of Exchange," mimeo, Boston University.
- [8] Curdia, Vasco and Michael Woodford, 2009, "Credit Frictions and Optimal Monetary Policy," mimeo, Columbia University.
- [9] Del Negro, Marco, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki, 2010, "The Great Escape? A Quantitative Evaluation of the Fed's Non-Standard Policies," mimeo, Princeton University.
- [10] Eggertsson, Gauti B. and Michael Woodford, 2003, "The Zero Interest-Rate Bound and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 1, 139-211.
- [11] Eggertsson, Gauti B., and Paul Krugman, 2011, "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo approach," mimeo, Federal Reserve Bank of New York.
- [12] Flodén, Martin and Jesper Lindé, 2001, "Idiosyncratic Risk in the U.S. and Sweden: Is there a Role for Government Insurance?," *Review of Economic Dynamics*, 4 (2), 406-437.
- [13] Geanakoplos, John, and Prandeeep Dubey, 2006, "Money and Production, and Liquidity Trap," *International Journal of Economic Theory*, 2 (3/4), 295-317.
- [14] Gertler, Mark and Peter Karadi, 2010, "A Model of Unconventional Monetary Policy," mimeo, New York University.
- [15] Gertler, Mark and Nobuhiro Kiyotaki, 2010, "Financial Intermediation and Credit Policy in Business Cycle Analysis," mimeo, New York University.

- [16] Goldberg, Jonathan, 2010, "A Model of Deleveraging," mimeo, MIT.
- [17] Guerrieri, Veronica, and Guido Lorenzoni, 2009, "Liquidity and Trading Dynamics," *Econometrica*, 77 (6), 1751-1790.
- [18] Hall, Robert, 2009, "The High Sensitivity of Economic Activity to Financial Frictions," mimeo, Stanford University.
- [19] Heathcote, Jonathan, Kjetil Storesletten and Gianluca Violante, 2009, "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, 1, 319-354.
- [20] Huggett, Mark, 1993, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17 (5-6), 953-69.
- [21] Jermann, Urban and Vincenzo Quadrini, 2009, "Macroeconomic Effects of Financial Shocks," NBER Working Paper no. 15338.
- [22] Kahn, Aubhik and Julia K. Thomas, 2010, "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity", mimeo, The Ohio State University.
- [23] Krugman, Paul R., 1998, "It's Back: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 2, 137-187.
- [24] Krusell, Per, Toshihiko Mukoyama, and Aysegul Sahin, 2010, "Labor-Market Matching with Precautionary Savings and Aggregate Fluctuations," *Review of Economic Studies*.
- [25] Mendoza, Enrique, Vincenzo Quadrini, and Jose-Victor Rios Rull, 2009, "Financial Integration, Financial Development and Global Imbalances," *Journal of Political Economy*, 117 (3), 371-416.
- [26] Midrigan, Virgiliu and Thomas Philippon, 2011, "Household Leverage and the Recession," mimeo, New York University
- [27] Nekarda, Christopher J. and Valerie A. Ramey, 2010, "The Cyclical Behavior of the Price-Cost Markup," mimeo, UCSD.

- [28] Ogaki, Masao and Carmen M. Reinhart, 1998, “Measuring Intertemporal Substitution: The Role of Durable Goods,” *Journal of Political Economy*, 106 (5), 1078-1098.
- [29] Ragot, Xavier and Edouard Challe, 2010, “Precautionary Savings in the Business Cycle,” mimeo, Paris School of Economics.
- [30] Shimer, Robert, 2005, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95 (1), 25-49.
- [31] Tauchen, G., 1986, “Finite State Markov-chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters* 20, 177–181.