

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: June 2011

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 outstanding research assistance. For helpful comments we thank Andy Atkeson, VV Chari, Bob Hall, John Leahy, Juan Pablo Nicolini, Thomas Philippon, Valery Ramey, Rob Shimer, Amir Sufi, Gianluca Violante, Pierre Yared, Mike Woodford and seminar participants at the Minneapolis Fed, the Milton Friedman Institute, the SED-Montreal, Columbia, NYU Stern, UCLA, UCSD, Richmond Fed, AEA meetings, ASU, MSU, Rice, Cambridge, LSE. Guerrieri thanks the Federal Reserve Bank of Minneapolis for its generous hospitality during part of this research. Lorenzoni thanks Chicago Booth 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 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 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 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 borrowing, 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, borrowing households are forced to deleverage and have to reduce consumption of both durable and non-durable goods. However, the precautionary motive induce lending households to save more, by accumulating both bonds and durables. Overall, with our calibration, lending households tend to favor bonds over durables, given that they are more liquid. This implies that the borrowers' contraction in durable purchases dominates the lenders' increase, leading to an aggregate drop in durable consumption and in aggregate 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 recent contributions that focus on transitional dynamics after different types of shocks. A good example is Mendoza, Rios Rull and Quadrini (2010) who look at the response of an economy opening up to international asset trade. Our treatment of durables is related to Carroll and Dunn (1997) and Fernandez-Villaverde and Kruger (2004), which incorporate durables in models 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), 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 households’ consumption, as we do here. Notable exceptions are Hall (2011a, 2011b), 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

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

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 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, especially when include durable goods in the model. Also, our analysis brings to attention the role of labor supply and 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 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. Section 5 studies the effects of simple fiscal 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

We consider a model of households facing idiosyncratic income uncertainty, who smooth consumption 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 can have negative bond holdings—i.e., they can borrow—up to an exogenous limit. We first analyze the steady state equilibrium of this economy, for a given borrowing limit. Then, we study the economy transitional dynamics following an unexpected shock that reduces this limit.

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

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

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

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

where θ_{it} is an idiosyncratic productivity shock which follows a Markov chain on the space $\{\theta^1, \dots, \theta^S\}$. We assume $\theta^1 = 0$ and we interpret this realization of the productivity shock as “unemployment.” For the moment, there are no aggregate shocks.

The household budget constraint is

$$q_t b_{it+1} + c_{it} + \tilde{\tau}_{it} \leq b_{it} + y_{it},$$

where b_{it} are the household bond holdings, q_t is the bond price and $\tilde{\tau}_{it}$ 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. The structure of tax payments is as follows: all households pay a lump sum tax τ_t and the unemployed receive the unemployment benefit v_t . That is, we set $\tilde{\tau}_{it} = \tau_t$ if $\theta_{it} > 0$ and $\tilde{\tau}_{it} = \tau_t - v_t$ if $\theta_{it} = 0$.² The household’s debt position is bounded below by the exogenous limit $\phi \geq 0$, that is, bond holdings must satisfy

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

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

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

The government chooses the aggregate supply of bonds B_t , the unemployment benefit v_t for all t and the lump sum tax τ_t so as to satisfy the budget constraint:

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

where $u = \Pr(\theta_{it} = 0)$ is the fraction of unemployed agents in the population. When we study transitional dynamics, we start from the assumption that the supply of government bonds and the unemployment benefit are kept constant at \bar{B} and \bar{v} , respectively, while the tax τ_t adjusts to ensure government budget balance. In Section 5, we consider the effects of alternative fiscal policies.

In our economy, the only supply of bonds outside the household sector comes from the government and is fixed at \bar{B} . When we calibrate our model, however, we'll interpret the bond supply \bar{B} more broadly as the sum of the liquid assets held by the household sector. In Section 6, we enrich the household portfolio choice by allowing households to accumulate both liquid bonds and a less liquid asset: durable consumption goods.

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 can borrow or lend at the same rate.

2.1 Equilibrium

Given a sequence of interest rates $\{r_t\}$ and of taxes $\{\tau_t\}$, let $C_t(b, \theta)$ and $N_t(b, \theta)$ denote the optimal levels of consumption and labor supply, at time t , for a household with bond holdings $b_{it} = b$ and current productivity $\theta_{it} = \theta$. 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 determines the distribution $\Psi_{t+1}(b, \theta)$, for a 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 consumption and labor supply policies $\{C_t(b, \theta), N_t(b, \theta)\}$, a sequence of taxes $\{\tau_t\}$, and a sequence of distributions for bond holdings and productivity levels $\{\Psi_t(b, \theta)\}$ such that, given the initial distribution $\Psi_0(b, \theta)$:

- (i) $C_t(b, \theta)$ and $N_t(b, \theta)$ are optimal given $\{r_t\}$ and $\{\tau_t\}$,
- (ii) $\Psi_t(b, \theta)$ is consistent with the consumption and labor supply policies,
- (iii) the tax satisfies

$$\tau_t = \bar{v}u + r_t\bar{B} / (1 + r_t),$$

- (iv) the bonds market clears:

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

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

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

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

$$\theta_{it} U_c(c_{it}, n_{it}) + U_n(c_{it}, n_{it}) = 0, \quad (3)$$

which holds for all households with $\theta_{it} > 0$.

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 is higher, by Jensen's inequality. This means that for a given level of interest rates, consumption today falls, as if there was a negative preference shock reducing the marginal utility of consumption today. This will be an important mechanism reducing consumption demand after a credit shock. In this sense, a model with precautionary savings provides a microfoundation for models that use preference shocks to push the economy in a liquidity trap, as, e.g., Christiano, Eichenbaum and Rebelo (2011).

2.2 Calibration

We analyze the model by numerical simulations, so we need to specify household preferences and calibrate the model parameters. In particular, we choose parameters that generate realistic levels of income uncertainty. Income uncertainty is the only motive for wealth accumulation in our model, as we are abstracting from life-cycle considerations and from other important drivers of household wealth dynamics, like durable goods purchases, health expenses, educational expenses, etc. However, our baseline calibration helps us identify some basic mechanisms which determine the responses of interest rates, consumption and output to a shock to credit access. These mechanisms extend to environments with richer motives for borrowing and lending. In particular, in Section 6, we extend the model to study the role of durable goods purchases.

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 our specification for the disutility of labor. The parameters are reported in Table 1. The parameter β is such that the yearly discount factor is $\beta^4 = 0.9103$ and is chosen to yield an interest rate of 2.5% in the initial steady state. This choice is meant to capture the low interest rates of the mid 2000's. The coefficient of risk aversion is $\gamma = 4$. Clearly, this coefficient is crucial in determining the consumers' precautionary behavior, so we also experiment with different values. The parameter ψ is chosen so that average hours worked for employed workers are 40% of their time endowment, normalized to 1 (following Nekarda and Ramey, 2010). The parameter η is chosen so that the average Frisch elasticity of labor supply is 1.

β	0.9777	ρ	0.967	ϕ	1.75
γ	4	σ_ε^2	0.017	\bar{B}	2.603
η	1.88	\bar{v}	0.168		
ψ	12.48				

Table 1. Baseline calibration

We calibrate the process for θ_{it} to capture wage and employment uncertainty. In particular, we assume that, when positive, θ_{it} follows an AR1 process in logs with autocor-

relation ρ and variance σ_ε^2 . We choose the parameters ρ and σ_ε^2 in line with the evidence in Floden and Lindé (2001), who use yearly panel data from the PSID to estimate the stochastic process for individual wages in the U.S. In particular, we focus on the persistent component of their wage process and choose parameters that yield a coefficient of autocorrelation of 0.9136 and a conditional variance of 0.0426 for yearly wages.³ The corresponding parameters are reported in Table 1. The wage process is approximated by a 13-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 quarterly transition probability from employment to unemployment at 0.057, and the quarterly transition from unemployment to employment at 0.882. We assume that when re-employed workers start at the average value of θ_{it} . For the unemployment benefit \bar{v} , we also follow Shimer (2005) and set it to 40% of average labor income.

Finally, we need to choose values for the bonds supply \bar{B} and for the borrowing limit ϕ . We choose these values in line with U.S. 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 broadly defined, namely the sum of their holdings of all deposits plus securities held directly by households which was approximately 170% 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

³See Table IV in Floden and Lindé (2001). The relation between quarterly and annual parameters (denoted by the subscript a) is $\rho = \rho_a^{1/4}$ and

$$\sigma_\varepsilon^2 = \frac{8(1-\rho^2)}{2+3\rho+2\rho^2+\rho^3} \frac{\sigma_{\varepsilon,a}^2}{1-\rho_a^2}.$$

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 almost linear in b . For lower levels of bond holdings, the consumption function is concave, as is common in precautionary savings models (see Carroll and Kimball, 1996). The optimality condition for labor supply implies that the relation between bond holdings and 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 most levels of b the substitution effect dominates the income effect and labor supply is higher for high wage earners. For lower levels of b , however, the income effect becomes stronger and dominates the substitution effect: very poor households want to work longer hours the lower is their hourly wage.




Figure 1: Optimal consumption and labor supply at the initial steady state (for $\theta = \theta^2$ and $\theta = \theta^{13}$).

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' . Recall that we chose $\phi' = 1.75$ to match the US households debt/GDP equal to 18%. We look at the effects of an unexpected shock leading to a decrease in the borrowing limit to $\phi'' = 0.98$. In particular, we choose the size of the shock so that the households debt over GDP ratio drops by 10 percentage points in the new steady state. 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 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 induce households to hold the same quantity of bonds.

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 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

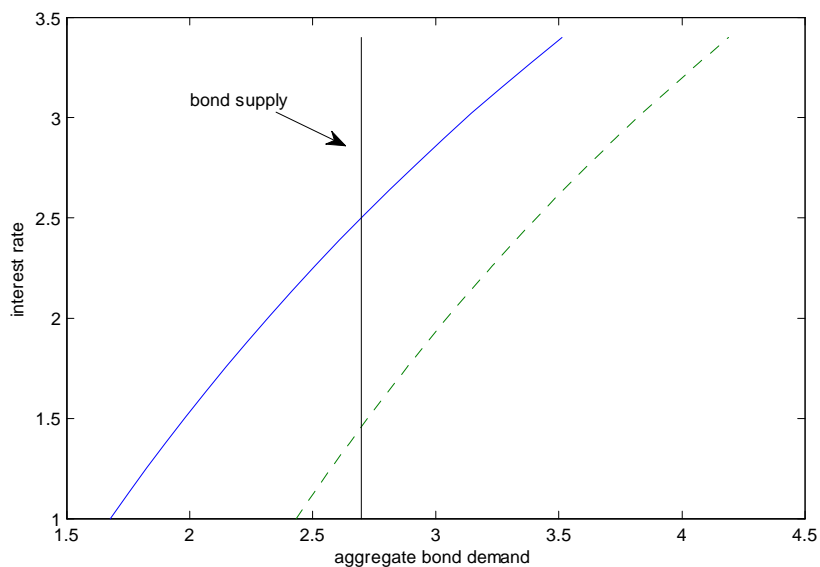


Figure 2: Aggregate bond demand in steady state

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 6 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 two panels of Figure 3, we plot the exogenous adjustment paths for the ratio of ϕ_t to GDP and for the debt/GDP. In the bottom two panels we plot the paths for the interest rate and the output level.

3.1 Interest rate dynamics

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 a 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 distribution 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

⁴We also experimented with different parameters configurations and the interest rate overshooting seems a very robust result.

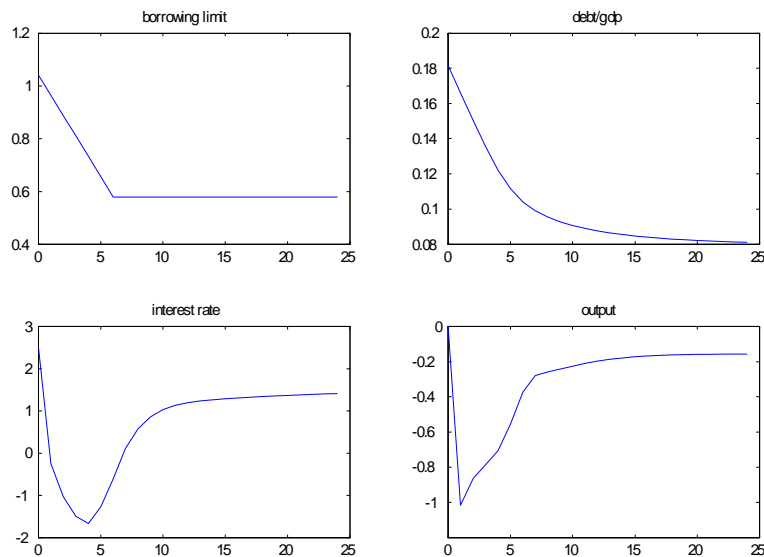


Figure 3: Interest rate and output dynamics after a drop in ϕ that reduces debt/gdp by 10 percentage points.

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 bottom panel of Figure 4 shows the densities associated, respectively, to F' (solid line) and to F'' (dashed line). 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. One can check formally that F' is a mean-preserving spread of F'' by showing that the integral $\int_{-\infty}^b (1 - F(\tilde{b})) d\tilde{b}$ for F'' is always above the one for F' . 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 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 $b_{it+1} - b_{it}$ (averaged over θ) as a function of the initial bond holdings b_{it} for the initial steady state (solid line) and the new one (dashed line). These functions are not exactly convex, but close to convex. 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




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

We are now ready to put the pieces together. Let us do a mental experiment and suppose the interest rate jumps immediately to its new steady state value at date 0. If the wealth distribution was already at the new steady state, average bond accumulation would be zero. In other words, the integral of the dashed function in the top panel weighted by the dashed density in the bottom panel is equal to zero. This implies that the integral of the dashed function weighted by the solid density is a positive number, because the dashed function is (approximately) convex and F' is a mean-preserving spread of F'' . Therefore, at the conjectured interest rate path, households want, on average, to accumulate bonds. Since the bond supply is fixed, this means that the conjectured interest rate path is not the equilibrium one, as it leads to an excess demand of bonds. To equilibrate the bonds market, we need a lower interest rate in the initial periods.

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).

3.2 Output response

Next, we want to understand what happens to output. Figure 3 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 approximately one percent reduction in output. We will consider below variations leading to larger output responses. 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 output overshooting. Again, let us make a mental experiment and suppose the interest rate jumps 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 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. 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 5 shows the behavior of interest rate and output for the same economy with $\gamma = 2$ (dashed lines) instead of $\gamma = 4$ (solid lines).⁶ According to 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 the economy takes longer to adjust.

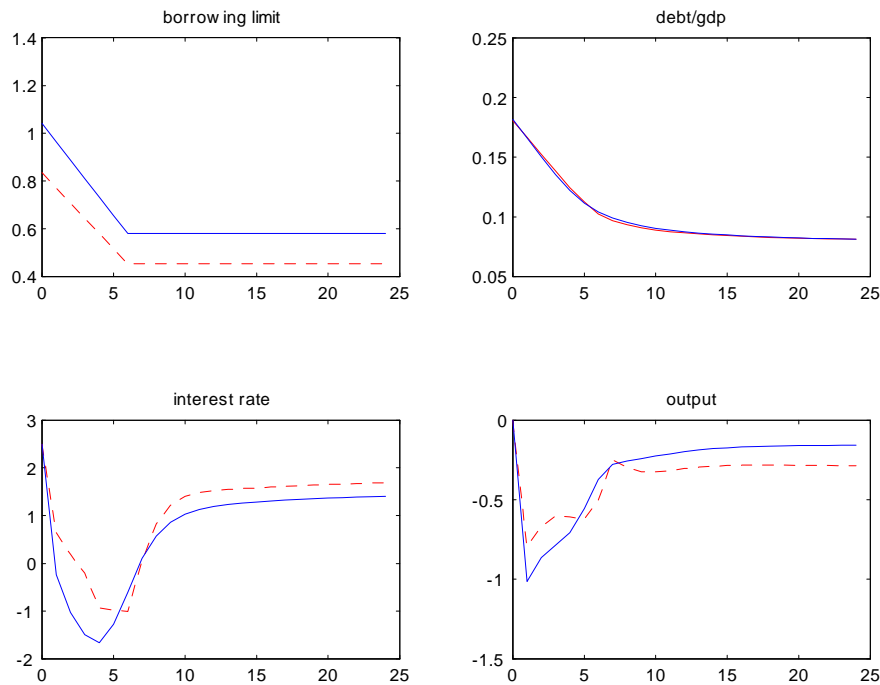


Figure 5: Changing the coefficient of relative risk aversion γ .

⁶The calibrated parameters, e.g. ϕ' , are re-calibrated when we change γ .

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 6: Output and employment for the baseline calibration and for an alternative calibration with lower ψ .

Finally, let us focus on the dynamics of aggregate employment. The top panel of Figure 6 shows the paths for output (solid line) and employment (dashed line). In our baseline calibration, aggregate employment increases in response to a credit crunch. The behavior of employment is determined by the shape of the labor supply policy rule. In

particular, when the labor supply is convex, the labor supply decision of poor households is sensitive to a credit crunch. As a result, total employment can increase, as it happens in our baseline calibration. While there is a lot of research, both theoretical and empirical, that suggests that the consumption function is concave in wealth, less is known on the second derivative of labor supply as a function of wealth. In the bottom panel of Figure 6 we show the paths of output (solid line) and employment (dashed line) for an alternative calibration with $\psi = .97$. In this case the labor supply policy function is concave enough so that aggregate employment drops by more than one percentage point. As a result, output drops much more than in the baseline calibration, by more than two percentage points. Notice that in our model, output always drops more than employment because the households who work more in response to a credit crunch are the poor ones, which typically have lower productivity.

4 Nominal Rigidities

Under flexible prices, the real interest rate is free to adjust to its equilibrium path to equilibrate the demand and supply of bonds, or—equivalently—the demand and supply of goods. In this section we explore what happens in a variant of the model with nominal rigidities. In presence of nominal rigidities, the central bank can affect the path of the real interest rate by setting nominal interest rates. However, the zero lower bound for nominal interest rates, together with nominal rigidities, implies that the central bank may not be able to achieve a real interest rate path that corresponds to the flexible price equilibrium. Therefore, a credit crisis which produces a large drop in real interest rates under flexible prices can drive the economy in a liquidity trap under sticky prices.

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

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

where $c_{it}(j)$ 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 θ_{it} as a shock to the efficiency of household i labor. Letting w_t denote the real wage rate per efficiency unit, the hourly wage rate for worker i is then $\theta_{it}w_t$.

Firms are owned by the consumers, so letting π_t denote total real profits, the budget constraint is

$$q_t b_{it+1} + c_{it} = b_{it} + \theta_{it} w_t n_{it} - \tilde{\tau}_{it} + \pi_t.$$

Monopolist j faces the demand

$$y_{jt} = (p_{jt}/P_t)^{-\varepsilon} C_t,$$

where C_t is aggregate consumption in the economy and P_t is the appropriate price index.

With flexible prices, 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

$$w_t = (\varepsilon - 1) / \varepsilon$$

and that households receive some 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. Therefore, the equilibrium path we describe is consistent with firms optimization with menu costs sufficiently large.

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. Notice that this model features nominal price rigidity but no wage rigidity. Therefore, when output contracts below its flexible equilibrium path the real wage needs to fall, to be consistent with the reduction in labor supply. An alternative interpretation is to assume that the labor market “wedge” adjusts endogenously so that a reduction in goods demand is translated in a reduction in labor inputs. Obviously, this is a crude theory of a demand-driven contraction in employment. But it works for our purpose here, which is to show how the zero

lower bound affects the required adjustment in consumer spending following a credit contraction.⁷

Suppose the central bank tries to replicate the flexible price path for the real interest rate, with the only constraint that the real interest rate cannot go negative. The last panel of Figure 7 shows the output response in this case (dashed line) together with the response in the flexible price case (solid line). With nominal rigidities the economy enters a liquidity trap and the output response is larger. As long as the economy is in the liquidity trap, output dynamics are fully determined by the demand side.

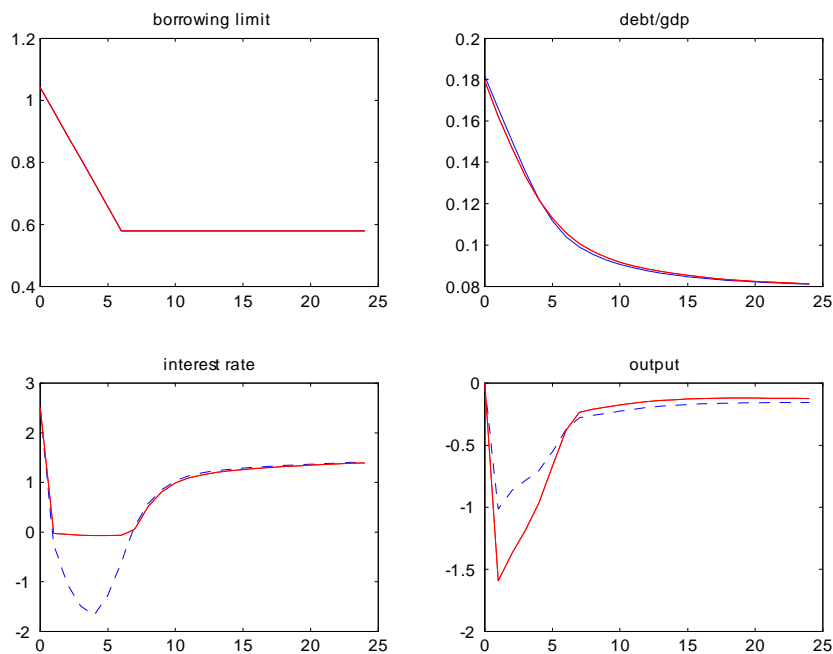


Figure 7: Interest rate and output responses with a binding zero lower bound on the interest rate (dashed lines: flexible price benchmark).

⁷As Hall (2011) points out, reconciling demand-driven contractions in employment with frictional models of the labor market is an important unresolved issue in macro.

5 Fiscal policy

We now explore the role of fiscal policy in mitigating the recession. In particular, we focus on simple policies in which the government changes the supply of government bonds by reducing the tax τ_t or by increasing the unemployment benefit v_t . 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. In our economy, this has a positive effect on spending given that Ricardian equivalence does not hold.

Since we assume lump sum taxation, an equivalence result holds between government supplied and privately supplied liquidity. Namely, an increase in the supply of government bonds B_t can exactly offset a change in the borrowing limit ϕ . In particular, the only thing that matters for the equilibrium is the sum $B_t + \phi$. This is a common result in this class of models, and it implies that in principle the government could completely neutralize the effect of a credit shock, by a sufficiently large increase in the supply of government bonds. However, for the sake of realism, here we look at the effects of policies that only partially offset the long run change in ϕ , possibly because of unmodelled 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 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 8 shows what happens to interest rate and output under these two policies. The blue lines represent the policy where the increase in B finances a temporary reduction in the tax τ_t ; the red lines represent the policy where the deficit goes partly to finance an increase in unemployment insurance.




Figure 8: Interest rate and output dynamics under alternative fiscal policies.

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.

6 Durable Goods

We now extend the model to include durable goods. The main reason for this extension is that credit access seems especially relevant for durable goods and that a large fraction of household borrowing takes the form of secured debt used to finance durable goods purchases.⁸

Since durable holdings are a form of savings, a model with durables enriches the

⁸Mian and Sufi (2011) document the importance of collateralized borrowing for consumers' spending both in the period preceding the recent crisis and during the crisis.

portfolio decision of the households. In particular, durables offer an alternative store of value, especially if we interpret durables to include housing. This means that when consumers' precautionary demand for assets increases, it can be directed not only towards bonds but also towards durables. This can potentially lead to an increase in durable accumulation as a result of an increase in precautionary savings. However, on the borrowers' side, opposing forces are at work: reduced credit access implies that borrowers need to sell durables in order to reduce their debt. This leads to durable goods decumulation. Whether the forces on the savers' side or on the borrowers' side dominate, depends on the model parameters. With realistic parameter values, we will see that the borrowers' side dominates and a credit shock leads to a reduction in durable purchases. As in the baseline model, these forces are especially strong in the short run, given that many borrowers find themselves with excess debt and are pushed to repay their debt faster.

A crucial parameter behind an aggregate contraction in durable purchases is the degree of liquidity of durable goods. In our model we assume that household face a discount when re-selling durables. A higher discount implies that durables are a less liquid form of savings. When households build up precautionary reserves following a credit shock, they tend to prefer more liquid assets, favoring bonds over durable goods. This reduces the increase in durable demand by savers and tends to generate an overall reduction in durable purchases. From a theoretical point of view, the interesting finding is that different degrees of liquidity can lead at the same time to an increased demand for bonds and to an overall reduction in the demand for less liquid assets. This captures a form of "flight to liquidity" on the households' side.

Households preferences are now represented by the utility function

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

where c_{it} is consumption of non-durable goods, k_{it} represent services from the stock of durables, and n_{it} is labor effort.

Each period durables depreciate at the rate δ and the household chooses whether to increase or decrease its stock of durables. If the household wants to accumulate durables, it needs to pay $k_{it+1} - k_{it}$ to purchase new durables plus δk_{it} to cover the depreciation of the existing stock. If it wants to decumulate durables, it faces real reselling

costs proportional to the capital sold, equal to $\zeta (k_{it} - k_{it+1})$. The parameter $\zeta > 0$ determines the illiquidity of durable goods. These assumptions are summarized in the adjustment cost function

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

We assume that $1 - \zeta > \delta$, so the household can always liquidate part of its durable stock to cover for depreciation.⁹

We also extend the model to allow for a spread between the interest rates on borrowing and lending. Let b_{it}^+ denote positive bond holdings and b_{it}^- denote negative holdings (borrowing). Also, let q_t be the price at which a household can buy bonds and \hat{q}_t the price at which a household can issue bonds. The household's budget constraint is then

$$q_t b_{it+1}^+ + \hat{q}_t b_{it+1}^- + g(k_{it+1}, k_{it}) + c_{it} + \tilde{\tau}_{it} \leq b_{it} + y_{it},$$

where the tax $\tilde{\tau}_{it}$ depends on the household's productivity θ_{it} as in the baseline model. A positive spread between interest rates for borrowers and lenders correspond to $\hat{q}_t < q_t$. We assume that there is a competitive banking sector that takes deposits, makes loans, and incurs monitoring costs proportional to the value of the funds intermediated. The intermediation cost per dollar of bonds issued is σ . This implies that $\hat{q}_t = (1 - \sigma) q_t$ and banks make zero profits. For interpretation, it is useful to express the intermediation cost in terms of the spread between borrowing and lending rates, which is equal to

$$\frac{1/\hat{q}_t}{1/q_t} - 1 = \frac{1}{1 - \sigma} - 1.$$

The production side of the model is as in the benchmark model, with the linear production function $y_{it} = \theta_{it} n_{it}$ and the exogenous Markov process for θ_{it} . Durable goods and non-durable goods are produced with the same technology.

The household's borrowing constraint is

$$b_{it+1} \geq -\phi - \phi_k k_{it+1}. \quad (4)$$

⁹An alternative assumption on transaction costs—made in Grossman and Laroque (1990) and Gruber and Martin (2003)—is to assume that 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 $g(k_{t+1}, k_t) = k_{t+1} - (1 - \zeta) k_t + \delta k_t$ if $k_{t+1} \neq k_t$ and $g(k_{t+1}, k_t) = \delta k_t$ otherwise. An advantage of our approach is that it keeps the household's problem concave.

The household now has access to both uncollateralized and collateralized debt. The maximum level of uncollateralized debt is ϕ . The parameter ϕ_k denotes the fraction of the value of the durable that can be used as collateral.

The government budget constraint is unchanged:

$$B_t + v_t u = q_t B_{t+1} + \tau_t.$$

and, as in the baseline, we fix the supply of government bonds and the unemployment benefit at the levels \bar{B} and \bar{v} , and let the tax τ_t adjust to ensure the government's budget balance.

6.1 Equilibrium and calibration

The main difference with the baseline model is that durable goods are now an additional state variable. The optimal decisions of the agents, for a 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. These three states fully determine the household's choice of non-durable and durable purchases, labor supply and the optimal level of borrowing or lending.

Let $\Psi_t(b, k, \theta)$ denote 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 the transition probability of the individual states (b, k, θ) , and, aggregating, a transition for the distribution Ψ_t . The definition of equilibrium is then the natural generalization of definition 1, where the bonds market clearing condition is now

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

To calibrate the model we adopt the utility function:

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

We choose a simple Cobb-Douglas specification to aggregate 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 ratio of non-durable consumption to total consumption. To compute this ratio we compute durables as the sum of durable consumption and consumption of housing services from NIPA. We take all other consumption (non-durables and non-housing services) as nondurables. This gives us $\alpha = 0.7$, looking at the average value for 2000-2010. 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 1. For the wage process and the transitions between employment and unemployment we proceed as in the benchmark calibration. However, we choose β to match a higher interest rate in our initial steady state, $r = 4.5\%$, to capture interest rates on longer maturity loans.

For the accumulation of durable goods, we need 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.

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 set ϕ_k to 0.8, which is in the range of loan-to-value ratios in mortgages and other durable financing. The value of ϕ is chosen assuming that borrowing households use secured and unsecured debt in proportion to the debt limits ϕ and $\phi_k k_t$. Then $\phi = 0.3$ delivers a ratio of unsecured debt to total debt equal to about 7%, in line with flow of funds evidence (we identify unsecured debt with revolving credit from the FRB table G.19). The parameters used are summarized in Table 2.

β	0.9713	ρ	0.967	ϕ	0.3
γ	4	σ_ε^2	0.017	ϕ_k	0.8
η	1.50	\bar{v}	0.160	δ	0.05
ψ	2.54	α	0.7	ζ	0.15

Table 2. Calibration with durables

In Figure 9 we represent the joint distribution of bond and durable holdings in the




Figure 9: The joint distribution of bonds and durable holdings in steady state.

initial steady state.¹⁰ The straight line represents the borrowing constraint. At low levels of total wealth (bonds plus durables) we find households who hold small durable stocks and small amounts of debt. If a household receives a positive, persistent income shocks, it responds by accumulating durables and taking on more debt. If households keeps receiving high income, it eventually pays off its debt and starts accumulating positive bond holdings. These wealth dynamics account for the curved shape of the durables/bond distribution. Moreover, for each level of durable holdings, there is a range of bond holdings such that households respond to transitory shocks by accumulating and decumulating bonds, keeping their durable holdings unchanged. This accounts for the fact that for each k , the distribution of bond holdings is fairly spread out.

6.2 Credit crunch

For the model with durables, we consider an alternative credit-tightening exercise, by looking at the effects of a permanent increase in spreads between borrowing and lending

¹⁰The contour lines correspond to the following percentiles: 0.1%, 1%, 30%, 40%, 80%.

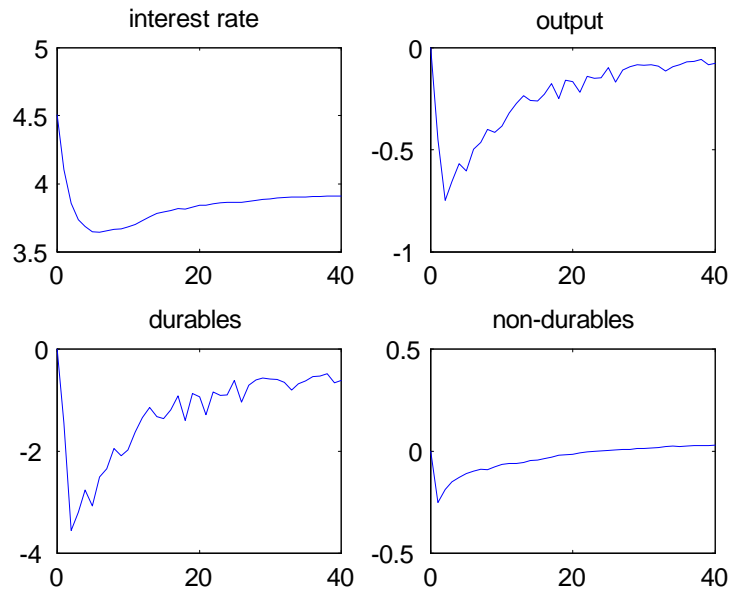


Figure 10: Interest rate and output dynamics in the model with durables: permanent increase in spread from 1% to 2.5%.

interest rates. We start the economy in steady state with the spread equal to $s = 1\%$ and we look at the effect of a permanent increase of the spread to $s = 2.5\%$. Such an increase determines a reduction in debt to GDP of 10% as in our baseline exercise.

Figure 10 shows the transitional dynamics for the real interest rate, output, and the production of durable and non-durable goods after the shock. Comparing this figure to Figure 3 shows that the patterns of real interest rate and output are qualitatively similar to our baseline exercise. From a quantitative point of view, however, both the interest rate drop and the output drop are smaller, but much more persistent (notice the difference in the number of time periods plotted in the two cases). The reason for the smaller effect on impact and for the increased persistence is due to the slower adjustment of the joint distribution of bonds and durables to its new steady state. A crucial difference between the model with durables and the baseline model is that the households most in debt are not the poorest households, but the households who had fairly good income prospects and an intermediate level of total initial wealth. These households can delever more gradually.

Looking at the two bottom panels of Figure 10, we see that the bulk of the adjustment takes place in durables, while non-durables experience a much smaller response than in the baseline. Here the crucial difference with the baseline is that households have an additional margin of adjustment when they are getting out of debt: instead of curtailing consumption, they can sell durables to repay their debt. Combining this with the observation above—that households more in debt are not the poorest households but households with intermediate wealth levels—helps to explain why deleveraging affects mostly durable purchases.

7 Concluding Remarks

We have proposed a model with uninsurable idiosyncratic risk to show how a credit crunch can generate a recession due to a combination of debt repayments and an increase in precautionary savings. This helps to explain why recessions driven by financial market trouble are more likely to drive the economy into a liquidity trap.

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 in which the labor market response to a drop in consumer demand leads to an endogenous increase in unemployment. The analysis in Hall (2011c) shows the open challenges that such an extension will face.

Finally, a missing element in the analysis is capital. Adding capital to the model would require a theory of why claims to physical capital cannot be costlessly transformed into assets equivalent to the bonds in our model. A model with spreads can be used to capture this illiquidity in a simple manner, while keeping the focus on household spending.

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, Gaudi 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, Gaudi B. and Michael Woodford, 2003, "The Zero Interest-Rate Bound and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 1, 139-211.
- [11] Eggertsson, Gaudi 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, 2011a, "The High Sensitivity of Economic Activity to Financial Frictions," *Economic Journal*, 121, 351-378.
- [19] Hall, Robert, 2011b, "The Long Slump," *American Economic Review*, 101(2), 431-469.
- [20] Hall, Robert, 2011c, "Clashing Theories of Unemployment," mimeo, Stanford University.
- [21] Heathcote, Jonathan, Kjetil Storesletten and Gianluca Violante, 2009, "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, 1, 319-354.
- [22] Huggett, Mark, 1993, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17 (5-6), 953-69.
- [23] Jermann, Urban and Vincenzo Quadrini, 2009, "Macroeconomic Effects of Financial Shocks," NBER Working Paper no. 15338.

- [24] Kahn, Aubhik and Julia K. Thomas, 2010, "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity", mimeo, The Ohio State University.
- [25] 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.
- [26] Krusell, Per, Toshihiko Mukoyama, and Aysegul Sahin, 2010, "Labor-Market Matching with Precautionary Savings and Aggregate Fluctuations," *Review of Economic Studies*.
- [27] 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.
- [28] Midrigan, Virgiliu and Thomas Philippon, 2011, "Household Leverage and the Recession," mimeo, New York University
- [29] Nekarda, Christopher J. and Valerie A. Ramey, 2010, "The Cyclical Behavior of the Price-Cost Markup," mimeo, UCSD.
- [30] Ogaki, Masao and Carmen M. Reinhart, 1998, "Measuring Intertemporal Substitution: The Role of Durable Goods," *Journal of Political Economy*, 106 (5), 1078-1098.
- [31] Ragot, Xavier and Edouard Challe, 2010, "Precautionary Savings in the Business Cycle," mimeo, Paris School of Economics.
- [32] Shimer, Robert, 2005, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95 (1), 25-49.
- [33] Sufi, Amir and Atif Mian, 2011, "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis," *American Economic Review*, forthcoming.
- [34] Tauchen, G., 1986, "Finite State Markov-chain Approximations to Univariate and Vector Autoregressions." *Economics Letters* 20, 177-181.