Credit Crises, Precautionary Savings, and the Liquidity Trap*

Veronica Guerrieri
University of Chicago and NBER

Guido Lorenzoni
Northwestern University and NBER

November 2016

Abstract

We study the effects of a credit crunch on consumer spending in a heterogeneous-agent incomplete-market model. After an unexpected permanent tightening in consumers' borrowing capacity, constrained consumers are forced to repay their debt and unconstrained consumers increase their precautionary savings. This depresses interest rates, especially in the short run, and generates an output drop, even with flexible prices. The output drop is larger with sticky prices, if the zero lower bound prevents the interest rate from adjusting downwards. Adding durable goods to the model, households take larger debt positions and the output response can be larger.

*Email addresses: vguerrie@chicagobooth.edu; guido.lorenzoni@northwestern.edu. For helpful comments, we are grateful to Robert Barro, Andrei Schleifer and two referees, as well as to Chris Carroll, V. V. Chari, Vasco Curdia, Gauti Eggertsson, Bob Hall, Greg Kaplan, John Leahy, Juan Pablo Nicolini, Thomas Philippon, Valery Ramey, Rob Shimer, Nancy Stokey, Amir Sufi, Gianluca Violante, Mike Woodford and numerous seminar participants. We thank Adrien Auclert and Iván Werning for generously sharing two analytical results presented in the Online Appendix. Adrien Auclert and Amir Kermani provided outstanding research assistance. Guerrieri thanks the Sloan Foundation for financial support and the Federal Reserve Bank of Minneapolis for its hospitality. Lorenzoni thanks the NSF for financial support and Chicago Booth and the Becker Friedman Institute for their hospitality.

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 a stationary state with a stable distribution of borrowing and lending positions. An unexpected shock hits the financial system and borrowing gets harder in terms of tighter borrowing limits and/or in terms of higher credit spreads. The most indebted consumers need to readjust towards lower levels of debt. Since the debtor position of one agent is the creditor position of another, this also means that lenders 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, focusing on the response of the household sector, using a workhorse Bewley (1977) model in which households borrow and lend to smooth transitory income fluctuations. Since the model cannot be solved analytically, our approach is to obtain numerical results under plausible parametrizations and to explore the mechanism behind them. The model captures two channels in the consumers' response to a reduction in their borrowing capacity. First, a direct channel, by which constrained borrowers are forced to reduce their debt. Second, a precautionary channel, by which unconstrained agents reduce their debt or increase their savings as a buffer against future shocks. Both channels increase net 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 interest rate dynamics and show that they are characterized by a sharp initial fall followed by a gradual adjustment to a new, lower steady state. The reason for the interest rate overshooting is that, at the initial asset distribution, the agents at the lower end of the distribution try to adjust towards a higher wealth target by increasing their net saving. To keep 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 responses of aggregate activity. In our baseline flexible price specification we obtain a mild output reduction of about 1.1%, in response to a shock that leads, in the long run, to a debt reduction of 10% of initial output.

We also study the economy's response under a simple form of nominal rigidities: fixed nominal wages. In the flexible price economy the nominal interest rate is negative in the short run following the credit tightening. Therefore, with nominal rigidities, the central bank reaches the zero lower bound and is unable to achieve the real interest rate that replicates the flexible price allocation. Therefore, with nominal rigidities the credit tightening leads to a larger contraction in output, of about 1.7%.[1]

We provide an interpretation of the output responses, with and without nominal rigidities, looking both at the demand and at the supply side of the model. On the demand side, the aggregate consumption response can be decomposed in two parts: a change due to the exogenous shift in the credit limit and a change due to the endogenous reduction in the interest rate. The first effect is large, about -4%, but is counteracted by a large consumption response to the endogenous drop in the interest rate. This decomposition shows that heterogeneous agent incomplete market models with standard preference parameters feature a fairly large interest rate elasticity of aggregate non-durable consumption. This decomposition also explains why the incomplete adjustment of the interest rate under nominal rigidities leads to a larger fall in consumption. The consumption predictions of the model are sensitive to the chosen calibration target for consumer savings. In our baseline calibration we target average liquid savings. When we target median savings instead of average savings, we obtain a calibration with a larger partial equilibrium response and a lower interest elasticity of aggregate consumption, leading to much larger responses of output both in the flexible price case and in the economy with nominal rigidities (respectively -1.8% and -5%).

We also look at the output response from the supply side. Here opposing forces are at work, since overly indebted agents would like to adjust by working more and increasing their current labor income. A composition effect tends to dampen the effect of this channel on output, as the more indebted households are also the one with worse employment opportunities (captured by lower labor productivity in the model). There-

[1]This was the first paper to combine nominal rigidities with a heterogeneous agent model of precautionary savings. Since this paper was circulated this combination of ingredients has proved useful to analyze other questions, most notably the effects of transfer spending in Oh and Reis (2012) and the effects of forward guidance in McKay, Nakamura, Steinsson (2016). Our working paper and the papers cited use nominal price rigidity, while here we use nominal wage rigidity. On page 23 we discuss the relative merits of these two approaches.

fore, in our baseline specification, the model predicts an increase in employment and a reduction in labor productivity following the credit shock. These predictions depend on the preference specification and we also present alternative calibrations where these effects are weaker and there is a reduction in employment following the credit shock.

Finally, we generalize the model to include durable consumption goods, which can be used as collateral. In this extension, households face a richer portfolio choice as they can invest in liquid bonds or in durable goods. To make bonds and durables imperfect substitutes, we assume a proportional cost of re-selling durables, so that durables are less liquid. After a credit crunch, net borrowers are forced to delevarage and have to reduce consumption of durable and non-durable goods. On the other hand, the precautionary motive induces net lenders to save more by accumulating both bonds and durables. Durable purchases may increase or decrease, depending on the strength of these two effects. In our calibration, the net effect depends on the nature of the shock. A pure shock to the credit limit affects only borrowers close to the limit, so the lenders' side dominates and durable purchases increase. A shock to credit spreads, on the other hand, affects a larger fraction of borrowers, leading to a contraction in durable purchases. Here the output effects can be large, leading to a 4% drop in consumption following a transitory shock that raises the spread on a one year loan from 1% to 3.8%. The consumption drop can be as large as 10% if prices are fixed and the zero lower bound is binding. However, the responses we obtain in this section are concentrated in durable consumption, with very small effects on non-durables.

Our paper focuses on households' balance sheets adjustment and consumer spending and is complementary to a growing literature that looks at the effects of credit shocks on firms' balance sheets and investment spending. [2] Hall (2011a, 2011b) argues that the response of the household sector to the credit tightening is an essential ingredient to account for the recent U.S. recession. Mian and Sufi (2011a, 2011b) use cross-state evidence to argue that the contraction in households' borrowing capacity, mainly driven by a de-

[2]Classic models of the role of firms' balance sheets are Kiyotaki and Moore (1997) and Bernanke, Gertler, and Gilchrist (1999). Recent contributions include Jermann and Quadrini (2009), Brunnermeier and Sannikov (2010), Gertler and Kiyotaki (2010), Khan and Thomas (2010), Buera and Moll (2011), Del Negro, Eggertsson, Ferrero, and Kiyotaki (2011), Cagetti, De Nardi, and Bassetto (2011). Goldberg (2011) is a model that combines financial frictions on both the firms' and the households' side, but focuses on steady states.

cline in house prices, was responsible for the fall in consumer spending and, eventually, for the increase in unemployment. Our model aims to capture the effects of a similar contraction in households' borrowing capacity in general equilibrium.

In modeling the household sector, we follow the vast literature on consumption and saving in incomplete market economies with idiosyncratic income uncertainty, going back to Bewley (1977), Deaton (1991), Huggett (1993), Aiyagari (1994), Carroll (1997).[3] Our approach is to compute the economy's transitional dynamics after a one-time, unexpected aggregate shock. This relates our paper to recent contributions that look at transitional dynamics after different types of shocks.[4] Much work on business cycles in economies with heterogenous agents and incomplete markets, follows Krusell and Smith (1998) and looks at approximate equilibria in which prices evolve as functions of a finite set of moments of the wealth distribution. Here, we prefer to keep the entire wealth distribution as a state variable at the cost of focusing on a one time shock, because our shock affects very differently agents in different regions of the distribution.[5] Midrigan and Philippon (2011) take a different (and complementary) approach to modeling the effects of a credit crunch on the household sector. They use a cash-in-advance model to explore the idea that credit access, as money, is needed to facilitate transactions. Finally, our model with durables is related to Carroll and Dunn (1997), an early paper that uses an heterogenous agent, incomplete market model with durable and non-durable goods to look at the dynamics of consumer debt and spending following a shock to unemployment risk.

The modern monetary policy literature has pointed out that at the roots of a liquidity trap there must be a shock that sharply reduces the "natural" interest rate, that is, the interest rate that would arise in a flexible price economy (Krugman, 1998; Woodford and Eggertsson, 2001). In representative agent models, the literature typically generates a liquidity trap by introducing a shock to intertemporal preferences, which mechanically increase the consumer's willingness to save (e.g., Christiano, Eichenbaum, and Rebelo,

[3]Heathcote, Storesletten, and Violante (2009) offer an excellent review.

[4]For example, Mendoza, Rios Rull and Quadrini (2010) look at the response of an economy opening up to international asset trade.

[5]Heathcote, Storesletten, and Violante (2009) point out that the nature of the shock is important in determining whether or not an heterogenous agent economy behaves approximately as its representative agent counterpart.

2011). Our model shows that in a heterogenous agent environment, shocks to the agents' borrowing capacity can be the underlying force that pushes down the natural rate, by reducing the demand for loans by borrowers and by increasing the supply of loans by lenders. This is consistent with the fact that, historically, liquidity trap episodes have always followed disruptions in credit markets. Two independent recent papers, Curdia and Woodford (2010) and Eggertsson and Krugman (2012), draw related connections between credit crises and the liquidity trap. The main difference is that they work with a representative borrower and a representative lender and mute wealth dynamics to aim for analytical tractability.[6] This implies that there are no precautionary effects, i.e., no direct responses for agents who are not at the debt limit, and that there is no internal dynamics associated to the wealth distribution. As we shall see, in our model the dynamics of the wealth distribution play an important role in generating large swings in the natural interest rates.

Two papers that explore the effects of precautionary behavior on business cycle fluctuations are Guerrieri and Lorenzoni (2009) and Challe and Ragot (2011). Both papers, derive analytical results under simplifying assumptions that eliminate the wealth distribution from the problem's state variables. In this paper we take a computational approach, to study 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 the precautionary motive in a monetary environment and focuses on the possibility of multiple equilibria.

Since this paper was first circulated, there has been a growing body of work on the effects of a credit crunch on a highly indebted household sector. Justinano, Primiceri, and Tambalotti (2013) take a quantitative perspective and evaluate the effects of a leveraging and deleveraging cycle in a stochastic business cycle model, but restrict attention to a model with only two types of households. Huo and Rios Rull (2014) start from a Bewley model like the one used here, but enrich it in many dimensions, mostly introducing a frictional labor market with search and matching so as to obtain more realistic implications for employment and hours worked. Rognlie, Shleifer, and Simsek (2014)

[6] Iacoviello (2005) is an early paper that studies monetary policy in a two-types model where households borrow to finance housing purchases, facing a collateral constraint similar to that in our durable section.

consider a model with an explicit treatment of housing and residential investment, in which the overinvestment in housing during the boom causes a slow recovery after a credit crunch. There has also been work on the role monetary policy in similar environments such as Buera and Nicolini (2014) and on the role of macroprudential policies in mitigating the effects of a debt-driven liquidity trap in work by Farhi and Werning (2016) and Korinek and Simsek (2016).

The paper is organized as follows. In Section 2, we present our model and characterize the steady state. In Section 3, we perform our main exercise, that is, we analyze the equilibrium transitional dynamics after a tightening of the borrowing limit. Section 4 introduces nominal rigidities. Section 5 presents some alternative calibrations. Section 6 studies the effects of fiscal policy. Section 7 presents the model with durable consumption goods. Section 8 concludes.

2 Model

Consider an infinite horizon economy populated by a continuum of households who face uninsurable idiosyncratic income risk. The only asset traded is a one-period risk-free bond. Households can borrow up to an exogenous limit. We first analyze the steady state equilibrium for a given borrowing limit. Then, we study transitional dynamics following an unexpected, one time shock that reduces this limit.

Households' preferences are represented by the utility function

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

where c_{it} and n_{it} are consumption and labor effort of household i and $\beta \in (0,1)$ is the discount factor. Each household produces consumption goods using the linear technology

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

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

The household's budget constraint is

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

where b_{it} are bond holdings, q_t is the bond price and $\tilde{\tau}_{it}$ are taxes. Tax payments are as follows: all households pay a lump sum tax τ_t and the unemployed receive the unemployment benefit ν_t, that is, $\tilde{\tau}_{it} = \tau_t$ if $\theta_{it} > 0$ and $\tilde{\tau}_{it} = \tau_t - \nu_t$ if $\theta_{it} = 0$. Household debt is bounded below by the exogenous limit ϕ, that is, bond holdings must satisfy[7]

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

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

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

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

where $u = \Pr(\theta_{it} = 0)$ is the fraction of unemployed agents in the population. For now, we assume that the supply of government bonds and the unemployment benefit are kept constant at B and ν, while the tax τ_t adjusts to ensure government budget balance. In Section 6, we consider alternative fiscal policies.

In the model, the only supply of bonds outside the household sector comes from the government. When we calibrate the model, we interpret the bond supply B broadly as the sum of all liquid assets held by the household sector. The main deviation from Aiyagari (1994) and most of the following literature is the absence of capital in our model.[8] The standard assumption in models with capital is that firms can issue claims to physical capital that are perfect substitutes for government bonds and other safe and liquid stores of value. This would not be a satisfactory assumption here, since we are trying to capture the effects of a credit crisis. A more general model of a credit crisis would have to include the effects of the crisis on the ability of firms to issue financial claims and on their accumulation of precautionary reserves, and it would have to allow for imperfect substitutability between different assets.[9] Here, we choose to focus on the household

[7] The presence of the unemployment benefit ensures that the natural borrowing limit is strictly positive. We always set ϕ to values smaller than the natural borrowing limit.

[8] Huggett (1993) studies an endowment economy with no capital and no outside supply of bonds.

[9] Along the lines of models such as those mentioned in footnote 2.

sector and we close the model by taking as given the net supply of liquid assets coming from the rest of the economy, B. In Section 7, we enrich the household portfolio choice by allowing households to accumulate both bonds and durable goods, which are a form of capital directly employed by the households. In that setup, we will introduce imperfect substitutability between the two assets.

In our baseline model, the only motive for borrowing and lending comes from income uncertainty. In particular, we abstract from life-cycle considerations and from other important drivers of household borrowing and lending dynamics, like durable purchases, health expenses, educational expenses, etc. Moreover, we assume that there is a single interest rate r_t, which applies both to positive and negative bond holdings, so that household can borrow or lend at the same rate. In Section 7, we address some of these limitations, by modeling durable purchases and introducing a spread between borrowing and lending rates.

2.1 Equilibrium

Given a sequence of interest rates $\{r_t\}$ and taxes $\{\tau_t\}$, let $C_t(b,\theta)$ and $N_t(b,\theta)$ denote the optimal consumption and labor supply at time t of a household with bond holdings $b_{it} = b$ and productivity $\theta_{it} = \theta$. Given consumption and labor supply, next period bond holdings are derived from 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 yields a transition probability for the individual states (b,θ). This transition probability determines the distribution Ψ_{t+1}, given the distribution Ψ_t. 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 $\{\Psi_t\}$ such that, given the initial distribution Ψ_0:*

(i) *$C_t(b,\theta)$ and $N_t(b,\theta)$ are optimal given $\{r_t\}$ and $\{\tau_t\}$,*

(ii) Ψ_t *is consistent with the consumption and labor supply policies,*

(iii) *the tax satisfies the government budget constraint,*

$$\tau_t = \nu u + r_t B / (1 + r_t),$$

(iv) *the bonds market clears,*

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

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

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

which holds with equality if the borrowing constraint $b_{it+1} \geq -\phi$ is slack. And the optimality condition for labor supply

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

which holds with equality if $n_{it} > 0$.

As we will see below, a tightening of the borrowing limit makes future consumption more responsive to income shocks, so that agents face higher future volatility. With prudence in preferences, this implies that the expected marginal utility on the right-hand side of (2) is higher, by Jensen's inequality. Therefore, 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. 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.

2.2 Calibration

We analyze the model by numerical simulations, so we need to specify preferences and choose parameter values. We assume the utility function is separable and isoelastic in consumption and leisure and we normalize the time endowment to 1, so we specify

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

Table 1: *Parameter Values*

Parameter	Explanation	Value	Target/Source
β	Discount factor	0.9711	Interest rate $r = 2.5\%$
γ	Coefficient of relative risk aversion	4	
η	Curvature of utility from leisure	1.5	Average Frisch elasticity = 1
ψ	Coefficient on leisure in utility	12.48	Average hours worked 0.4 of endowment (Nekarda and Ramey, 2010)
ρ	Persistence of productivity shock	0.967	Persistence of wage process in Floden and Lindé (2001)
σ_ε	Variance of productivity shock	0.017	Variance of wage process in Floden and Lindé (2001)
$\pi_{e,u}$	Transition to unemployment	0.057	Shimer (2005)
$\pi_{u,e}$	Transition to employment	0.882	Shimer (2005)
ν	Unemployment benefit	0.10	40% of average labor income
B	Bond supply	1.6	Liquid assets (Flow of funds)
ϕ	Borrowing limit	0.959	Total gross debt (Flow of funds)

Note: See the text for details on the targets.

Our baseline parameters are reported in Table 1. The time period is a quarter. The discount factor β is chosen to yield a yearly interest rate of 2.5% in the initial steady state. The coefficient of risk aversion is $\gamma = 4$. Clearly, this coefficient is crucial in determining precautionary behavior, so we will experiment with different values. The parameter η is chosen so that the average Frisch elasticity of labor supply is 1. The parameter ψ is chosen so that average hours worked for employed workers are 40% of their time endowment, in line with the evidence in Nekarda and Ramey (2010).[10] As we shall see, the value of ψ is relevant in determining the shape of wealth effects on labor supply and thus the model's implications for employment.

The average level of θ is chosen so that yearly output in the initial steady state is normalized to 1. The remaining moments of the θ process are chosen to capture wage and employment uncertainty. We assume that, when positive, θ follows an AR1 pro-

[10]See Prescott (2004) for a similar calibration of that parameter. Figure 1 in Nekarda and Ramey (2010) shows about 39 weekly hours per worker in 2000-2008. Subtracting 70 hours per week for sleep and personal care from a time endowment of 168 hours, we obtain $39/98 = 0.40$.

cess in logs with autocorrelation ρ and variance σ_ε^2. The parameters ρ and σ_ε^2 are chosen to match 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, our parameters yield a coefficient of autocorrelation of 0.9136 and a conditional variance of 0.0426 for yearly wages, matching the same moments of the persistent component of their wage process.[11] The wage process is approximated by a 12-state Markov chain, following the approach in Tauchen (1986). 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 then choose transition probabilities equal to 0.057 from employment to unemployment and equal to 0.882 from unemployment to employment. When first employed workers draw θ from its unconditional distribution. For the unemployment benefit ν, we also follow Shimer (2005) and set it to 40% of average labor income.

Finally, we choose values for the bonds supply B and the borrowing limit ϕ to reflect U.S. households' balance sheets in 2006, before the onset of the financial crisis. Defining liquid assets broadly as the sum of all deposits plus securities held directly by households, the liquid assets to GDP ratio in 2006 was equal to 1.78.[12] We choose B to match this ratio, computing liquid assets as the sum of households' *positive* bond holdings.[13] Second, we match debt in our model to consumer credit, which was 18% of GDP in 2006.[14] We choose ϕ to match this ratio, computing debt as the sum of households' negative bond holdings. The value of ϕ that we obtain in this way is equal to about one year of average income.

[11] See Table IV in Floden and Lindé (2001). Since our wage process is quarterly, we use the fact that the variance and autocovariance of the yearly-average of a quarterly AR1 process are given by the following expressions

$$\frac{1}{4^2}\left(4+6\rho+4\rho^2+2\rho^3\right)\frac{\sigma_\varepsilon^2}{1-\rho^2},$$

$$\frac{1}{4^2}\left(\rho+2\rho^2+3\rho^3+4\rho^4+3\rho^5+2\rho^6+\rho^7\right)\frac{\sigma_\varepsilon^2}{1-\rho^2},$$

and match them to the corresponding yearly moments.

[12] Federal Reserve Board Flow of Funds (Z.1) table B.100, sum of lines 9, 16, 19, 20, 21, 24, and 25.

[13] Since gross debt is calibrated at 0.18, setting $B = 1.6$ yields gross positive asset holdings equal to 1.78.

[14] Also in table B.100, line 34, which essentially corresponds to total household liabilities minus mortgage debt.

Figure 1: *Optimal Consumption and Labor Supply in Steady State*

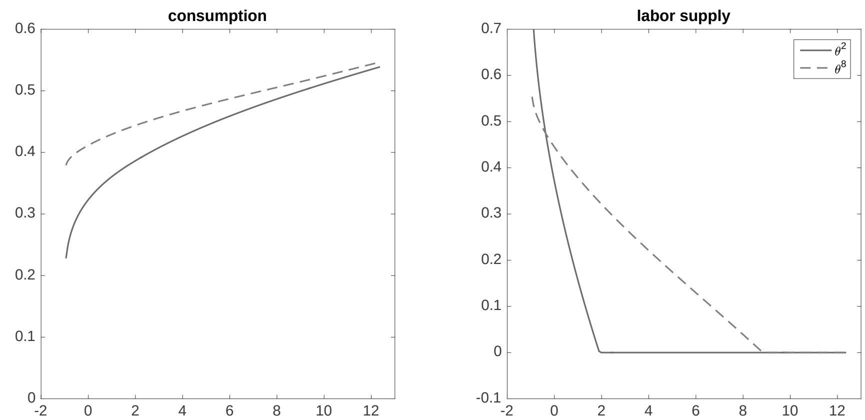

2.3 Steady state

To conclude this section, we briefly describe the household policies in steady state. Figure 1 shows the optimal values of consumption and labor supply as a function of the initial level of bond holdings, for two productivity levels, the lowest positive productivity level θ^2 (solid line) and the average productivity level θ^8 (dashed line).

Different responses at different levels of bond holdings are apparent. At high levels of b, consumer 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 (Carroll and Kimball, 1996). The optimality condition for labor supply implies that labor supply is a non-increasing function of consumption. So the relation between labor supply and bond holdings is non-increasing and the values of ψ and η determine the shape of this relation. Our baseline calibration, yields a convex labor supply function. So labor supply is steeply decreasing in b for low levels of b. For b large enough labor supply hits a corner at 0. As we will see, the shape of this function matters for the model's predictions regarding the aggregate response of employment to a credit crunch. Finally, the comparison between labor supply curves for different θs reflects both substitution and income effects at work. For most levels of b the substitution effect dominates the income effect and higher wages are associated to higher labor supply. For very low levels of

b, however, the income effect dominates and low wage households supply more hours than high wage households.

3 Credit Crunch

We now explore the response of our economy to a credit crunch. We consider an economy that starts at $t = 0$ in steady state with the borrowing limit $\phi = 0.959$. We then look at the effects of an unexpected shock at $t = 1$ that gradually and permanently decreases the borrowing limit to $\phi' = 0.525$. The size of the shock is chosen so that the debt-to-GDP ratio drops by 10 percentage points in the new steady state.

Starting at $t = 1$, the borrowing limit ϕ_t follows the linear adjustment path

$$\phi_t = \max\left\{\phi', \phi - \Delta\phi \cdot t\right\},$$

and households perfectly anticipate this path. We choose $\Delta\phi$ so that the adjustment lasts 6 quarters. Since all debt in the model has a one-quarter maturity, a sudden adjustment in the debt limit would require unrealistically large repayments by the most indebted households. An assumption of gradual adjustment of the debt limit is a simple way of capturing the fact that actual debt maturities are longer than a quarter, so that after a credit crunch households can gradually pay back their debt. An adjustment period of 6 quarters ensures that no household is forced into default. Default and bankruptcy are clearly an important element of the adjustment to a tighter credit regime, but are beyond the scope of this paper.

Before looking at transitional dynamics, let us briefly compare steady states. In Figure 2 we plot the aggregate bond demand in the initial steady state (solid line) and in the new steady state (dashed line). Two effects contribute to shifting the demand curve to the right. First there is a mechanical effect, as households with debt larger than ϕ' need to reduce their debt. Second there is a precautionary effect, as households accumulate more wealth to stay away from the borrowing limit. As the supply of bonds is fixed at B, the shift in bond demand leads to a lower equilibrium interest rate.

Figure 2: *Bond Market Equilibrium in Steady State*

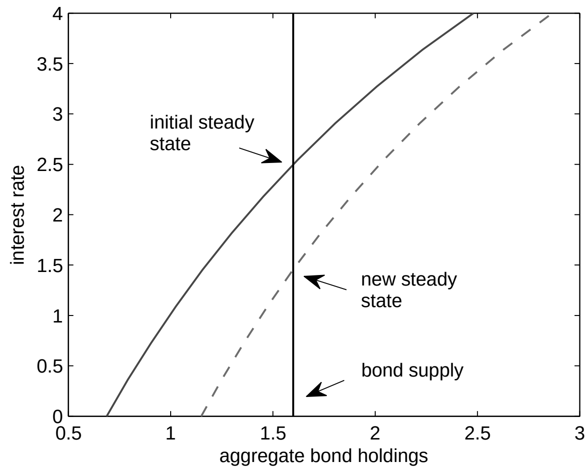

Note: Interest rate is in annual terms.

3.1 Transitional dynamics: interest rate

Figure 3 illustrates the economy's response to the debt limit contraction. In the top left panel, we plot the exogenous adjustment path for ϕ_t. The remaining panels show the responses of the debt-to-GDP ratio (top right panel), the interest rate (bottom left panel), and output (bottom right panel).

The interest rate drops sharply after the shock, going negative for the first 5 quarters. The interest rate overshooting after a debt contraction is our first main result. From numerical experiments, this result seems a fairly general qualitative outcome of this class of models and not just the consequence of our choice of parameters. To provide some intuition, we look at some properties of the household policy functions and of the steady state distributions that help explain the result.

Let us first look at the policy functions. The top panel of Figure 4 plots the optimal bond accumulation $b_{it+1} - b_{it}$ (averaged over θ) as a function of the initial bond holdings b_{it}, at the initial steady state (solid blue line) and at the new steady state (dashed red line). The function is decreasing and convex. The steeper portion, for low levels of b, reflects the strong incentives to save for households at the left tail of the distribution, who want to move away from their borrowing limit. Notice that the convexity of the

Figure 3: *Interest Rate and Output Responses*

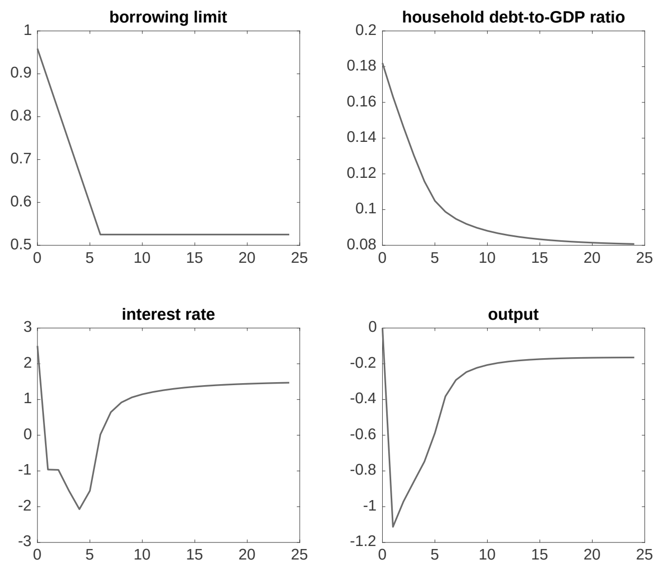

Note: Interest rate is in annual terms. Output is in percent deviation from initial steady state.

bond accumulation function follows from the budget constraint, the concavity of the consumption function and the convexity of the labor supply function (see Figure 1).

Consider next the stationary bond distributions. The bottom panel of Figure 4 shows the marginal density of bond holdings at the initial steady state (solid blue line) and at the new steady state (dashed red line). The two distributions have the same average, as the bond supply is the same in the two steady states, but the new distribution is more concentrated.[15] A comparison of the policies in the top panel helps to explain why. 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.

We are now ready to put the pieces together. In equilibrium, aggregate net bond accumulation must be zero as the bond supply is fixed. In steady state, this means the

[15]Formally, the initial distribution is a mean-preserving spread of the new distribution. We checked this property numerically plotting the integral of the CDF of b for the two distributions and comparing them at each value of b.

Figure 4: *Bond Accumulation and Distributions in the Two Steady States*

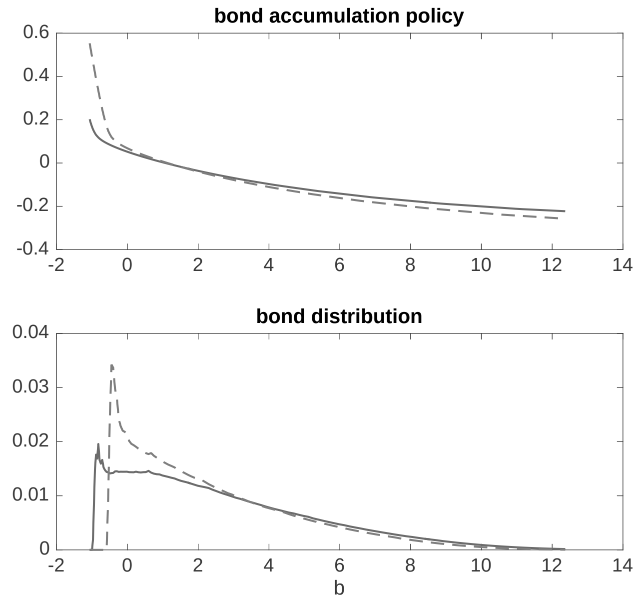

integral of the solid (dashed) function in the top panel weighted by the solid (dashed) density in the bottom panel is equal to zero. Let us make a "disequilibrium" experiment: suppose that instead of following its equilibrium transition path the interest rate jumped directly to its new steady state value at $t = 1$ and stayed there from then on. Average bond accumulation could then be computed by integrating the dashed function in the top panel weighted by the solid density in the bottom panel. This gives a positive number, because the bond accumulation function is convex and the solid distribution is a mean-preserving spread of the dashed one. Therefore, at the conjectured interest rate path, there is excess demand of bonds and we need a lower interest in the initial periods to equilibrate the bonds market. Intuitively, the economy begins with too many households at low levels of debt, with a strong incentive to save. This is not compensated by the presence of households at high wealth levels, who tend to decumulate assets, due to the convexity of the bond accumulation function. Therefore, the net effect is to push down equilibrium interest rates. As the economy reaches its new steady state, the lower tail of the distribution converges towards higher levels of bond holdings, the saving pressure subsides, and the interest rate goes back up.

3.2 Transitional dynamics: output

Next, we want to understand what happens to output. The bottom right panel of Figure 3 shows that output contracts by 1.1% on impact and then recovers, converging towards a level below the initial steady state.

The output response depends on the combination of consumption and labor supply decisions, with the interest rate acting as the equilibrating price. In partial equilibrium, if the interest rate does not adjust, a contraction in the credit limit leads to lower consumption demand and to higher output supply, as households adjust to the tighter limit by spending less and working more. The drop in the interest rate equilibrates the goods market by increasing consumption and by lowering labor supply, via intertemporal substitution channels. The market clearing level of output can then, in general, be above or below its new steady state level, depending on whether the adjustment is more on the consumption side or on the labor supply side. Given our chosen parameters, the consumption side dominates, leading to a contraction in output.

In the rest of this section, we explore in more detail the consumption and employment responses.

3.2.1 Consumption dynamics

It is useful to decompose the consumption response in two parts: the partial equilibrium response to the debt limit shock—keeping the interest rate at the initial steady state level—and the response to the endogenous change in the interest rate path. The decomposition is presented in Figure 5. The solid line replicates the total response in Figure 3, the dashed line is the partial equilibrium response to the debt limit shock and the dotted line is the response to the interest rate changes.

Let us look first at the partial equilibrium response, to understand the forces at work. Simulations show that, for a given interest rate, a reduction in the debt limit leads to an approximately uniform horizontal shift of the consumption function to the right, by an amount approximately equal to the reduction in the debt limit. That is, all agents, not only those at the constraint, behave as if they had experienced a wealth reduction equal to the contraction in the debt limit. The horizontal shift in the consumption function implies that the partial equilibrium reduction in consumption driven by a $d\phi$ change in the

Figure 5: *Consumption Response Decomposition*

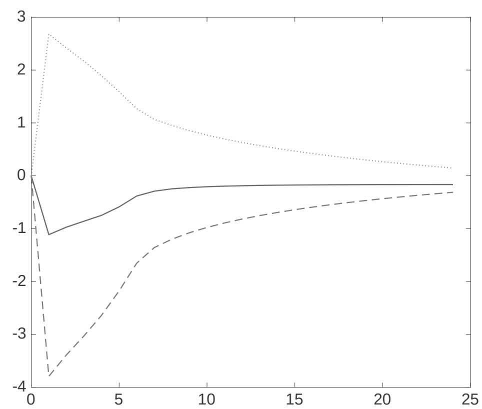

Note: Percent deviations from initial steady state. Solid line: general equilibrium response. Dashed line: partial equilibrium response to debt limit reduction. Dotted line: response to the equilibrium sequence of interest rate changes.

debt limit is approximately equal to $MPC \cdot d\phi$, where MPC is the propensity to consume out of a one time transfer.[16] In our experiment the total change in ϕ is equal to 0.43 and the MPC is equal to 0.023. The partial equilibrium change in consumption is equal to $0.0095 \approx 0.023 \times 0.43$. Since the initial value of quarterly consumption is equal to 0.25 (as we normalize output in the initial steady state to 1), we have a partial equilibrium contraction in consumption of $0.0095/0.25 = 3.8\%$. The fact that all households, not just those at the debt limit, respond to the shock, is a distinctive feature of our modeling approach, relative to more stylized models of household deleveraging that simply assume two groups of households, one of which is exactly at the constraint, as in Eggertsson and Krugman (2012). This also has the advantage that the calibration is much less sensitive to assumptions about the mass of agents who are at the constraint.[17]

The dotted line in Figure 5 shows that the general equilibrium effect of lower interest rates is strong and dampens substantially the effect of the credit crunch on consumption. A drop in the interest rate of 4.5 percentage points in the short run and of 1 percentage

[16] We thank Adrien Auclert for pointing out this relation. An analytical result that explains this relation is Proposition 1 in the Online Appendix.

[17] In fact, in our baseline calibration only 1% of agents start exactly at the constraint.

point in the long run, leads to an increase in aggregate consumption of about 2.7 percentage points on impact.[18]

This discussion highlights that the *MPC*s and the interest elasticity of consumption are important elements to determine the quantitative impact of a credit tightening. In particular, our baseline *MPC* is very low when compared to empirical estimates such as in Johnson, Parker and Soulelels (2006).[19] In Section 5 below, we experiment with alternative calibrations that feature higher *MPC*s and lower interest elasticities and show that, combined with nominal rigidities and the zero lower bound, they produce much larger consumption contractions.

Turning to the cross-sectional predictions of the model, let us look at the responses of consumers who start with different liquid wealth holdings. In Figure 6, we plot consumption responses for five groups. The first group includes only consumers at the debt limit in the initial steady state, which corresponds to the first 1% of the initial distribution. The other three groups are the 10th, 20th and 50th percentiles of the initial wealth distribution. The concavity of the consumption implies that *MPC*s are higher for consumers with lower initial wealth. Therefore, the partial equilibrium response is larger for those consumers. Our simulations also show that the response to the endogenous interest rate reduction is stronger for the consumers with higher wealth.[20] The net effect of these differential responses is that lower wealth consumers experience a large reduction in consumption, while higher wealth consumers experience a moderate increase. These cross-sectional predictions of the model are qualitatively in line with evidence by Heathcote and Perri (2015). They use Consumer Expenditure Survey data to show that consumers with lower ratios of wealth to permanent income did experience a larger contraction in consumer expenditure during the 2008 recession.

[18] We also performed related exercises, by looking at the effects of a temporary monetary policy shocks in the context of the sticky wages model of Section 4, and obtained large elasticities of aggregate consumption to temporary interest rate shocks (in the range of 0.35-0.5, depending on the initial condition). These responses seem large, but it is useful to remark that they not only embed the response of consumers to interest rate changes, but also the endogenous responses of income. There is a growing literature on the effect of interest rates on consumer spending in heterogeneous agents economies, including Auclert (2015) and Wong (2016). Werning (2016) emphasizes the importance of taking into account endogenous income responses.

[19] Estimates in the recent empirical literature all range near 0.2.

[20] The decomposition by percentile is not reported for reasons of space.

Figure 6: *Consumption Response by Percentile in Initial Wealth Distribution*

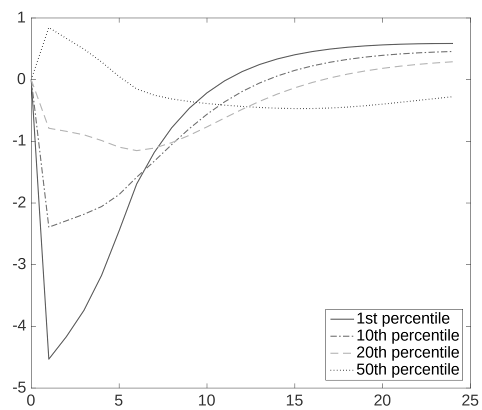

Note: Percent deviations from steady state path conditional on initial wealth being in the reported percentile.

3.2.2 Employment and output dynamics

The response of employment is also driven by partial equilibrium and general equilibrium effects. However, to connect the labor supply response to the output response, we also need to consider compositional effects, namely how labor supply responses are distributed across workers with different productivity.

Symmetrically to what happens to consumption, the partial equilibrium effect of the reduction in the debt limit is to increase labor supply, as workers increase work effort to reduce their debt or increase their savings. The reduction in the interest rate has an opposing effect, as it leads to intertemporal substitution leading to a reduction in labor effort today. In our baseline calibration, the first effect dominates and total hours go up, as illustrated in Figure 7. However, the compositional effect is sufficiently strong that the total increase in hours is actually associated to a decrease in total output as seen in Figure 3. This is due to the fact that hours worked increase for low-productivity workers at the bottom end of the bond distribution, who are closer to the borrowing limit and are least interest-sensitive, while hours worked drop for high-productivity workers with high bond holdings, who are farther from the debt limit and are more interest-sensitive. So behind the drop in output there is a compositional effect and a drop in average labor

Figure 7: *Employment Response*

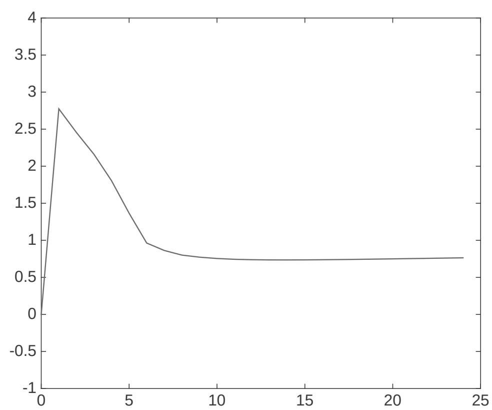

Note: Percent deviations from initial steady state.

productivity.[21]

The prediction of an aggregate increase in hours worked is in part due to the fact that we have introduced no frictions on the supply side of the model. A first step in this direction will be to introduce nominal wage rigidities and the zero lower bound in the next section. As we shall see, that will not be enough to produce a contraction in employment in our baseline calibration, but it will once we experiment with alternative calibrations in Section 5.

4 Nominal Rigidities and the Zero Lower Bound

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 the nominal interest rate. However, the presence of the zero lower bound implies that the central bank may not be able to replicate the real interest rate path corresponding to the flexible price equilibrium. Therefore, a credit crisis which

[21]This compositional effect is closely related to the steady state labor misallocation analyzed in Heathcote, Storesletten, and Violante (2008).

produces a large drop in real interest rates under flexible prices can drive the economy into a liquidity trap and into a recession under sticky prices.

To introduce nominal rigidities, we consider a simple model with nominal wage rigidity. Namely, we assume that nominal wages are fixed at W. We interpret the shock θ_{it} as a shock to the efficiency of household i's labor and assume that workers are hired by competitive firms that produce consumption with a linear technology. Therefore, constant nominal wages translate into a constant nominal price level. To clear the labor market we introduce a wedge in labor supply decisions, which is a simple way of capturing a labor market friction that reduces the labor input in response to low aggregate demand for goods. In particular, we denote the wedge with ω_t and replace the optimality condition (3) with the optimality condition:

$$(1-\omega_t)\theta_{it}U_c(c_{it}, n_{it}) + U_n(c_{it}, n_{it}) \leq 0. \tag{4}$$

The budget constraint and the optimality condition for bond holdings are unchanged.

How does an equilibrium with fixed wages work? Since the price level is constant, the nominal interest rate is equal to the real interest rate. The central bank, by choosing a sequence of nominal interest rates, chooses a sequence of real rates r_t. We assume that the central bank sets the interest rate r_t to replicate a flexible price allocation whenever possible, that is, to reach an allocation with $\omega_t = 0$. The only constraint for the central bank is the zero lower bound, i.e., the interest rate cannot go negative. Therefore, we define an equilibrium in this section as given by two sequences $\{r_t, \omega_t\}$ such that $r_t \geq 0$, $\omega_t \geq 0$, and at least one of these two conditions hold as an equality.[22] The remaining equilibrium conditions are as in Definition 1.

The assumption of fixed wages is clearly an extreme form of nominal rigidity and it is made here to simplify the analysis. Notice however that at the zero lower bound, the assumption of fixed prices is actually a conservative assumption, as the deflation triggered by output below its natural level has an amplifying effect in the standard new Keynesian model, given that deflation leads to a lower real interest rate (as the nominal rate is unchanged at zero). This amplifying effect is well understood in the recent literature on the zero lower bound and we simply keep it muted here by having constant

[22] The reasoning behind this definition is that whenever $\omega_t > 0$ and $r_t > 0$ the central bank can lower the rate r_t, increase output and employment, and thus decrease ω_t.

prices.[23] Another amplifying effect of deflation ignored here has to do with Fisher's debt deflation channel, which we analyze in Section 6.

There are alternative ways of incorporating nominal rigidities in the model. In the working paper version of this paper (Guerrieri and Lorenzoni, 2011), we use monopolistic competition and sticky prices. The reason why we use sticky wages here is because under sticky prices the presence of firms' monopoly profits introduce firms' ownership shares as an additional asset. Moreover, since real wages need to fall in a recession, the value of this asset increases automatically in recessions, a mechanism that we don't find plausible. As it turns out, for the exercises conducted here, the choice of the form of nominal rigidity does not affect the results. There are also different ways of dealing with rationing in the labor market in a demand-determined model. Here we assume that labor is reallocated so that all workers face the same wedge.[24]

Figure 8 shows what happens to interest rates and output under fixed wages. The solid line is the flexible price baseline. The dashed line is the equilibrium with fixed wages. The presence of the zero lower bound implies that consumption, and thus output, drop more when the shock hits. This also implies that the adjustment of the wealth distribution is slower, since incomes are lower in the short run, which slows down bond accumulation for poorer households. We can see the effects of this slower wealth adjustment in the interest rate dynamics: in period 7 the interest rate would be positive in the flexible price regime, but is still stuck at zero in the fixed wage regime. The presence of the zero lower bound slows down the deleveraging process for the household sector, but this slow adjustment is not a positive symptom, since it comes from depressed incomes, not from a milder consumption drop.[25]

Notice that in our baseline calibration the effects of the zero lower bound are not large. Moreover, the response of aggregate employment is still positive, as in the flexible price case. In the next section we explore alternative calibrations that yield larger consumption contractions and also employment reductions.

[23] See, e.g., Eggertsson and Woodford (2003).

[24] Werning (2016) uses a proportional rationing rule.

[25] The slow adjustment of the wealth distribution as a source of propagation also distinguishes our exercise from Eggertsson and Krugman (2012), where the adjustment takes place in one period.

Figure 8: *Responses with Fixed Wages*

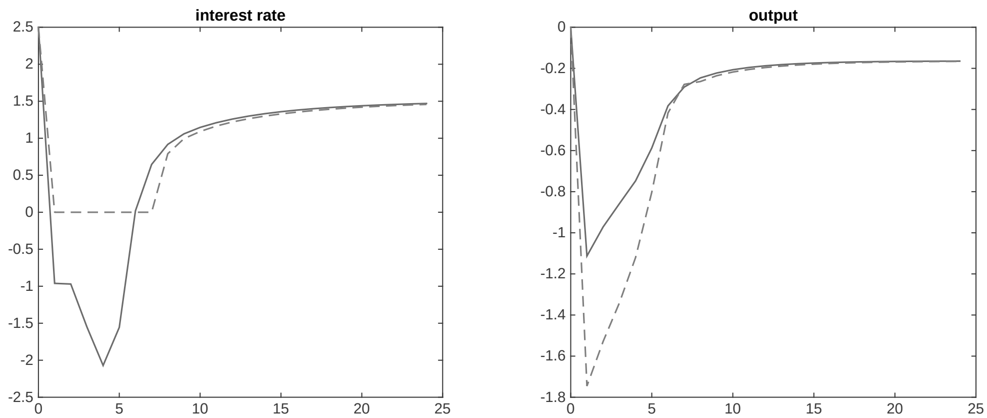

Note: Solid line: flexible price economy. Dashed line: economy with fixed wages. Interest rate in annual terms. Output in percent deviation from initial steady state.

5 Alternative Calibrations

We now consider some alternative calibrations that illustrate how the model's aggregate implications are sensitive to parameter choices. The parameters of all the calibrations in this section are in the Online Appendix (Tables 1 and 2).

5.1 Targeting median liquid wealth

First, we consider what happens when we target a different moment of the wealth distribution. The baseline model, as is common in this class of models, does not match well the fat right tail of the empirical wealth distribution.[26] An implication of this is that if we calibrate B to target average wealth, median households wealth in the model is unrealistically large. Therefore, our first experiment is to target median liquid wealth instead of average liquid wealth. Namely, we match median holdings of liquid assets of U.S. households from the 2001 Survey of Consumer Finances (SCF). We define liquid assets as the sum of all deposits plus securities held directly by households. Median liquid assets holdings are \$2,726 and taking the ratio of this value to average labor income in

[26]Carroll (1997) and Quadrini (1999) discuss this issue focusing on the distribution of total (liquid and illiquid) wealth. A similar issue arises if we restrict attention to the distribution of liquid wealth.

Figure 9: *Interest Rate and Output: Median Wealth Calibration*

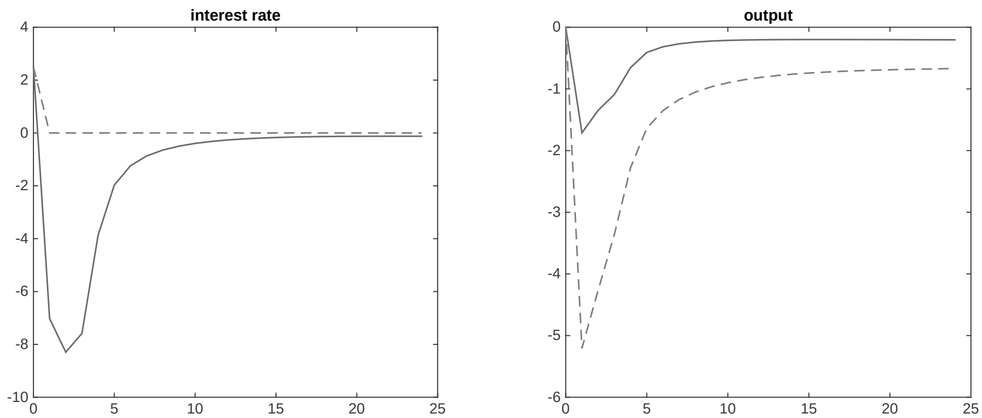

Note: Solid line: flexible price economy. Dashed line: economy with fixed wages. Interest rate in annual terms. Output in percent deviation from initial steady state.

the same sample yields 5.2%. The value of median bond holdings over average output in the model is matched to this number.[27]

The results for this alternative calibration, under flexible prices (solid lines) and under fixed wages (dashed lines), are reported in Figure 9. The flexible price interest path drops now much more sharply below zero after the shock and converges to a negative number close to zero in the long run.[28] This implies that the zero lower bound has more dramatic effects on consumption dynamics, leading to a much larger and persistent output drop compared to our baseline calibration. An important difference between this calibration and the baseline is that, by targeting lower wealth levels, this calibration implies that households are more impatient (lower β). This yields a higher average MPC equal to 0.0422, getting us closer to available empirical estimates.[29] The discussion in Section 3.2.1 helps to explain why larger MPCs lead to larger consumption responses, especially when the interest rate is not allowed to adjust downard.

[27] A similar calibration strategy for liquid wealth is adopted in Kaplan and Violante (2014).

[28] In this simulation the long run interest rate is close enough to zero that the long run implications with or without nominal rigidities are indistinguishable. On negative rates in the long run, see discussion on page 27.

[29] This is still lower than available estimates. Kaplan and Violante (2014) propose a model that reconciles Beweley-Ayiagari models with high observed MPCs.

Figure 10: *Interest Rate, Output, Employment: Low ψ Calibration*

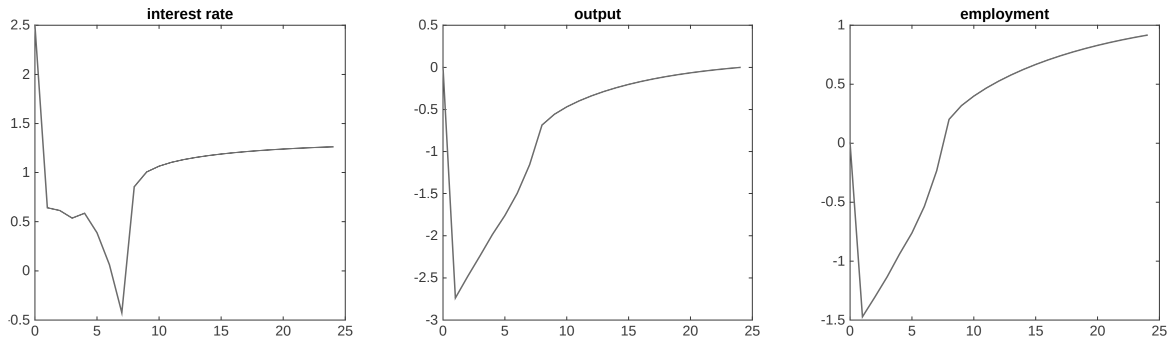

Note: Interest rate is in annual terms. Output is in percent deviation from initial steady state.

5.2 Exploring employment responses

Next, we want to better understand the determinants of the labor supply response reported in Figure 7. In particular, we focus on the shape of the labor supply policy function. In our baseline calibration this function is convex (see Figure 1), which implies that households at low levels of wealth respond to a negative wealth shock with a large increase in labor supply. In Figure 10 we plot the interest rate, output and employment for an alternative calibration in which the labor supply policy is concave at low levels of wealth. The parameters for this calibration are obtained reducing the maximum time endowment used to calibrate ψ. In particular, we assume that the weekly endowment is of 50 hours instead of 98 hours, so our target for hours worked is 0.8 instead of 0.4 of the time endowment (see footnote 10). Figure 1 in the Online Appendix shows that this calibration delivers a concave labor supply policy function for low levels of wealth. The reason is simply that poor agents are closer to the maximum time endowment, so they have limited room to respond with an increase in hours.

Figure 10 shows that this calibration delivers an employment reduction together with the output reduction. The compositional effect is still present, so average productivity decreases, although less than in the baseline. Notice however that this calibration also affects the model implications for how consumption and output respond to interest rates, so in equilibrium the model delivers a much lower interest rate response and the natural interest rate only goes below zero for one period. Therefore, it is interesting

Figure 11: *Interest Rate, Output, Employment: Median Wealth and Low ψ Calibration*

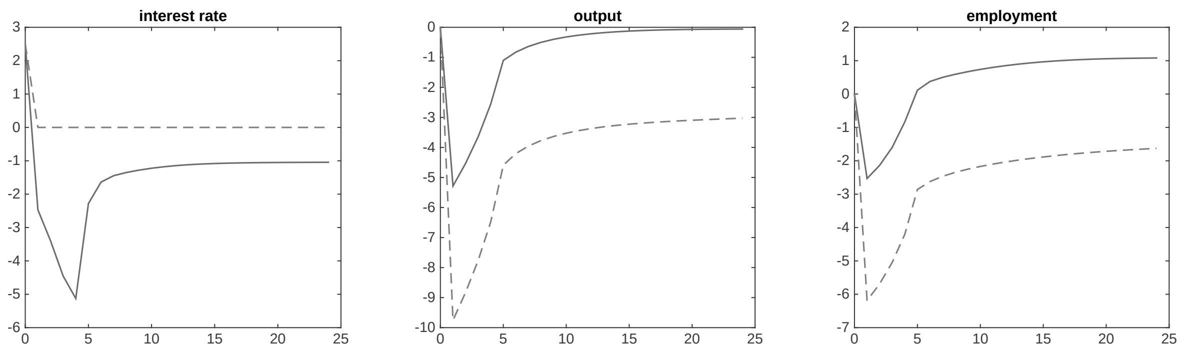

Note: Solid line: flexible price economy. Dashed line: economy with fixed wages. Interest rate is in annual terms. Output is in percent deviation from initial steady state.

to also consider a calibration that combines the two ingredients discussed so far in this section, both targeting a lower wealth level and a lower value of ψ. The results for this calibration are plotted in Figure 11.

In this calibration the permanent credit contraction leads to a natural interest rate that is permanently negative, so this economy is affected by what Summers recently dubbed "secular stagnation." A precautionary saving model without capital can easily produce a negative natural rate in the long run, but the full analysis of its implications for the secular stagnation hypothesis is outside the scope of this paper.[30]

5.3 Risk aversion

To further investigate the consumption response, it is useful to experiment with different values of γ. Figure 12 shows the behavior of the interest rate and output in the baseline with $\gamma = 4$ (solid lines) and in an alternative calibration with $\gamma = 6$ (dashed lines).[31] Different effects are at work here. On the one hand, the effect of higher risk aversion is to make the precautionary effect weaker and thus the consumption policy more concave. On the other hand, higher risk aversion also implies that consumers tend to borrow less in the initial steady state, so the initial distribution displays fewer agents

[30]Eggertsson and Mehrotra (2014) is the first paper to model the secular stagnation hypothesis using an overlapping generations model.

[31]Calibrated parameters are reported in Table 2.

Figure 12: *Interest Rate and Output Responses:* $\gamma = 6$ Calibration

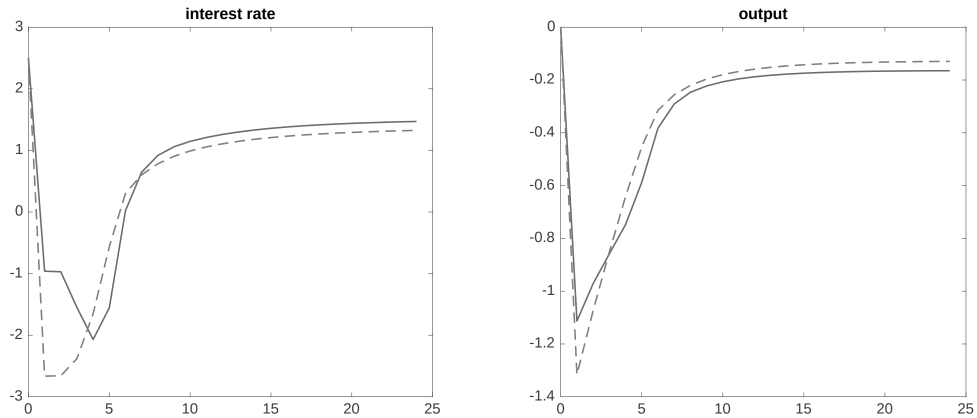

Note: Interest rate is in annual terms. Output is in percent deviation from initial steady state.

near the borrowing limit. These two effects go in opposite directions, the first increasing and the second decreasing the initial shift in consumer demand. Finally, an increase in γ also implies a lower elasticity of intertemporal substitution, which implies that consumer demand is less interest rate elastic. The net effect is that we obtain a larger drop in the interest rate and output. However, since opposing effects are present the relation between γ and the initial drop in the interest rate and output is in general non-monotone.

6 Fiscal Policy

We now explore the role of fiscal policy in mitigating the recession. In particular, we focus on the effects of simple transfers financed by an increase in government debt. Such transfers have an effect in the economy analyzed here because Ricardian equivalence does not hold, due to heterogeneity, uninsurable income risk, and borrowing constraints.

Since we have lump sum taxation, in our environment an equivalence result holds between public and private supply of liquidity. Namely, there exists a sequence of lump sum taxes τ_t and government bond supplies B_t that exactly offsets changes in the bor-

Figure 13: *Fiscal Policy*

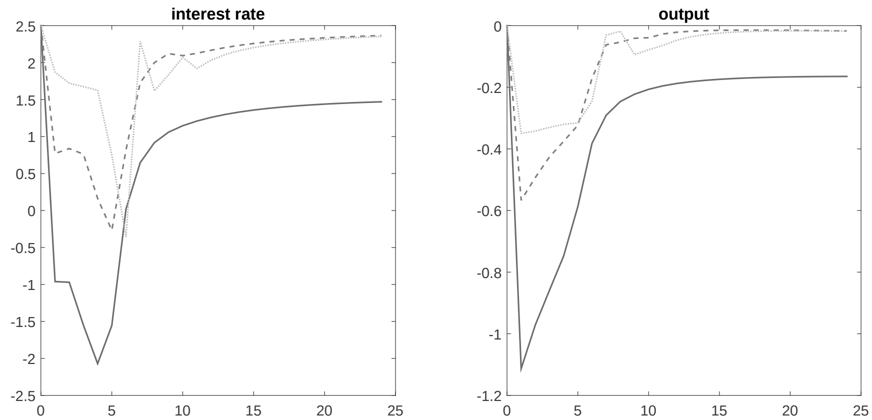

Note: Solid line: baseline. Dashed line: temporary reduction in lump sum tax. Dotted line: temporary increase in unemployment benefits. Interest rate is in annual terms. Output is in percent deviation from initial steady state.

rowing limit ϕ_t.[32] Basically, the government can neutralize the effects of a credit shock, by transferring to the agents resources equivalent to their loss in borrowing capacity and financing the transfer through government borrowing. This is a common result in this class of models and clearly derives from a number of simplifications, like abstracting from distortionary taxation and assuming government debt is always default-free. Therefore, 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 costs of a higher debt level in the long run.[33] The main objective of these experiments is to show that the way in which the transfers are implemented matters greatly for their impact.

In Figure 13, we compare two transfer policies, one in which the government temporarily reduces the lump sum tax τ for all households and the other in which the government temporarily raises the unemployment benefit. We fix the sequence of govern-

[32]The result is stated formally in Proposition 3 in the Online Appendix.

[33]Aiyagari and McGrattan (1998) study the trade off between distortionary taxation and the self-insurance benefits of government bonds in steady state.

Figure 14: *Deflation: Fisher Effect*

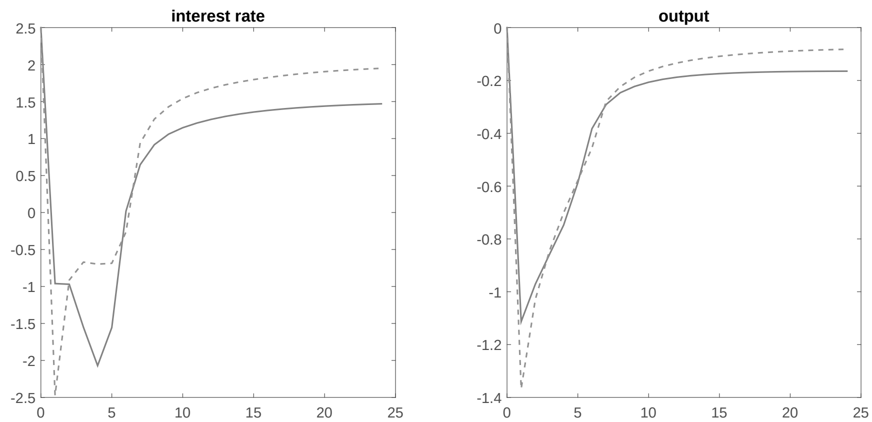

Note: Solid line: baseline. Dashed line: 10% deflation at $t = 0$. Interest rate is in annual terms. Output is in percent deviation from initial steady state.

ment deficits by assuming the bond supply follows the path

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

where B' is the new long run level of B_t and is 20% higher than in the initial steady state and where $\rho_b = 0.95$ determines the speed of adjustment of B_t. We then consider two different ways of spending the associated deficit. First, we consider a policy where the unemployment benefit is kept constant and the tax τ_t adjusts. Second, we consider an increase of the unemployment benefit by 50% for the first two years after the shock. Afterwards, the unemployment benefit reverts to its initial value and throughout the transition path τ_t adjusts to satisfy the government budget constraint. The solid lines in Figure 13 plot the baseline with no fiscal policy response, the dashed lines represent the case of a temporary transfer to all households and the dotted lines the case of a temporary increase in the unemployment benefit ν_t.

The figure shows that increasing the supply of government bonds dampens the responses of both interest rates and output. Moreover, increasing the unemployment benefit in the short run has larger effects than reducing the lump-sum tax, because it is a policy targeted towards households that are more likely to be credit constrained.

It is useful to notice that an increase in the stock of government debt per se is not necessarily helpful in this economy. A particularly bad way of increasing the real stock of government debt in this economy is the combination of nominal debt and deflation. To show this, in Figure 14 we present a simple experiment in which we assume that all debt (private and public) is in nominal terms and there is an unexpected 10% reduction in the price level at date $t = 0$. After that, the price level is constant. The effects of this deflation is to increase the real debt burden of the indebted households and revalue the bond holdings of the household with positive savings. This leads to a sharper contraction in the interest rate and output in the short run. The logic behind this figure is in line with Fisher's (1933) theory of debt deflation.

To conclude on monetary and fiscal policy, let us add that in this paper we do not analyze the effects of quantitative easing as a possible intervention. It would not be too difficult to add an explicit model of money demand to the model, which would provide an explicit justification for the zero lower bound. However, at the zero lower bound, exchanges of money for short term government bonds would be perfectly neutral, as the two assets would be perfect substitutes. Therefore, to introduce quantitative easing requires the introduction of additional assets that are less liquid than the one period safe bonds in the model and let the central bank exchange money (or short term bonds) for these less liquid assets. Such an extension is beyond the scope of this paper.

7 Durable Goods

In this section, we extend the model adding durable goods. A large part of household borrowing is associated to durable purchases and takes the form of secured debt, with durables as collateral. Therefore, a model with durables is more realistic in capturing both the motive for borrowing and the nature of the credit limit. The extended model is discussed briefly here and treated in detail in the Online Appendix.

A model with durables enriches the household portfolio decision. As durables offer an alternative store of value, when the 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. An opposing force is at work on the borrowers' side: reduced credit access implies that borrowers need to sell durables in order to reduce their debt. This leads to durable goods decumulation. Whether the force on the savers' side or on the borrowers' side dominates, depends on the model parameters and on the nature of the shock hitting the economy, as we will see shortly.

Households' portfolio decisions are also affected by the fact that durables are a less liquid form of savings than bonds. To capture the illiquidity of durable goods, we assume that households face a discount when re-selling durables. 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. The interesting finding here is that in a model with liquid and illiquid assets a credit shock can lead, at the same time, to an increase in demand for the liquid asset and to a reduction in demand for the illiquid asset.

In the extended model, households consume non-durables c_{it} and receive services, one for one, from the stock of durables k_{it}. Durables depreciate at the rate δ and the household incurs a proportional reselling cost $\zeta \cdot (k_{it} - k_{it+1})$ when it decides to reduce its durable holdings. The parameter ζ controls the degree of illiquidity of durable holdings. The borrowing constraint is now

$$b_{it+1} \geq -\phi_k k_{it+1}, \tag{5}$$

so all household debt is collateralized by durables. The parameter ϕ_k is the fraction of the value of the durable that can be used as collateral.

The extended model also features a spread between borrowing and lending interest rates. Specifically, if the household is a net seller of bonds ($b_{it+1} < 0$), the household needs to buy intermediation services from a competitive banking sector. The banking sector faces linear intermediation costs, at a rate of χ per unit of bonds issued. This implies that households receive a net price $\hat{q}_t = (1 - \chi)\, q_t$ per bond issued and banks make zero profits.

The production side of the model and the supply of bonds by the government are as in the benchmark model. Durable and non-durable goods are produced with the same technology, so the relative price of durables is 1.

Figure 15: *Responses to a Shock to the Borrowing Limit ϕ_k*

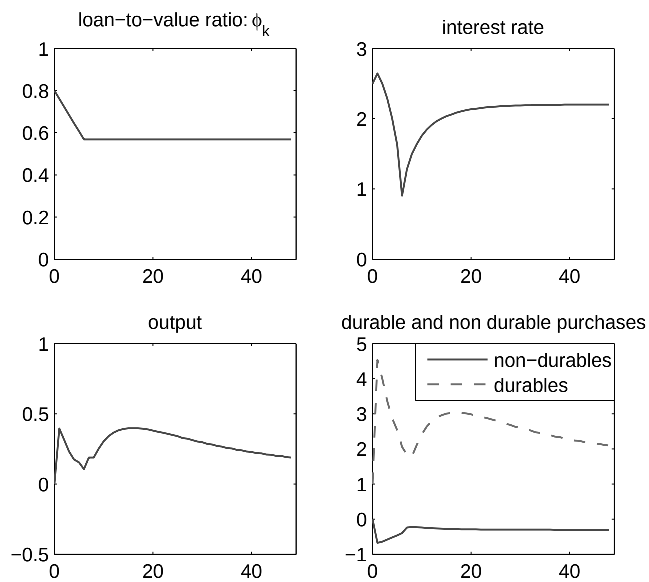

Note: Interest rate is in annual terms. Quantities are in percent deviation from initial steady state.

7.1 Credit crunch

The calibration parameters are presented in the Online Appendix. Here we show the outcomes of two different credit-tightening exercises, by looking at the effects of a permanent reduction in the borrowing limit ϕ_k and of a transitory increase in the spread χ. As in our baseline exercise, the aggregate shocks are unexpected and hit the economy in steady state.

Figure 15 shows the response to a permanent contraction in the borrowing limit ϕ_k from 0.8 to 0.56, which yields a 10 percentage points reduction in the household debt-to-GDP ratio from 54% to 44%. The contraction in ϕ_k is gradual and follows a linear path that lasts 6 quarters. The contraction in the interest rate is less strong than in our baseline exercise and output actually increases by a 0.4%. The reason behind these results is that durable purchases are very interest elastic. So a smaller interest rate reduction is sufficient to equilibrate the goods market and yields higher total spending in equilibrium. This is confirmed by the bottom right panel of Figure 15, which shows that there is a contraction in non-durable spending, similar in size to the contraction obtained in our baseline, but the contraction is more than compensated by a 4% increase in durable spending.

Numerical experiments show that this increase in durable spending is due to the endogenous reduction in the interest rate. A simple disequilibrium exercise shows that durable spending would drop by about 18% if the interest rate adjusted immediately to its new long run level, which is 2.2%. A short lived drop in the interest rate to 0.9% is sufficient to turn a 18% contraction in durable spending into a 4% increase. This may seem an unrealistically large interest elasticity of durable spending which indicates that in our model bonds and durables remain very good substitutes, despite the illiquidity cost. This points in the direction of extending the model using alternative specifications of the durables adjustment cost or accounting explicitly for the price risk associated with durable purchases (especially of housing), to reduce the substitutability between the two assets.[34] We leave these developments to future work.

Our second experiment is to look at a transitory increase in the intermediation cost by 6 percentage points at an annual rate. We assume the shock decays geometrically with a rate of decay of 0.6.[35] This implies that the rate on a 1 year loan goes up by about 3.9% in the first quarter after the shock. In constructing this shock we follow Hall (2011a) who argues that it is a reasonable representation of the credit shock in U.S. financial markets in 2008-2009. The responses are in Figure 16. The shock has a much larger, but short lived effects on quantities, with a 3.5% output drop. Unlike in the case of a shock to the credit limit, the adjustment is now all in durables (-17.2%), while non-durables are essentially unchanged. The crucial difference is that a spread shock is more pervasive, as it affects all borrowers and not just those near the borrowing limit. This explains both why it leads to a contraction in durable purchases (as also borrowers far from the borrowing limit find it more costly to finance durable purchases) and why we have a smaller response of non-durables (as borrowers far from the borrowing limit have lower marginal propensities to consume). This shock is sufficiently large to drive the interest rate into negative values. In the Online Appendix we show that introducing

[34] In particular, models with fixed costs of adjustment can potentially introduce more sluggish responses of durable spending to changes in the interest rate, as shown by Berger and Vavra (2014).

[35] That is, we have $\chi_t = 0.0025 + 0.015 \cdot 0.6^{-(t-1)}$, for $t = 1, 2, ...$

Figure 16: *Responses to a Temporary Shock to the Intermediation Cost χ*

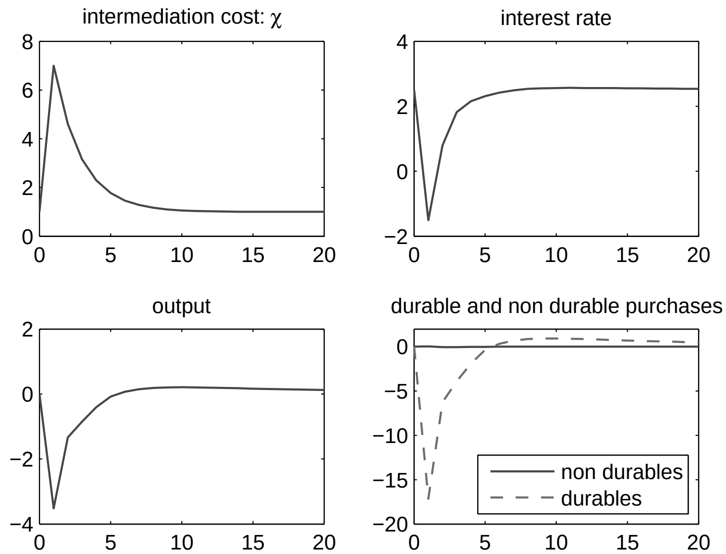

Note: Interest rate is in annual terms. Quantities are in percent deviation from initial steady state.

nominal rigidities leads to a larger output contraction, of 9.7%.

8 Concluding Remarks

We have proposed a model with uninsurable idiosyncratic risk to show how a credit crunch can generate a recession with low interest rates, 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.

A 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 an explicit treatment of labor market frictions, in which the labor market response to a drop in consumer demand leads to an endogenous increase in unemployment.[36]

[36]Krusell, Mukoyama, and Sahin (2010) introduce search frictions in an Ayiagari (1994) incomplete

Finally, a missing element in the analysis is capital. Adding capital to the model requires a theory of why claims to physical capital cannot be costlessly transformed into perfectly liquid assets like the bonds of our model. A way to move in this direction would be to combine our analysis of the household sector with financial frictions on the firms' side or a richer model of intermediation.

markets environment. Hall (2011c) discusses ways of adding search frictions to monetary environments with a binding zero lower bound on the nominal interest rate.

References

[1] Aiyagari, Rao, 1994, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109 (3), 659-84.

[2] Aiyagari, Rao and Ellen R. McGrattan, 1998. "The optimum quantity of debt,"*Journal of Monetary Economics*, 42(3), 447-469.

[3] Auclert, Adrien, 2015. "Monetary Policy and the Redistribution Channel", mimeo, Stanford.

[4] Auclert, Adrien and Matt Rognlie, 2016. "Inequality and Aggregate Demand", mimeo, Stanford.

[5] Berger, David, and Joe Vavra, 2015. "Consumption Dynamics During Recessions ", forthcoming, *Econometrica*.

[6] Bernanke, Ben, Mark Gertler, and Simon Gilchrist, 1999. "The financial accelerator in a quantitative business cycle framework,"*Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford (ed.), vol. 1, chap. 21, 1341-1393.

[7] Bewley, T., 1977, "The Permanent Income Hypothesis: A Theoetical Formulation," *Journal of Economics Theory*, 16(2), 252-92.

[8] Brunnermeier, Markus and Yuliy Sannikov, 2010, "A Macroeconomic Model of the Financial Sector," mimeo, Princeton University.

[9] Buera, F. and B. Moll, 2011, "Aggregate Implications of a Credit Crunch," mimeo, Princeton University.

[10] Buera, Francisco and Juan Pablo Nicolini, 2014, "Liquidity Traps and Monetary Policy: Managing a Credit Crunch," Federal Reserve Bank of Chicago WP 2014-14.

[11] Caballero, Ricardo, 1990. "Expenditure on Durable Goods: A Case for Slow Adjustment," *Quarterly Journal of Economics*, 422(3), 727-744.

[12] Cagetti, M., M. De Nardi, and M. Bassetto, 2011. "Credit Crunches and Credit Allocation in a Model of Entrepreneurship,"mimeo, Federal Reserve Bank of Chicago.

[13] Carroll, Christopher D., 1997, "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112(1), 1-55.

[14] Carroll, Christopher D., 2006, "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters* 91 (3), 312-320.

[15] Carroll, Christopher D. and Miles S. Kimball, 1996, "On the Concavity of the Consumption Function," *Econometrica*, 64(4), pp. 981-92.

[16] Carroll, Christopher D. and Wendy E. Dunn, 1997, "Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets," *NBER Macroeconomics Annual*, 12, 165-217

[17] Challe, Edouard and Xavier Ragot, 2010. "Precautionary Savings in the Business Cycle," mimeo, Paris School of Economics.

[18] Chamley, Christopher, 2010. "Self-fulfilling Liquidity Traps with Money as a Medium of Exchange," mimeo, Boston University.

[19] Christiano, Lawrence, Martin Eichenbaum and Sergio Rebelo, 2009. "When is the Government Spending Multiplier Large?," NBER WP 15394.

[20] Curdia, Vasco and Michael Woodford, 2010. "Credit Spreads and Monetary Policy," *Journal of Money, Credit and Banking*, 42(s1), 3-35.

[21] Deaton, A., 1991. "Saving and Liquidity Constraints." *Econometrica* 59(5):1221—48.

[22] Del Negro, Marco, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki, 2011, "The Great Escape? A Quantitative Evaluation of the Fed's Non-Standard Policies," mimeo, Federal Reserve Bank of New York.

[23] Eggertsson, Gauti B. and Michael Woodford, 2003, "The Zero Interest-Rate Bound and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 1, 139-211.

[24] Eggertsson, Gauti B., and Paul Krugman, 2012, "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo approach, " *Quarterly Journal of Economics*, 127(3), 1469-1513.

[25] Farhi, Emmanuel and Ivan Werning, 2016. "A Theory of Macroprudential Policies in the Presence of Nominal Rigidities," *Econometrica*, 84 (5), 1645-1704.

[26] 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.

[27] Geanakoplos, John, and Prandeep Dubey, 2006, "Money and Production, and Liquidity Trap," *International Journal of Economic Theory*, 2 (3/4), 295-317.

[28] Gertler, Mark and Peter Karadi, 2010, "A Model of Unconventional Monetary Policy," mimeo, New York University.

[29] Gertler, Mark and Nobuhiro Kiyotaki, 2010, "Financial Intermediation and Credit Policy in Business Cycle Analysis," mimeo, New York University.

[30] Goldberg, Jonathan, 2011, "A Model of Deleveraging," mimeo, MIT.

[31] Grossman, Sanford J. and Guy Laroque, 1990. "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods,"*Econometrica*, vol. 58(1), 25-51.

[32] Gruber, Joseph and Robert Martin, 2003. "Precautionary Savings and the Wealth Distribution with Illiquid Durables,"*International Finance Discussion Papers* 773, Board of Governors of the Federal Reserve System.

[33] Guerrieri, Veronica, and Guido Lorenzoni, 2009. "Liquidity and Trading Dynamics," *Econometrica*, 77 (6), 1751-1790.

[34] Hall, Robert, 2011a. "The High Sensitivity of Economic Activity to Financial Frictions," *Economic Journal*, 121, 351-378.

[35] Hall, Robert, 2011b. "The Long Slump," *American Economic Review*, 101(2), 431-469.

[36] Hall, Robert, 2011c. "Clashing Theories of Unemployment," mimeo, Stanford University.

[37] Heathcote, Jonathan, Fabrizio Perri, 2015. "Wealth and Volatility," mimeo, Minneapolis FRB.

[38] Heathcote J., K. Storesletten, and G. Violante, 2008. "Insurance and Opportunities: a Welfare Analysis of Labor Market Risk." *Journal of Monetary Economics*, 55(3), 501-25.

[39] Heathcote, Jonathan, Kjetil Storesletten, and Gianluca Violante, 2009. "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, 1, 319-354.

[40] Hintermaier, Thomas and Koeniger, Winfried, 2010. "The Method of Endogenous Gridpoints with Occasionally Binding Constraints among Endogenous Variables," *Journal of Economic Dynamics and Control*, 34(10), 2074-2088.

[41] Huggett, Mark, 1993. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17 (5-6), 953-69.

[42] Huo, Zhen and Jose-Victor Rios Rull. "Tightening Financial Frictions on Households,Recessions, and Price Reallocations, "*Review of Economic Dynamics* forthcoming.

[43] Iacoviello, Matteo, 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle,"*American Economic Review* 95 (3): 739–764.

[44] Jermann, Urban and Vincenzo Quadrini, 2009, "Macroeconomic Effects of Financial Shocks," NBER Working Paper no. 15338.

[45] Justiniano, Alejandro, Primiceri, Giorgio E., and Andrea Tambalotti. "Household Leveraging and Deleveraging, "*Review of Economic Dynamics* forthcoming.

[46] Khan, Aubhik and Julia K. Thomas, 2010. "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity", mimeo, The Ohio State University.

[47] Kaplan, Greg and Gianluca Violante, 2011. "A Model of the Consumption Response to Fiscal Stimulus Payments," mimeo, New York University.

[48] Kiyotaki, Nobuhiro and John Moore, 1997. "Credit Cycles," *Journal of Political Economy* 105 (2): 211–248.

[49] Korinek, Anton and Alp Simsek, 2016. "Liquidity Trap and Excessive Leverage," *AER*.

[50] Krueger, Dirk and Jesus Ferndandez-Villaverde, 2011. "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?"*Macroeconomic Dynamics*, forthcoming.

[51] Krugman, Paul R., 1998, "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 2, 137-187.

[52] Krusell, Per, Toshihiko Mukoyama, and Aysegul Sahin, 2010, "Labor-Market Matching with Precautionary Savings and Aggregate Fluctuations," *Review of Economic Studies*, 77(4), 1477-1507.

[53] Lazear, E. P., Kathryn L. Shaw, Christopher Stanton, 2013. "Making Do With Less: Working Harder During Recessions, "NBER Working Paper 19328.

[54] Leahy, John and Joseph Zeira, 2005. "The Timing of Purchases and Aggregate Fluctuations, " *Review of Economic Studies*, December 2005, 1127-1151.

[55] McKay, A., E. Nakamura, and J. Steinsson, 2016. "The Power of Forward Guidance Revisited,"*American Economic Review,* 106(10), 3133-3158.

[56] 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.

[57] Mian, A. and A. Sufi, 2011a. "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis," *American Economic Review*, forthcoming.

[58] Mian, A. and A. Sufi, 2011b. "What Explains High Unemployment? The Deleveraging–Aggregate Demand Hypothesis," mimeo, University of Chicago Booth School of Business.

[59] Midrigan, Virgiliu and Thomas Philippon, 2011, "Household Leverage and the Recession," mimeo, New York University

[60] Nekarda, Christopher J. and Valerie A. Ramey, 2010, "The Cyclical Behavior of the Price-Cost Markup," mimeo, UCSD.

[61] Ogaki, Masao and Carmen M. Reinhart, 1998, "Measuring Intertemporal Substitution: The Role of Durable Goods," *Journal of Political Economy*, 106 (5), 1078-1098.

[62] Oh, Hyunseung and Ricardo Reis, 2012. "Target Transfers and the Fiscal Response to the Great Recession "*Journal of Monetary Economics*, 59(S), S50-S64.

[63] Prescott, Edward C., 2004, "Why do Americans work so much more than Europeans?," *Quarterly Review, Federal Reserve Bank of Minneapolis*, July, pages 2-13.

[64] Rognlie, Matthew, Andrei Shleifer, and Alp Simsek, 2014 "Investment Hangover and the Great Recession," mimeo, Harvard University.

[65] Shimer, Robert, 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95 (1), 25-49.

[66] Tauchen, G., 1986, "Finite State Markov-chain Approximations to Univariate and Vector Autoregressions." *Economics Letters* 20, 177–181.

[67] Werning, I., 2016, "Incomplete Markets and Aggregate Demand" mimeo, MIT