HARD VS. SOFT FINANCIAL CONSTRAINTS IMPLICATIONS FOR THE EFFECTS OF A CREDIT CRUNCH

Cristian Alonso*

Princeton University

First Version: December 2015

This Version: June 2016

Abstract

In the aftermath of the Great Recession, understanding how households' consumption responds to a credit crunch has been a central goal of macroeconomics. Most of the recent research has explored this question using a "hard constraint" modeling device, where households can borrow at the risk-free rate only up to an exogenous amount. An alternative, and more realistic, way to model financial frictions is to allow households to borrow as much as they want but at an interest rate that depends on the level of debt. I refer to the latter as the "soft constraint" model. In a Standard Incomplete Markets framework with heterogeneous agents, I calibrate two economies differing only in the type of financial constraint that households face and I show that a credit crunch in the hard constraint economy (i.e. decrease in the exogenous borrowing limit) produces a drop in consumption significantly more severe than an equivalent crunch in the soft constraint version (i.e. increase in the borrowing interest rate). I conclude that the quantitative consequences of a credit crunch largely depend on the modeling approach.

JEL Codes: E21, E44

Keywords: Credit Crunch, Borrowing Constraints, Consumption

*calonso@princeton.edu. I thank Greg Kaplan for the valuable advice he has provided at each stage of the project. I am grateful for comments and suggestions from Richard Rogerson, Mike Golosov, and Nobu Kiyotaki. The project has also benefited from comments from participants at the Macro Student Workshop at Princeton University and at EconoCon 2015.

1 Introduction

After the Great Recession, research has flourished in an attempt to enhance our understanding of the effects of a credit crunch, i.e. a decrease in the availability of credit in the economy, on real variables such as consumption and employment[1]. Most recent work has modeled credit conditions as a hard borrowing constraint, i.e. an exogenous limit on the amount that households can borrow. In this context, a credit crunch is defined as a reduction in the borrowing limit, which forces households to reduce their consumption until they satisfy the new, lower limit. An alternative approach is to assume what I will call a soft borrowing constraint. In this setup, households can borrow as much as they want, up to their natural borrowing limit, but at an interest rate that is higher than the saving rate and potentially increases with the amount borrowed. In this context, a credit crunch is modeled as an increase in the borrowing interest rate. Both models can easily accommodate loose financial conditions (loose borrowing limit, low borrowing rate) as well as tight ones (tight borrowing limit, high borrowing rate). In this paper, I demonstrate that the choice to describe credit market conditions in terms of hard or soft constraints is not inconsequential. Rather it has important implications for inference on the effects of a credit crunch in the macroeconomy.

I compare these two alternative specifications of the financial constraints in a life-cycle Standard Incomplete Markets framework with heterogeneous agents in partial equilibrium. Consumers receive a stochastic and idiosyncratic income shock every period and decide how much to consume and how much to save or borrow of a risk free asset given the credit constraints in place, i.e. a borrowing limit in the hard constraint economy and a borrowing interest spread in the soft constraint case. The setting is deliberately simple. I abstract away from many important aspects of debt accumulation such as mortgages, default, and endogenous labor market decisions, because the goal of this exercise is to explore the implications of financial conditions in the simplest, but still realistic, framework for consumption. I calibrate the discount factor and borrowing parameters in these two economies using the Method of Simulated Moments to match the levels of aggregate wealth and debt in the US economy in 2006 according to the Survey of Consumer Finances.

While the two models predict very similar life-cycle patterns for asset accumulation, the soft constraint model matches the empirical debt distribution better than the hard constraint model in the baseline specification. Because the income profile grows with age, households have incentives to borrow early in life to smooth consumption. Debt is also useful to smooth transitory shocks to income, but as households grow older and richer, this role vanishes. Relative to the data, both models overstate the amount of debt contracted early in life and predict no

[1]Eggertsson and Krugman (2012), Guerrieri and Lorenzoni (2011), Buera and Moll (2015), Midrigan and Philippon (2011), Huo and Ríos-Rull (2016), and Justiniano, Primiceri, and Tambalotti (2015)

borrowing after middle age. But in terms of assets and debt distributions, the soft constraint model obtains a significantly better fit to the empirical evidence in the Survey of Consumer Finances. The debt distribution, i.e. the distribution of assets conditional on being negative, is better captured by the soft constraint model since the hard constraint model misses all the households with debts above the calibrated borrowing limit and predicts a counterfactual mass point at that level. Nevertheless, this better fit is not an intrinsic characteristic of the model and alternative specifications of the soft constraint, a convex borrowing cost for instance, yield a debt distribution closer to the one derived from the hard constraint model[2].

My main quantitative result is that a tightening of borrowing conditions induces a much more severe drop in consumption in the hard constraint economy than in the soft constraint economy. I follow Guerrieri and Lorenzoni (2011) in the definition of a credit crunch: a change in the borrowing parameter (i.e. a decrease in the borrowing limit in the hard constraint model or an increase in the borrowing spread in the soft constraint model) that produces a decrease of the debt to GDP ratio of 56% in the new long run equilibrium of the economy. Following a credit crunch, the drop in consumption in the hard constraint economy is more than double the drop in consumption in the soft constraint economy. The reason for this difference is that in the hard constraint economy, the credit crunch induces a large drop in consumption by forcing households to deleverage immediately or soon after, whereas in the soft constraint economy, the incentives to deleverage are provided through the interest rates, in response to which households choose to optimally reduce their debt at a lower pace. A number of robustness exercises regarding the calibration targets, elasticity of intertemporal substitution, and initial distribution of assets, confirm that the milder response in the soft constraint setting is a very general result.

In this paper, I adopt the most basic form of the soft constraint, a constant spread between the saving and the borrowing rate. The assumption is driven by its simplicity and parsimony, because it reduces the calibration to only one parameter. In Section 6, I introduce some heterogeneity by allowing the borrowing spread to differ across households. I choose to model such heterogeneity as a fixed financial type to be consistent with the lack of life-cycle features of the reported borrowing interest rate. I take the borrowing spreads directly from the data and I use the discount factor to match the amount of debt in the economy. In this context, I find that a sevenfold increase in the borrowing spread is required for the soft constraint to deliver drops in consumption similar to those in the hard constraint. An alternative approach would have been to estimate the soft constraint directly in the data using information on interest rates and debt levels. Unfortunately, the Survey of Consumer Finances does not include data on credit scores and so, a regression of interest rates on debt levels would suffer from omitted variable bias. In the Appendix, I also consider the case of a convex borrowing cost and show that still

[2]This case is shown in the Appendix.

the soft constraint exhibits a milder consumption response.

The idea of a soft constraint, a setting in which the interest rate that borrowers have to pay depends on the amount they want to borrow, has already been explored in the literature but not in the context of a credit crunch. An early study of an increasing and convex interest rate schedule on assets emerging from the default risk is found in Eaton and Gersovitz (1981) in the context of sovereign debt. More recently, Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007) and Livshits, MacGee, and Tertilt (2007) incorporate default in a life-cycle model with incomplete markets and endogenously derive the price of consumer loans in equilibrium[3]. Even in the absence of default, a positive borrowing spread can be understood as an intermediation cost in the spirit of Bernanke (1983). Since I do not model default explicitely, I prefer the latter interpretation of the borrowing spread. Finally, agnostic about the origin of the constraint, Fernández-Corugedo (2002) studies the consumption-savings problem and finds that precautionary savings are higher in the hard constraint model than in the soft constraint. However, since I calibrate impatience and borrowing parameters to match the same amounts of wealth and debt in equilibrium, such difference will be absorbed by the discount factor.

The main contribution of this paper is to the literature on consumption and credit conditions. Building on the influential work by Guerrieri and Lorenzoni (2011), I show that the magnitude of the macroeconomic consequences of a credit crunch largely depends on the modeling strategy employed. If the tightening of credit market conditions occurs through interest rates rather than borrowing limits, my results cast doubts on the quantitative importance of this mechanism to explain the drop in consumption during the Great Recession. Other mechanisms such as the increase in the risk of unemployment or the negative wealth effect of lower housing prices are then more likely explanations of the last recession. But the question of whether the credit crunch affected households via borrowing limits or spreads is ultimately an empirical question that this paper cannot address because of data limitations. An ideal dataset to answer such question will include household's borrowing conditions before and after the credit crunch. To the best of my knowledge there is not public dataset with that information[4].

The rest of the paper proceeds as follows. Section 2 describes the model, which is calibrated in Section 3. Section 4 presents the main results in the initial steady state of each economy and Section 5 explores the consumption response to a worsening of the financial conditions. Section 6 extends the basic framework to allow for heterogeneous financial conditions and provides additional robustness checks. Finally, Section 7 concludes.

[3]Athreya, Tam, and Young (2012) is also concerned with unsecured debt, although not in a life cycle framework.

[4]The 2007-09 Panel Survey of Consumer Finances provides information on households' changes in wealth during the Great Recession, but the dataset does not include interest rates on unsecured debt, nor credit scores.

2 Model

In this section, I describe the model that I use to study the differences between a hard constraint and a soft constraint economy.

I compare the aggregate implications of the different models of credit conditions in a Standard Incomplete Markets framework with heterogeneous agents in the tradition of Aiyagari (1994). Agents in this model are consumers/households who receive an stochastic income every period and decide how much to spend that period and how much to save or borrow in a risk-free asset. The stochastic income is idiosyncratic and cannot be insured against because financial markets are incomplete. In this framework, financial conditions are introduced as limits to the agent's borrowing ability. In the hard constraint economy, the agent can borrow at the risk-free rate only up to an exogenous amount. There is no such limit in the soft constraint economy so that the consumer can borrow as much as she wants up to the natural borrowing limit, but the interest rate she will have to pay on her debt is higher than the rate she would received if she had positive assets.

Time is discrete. The economy is populated by a continuum of households of measure one. Each agent works and consumes for a finite but uncertain number of periods. She starts working immediately after she is born and retires at age T_w. After retiring from the workforce, the agent starts facing a risk of death. A consumer of age t survives the period with probability ζ_t and no agent grows older than $\bar{T}$, i.e. $\zeta_{\bar{T}} = 0$. New households are born every period to replace the ones who die keeping the population size constant.

The household has preferences over consumption in different periods and states of the world. These preferences are represented by the utility function:

$$E\left[\sum_{t=1}^{\bar{T}} \beta^t U\left(C_t\right)\right]$$

with

$$U\left(C\right) = \begin{cases} \frac{C^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1 \\ \ln C & \text{if } \gamma = 1 \end{cases}$$

where C_t is the consumption level when the agent is t periods old, β is the discount factor, γ is the inverse of the elasticity of intertemporal substitution and $E\left[\cdot\right]$ is the expectation operator.

During her working life, when the agent is younger than T_w, she receives stochastic earnings every period in compensation for her work, which she supplies inelastically. Let $Y_{i,t}$ be earnings before taxes for agent i of age t. Log earnings, $y_{i,t}$, are assumed to be decomposed into a common deterministic experience profile, κ_t; an individual fixed effect, μ_i; a persistent shock, $z_{i,t}$; and a transitory shock, $\epsilon_{i,t}$.

$$y_{i,t} = ln(Y_{i,t}) = \kappa_t + \mu_i + z_{i,t} + \epsilon_{i,t}$$

$$z_{i,t} = z_{i,t-1} + \eta_{i,t}$$

where $\{\eta_{i,t}\}_{t=1}^{T_w}$ is a sequence of independent random variables with normal distribution, zero mean and variance σ_η^2 and $z_0 = 0$. $\{\epsilon_{i,t}\}_{t=1}^{T_w}$ is a sequence of independent random variables with normal distribution, zero mean and variance σ_ϵ^2. The individual fixed effect exhibits also a normal distribution across agents with zero mean and variance σ_μ^2. All these shocks are independent across agents in the economy. Then, a version of the Law of Large Numbers apply in this economy and there is no aggregate risk. To solve the model by simulation, the income process is discretized and so, the Natural Borrowing Limit is not zero.

During retirement, the agent receives a constant pension benefit per period until her death. This benefit is a function of all the history of individual earnings before taxes during her working life, i.e. $P\left(Y_{i,g,1}, ..., Y_{i,g,T_w-1}, Y_{i,g,T_w}\right)$, which I approximate keeping track of average individual gross earnings.

Both earnings and pension benefits are taxed using a non-linear schedule on pre-tax income, $\tau(\cdot)$. After-tax income $\bar{Y}_{i,g,t}$ is then given by $\bar{Y}_{i,g,t} = Y_{i,g,t} - \tau(Y_{i,g,t})$.

Unsecured debt is more prevalent among young households, but also retired agents borrow in the data, as shown in Figure 5. Since in the model retired agents face death risk, their natural borrowing limit in absence of life insurance would be zero, and so they would be unable to borrow. To match the empirical fact that even old households borrow, I assume that there exist perfect, actuary fair, annuities markets that allow retired agents to purchase insurance against the death shock and so, to borrow.

Other than life insurance, markets are incomplete. Households are allowed to borrow and save only through a risk-free, one period asset. This asset pays a deterministic interest rate given by the function $R(A)$, where A is the level of assets held by the agent. Also, there is a limit $\bar{A}_t$ on the amount that can be borrowed as in Aiyagari (1994). There is no default option in the model. Certainly allowing for agents to default would be an interesting extension, but it is beyond the scope of this paper.

Therefore, the agent's budget constraint per period can be summarized as follows:

$$C_{i,t} + A_{i,t+1} = R\left(A_{i,t}\right) A_{i,t} + \bar{Y}_{i,t} + TR_{i,t} \qquad \text{if } t \leq T_w$$

$$C_{i,t} + \tfrac{1}{\zeta_{t+1}} A_{i,t+1} = R\left(A_{i,t}\right) A_{i,t} + \bar{P}\left(Y_{i,1}, ..., Y_{i,T_w-1}\right) + TR_{i,t} \qquad \text{if } t > T_w$$

where $TR_{i,t}$ is a government transfer to ensure a consumption floor $\bar{c}$ for all households. $\bar{c}$ will be calibrated to a very small level and no positive transfers will take place in these economies at the initial steady state. However, after the impact of a credit crunch, some households could be forced into negative consumption if their income is not enough to cover the required deleverage, and in those cases positive transfers will emerge. These transfers can then be thought as a rudimentary safety net.

Finally, initial assets are drawn from a distribution $H(A_0)$.

Financial Conditions in the Hard Constraint Model. Agents can borrow only up to a certain level at a constant interest rate. There is no spread between the borrowing and saving rates. The borrowing limit $\bar{A}_t$ in period t is the minimum of the natural borrowing limit for that period (NBL_t) (i.e. the maximum amount that could be fully repaid with probability one) and an exogenous amount $\bar{\phi}$, assumed constant over the life cycle. Then, $\bar{A}_t = min\left\{\bar{\phi}, NBL_t\right\}$. The interest rate function is simply $R(A) = R^f$ where R^f denotes the gross risk-free rate. In Section 6, heterogeneity in the borrowing limit will be allowed and the distribution will be inferred from the data, but for the rest of the paper all the households will face the same borrowing limit regardless of their income. In the Appendix, I study an alternative specification for the hard constraint, where the borrowing limit is a fraction of the natural borrowing limit, i.e. $\bar{A}_t = \bar{\phi}_{NBL} NBL_t$. The results do not vary significantly.

Financial Conditions in the Soft Constraint Model. Agents can borrow as much as they want up to the natural borrowing limit ($\bar{A}_t = NBL_t$), but the interest rate is no longer constant. In particular, a constant borrowing spread is assumed such that:

$$R\left(A\right) = \begin{cases} R^f & \text{if } A \geq 0 \\ R^f + \phi & \text{if } A < 0 \end{cases}$$

For parsimony, the constant borrowing spread assumption is maintained in the baseline specification. Heterogeneity in the borrowing spread will be allowed in Section 6. In the Appendix, I consider the case of a convex function for the borrowing spread.

Figure 1 illustrates the difference between these two alternative approaches to modeling financial conditions in terms of the shape of the resulting budget constraint in a simple two-period model. Both can accommodate very loose and very tight borrowing conditions, although the soft constraint seems more realistic as we do observe borrowing spreads in the data, whereas borrowing limits are more difficult to identify.

Figure 1: Budget and Credit Constraints in a Simple Two-Period Model

Panel A. Hard constraint — Panel B. Soft constraint

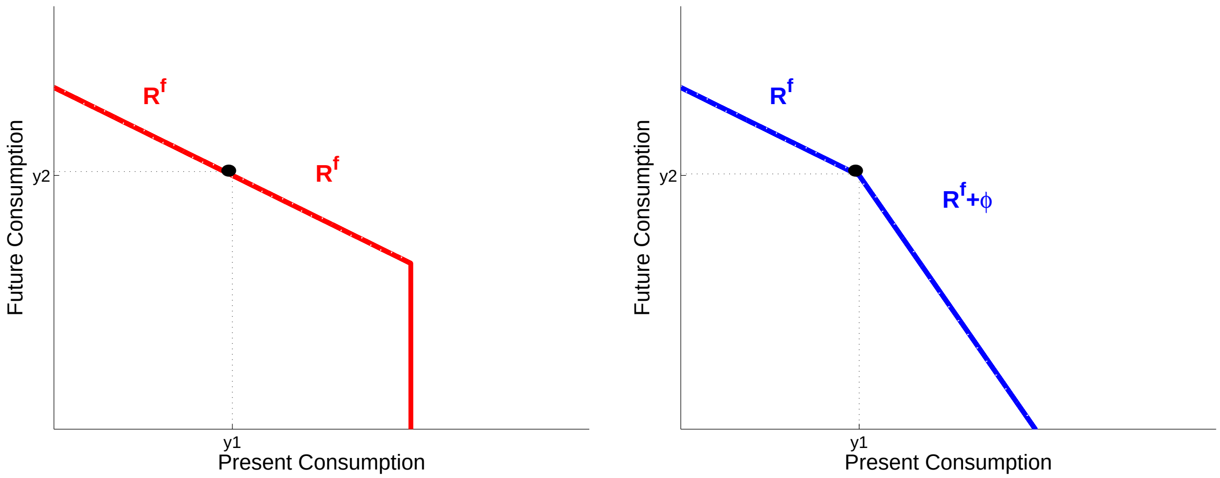

Consumption possibilities differ only for borrowers
y_1 and y_2 are income in the present and future respectively, i.e. (y_1 , y_2) is the autarky point. Households whose optimal consumption bundle is to the left of y_1 are saving and those to the right of y_1 are borrowing.
In Panel A the slope of the budget constraint is R^f until the hard borrowing starts to bind and the slope becomes infinite. In Panel B, the slope is R^f to the left of the autarky point (household saving) and $R^f + \phi$ to the right (household borrowing).

3 Calibration

In this section I discuss the calibration of the model. The strategy largely follows Kaplan and Violante (2010), although my calibration is at the quarterly level and theirs is annual. I match aggregate moments of the US economy in 2006, before the beginning of the Great Recession, using the 2007 Survey of Consumer Finances (SCF 2007 henceforth).

Assets are defined as net liquid financial wealth[5]. This definition follows Guerrieri and Lorenzoni (2011) not only for comparability of the results, but also to give unsecured debt a meaningful role as an instrument to smooth consumption. The definition excludes housing, mortgages, and other types of secured debt because the model does not capture the main features of these assets.

Demographics. Households join the labor market at age 25 ($t = 1$ in the model) and retire at age 65 ($T_w = 160$). Survival probabilities are obtained from population data from the

[5]Checkings, savings and money market accounts, stocks, bonds, and certificates of deposits minus revolving credit card debt, consumer and educational loans.

2010 Census[6]. Agents die with certainty after they turn 95 ($\bar{T} = 280$). Therefore, households work for 40 years (160 quarters) and live on their pension benefits and accumulated assets for, at most, 30 years (120 quarters).

Initial Assets Distribution. All agents start their economic life with zero assets. In Section 6, results are shown to be robust to including a non-degenerate initial distribution of assets, $H(A_0)$, estimated from the SCF 2007 as the distribution of financial wealth for households younger than 25.

Earnings Before Tax. The age profile of labor income is estimated using data from the Panel Study of Income Dynamics (PSID henceforth) between 1967 and 2002 as suggested in Heathcote, Perri, and Violante (2010). The deterministic experience profile, $\{\kappa_t\}_{t=1}^{T_w}$, is computed with a fourth-order polynomial on potential experience[7] over log earnings for households with age between 25 and 64. For comparability between assets data obtained from the SCF and earnings figures obtained from the PSID, the intercept is then adjusted to match average pre-tax annual earnings before retirement in the SCF 2007: \$61,521.

The variance of the residuals of earnings after subtracting the deterministic component and controlling by time fixed effects[8] rises almost linearly, a point previously noted by Kaplan and Violante (2014), which justifies the assumption of a unit root earnings process. Following Kaplan and Violante (2014), I set the variance of the individual fixed effect to 0.18 to match the initial dispersion in earnings and the variance of the permanent shock to 0.003 to replicate the rise in dispersion over the life cycle. The quarterly variance of the transitory shock is set equal to 0.19 to reproduce an annual variance of 0.05 as used in Kaplan and Violante (2010). In the Appendix, I show that my results are robust to an annual specification of the transitory shock instead of a quarterly one.

Pension Benefits and Tax System. Following Kaplan and Violante (2010), pension benefits and the tax schedule are computed such that they resemble the actual US systems. Social security payments are equal to 90% of average individual gross earnings up to a first bend point (18% of cross-sectional average), 32% up to a second bend point (110%) and 15% from there on. Then, the payments are scaled to get an average replacement rate of 45%.

Gross earnings are taxed through the nonlinear function estimated by Gouveia and Strauss (1994):

$$\tau(Y) = \tau^b \left[Y - \left(Y^{-\tau^p} - \tau^s \right)^{-\frac{1}{\tau^p}} \right]$$

with $\tau^b = 0.258$, $\tau^p = 0.768$ and τ^s chosen to match the ratio of personal current tax receipts on labor income to total labor income in the US economy, 25%. As in the data, 85% of the

[6]Source: U.S. Census Bureau, Population Division.

[7]Mincer (1958)

[8]Similar results are obtained after controlling for cohort fixed effects.

social security benefits are taxed in the model.

The consumption floor is defined as the quarterly average SNAP benefit per person in 2007.

Preferences. For the baseline specification, logarithmic preferences, $\gamma = 1$, are assumed. Section 6 shows that results are robust to other sensible values of $\gamma \in [0.5, 4]$. The calibration of the discount factor is described below.

Saving Interest Rate. As in Telyukova (2013) and because of the lack of other high-yield assets in the model, the annual interest rate of financial assets to 4%.

Discount Factor and Borrowing Parameters. The discount factor β and the borrowing parameters (the exogenous borrowing limit in the hard constraint version and the borrowing spread in the soft constraint) are jointly estimated to match two significant moments of the data through the Method of Simulated Moments. Since I am interested in studying the dynamics of unsecured debt over the life cycle, I match the ratios of aggregate gross financial wealth and debt to income from the SCF 2007. Guerrieri and Lorenzoni (2011) also target these moments, but they compute them from national accounts data. Instead, I compute them from micro data from where the age profiles can also be recovered. In the SCF 2007, the ratio of aggregate financial wealth to income is 1.4654, whereas the debt to income ratio is 0.0556. In Section 6, I show that my results are even stronger when calibrating the models to match the moments as computed from the national accounts data.

For the hard constraint, this approach results on $\beta = 0.9874$ (equivalent to an annual discount factor of 0.9505) and a borrowing limit of $\bar{\phi} = -20,952$. For the soft constraint, I obtain $\beta = 0.9873$ (equivalent to an annual discount factor of 0.9503) and a borrowing wedge of $\phi = 0.0079$, which implies an annual borrowing interest rate of 7.3%.

The calibration is summarized in Table 1. The discount factor, β, is identical for agents in both economies and it is not implausible low. The borrowing limit, $20,952$, does not seem unreasonable considering I am matching only unsecured debt. According to the SCF 2007, less than a quarter of the debtors had unsecured debt for amounts greater than this borrowing limit. Also, in terms of self-reported borrowing limit on credit cards, around 75% of the households indicated a limit below my calibration. In the model, 0.5% of the households are borrowing at the limit. On the other hand, the borrowing spread is below what is observed in the data by implying an annual borrowing interest rate of 7.3%. The median annual interest rate on unsecured debt in the SCF 2007 was 8.0%, whereas the mean was 9.4%. Certainly a limitation of the baseline model is the assumption that the same financial conditions apply to everyone. This issue is addressed in Section 6, where the distribution of borrowing rates and credit limits are taken directly from the data.

Table 1: Calibration

Parameter	Description	Value
γ	Coefficient of relative risk aversion	1
R^f	Annual net risk-free rate	4%
σ_η^2	Quarterly variance of the persistent shock	0.003
σ_μ^2	Variance of the individual fixed effect	0.18
σ_ϵ^2	Quarterly variance of the transitory shock	0.19
$\bar{c}$	Consumption Floor	289
Hard Constraint Economy		
β	Annual discount factor	0.9505
$\bar{\phi}$	Borrowing limit	$20,952$
Soft Constraint Economy		
β	Annual discount factor	0.9503
ϕ	Annual borrowing interest rate	7.3%

The interest rates and discount factors are reported in annual levels.
Discount factor and borrowing parameters chosen to match aggregate debt and aggregate wealth to GDP ratios in each economy. Refer to the text for details on the calibration of the remaining parameters.

4 Steady State Results

The model is solved numerically. I use Carroll (2005)'s Endogenous Grid Points Method[9] to find the policy functions for each model and then I simulate economies with $100,000$ households. In this section, I compare the long run equilibrium outcomes of the hard constraint and soft constraint economies, in particular, the household policy functions, life-cycle paths, and asset distributions.

The credit conditions affect consumption decisions only when the level of assets is low. Figure 2 plots the consumption policy functions resulting from the models calibrated as indicated in the previous section for agents at two different ages (29 and 62, very early and very late in their working life). As noted in the consumption literature[10], consumption is linear in wealth (or cash-in-hand) provided the amount of assets held by the household is positive and large enough, as the financial constraints are then irrelevant. In fact, the policy functions for the soft

[9] I use 62 grid points for assets, 30 for negative values and 32 for positive. I space the points with a polynomial of exponent 0.4 such that more grid points are obtained close to the borrowing limit and to zero. I use 11 grid points for the permanent component, 5 for the transitory shock, 5 for the individual fixed effect, and 5 for lifetime average earnings

[10] Deaton (1991), Carroll and Kimball (1996)

Figure 2: Consumption Policy Functions

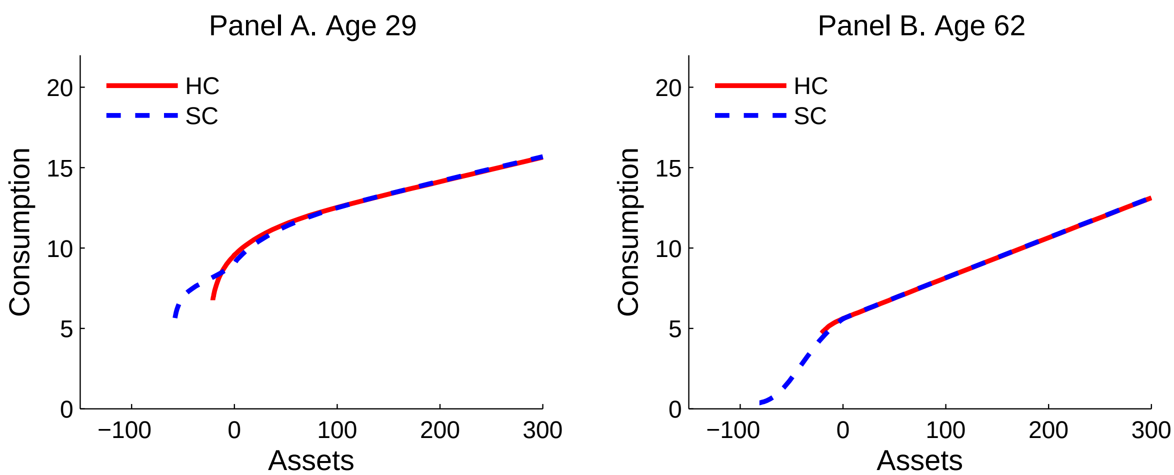

Consumption decisions only differ significantly between the two economies when the agent is borrowing or has a positive but small stock of assets. Quarterly consumption level as a function of the assets held by an agent of age 29 and 62 (15 and 150 in the model) with average earnings history, average permanent shock, average individual fixed effect, and average transitory shock. Consumption and assets are expressed in thousands of US dollars.
HC: Hard Constraint Model. SC: Soft Constraint Model.

and hard constraint versions are very similar in this area of the state space.

As the amount of wealth decreases and becomes close to zero, the policy functions start to differ significantly. For the hard constraint, the well-known concave shape emerges from the combination of the borrowing limit and the stochastic earnings process. Uncertainty vanishes for an agent close to retirement and so, the policy function is almost linear for all the domain for the older agent. On the other hand, the soft constraint allows for an extended domain as households can borrow greater amounts. Around the Natural Borrowing Limit, the consumption policy function is concave for the young agent. When the amount of debt is large but far from the limit, the policy function is again linear as the savings rate does not factor into the problem in the short and medium term. Around zero, the discontinuity in the interest rate induces another non-linearity in the consumption function. The household chooses a level of consumption below the hard constraint amount acknowledging the greater borrowing cost. A different non-linearity emerges in the policy function of the old agent in the soft constraint model. When the amount of debt is large enough, the agent chooses to default into the consumption floor and so, the policy function is flat in that region.

The two models predict a similar evolution of mean consumption and wealth over the life cycle. Figure 3 summarizes them. The combination of financial constraints and a stochastic

Figure 3: Mean Consumption, Income and Assets over the Life Cycle

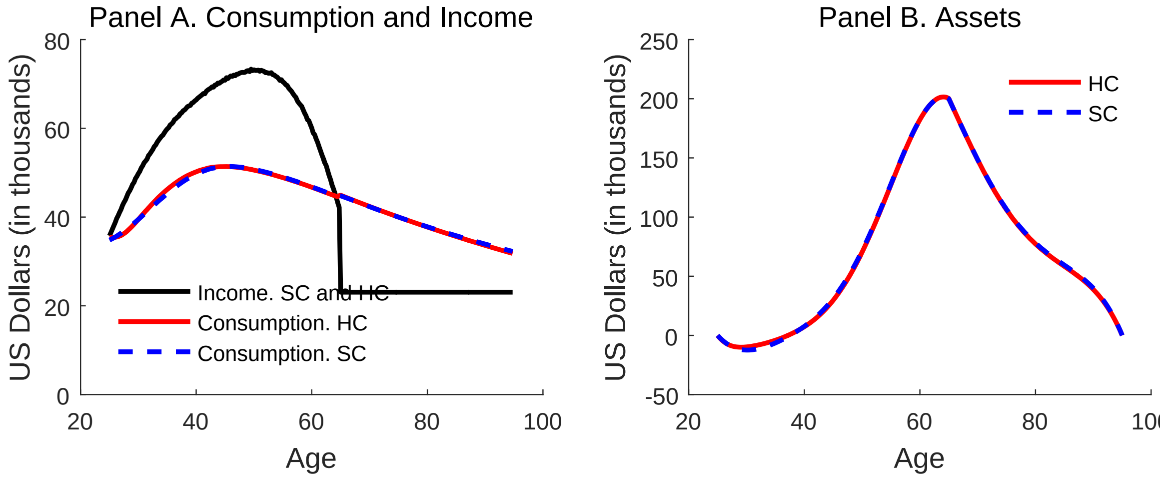

On average, consumption and assets show a similar evolution over the life cycle in both models.

Means computed for the simulation of two economies with 100,000 households each. Consumption and income are annualized.

HC: Hard Constraint Model. SC: Soft Constraint Model.

labor income produces a hump-shaped consumption. Mean consumption attains a maximum at age 45 in the hard constraint model and 46 in the soft constraint. The timing of the peak is consistent with the empirical evidence documented by Fernández-Villaverde and Krueger (2007), who find that both total and non-durable consumption peak in the late forties. The soft constraint presents a slightly steeper consumption growth early in life, with consumption at the peak being 48% higher than at age 25, versus 45% in the hard constraint. In both cases, consumption grows over the first years of the working life as the household moves away from the financial constraints, and then it decreases as retirement approaches because the household is more impatient than the market and the stochastic component of income starts to vanish. Other features traditionally associated with the hump-shape of consumption such as changes in family demographics, housing and other durable goods, and non-separability between consumption and leisure[11], are absent in my model.

The degree of consumption smoothing is similar in the aggregate for the two models, but there are life cycle differences early in life. Following Kaplan and Violante (2010), I compute the insurance coefficients[12] in both economies using the observed permanent and transitory

[11] Fernández-Villaverde and Krueger (2007), Attanasio, Banks, Meghir, and Weber (1999)

[12] $\phi^x = 1 - \dfrac{cov(\Delta c_{it}, x_{it})}{var(x_{it})}$ where x_{it} is either the transitory or the permanent shock. With this definition, an insurance coefficient of 1 implies that consumption does not react to the shock at all (full insurance), whereas an insurance coefficient of 0 indicates complete absence of insurance.

Figure 4: Insurance Coefficients during the Working Life

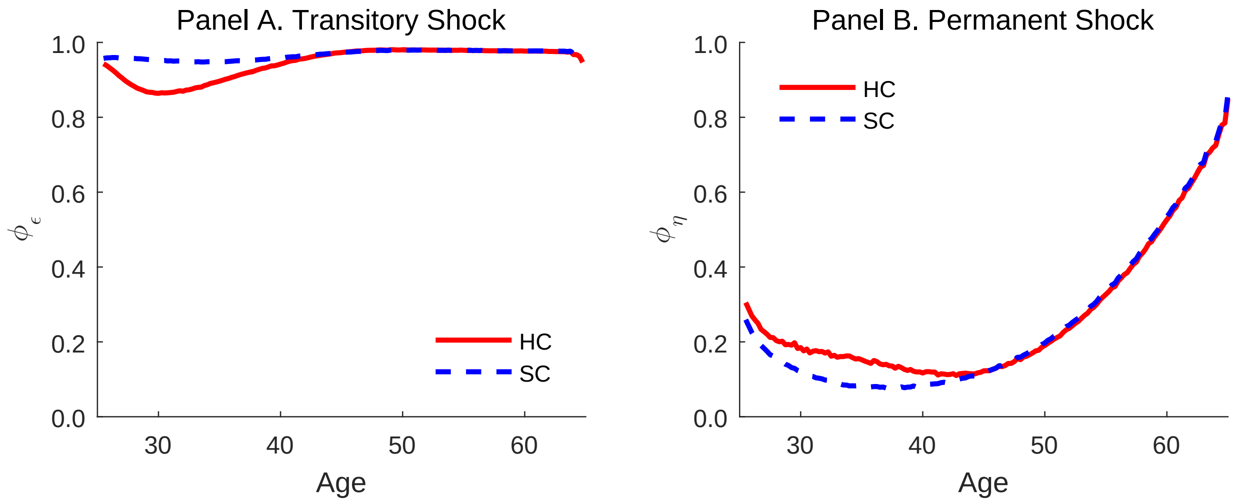

The models differ on the insurance options early in life. In the soft constraint economy it is easier to insure transitory shocks, but it is more difficult to insure permanent shocks.
Insurance coefficients are computed following the procedure described in Kaplan and Violante (2010) and employing the true values of the permanent and transitory shocks.
HC: Hard Constraint Model. SC: Soft Constraint Model. Refer to the text for further details.

shocks. On the one hand, the insurance coefficient for the permanent shock is 0.21 in the hard constraint economy and 0.18 in the soft constraint economy. As in Kaplan and Violante (2010), both models underestimate the amount of insurance with respect to what Blundell, Pistaferri, and Preston (2008) find in the data, 0.36, and they yield a U-shaped profile over the life cycle (Figure 4). However, the soft constraint economy exhibits less insurance early in life against the permanent shock because of the higher borrowing interest rate. On the other hand, both models predict high insurance for transitory shocks: 0.95 in the hard constraint and 0.97 in the soft constraint, consistent with Blundell, Pistaferri, and Preston (2008)'s empirical finding of 0.95. In the soft constraint version, the age profile of the transitory insurance coefficient is fairly flat around 0.97. In contrast, in the hard constraint the insurance coefficient decreases early in life as households accumulating debt are less able to self-insure against the shock when they approach the borrowing limit. As households age and save, the insurance coefficient grows and converges to the level of the soft constraint economy.

The predictions of both models in terms of the ratios of debt and wealth to income over the life cycle are fairly similar. Relative to the data, both models overstate the amount of borrowing early in life, while they are unable to explain any borrowing between age 45 and retirement. In a similar fashion, the two models miss the continuously increasing wealth to

Figure 5: Debt to Income and Wealth to Income Ratios. Models vs. Data

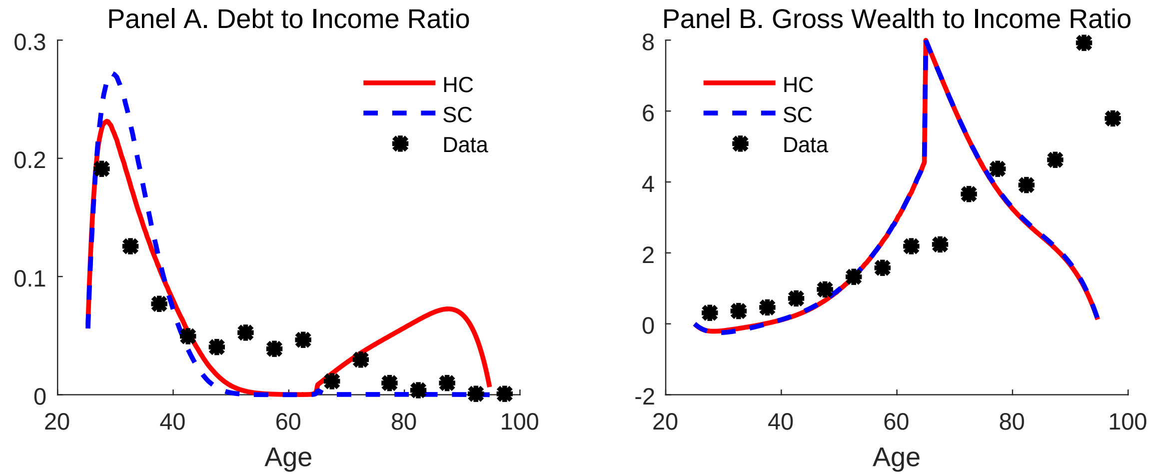

Both models predict a similar pattern for the ratios of wealth and debt to income, overstating the amount of debt early in life and understating it in middle age

Debt to income is the ratio of the absolute value of aggregate negative assets to aggregate annual income. Gross wealth to income is the ratio of aggregate positive assets to aggregate annual income.

HC: Hard Constraint Model. SC: Soft Constraint Model. Data: SCF 2007. Refer to the text for further details.

income ratio, especially after retirement[13]. The failure of the models in those dimensions could be corrected by extending them to include durable consumption, changes in family size, or altruistic inheritances which would enhance the use of unsecured debt late in the working life and provide incentives to keep accumulation wealth after retirement. However, the goal of this paper is to study the implications of financial conditions in the simplest consumption framework. To that extent, the present model offers a reasonable first pass.

By construction, both models produce the same amount of aggregate wealth and debt. However, the distributions differ. In the hard constraint version, 30.7% of the households are borrowing, with 1.6% of them borrowing at the limit. In the soft constraint, the ratio of households borrowing is 25.3%, closer to the fraction of borrowers observed in the data, 25.7%.

It is a well established fact that this class of models struggles to match the entire wealth distribution, even when targeting multiple moments of the distribution[14]. It is then no surprise that none of the models does a great job at matching the wealth or debt distributions, although the soft constraint version does better as shown in Figure 6. In terms of the distribution

[13]Note that the sharp increment in the gross wealth to income ratio at retirement is generated in both models mechanically by the sudden (but predictable and expected) decrease in labor income associated with the end of the working life, rather than by additional savings.

[14]Huggett (1996), Castañeda, Ana, Díaz-Gimenez, and Ríos-Rull (2003).

Figure 6: Assets and Debt Distribution during Working Life

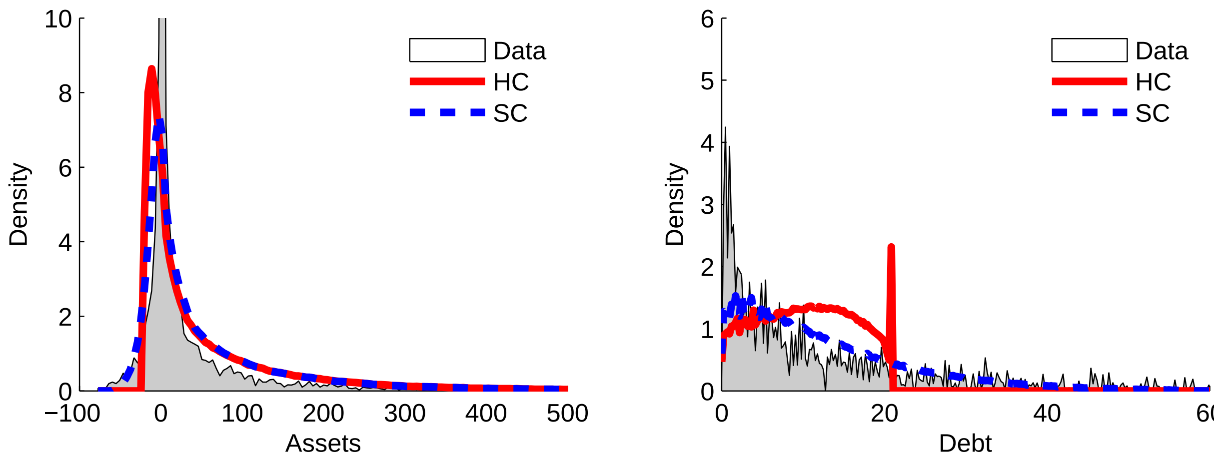

The soft constraint model produces assets and debt distributions closer to the data

Assets and debt are in thousands of US dollars. The axis have been truncated to facilitate visualization.

HC: Hard Constraint Model. SC: Soft Constraint Model. Data: SCF 2007. Refer to the text for further details.

of assets, both models fail to replicate the large proportion of households with no wealth by overstating the asset accumulation. However, the soft constraint does a better job with the left tail and, as a result, the distribution looks more symmetric than the one generated by the hard constraint. As shown in the Appendix, the better match of the soft constraint to the assets distribution remains when the hard constraint is modeled as a fraction of the natural borrowing limit, instead of a dollar amount. But when the soft constraint is modeled as a convex borrowing cost on the amount of debt, the models yield very similar assets distributions.

In terms of the debt distribution, it is clear that the soft constraint produces a better match. The Jensen-Shannon divergence[15] between the model-derived debt distribution and its empirical counterpart is 0.1556 in the hard constraint economy, whereas it is 0.0718 in the soft constraint economy. By unrealistically offering cheap credit, a significant number of households end up at the borrowing limit and so, the distribution is not monotonically decreasing in the amount of debt, but actually has a mass point at the exogenous borrowing limit. As it will be discussed in the next section, these households borrowing at the limit will be the ones responding the most to a credit crunch just because mechanically they will be forced to delever. But there are relatively so few of these constrained borrowers that the aggregate response will be driven

[15]The Jensen-Shannon divergence is a measure of the similarity between two probability distributions. It is symmetric and bounded between 0 and log(2). In this paper, I discretize the model and empirical distributions using the same 500 bins in both and I use the natural logarithm when computing the Jensen-Shannon divergence.

instead by the unconstrained borrowers, i.e. the households borrowing, but not at the limit.

To summarize, the two models have similar predictions for the aggregates in the economy, but the soft constraint produces a better fit in terms of the wealth and debt distributions in the baseline framework. In the Appendix, I will show that alternative specifications of the borrowing conditions still deliver better fits in the soft constraint economy, but the difference becomes smaller. The better match to the empirical distributions by the soft constraint model should then be interpreted as a strength of the baseline framework, but not as a general result.

5 Credit Crunch

A credit crunch is a tightening of the financial conditions faced by the households. In this section, the credit crunch exercise in Guerrieri and Lorenzoni (2011) is replicated, but within a partial equilibrium framework to focus on the response of aggregate consumption. These results can be thought as an upper bound on the general equilibrium effects because if the interest rate was allowed to adjust after the crunch, it would decrease, lessening the initial drop in consumption[16].

I follow Guerrieri and Lorenzoni (2011) and I define a credit crunch as a change in borrowing parameters that leads to a new steady state with lower amount of debt. In their section focusing on unsecured debt, Guerrieri and Lorenzoni (2011) assume an initial debt to GDP ratio of 18% and model the tightening of financial conditions as the decrease in the exogenous borrowing limit necessary for the debt to GDP ratio to drop by 10 percentage points in the new steady state. Because I calibrate the models to match the evidence in the micro data on debt to income averages by age, rather than national accounts figures, debt to GDP in the initial steady state is set to 5.56% and so it cannot decrease by 10 percentage points. Instead, I re-calibrate the borrowing parameters to match a decrease of 55.56% (1-0.08/0.18) in the debt to income ratio in the new steady state as in Guerrieri and Lorenzoni (2011). Thus, I find the borrowing limit in the hard constraint and the borrowing spread in the soft constraint to match a debt to income of 2.47% in the new steady state. In the next section, I show that the results hold when the

[16]In fact, if the economy was closed and there was not an exogenous supply of bonds, the net wealth in the economy would have to always be zero, and aggregate consumption would always be equal to aggregate income. Aggregate income is the sum of compensation of employees and the interest income from net wealth. I do not count the borrowing spread in the soft constraint economy as part of income. Instead, I assume that borrowing is conducted by a competitive financial industry using only labor, whose compensation is already counted in GDP. I think of a credit crunch in the soft constraint economy as a negative shock to labor productivity in the sector. The shock would then displace some workers from the financial industry, but they would immediately obtain equivalent jobs in the rest of the economy. Then, aggregate income would not react to a credit crunch and neither would aggregate consumption. In this case, the interest rate would drop inducing savers to consume more to make up for the reduced consumption of borrowers. Debt and gross wealth will decrease by the same amount, so that net wealth will remain constant at zero. Thus, the consumption response to a credit crunch in general equilibrium for a closed economy with zero aggregate net wealth would be zero.

economies are calibrated to match the moments in Guerrieri and Lorenzoni (2011).

I consider two alternatives for the timing of the credit crunch. First, an unexpected worsening of the conditions that hits the economy at time 1 and forces households to adjust either immediately or over a period of time. Second, I allow for an expected worsening in which households learn about the future credit crunch a few quarters before it actually occurs. Since in the former the households are surprised by the shock, the drop in consumption there can be understood as an upper bound; whereas in the latter, I try to capture the fact that households can anticipate trouble and start adjusting before the shock actually hits, so that the decrease is smaller.

Finally, I consider an alternative definition of a credit crunch, one in which the shock to the amount of debt is the same not in the long run, but in the very short run. In the previous exercises, debt to GDP would converge to the same, lower level in the long run, but implying different paths immediately after the shock. An alternative definition could require debt to GDP (and consumption) to fall by the same amount in both economies in the first period and then evaluate the long run consequences of the credit crunch. I calibrate the financial parameters to produce the same initial drop in consumption after the shock hits, so that the initial decrease in the supply of loanable funds is the same in both economies. And I study the implications of this calibration for the level of debt in the long run.

5.1 Unexpected Credit Crunch

In the unexpected credit crunch, households learn at time 1 that financial conditions will start to deteriorate immediately. For the hard constraint, the calibration of tightening of financial conditions yields a new borrowing limit of \$12,209, a decrease of 8,720 dollars or 41.7%, similar to the drop obtained in Guerrieri and Lorenzoni (2011), 44.2%. For the soft constraint, the credit crunch will be modeled as an increase in the borrowing interest rate spread to 0.0134, an increase of almost 70% and equivalent to an annual borrowing rate of 9.6%. The new borrowing limit of \$12,209 and the new borrowing interest rate spread of 0.0134 are such that in the final steady state the debt to income ratio will be 2.47% in both economies.

I consider two alternative dynamics for the credit shock. First, I assume households are forced to adjust immediately after the shock hits. Because of the size of the shock, around a sixth of average annual income, the deleverage does not appear implausible. In addition, the existence of transfers guarantees that no household will be lead into negative consumption. Then, I follow Guerrieri and Lorenzoni (2011) and allow the adjustment of financial conditions to take place over six quarters. Comparing the differential dynamics will be instructive as the six-quarter adjustment period decreases the importance of forced-deleverage in the aggregate response.

When the credit crunch shock requires immediate adjustment, the initial drop in consumption in the soft constraint economy is only a third of the observed in the hard constraint. As illustrated in Figure 7, aggregate consumption drops by 4.8% when the credit crunch hits the hard constraint economy, but only by 1.6% in the soft constraint one. The decrease in the hard constraint is lessened by the presence of transfers which benefit 2.6% of the households when the shock hits, compared with only 0.05% of the households in the soft constraint. Once the new borrowing limit settles in, the drop in consumption decreases in magnitude until it eventually turns positive after 42 quarters. In the soft constraint, the change in consumption also becomes positive after 42 quarters.

The right hand side panel of Figure 7 shows how the deleverage takes place almost immediately in the hard constraint version, with the debt to income ratio dropping by 34% within the first year after the shock. On the other hand, the speed of adjustment is much slower in the soft constraint as it takes more than 12 quarters to produce the deleverage that the decrease in the borrowing limit achieved in a year. The mechanism behind these different responses is not surprising. In the hard constraint economy, a quick deleverage is fabricated by forcing households borrowing at the limit or close, to delever immediately, while also inducing greater precautionary savings among households not mechanically affected by the decrease in the limit. In the soft constraint version, the incentives to delever are provided through an increase in the borrowing rate, but highly-indebted households are still optimizing in an interior point of their budget constraint according to their Euler equation and so, it is optimal for them to follow a smoother adjustment path.

Figure 8 presents the differential responses to the shocks of groups of households defined in terms of their assets position when the shock hits. With a drop of 21% in their consumption, borrowers drive the adjustment in the hard constraint economy, while savers virtually do not react. Households borrowing at the limit when the shock hits reduce their consumption by more than 50%, but since there is relatively few of them, most of the change in borrowers' consumption is explained by unconstrained borrowers' behavior. In a similar fashion, the adjustment in the soft constraint is driven by the borrowers' response, although their consumption goes down by only 7% when the shock hits, and savers are again almost completely unaffected.

Thus, the drop in aggregate consumption is three times more severe in the hard constraint economy than in the soft constraint when new financial conditions take place immediately. In both cases the adjustment is led by borrowers, but it is more rapid in the hard constraint as these households need to delever faster to avoid the risk of hitting the borrowing limit.

Next, I use the time path in Guerrieri and Lorenzoni (2011) and I split the shock to financial conditions linearly along six quarters. Then, in the hard constraint economy at time 1, when the shock hits, households learn that borrowing conditions will worsen over the following quarters: that quarter they will not be able to borrow more than 19,476, the next one the limit will be

Figure 7: Unexpected Tightening of the Borrowing Conditions. Immediate Adjustment

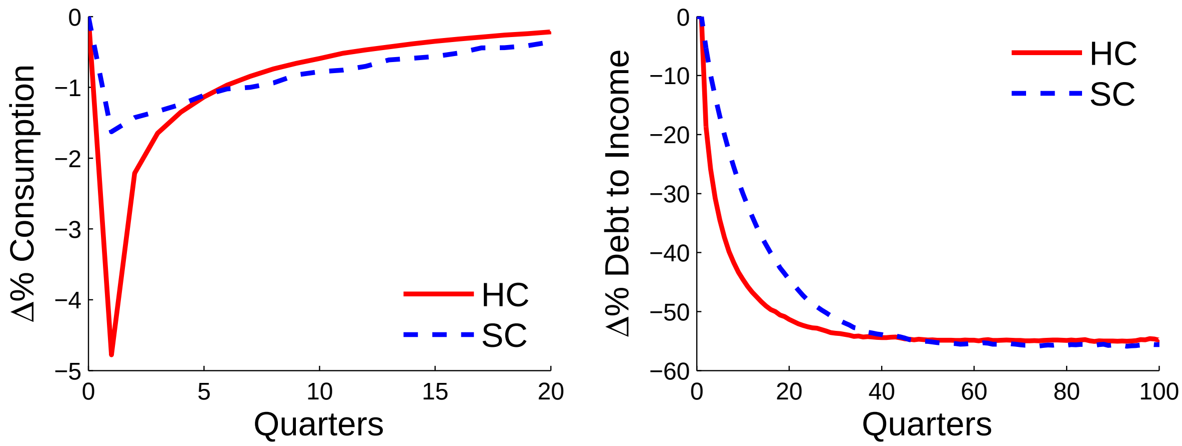

Consumption and debt decrease in both economies after the credit crunch, but in the hard constraint economy the drop is much more severe.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.
HC. Initial Borrowing Limit: -20,952. Final Borrowing Limit: -12,209.
SC. Initial Borrowing Spread: 0.0079. Final Borrowing Spread: 0.0134.

18,022, and so on. Finally, by the sixth quarter, the credit limit will settle into the new steady level, 12,209. I assume the same linear path also applies to the borrowing spread parameter in the soft constraint economy.

The consumption response is again much milder in the soft constraint (-1.2%) than in the hard constraint economy (-2.1%). As shown in Figure 10, borrowers lead the contraction in both economies, although in the soft constraint the decrease is significantly milder. Unlike the previous specification where constrained borrowers' consumption dropped by 40%, here their response is similar to that of the unconstrained borrowers as the adjustment path allows for a more gradual deleverage.

Figure 8: Unexpected Tightening of the Borrowing Conditions. Immediate Adjustment

Panel A. Hard Constraint. Savers and Borrowers

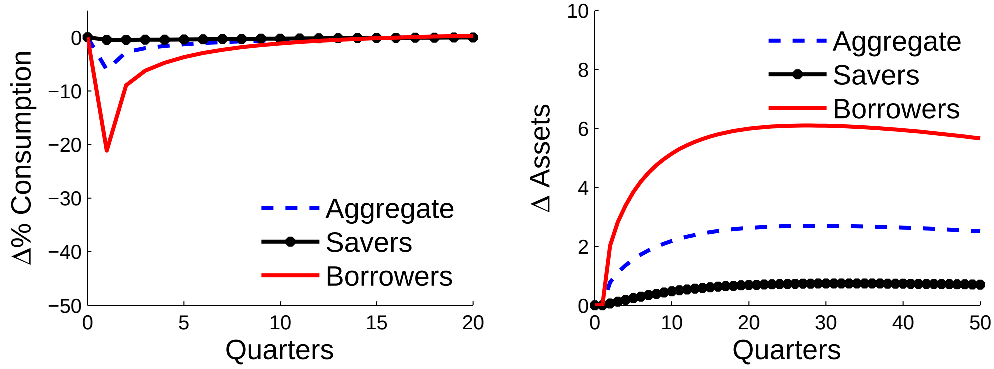

Panel B. Hard Constraint. Constrained Borrowers and Unconstrained Borrowers

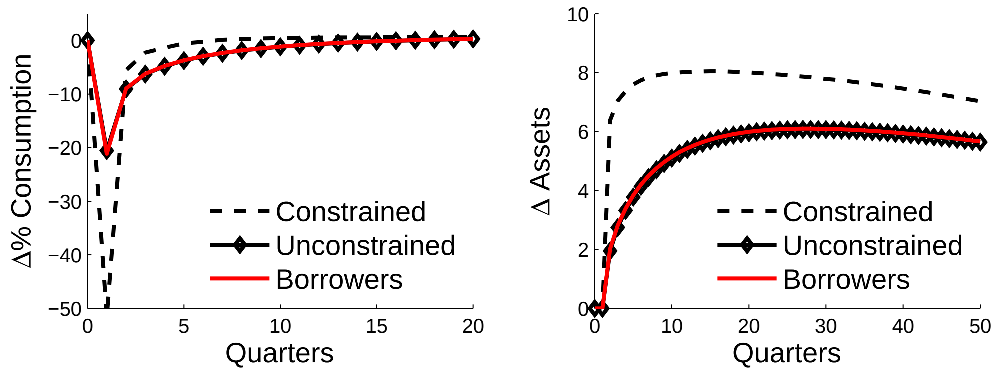

Panel C. Soft Constraint. Savers and Borrowers

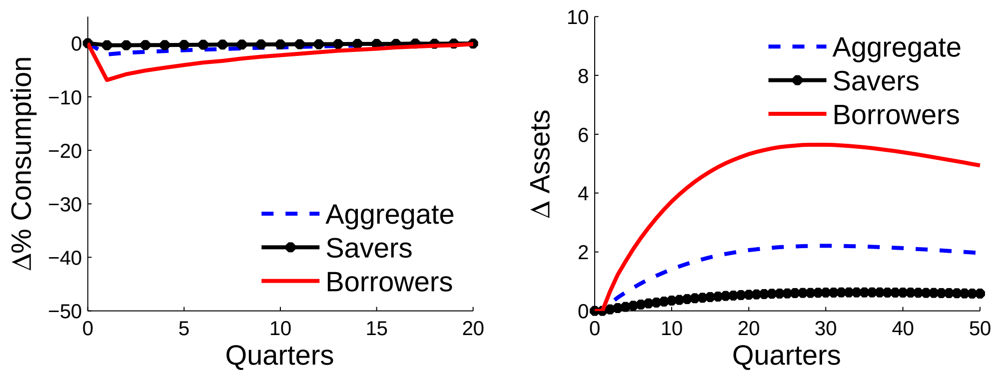

In both economies, the aggregate response is driven by borrowers reaction. But in the soft constraint, that response is significantly milder.

Borrowers (Savers) refers to the average response of households that had a negative (positive) level of assets when the shock hit. Aggregate refers to the average response of the economy as a whole. Constrained (Unconstrained) refers to the average response of households that had a negative level of assets when the shock hit and were (were not) at their borrowing limit.

Consumption and assets responses by group after an unexpected tightening of the credit conditions decreases the debt to income ratio in the new steady state by 55.56%. Consumption response is the percentage change with respect to the initial steady state group level. Assets response is expressed in thousands of US dollars with respect to the initial steady state group level. The x-axis measures the number of quarters after the credit crunch hits the economy.

Figure 9: Unexpected Tightening of the Borrowing Conditions. Adjustment in 6 quarters

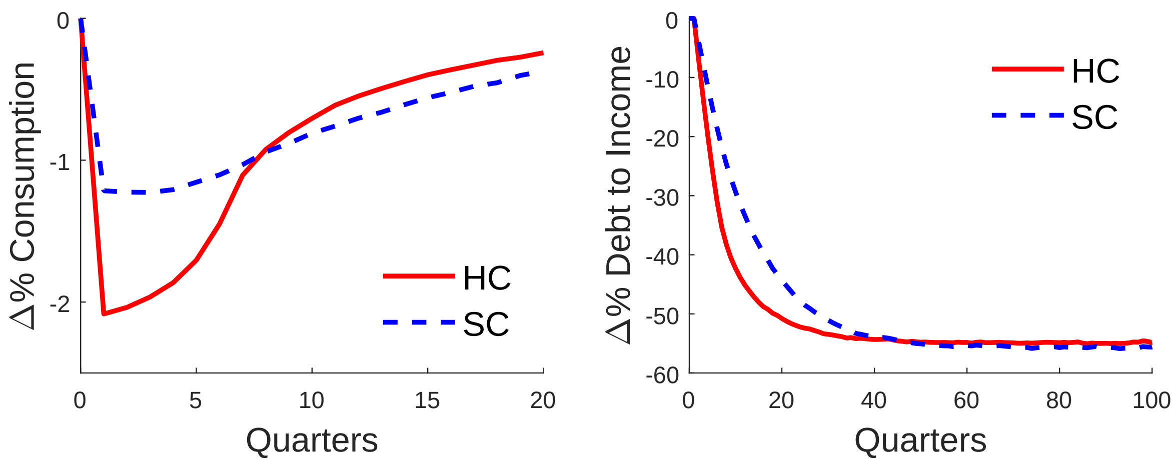

When the credit crunch is gradual, the consumption response is again milder in the soft constraint economy.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions decreases the debt to income ratio in the new steady state by 55.56%. Borrowing parameters are assumed to adjust linearly over six quarters. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.
HC. Initial Borrowing Limit: -20,952. Final Borrowing Limit: -12,209.
SC. Initial Borrowing Spread: 0.0079. Final Borrowing Spread: 0.0134.

Figure 10: Unexpected Tightening of the Borrowing Conditions. Adjustment in 6 quarters

Panel A. Hard Constraint. Savers and Borrowers

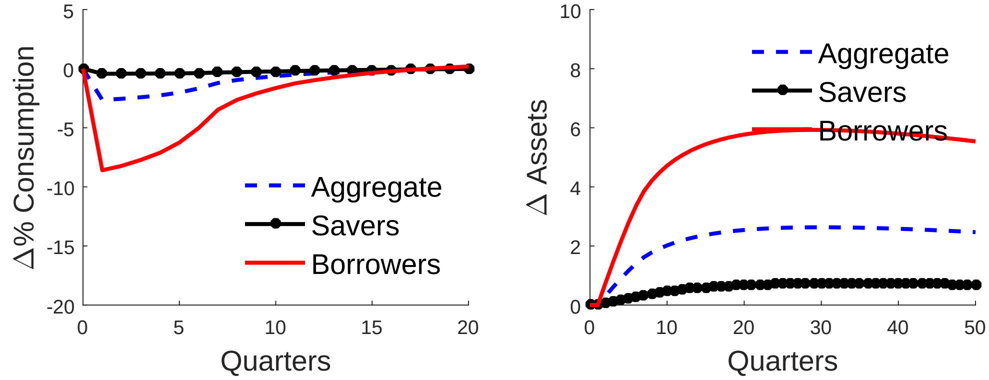

Panel B. Hard Constraint. Constrained Borrowers and Unconstrained Borrowers

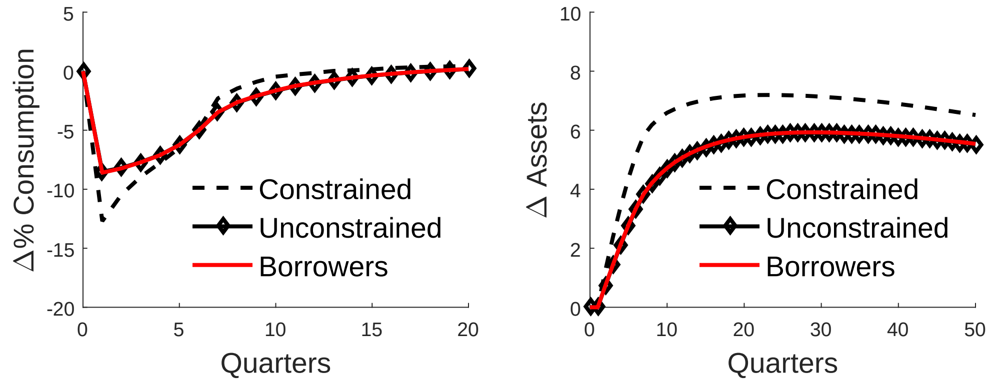

Panel C. Soft Constraint. Savers and Borrowers

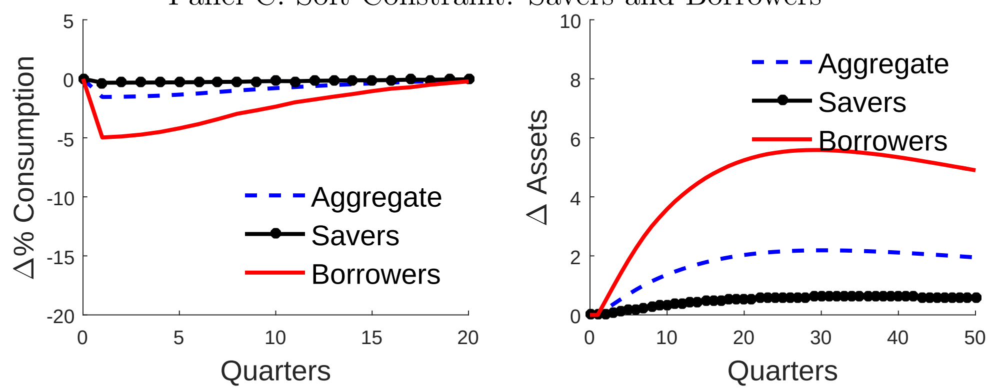

In both economies, the aggregate response is driven by borrowers reaction. But in the soft constraint, that response is significantly milder.
Borrowers (Savers) refers to the average response of households that had a negative (positive) level of assets when the shock hit. Aggregate refers to the average response of the economy as a whole. Constrained (Unconstrained) refers to the average response of households that had a negative level of assets when the shock hit and were (were not) at their borrowing limit.
Consumption and assets responses by group after an unexpected tightening of the credit conditions decreases the debt to income ratio in the new steady state by 55.56%. Borrowing parameters are assumed to adjust linearly over six quarters. Consumption response is the percentage change with respect to the initial steady state group level. Assets response is expressed in thousands of US dollars with respect to the initial steady state group level. The x-axis measures the number of quarters after the credit crunch hits the economy.

5.2 Expected Credit Crunch

An alternative experiment to account for the fact that the agents could anticipate the arrival of the shock is to allow households to find out about the credit crunch before it actually takes place. Such a setting will permit the evaluation of the adjustment path chosen by households reducing the concern of the result being driven mechanically by a forced deleverage of those households at the borrowing limit. In this case, at time 1, households learn that within a year borrowing conditions will worsen and new financial conditions will apply. I assume that the magnitude of the shock (i.e. a new borrowing limit of \$12,209 and a new borrowing interest rate spread of 0.0134 such that in the final steady state the debt to income ratio will be 2.47% in both economies) is exactly as before so the only difference between the credit crunch with immediate adjustment studied in the previous section and this one is that here households find out about the forthcoming tightening one year before it actually happens.

The arrival of news about a forthcoming credit crunch reduces the speed of the adjustment but does not alter the dynamics significantly. As soon as households learn about the shock, they reduce their consumption and start to delever. Since the incentives to delever in this case are weaker, so is the consumption response. Aggregate consumption decreases by 2.4% in the hard constraint economy and by 1.2% in the soft constraint. Thus, anticipating the arrival of a credit crunch reduces significantly the magnitude of the aggregate response to the shock in both economies. The aggregate response is still greater in the hard constraint model.

Figure 11: Expected Tightening of the Borrowing Conditions. News arrive 1 year earlier

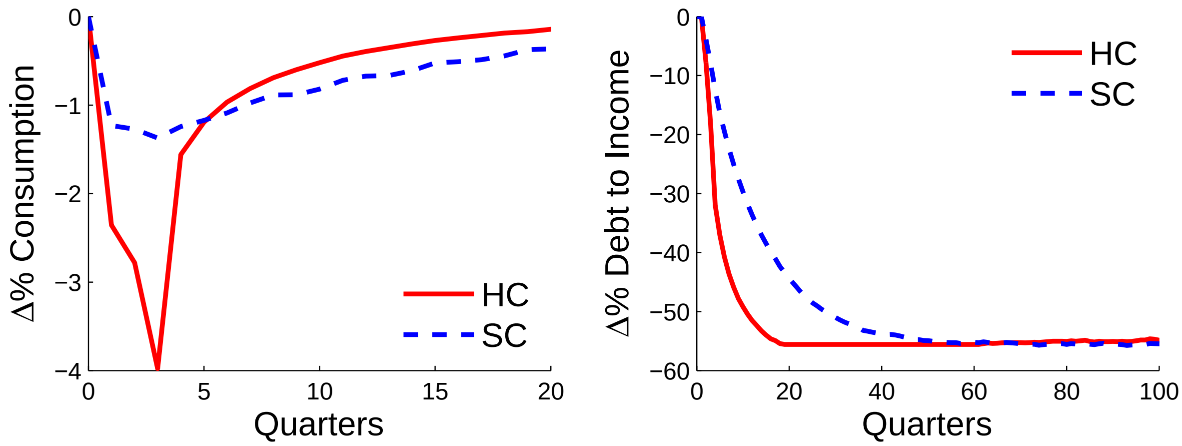

The consumption response to an expected credit crunch is milder in the soft constraint economy.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after the arrival of news of a tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. The news arrive four quarters before the actual tightening takes place. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.
HC. Initial Borrowing Limit: -20,952. Final Borrowing Limit: -12,209.
SC. Initial Borrowing Spread: 0.0079. Final Borrowing Spread: 0.0134.

5.3 Short-Run Credit Crunch

The previous exercises defined a credit crunch as a change in borrowing parameters that reduces the amount of debt in the economy in the long run and examine the aggregate implications in the short run. However, a credit crunch could also be defined in terms of a sudden decrease in the amount available to be borrowed in the very short run. In this section, I try to capture this idea. I pick the new financial parameters, the borrowing limit in the hard constraint and the borrowing spread in the soft constraint, such that they produce the same drop in consumption (implying the same decrease in loanable funds) in the period when the shock hits. I assume the shock requires immediate adjustment. From there, the long-run implications will be compared. A more severe tightening of borrowing conditions will show up as a greater drop in the amount of debt in the final steady state.

As a target for calibration, I look at the consumption response at the beginning of the Great Recession. Obviously not all of that drop in consumption is attributable to the credit crunch. Many other explanations are also possible: increasing unemployment risk, the collapse of the housing bubble, a change in expectations, and so on. The question here is how large would the

Figure 12: Unexpected Tightening of the Borrowing Conditions. Immediate Adjustment

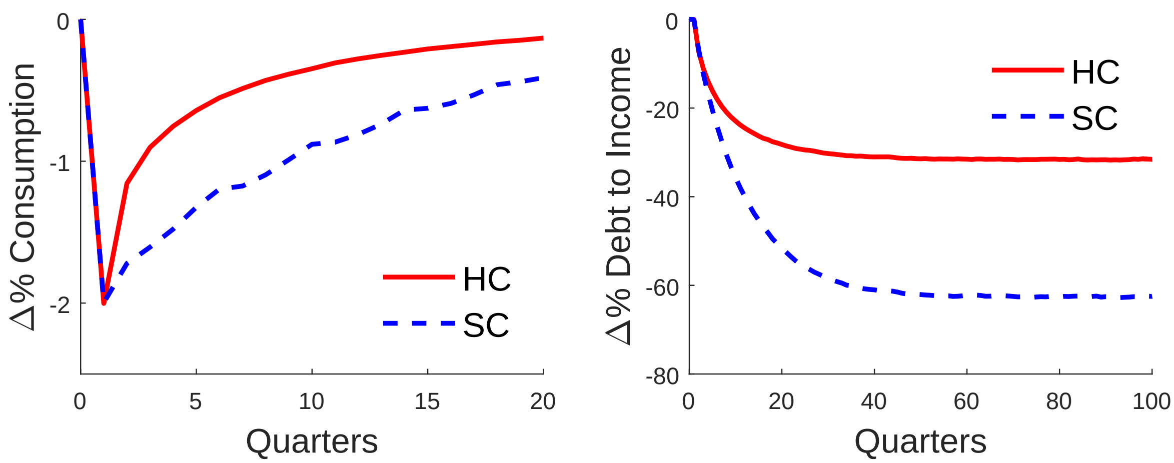

To produce the same initial drop in consumption, a much more severe reduction of the level of debt must take place in the soft constraint economy.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that produces an immediate drop in aggregate consumption of 2%. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.
HC. Initial Borrowing Limit: -20,952. Final Borrowing Limit: -16,178.
SC. Initial Borrowing Spread: 0.0079. Final Borrowing Spread: 0.0149.

credit crunch have needed to be if all those other factors had been muted. Following Blinder (2013) I define the beginning of the credit crunch to be the bankruptcy of Lehman Brothers on September 15, 2008. I target the average drop in consumption per quarter in the following year, which was 2%.

To produce an immediate drop in consumption of 2%, the new borrowing limit in the hard constraint is 16,178, a 22.7% reduction from the initial borrowing limit. In the soft constraint model, the interest rate needs to increase from 7.3% to 10.3%, a 87.0% increment in the borrowing spread. Figure 12 shows the aggregate dynamics that follow the shock. The right panel shows that the shock in the soft constraint economy adversely affects the long run twice as much as it does in the hard constraint economy. The credit crunch reduces equilibrium debt in the economy by 62% in the soft constraint economy, but only 32% in the hard constraint economy. The tightening of financial conditions is significantly more severe in the hard constraint economy.

6 Robustness

6.1 Heterogeneity in Borrowing Conditions

In the data, there is a great deal of heterogeneity in the financial conditions that households face (Table 2). In this section, I take this heterogeneity as given and study its implications for the consumption change induced by a credit crunch. Since the model does not include default, I assume agents are born with a credit type which they will carry out for life. Considering the lack of life-cycle features of the interest rates on unsecured debt, this assumption does not seem completely implausible. This credit type will determine the agent's borrowing limit in the hard constraint economy or her borrowing interest rate in the soft constraint version.

I define five financial types. Each household draws her type at the beginning of her working life from a uniform distribution independent from any of her other characteristics. This assumption is motivated by the lack of predictive power displayed by observables in the SCF 2007[17].

I use the empirical distribution of borrowing conditions in the SCF 2007 to calibrate the financial types in each model. In the hard constraint model, I define the borrowing limit of the n^{th} type to be the mean of the credit card borrowing limit of the n^{th} quintile in the empirical distribution of the borrowing limit from the SCF 2007. The credit card limit is not the ideal target as households can usually obtain credit from other sources such as consumer and educational loans, secured credit or informal borrowing from family and friends. But it is the only measure available in the data. On the other hand, the target for the soft constraint model, the borrowing spread, is observable in the data. For consistency, I assume the borrowing rate of the n^{th} type to be the mean of the credit card interest rate of the n^{th} quintile. Table 2 presents the levels used for the hard constraint borrowing limit (Column 1) and the soft constraint borrowing interest rate (Column 3).

After having imputed directly from the data the borrowing conditions, there is only one free parameter to pick, β, which I set to match the aggregate debt to income ratio in the economy, 5.56%. In the hard constraint, this yields a $\beta = 0.9886$ (annual discount factor of 0.9552), whereas in the soft constraint we obtain $\beta = 0.9882$ (annual discount factor of 0.9534). Agents are slightly more impatient in the soft constraint model.

This relative impatience of agents in the soft constraint economy slightly reduces the amount of aggregate wealth in the initial equilibrium relative to the hard constraint case. The ratio

[17] Results are not shown, but are available upon request. Observables include demographic characteristics, income, and net wealth. The SCF does not include credit scores, which are probably key to explain the cross-sectional variance in borrowing conditions. The lack of this variable prevents me from estimating a function of the borrowing spread on the amount of debt directly in the data as such estimation will likely suffer from omitted variable bias.

Table 2: Credit Card Financial Conditions by Quintile. SCF 2007

	Borrowing Limit		Interest Rate		Credit Crunch	
	Mean	Median	Mean	Median	Limit	Rate
1^{st} Quintile	1,495	1,100	4.16	4.90	751	5.16
2^{nd} Quintile	6,288	6,000	9.22	9.00	3,159	46.34
3^{rd} Quintile	13,201	13,000	12.75	12.50	6,631	80.40
4^{th} Quintile	27,628	26,000	16.87	17.00	13,879	126.03
5^{th} Quintile	72,593	60,000	22.27	21.00	36,467	195.93

Columns 2 to 5 are obtained from the SCF 2007. Borrowing limit refers to the reported total limit on the household's credit cards. Interest rate refers to the rate paid on credit card debt.
Columns 6 and 7 represent the values that these variables need to take after an unexpected and proportional tightening of the credit conditions decreases the debt to income ratio in the new steady state by 55.56%.

of gross wealth to income is 1.74 in the soft constraint model, whereas in the hard constraint model it is 1.87. Both are well above the gross wealth to income ratio in the data: 1.47.

I now repeat the credit crunch exercise assuming a homogeneous shock among financial types that leads to a reduction of 55.56% in the amount of debt in the new steady state. First, I compute the ratio by which the five borrowing limits should decrease simultaneously for the final steady state to have a debt to income ratio of 2.47%. The result is 0.502, i.e. the borrowing limit for the first type will go down from 1,495 to 751, for the second group it will decrease from 6,288 to 3,159, and so on. The complete list is in Column 5 of Table 2. Second, in the soft constraint model, I look for the homogeneous increase in the quarterly borrowing spread which produces the same drop in the level of debt to GDP in the economy. In this calibration, the five spreads must simultaneously increase by 624%. For the first financial type this would mean that the annual borrowing rate increases from 4.16% to 5.16%, whereas for the fifth group it will soar from 22.27% to 195.93%.

Only under this unrealistic and massive increment in the borrowing interest rates the severity of the consumption drop in the soft constraint economy becomes comparable to that in the hard constraint. Figure 13 shows the effects of this credit crunch in consumption and debt when immediate adjustment is assumed. The response of aggregate consumption in the soft constraint is still milder: 2.9% versus 3.7% in the hard constraint. It is only when the adjustment is allowed to take place over six periods, with a linear, gradual adjustment of the borrowing parameters, shown in Figure 14, that the drop in consumption is smaller in the hard constraint economy: 1.6% versus 2.0%. Thus, for the drop in consumption in a soft constraint economy to be

Figure 13: Unexpected Tightening of the Borrowing Conditions. Heterogeneous Financial Types. Immediate Adjustment

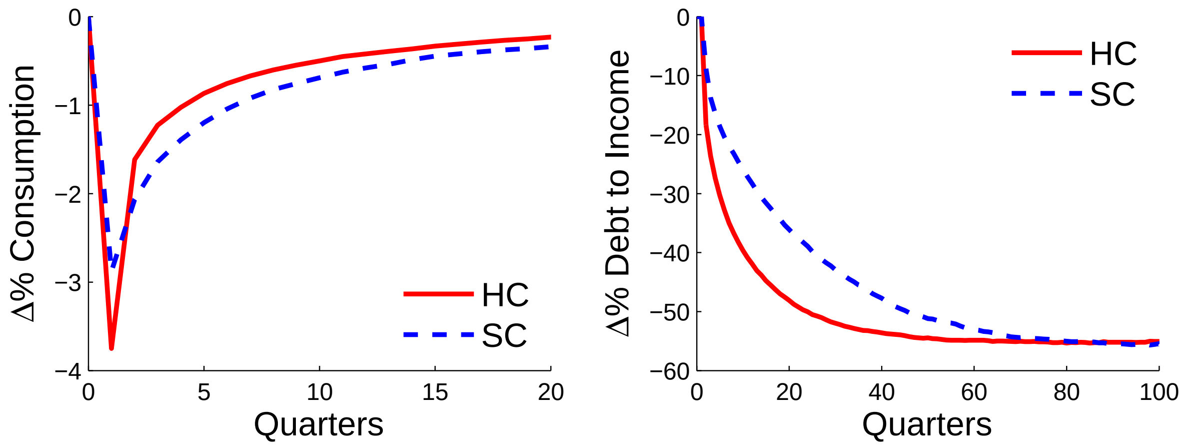

The credit crunch is less severe in the soft constraint economy
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Percentages are with respect to the initial steady state levels.
Initial and Final Borrowing Parameters as indicated in Table 2. HC: Columns 2 and 6. SC: Columns 4 and 7.

severe enough to exceed the one in the hard constraint, borrowing interest rates must increase massively, whereas borrowing limits just need to decrease by half.

Figure 14: Unexpected Tightening of the Borrowing Conditions. Heterogeneous Financial Types. Adjustment in 6 periods

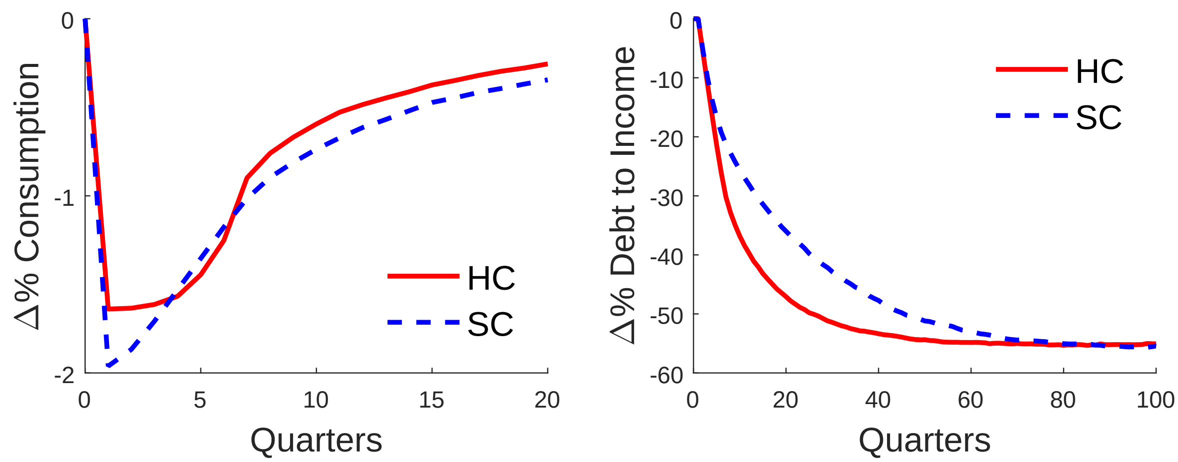

The credit crunch is more severe in the soft constraint economy when borrowing spreads soar by more than 600%
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Borrowing parameters are assumed to adjust linearly over six quarters. Percentages are with respect to the initial steady state levels.
Initial and Final Borrowing Parameters as indicated in Table 2. HC: Columns 2 and 6. SC: Columns 4 and 7.

6.2 Other Robustness Checks with Homogeneous Borrowing Conditions

Finally, I return to the case of homogeneous financial conditions to verify that the main result of the paper, that a credit crunch in a hard constraint economy is much more severe than in the soft constraint, is robust to alternative specifications.

Table 3 summarizes the robustness checks. First, I consider a non-degenerate initial distribution of assets estimated empirically from the SCF among agents younger than 25, instead of the assumption of no initial assets in the baseline model. Second, I calibrate the model to match the NIPA moments used in Guerrieri and Lorenzoni (2011): 1.78 for gross wealth to GDP and 0.18 for debt to GDP. Finally, I consider two other values for γ, the preference parameter, 0.5 and 4.

The main result of the paper is robust to all these exercises. The contraction in consumption when the credit crunch is modeled within a hard constraint framework is much more severe than when it happens through interest rates.

Table 3: Robustness Checks

	Initial Assets		NIPA Target		$\gamma = 0.5$		$\gamma = 4$	
	HC	SC	HC	SC	HC	SC	HC	SC
Parameters								
β	0.950	0.950	0.957	0.957	0.956	0.956	0.897	0.897
Borrowing Limit	(20,539)		(100,546)		(20,095)		(26,764)	
Borrowing Interest Rate		7.61	3.99			5.90		9.17
Initial Steady State								
Wealth to Income	1.465	1.465	1.780	1.780	1.465	1.465	1.465	1.465
Debt to Income	0.056	0.056	0.180	0.180	0.056	0.056	0.056	0.056
Initial Drop in Consumption (Δ%) - Immediate Adjustment								
Aggregate	6.306	2.213	11.570	3.110	6.344	1.882	4.574	2.291
Borrowers	22.260	7.680	30.418	7.953	21.558	6.030	18.806	9.193
Savers	0.436	0.369	0.126	0.174	0.480	0.377	0.363	0.285
Initial Drop in Consumption (Δ%) - Adjustment in 6 quarters								
Aggregate	2.650	1.634	4.792	2.411	2.635	1.501	2.657	1.733
Borrowers	8.739	5.489	12.486	6.120	8.310	4.708	10.442	6.775
Savers	0.410	0.333	0.120	0.162	0.447	0.338	0.353	0.267
Final Steady State								
Wealth to Income	1.524	1.520	1.933	1.941	1.523	1.517	1.528	1.524
Debt to Income	0.025	0.025	0.080	0.080	0.025	0.025	0.025	0.025

HC: Hard Constraint. SC: Soft Constraint.
NIPA Target: Guerrieri and Lorenzoni (2011)'s moments for debt and wealth to income (0.18 and 1.78) from national accounts.
The interest rates and discount factors are reported in annual levels.

7 Conclusions

The Great Recession raised once again the question of whether credit conditions can affect the real economy. The literature has modeled credit conditions in the form of a borrowing limit (hard constraint) or a borrowing spread (soft constraint). In this paper, I have built on the work by Guerrieri and Lorenzoni (2011) and I have shown that the macroeconomic implications of a worsening in the financial conditions largely depend on the modeling approach. I have argued that the standard approach to model financial constraints, the hard constraint, mechanically creates a significant drop in consumption by forcing households to deleverage. Modeling financial frictions with a soft borrowing constraint only yields a minor decrease in consumption to a credit crunch. My results highlight the importance of examining alternative explanations for the phenomena observed in the recent crisis, while cast doubt on the use of the soft constraint as an appropriate modeling device.

Further work should attempt to endogeneize the borrowing spread by introducing default in the model as in Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007) and Livshits, MacGee, and Tertilt (2007). In such a model, an aggregate shock would increase the probability of default, which in turn would increase the borrowing spread, amplifying the initial shock. Finally, the inclusion of an illiquid asset, e.g. housing, as in Kaplan and Violante (2014) would provide with a complete framework to test the alternative explanations to the recent crisis and compare the results with the empirical evidence.

References

Aiyagari, S. R. (1994, August). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Athreya, K., X. S. Tam, and E. R. Young (2012). A Quantitative Theory of Information and Unsecured Credit. *American Economic Journal: Macroeconomics 4*(3), 153–183.

Attanasio, O., J. Banks, C. Meghir, and G. Weber (1999). Humps and Bumps in Lifetime Consumption. *Journal of Business and Economic Statistics 17*(1), 22–35.

Bernanke, B. S. (1983). Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression. *American Economic Review 73*(3), 257–76.

Blinder, A. (2013). *After the Music Stopped: The Financial Crisis, the Response, and the Work Ahead*. Penguin Press.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption Inequality and Partial Insurance. *American Economic Review 95*(5), 1887–1921.

Buera, F. and B. Moll (2015). Aggregate Implications of a Credit Crunch: The Importance of Heterogeneity. *American Economic Journal: Macroeconomics 7*(3), 1–42.

Carroll, C. (2005). The Method of Endogenous Grid Points for Solving Dynamic Stochastic Optimal Problems. *Economic Letters 91*, 312–20.

Carroll, C. and M. Kimball (1996). On the Concavity of the Consumption Function. *Econometrica 64*(4), 981–992.

Castañeda, Ana, J. Díaz-Gimenez, and J. V. Ríos-Rull (2003). Accounting for the U.S. Earnings and Wealth Inequality. *Journal of Political Economy 111*(4), 818–857.

Chatterjee, S., D. Corbae, M. Nakajima, and J. V. Ríos-Rull (2007, November). A Quantitative Theory of Unsecured Consumer Credit with Risk of Default. *Econometrica 75*(6), 1525–1589.

Deaton, A. S. (1991). Saving and Liquidity Constraints. *Econometrica 59*(5), 1221–1248.

Eaton, J. and M. Gersovitz (1981, April). Debt with Potential Repudiation: Theoretical and Empirical Analysis. *The Review of Economic Studies 48*(2), 289–309.

Eggertsson, G. B. and P. Krugman (2012, August). Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. *The Quarterly Journal of Economics 127*(3), 1469–1513.

Fernández-Corugedo, E. (2002). Soft Liquidity Constraints and Precautionary Saving. Bank of England.

Fernández-Villaverde, J. and D. Krueger (2007). Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data. *The Review of Economics and Statistics 89*(3), 552–565.

Flodén, M. and J. 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.

Gouveia, M. and R. P. Strauss (1994). Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis. *National Tax Journal 47 (2)*, 317–39.

Guerrieri, V. and G. Lorenzoni (2011). Credit Crises, Precautionary Savings, and the Liquidity Trap. NBER Working Paper No. 17583.

Guvenen, F. (2009). An Empirical Investigacion of Labor Income Processes. *Review of Econ 12*, 58–79.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2015). What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk? Mimeo.

Heathcote, J., F. Perri, and G. L. Violante (2010). Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006. *Review of Economic Dynamics 13*, 15–51.

Huggett, M. (1996). Wealth Distribution in Life-Cycle Economies. *Journal of Monetary Economics 38*, 469–494.

Huo, Z. and J.-V. Ríos-Rull (2016). Financial Frictions, Asset Prices, and the Great Recession. Mimeo.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2015). Household Leveraging and Deleveraging. *Review of Economic Dynamics 18*(1), 3–20.

Kaplan, G. and G. L. Violante (2010). How Much Consumption Insurance Beyond Self-Insurance? *American Economic Journal: Macroeconomics 2*, 53–87.

Kaplan, G. and G. L. Violante (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica 82*(4), 1199–1239.

Livshits, I., J. MacGee, and M. Tertilt (2007). Consumer Bankruptcy: A Fresh Start. *American Economic Re 97*(1), 402–418.

Midrigan, V. and T. Philippon (2011, November). Household Leverage and the Recession. Mimeo.

Mincer, J. (1958, August). Investment in Human Capital and Personal Income Distribution. *Journal of Political 66*(4), 281–302.

Telyukova, I. A. (2013). Household Need for Liquidity and the Credit Card Debt Puzzle. *Review of Economic Studies 80*(3), 1148–1177.

A Appendix

A.1 Annual Transitory Shock

In this Appendix, I explore the implications of an annual transitory shock, instead of the quarterly shock employed in the main section of the paper. I show that the credit crunch still leads to a more severe drop in consumption in the hard constraint economy than in the soft constraint one.

The model I have used in this paper is at a quarterly frequency, however data on earnings is available only at the annual level. Estimations of the earnings process have typically used individual longitudinal data on annual labor income or wages as found in the Panel Study of Income Dynamics (Flodén and Lindé (2001), Guvenen (2009)) or, more recently, using U.S. administrative records (Guvenen, Karahan, Ozkan, and Song (2015)). And so, no information is available regarding the nature of the transitory shock at higher frequencies than a year. Guerrieri and Lorenzoni (2011) opt to introduce transitory risk at the quarterly level through an unemployment shock. In the main section of this paper, I assumed households draw a new transitory shock every quarter with the variance of this quarterly shock chosen to match, when aggregated, the variance of the annual transitory shock as in Kaplan and Violante (2010). Telyukova (2013) follows a similar approach when calibrating the income process at the monthly level. In this Appendix, I explore an annual specification for the transitory shock, which delivers the same variance at the annual frequency, but has very different implications for risk insurance.

I modify the baseline model described in Section 2 only in what respects to the transitory shock, $\epsilon_{i,t}$. I assume households now find out about the new realization of the transitory shock at the beginning of the calendar year and that realization lasts for four quarters[18]. I set the variance of this annual transitory shock to 0.05 as in Kaplan and Violante (2010). And I recalibrate the discount factor and the borrowing parameters to match debt and wealth to GDP ratios.

Figures A.1 and A.2 show that, after recalibrating the models, the frequency of the transitory shock does not produce major differences in the consumption policy functions, nor in the distribution of assets. The only noticeable difference in Figure A.1 is that in the soft constraint economy the Natural Borrowing Limit is greater in the case of annual shocks because the smallest value is not as small as the smallest value in the quarterly specification, which allows young workers to borrow more. In Figure A.2, the main difference between transitory shock specifications is in the debt distribution for the hard constraint economy where a lower calibrated borrowing limit produces a greater spike. However, the debt distributions even for

[18] I achieved this with a time-varying transition matrix for the Markov process, which implies an i.i.d. shock in quarters $1, 5, 9, ...$ and full persistence for the rest.

Figure A.1: Consumption Policy Functions

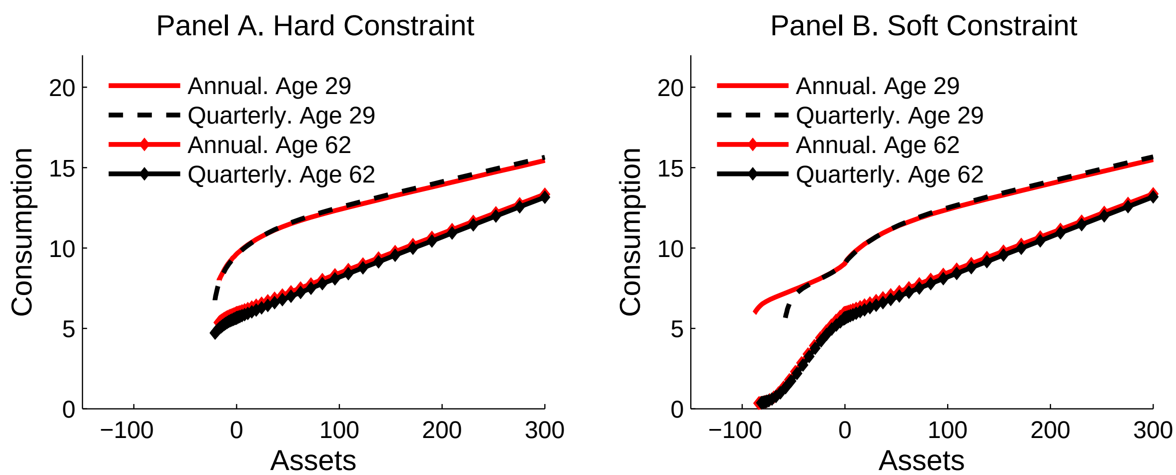

Consumption decisions are very similar regardless of the frequency of the transitory shock.
Quarterly consumption level as a function of the assets held by an agent of age 29 and 62 (15 and 150 in the model) with average earnings history, average permanent shock, average individual fixed effect, and average transitory shock. Consumption and assets are expressed in thousands of US dollars.
Annual corresponds to the annual specification for the transitory shock as described in the Appendix. Quarterly refers to the quarterly transitory shock employed in the main section of the paper.

that model look very similar below $15,000. The life cycle patterns for mean consumption and insurance coefficients[19] are also very similar across specifications (not shown).

Finally, I recalibrate the borrowing parameters in each economy to reproduce the credit crunch exercise. I find that the consumption response to a credit crunch (Figure A.3) is stronger in the hard constraint economy, but milder in the soft constraint when the transitory shock is annual rather than quarterly. So, the difference between the hard constraint and soft constraint is even magnified with respect to the results reported in the main section of the paper.

[19]The insurance coefficient for the transitory shock exhibits a bumpier evolution in the quarterly specification, being lower in the quarters where the new realization arrives but jumping up to 1 in the following quarters as the household smooths consumption. Nevertheless, this bumpier behavior oscillates around the insurance coefficient in the quarterly specification.

Figure A.2: Assets and Debt Distribution during Working Life

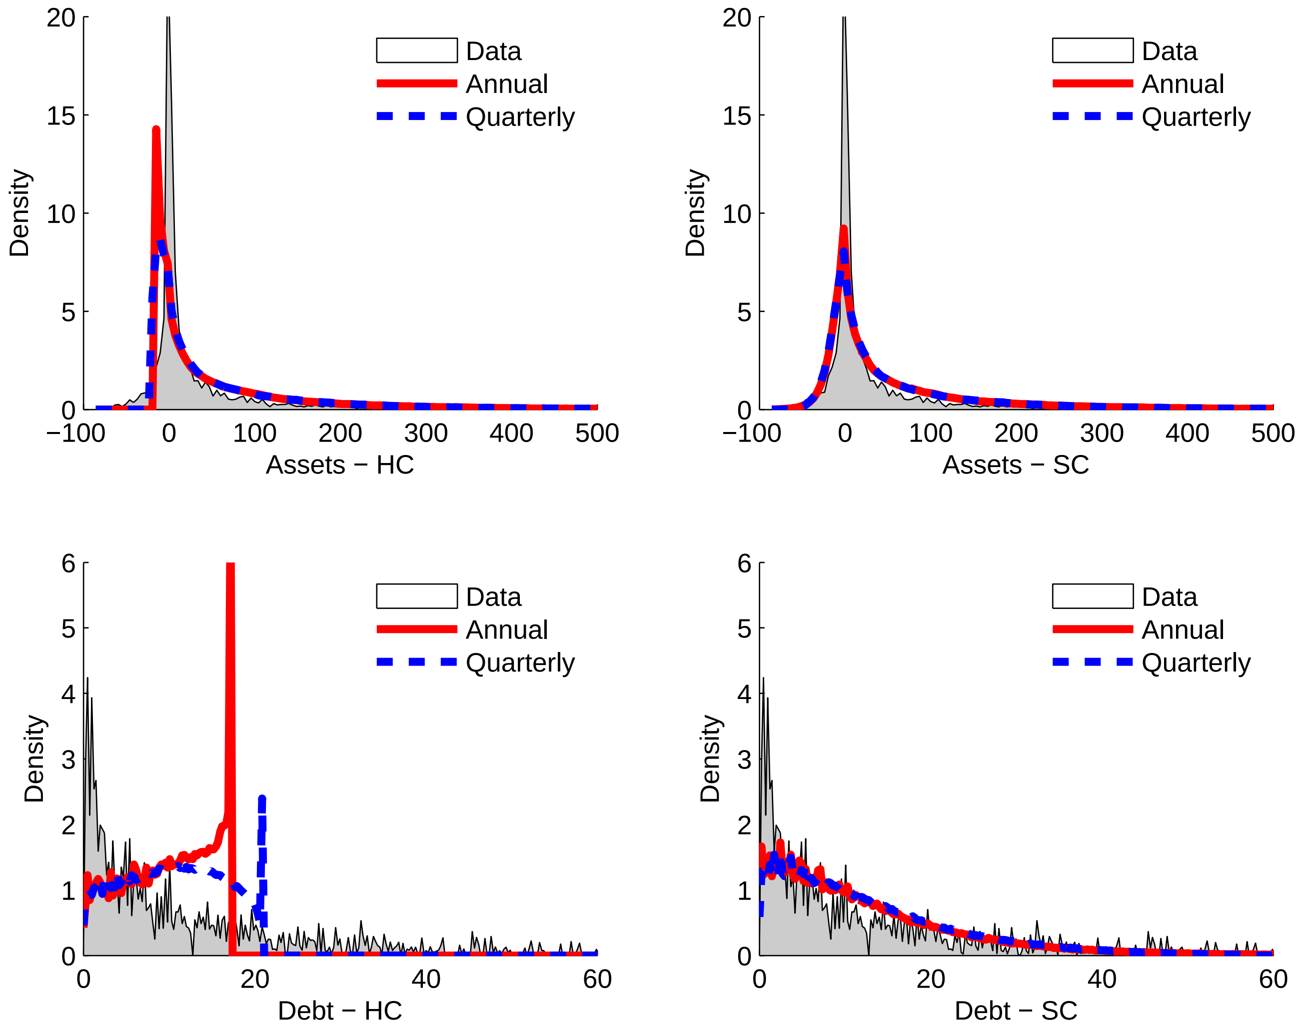

Assets and debt distributions are very similar regardless of the frequency of the transitory shock.
Assets and debt are in thousands of US dollars. The axis have been truncated to facilitate visualization.
Annual corresponds to the annual specification for the transitory shock as described in the Appendix. Quarterly refers to the quarterly transitory shock employed in the main section of the paper.
HC: Hard Constraint Model. SC: Soft Constraint Model.

Figure A.3: Unexpected Tightening of the Borrowing Conditions for Annual and Quarterly Transitory Shocks

Panel A. Hard Constraint

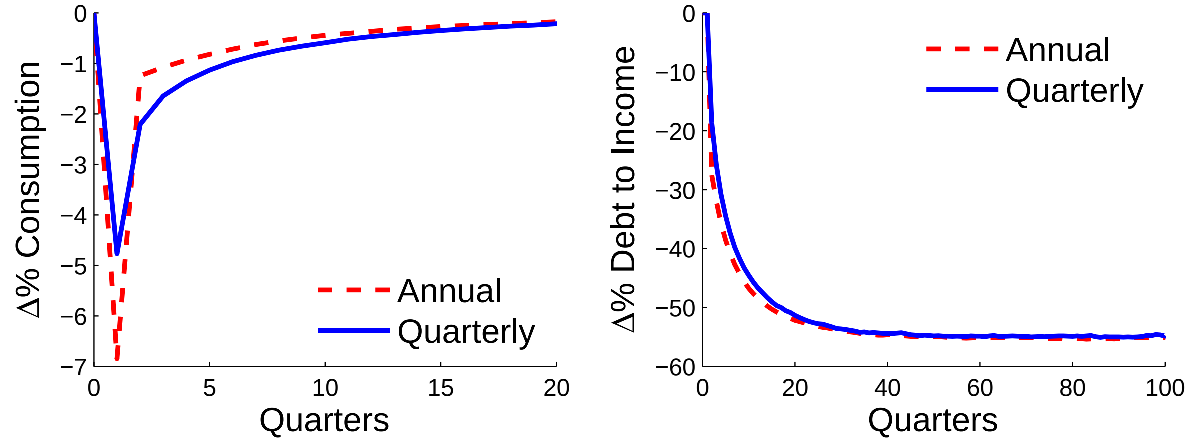

Panel B. Soft Constraint

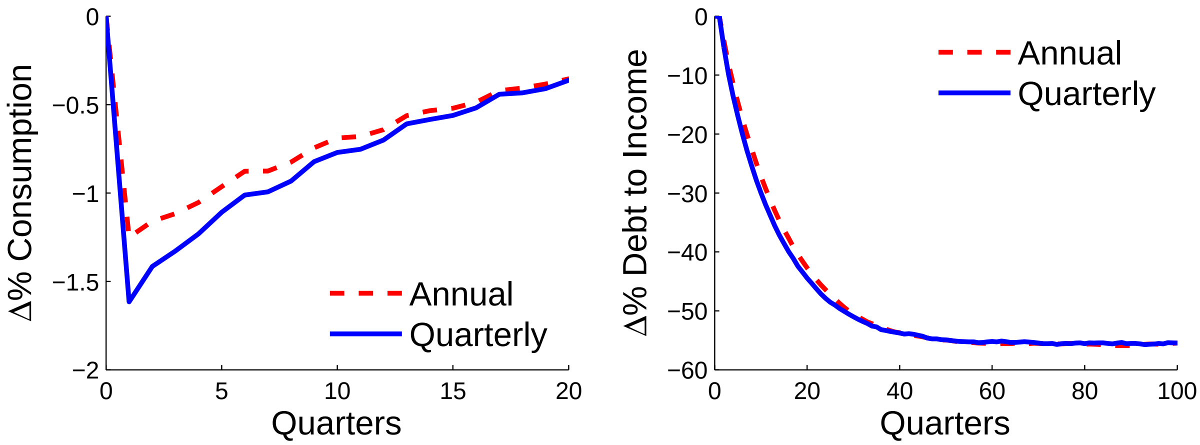

The credit crunch is much less severe in the soft constraint economy
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Percentages are with respect to the initial steady state levels.
Annual corresponds to the annual specification for the transitory shock as described in the Appendix. Quarterly refers to the quarterly transitory shock employed in the main section of the paper.

A.2 Hard Constraint as a Fraction of Natural Borrowing Limit

In this Appendix, I show that my results are robust to modeling the hard constraint as a fraction of the natural borrowing limit, rather than a constant amount across the lifecycle.

In the main section of the paper, I modeled the hard constraint as a limit on the amount households could borrow that was given exogenously and did not depend on household characteristics, i.e. $\bar{\phi}$. In that setting, households were able to borrow up to the minimum between the exogenous limit $\bar{\phi}$ and their natural borrowing limit NBL_t:

$$\bar{A}_t = min \left\{ \bar{\phi}, NBL_t \right\}$$

In this Appendix, I instead model the maximum amount that households can borrow as a fraction of their natural borrowing limits.

$$\bar{A}_t = \bar{\phi}_{NBL} NBL_t$$

where $\bar{\phi}_{NBL}$ is a number between 0 and 1. The smaller $\bar{\phi}_{NBL}$, the tighter the borrowing constraint because households can borrow up to a lower fraction of their natural borrowing limit.

As before, I recalibrate the discount factor and borrowing limit to match the wealth and debt to GDP ratios. I find that, under this new parametrization of the hard constraint, $\beta = 0.9874$ (equivalent to an annual discount factor of 0.9506) and $\bar{\phi}_{NBL} = 0.2215$.

Both in terms of lifecycle averages and distributions, the resulting model does not differ substantially from the baseline specification of the hard constraint discussed in the main body of the paper. Figure A.4 shows that the hard constraint does a slightly better job at matching the debt to income ratio early in life, but it overestimates the amount of debt later. On the other hand, the debt distribution is better matched for this new version of the hard constraint (Figure A.5) as there is not a single borrowing limit inducing a mass point. However, the soft constraint model still yields a better match to the distribution of both assets and debt. The Jensen-Shannon divergence between the model-derived debt distribution and its empirical counterpart is 0.1500 in the hard constraint economy, better than the 0.1556 in the baseline model, but still much greater than the 0.0718 in the soft constraint economy.

The two versions of the hard constraint produce very similar results in the initial steady state because in this model the main driver for borrowing is consumption smoothing early in life, when income is very low and is expected to grow significantly. As most of the debt is contracted in that period of life, rather than being distributed more smoothly over the lifetime, there is not much difference on effective borrowing limits in the cross section for agents willing to borrow heavily, because most young agents will have fairly similar natural borrowing limits.

Figure A.4: Debt to Income and Wealth to Income Ratios. Models vs. Data

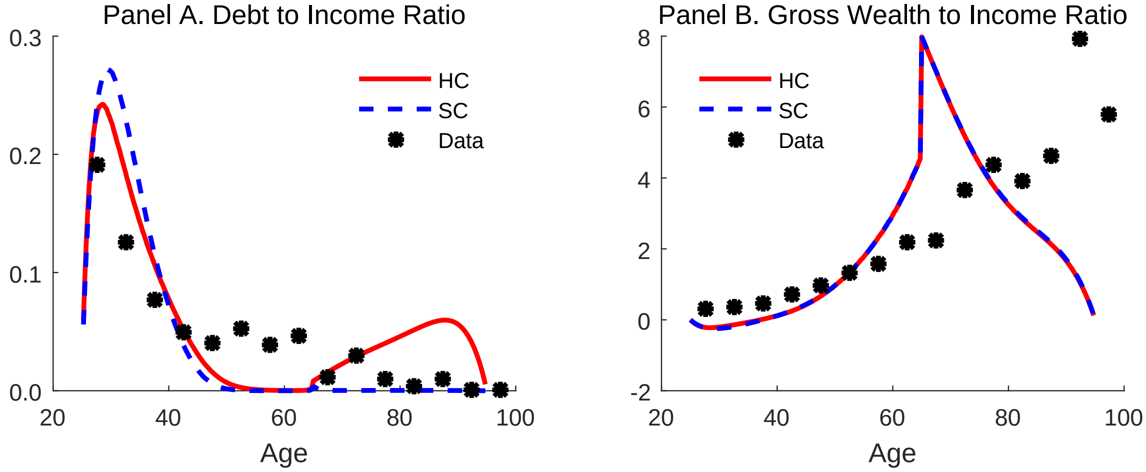

Both models predict a similar pattern for the ratios of wealth and debt to income, overstating the amount of debt early in life and understating it in middle age

Debt to income is the ratio of the absolute value of aggregate negative assets to aggregate annual income. Gross wealth to income is the ratio of aggregate positive assets to aggregate annual income.

HC: Hard Constraint Model. SC: Soft Constraint Model. Data: SCF 2007. Refer to the text for further details.

Then, a hard constraint expressed as a constant amount is not too different from one expressed as a fraction of the natural borrowing limit.

Finally, the milder response of consumption to a credit crunch in the soft constraint with respect to the hard constraint still holds for the new specification of the hard constraint. As before, I calibrate the tighter borrowing limit to obtain of drop of 55.56% in debt to income in the new steady state. The new value of the borrowing parameter is $\bar{\phi}_{NBL} = 0.1308$, a contraction of 41.0%, very close to the 41.7% found for the baseline specification of the hard constraint. Figure A.6 shows that when adjustment to the credit crunch is required to be immediate, the consumption drop is three times greater in the hard constraint. And when the adjustment is allowed to extend for six periods (Figure A.7), the consumption drop in the hard constraint is twice the response in the soft constraint model.

Figure A.5: Assets and Debt Distribution during Working Life

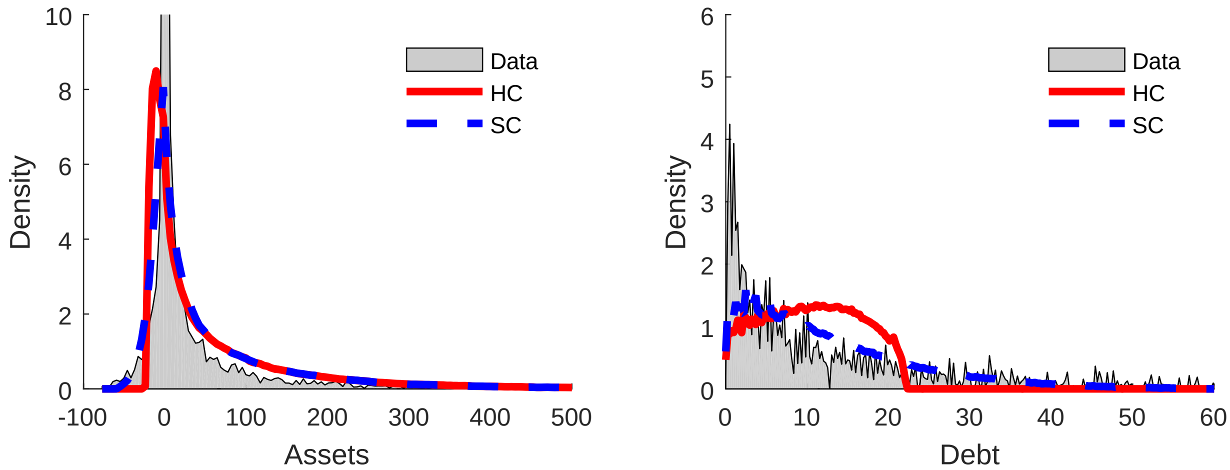

The soft constraint model produces assets and debt distributions closer to the data
Assets and debt are in thousands of US dollars. The axis have been truncated to facilitate visualization.
HC: Hard Constraint Model. SC: Soft Constraint Model. Data: SCF 2007. Refer to the text for further details.

Figure A.6: Unexpected Tightening of the Borrowing Conditions. Immediate Adjustment

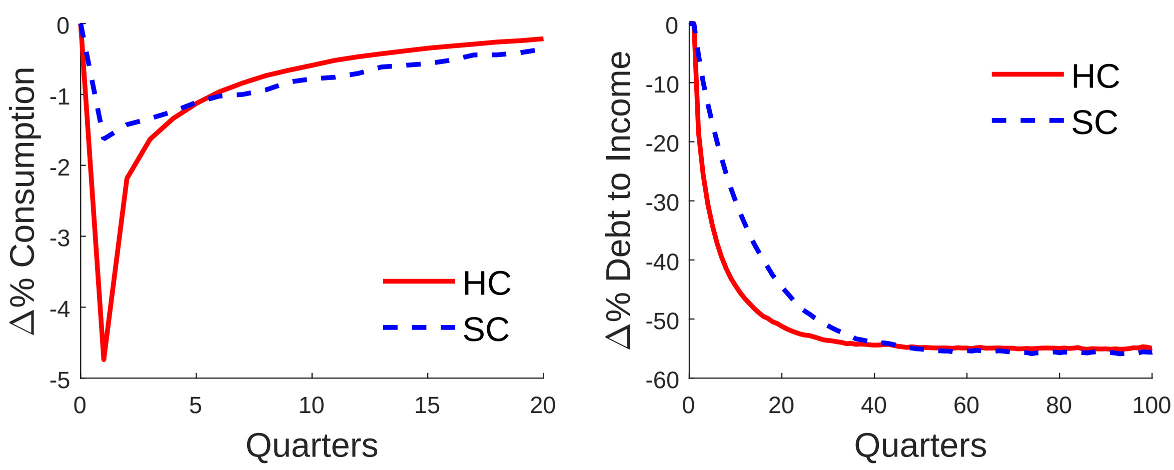

Consumption and debt decrease in both economies after the credit crunch, but in the hard constraint economy the drop is much more severe.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.

Figure A.7: Unexpected Tightening of the Borrowing Conditions. Adjustment in 6 quarters

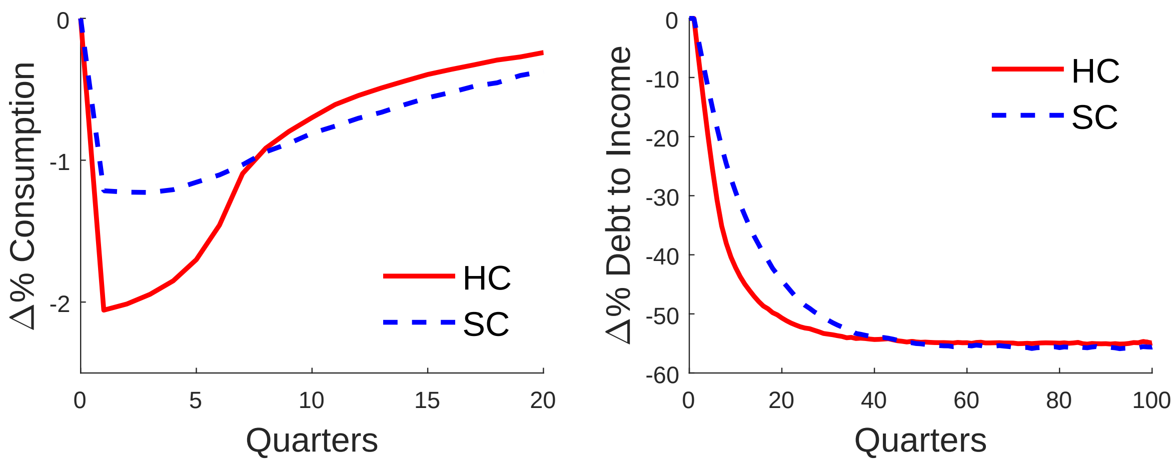

Consumption and debt decrease in both economies after the credit crunch, but in the hard constraint economy the drop is much more severe.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Borrowing parameters are assumed to adjust linearly over six quarters. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.

A.3 Soft Constraint with an Interest Rate Increasing in Debt

In this Appendix, I show that my results are robust to modeling the soft constraint as a strictly increasing function on the level of debt, rather than a constant spread.

In particular, I assume that the interest rate is given by:

$$R(A) = \begin{cases} R^f & \text{if } A \geq 0 \\ R^f + \phi(A) & \text{if } A < 0 \end{cases}$$

where:

$$\phi(A) = \phi_1 + \left(-\frac{A}{Amax}\right)^{\phi_2}$$

I use $Amax$ to normalize the amount of debt. By setting it equal to $300,000, the ratio $-\frac{A}{Amax}$ will be positive but lower than 1 for the relevant debt levels[20]. ϕ_1 allows for the existence of a discontinuity in the interest rate when the household becomes a debtor. If ϕ_2 is greater than 1, the borrowing cost will be increasing and convex on the amount of debt. Figure A.8 shows the borrowing interest rate for different values of the parameters used in the credit crunch exercises in this Appendix.

To keep the calibration approach as close as possible to the one in the main text, I start by assuming that in the initial steady state $\phi_1 = 0$. The borrowing cost function is then parametrized only by ϕ_2. As before, I use the Method of Simulated Moments to match the ratios of debt and wealth to GDP with the discount factor, β, and the borrowing parameter, in this case ϕ_2. I obtain $\beta = 0.9874$ (equivalent to an annual discount factor of 0.9507) and a borrowing parameter of $\phi_2 = 1.9978$, which implies an annual borrowing interest rate of 5.9% for a debt of $20,000 and a rate of 16.0% for a debt of $50,000.

The use of a strictly convex borrowing cost makes the soft constraint economy more similar to the hard constraint one. Figure A.9 shows that both models have now almost identical predictions in terms of average debt and wealth accumulation during the life cycle. The cross-sectional distribution of debt and assets (Figure A.10) is also now very similar between the two models, although the soft constraint does not exhibit a mass point. The Jensen-Shannon divergence between the model-derived debt distribution and its empirical counterpart in the soft constraint economy is now 0.1087, greater than the baseline model, but still better than the 0.1556 divergence in the hard constraint economy.

Next, I consider two versions of a credit crunch in this soft constraint environment. A credit crunch can take place either through an increase in ϕ_1, a shock that hits every borrower in the same magnitude, or through a decrease in ϕ_2, a shock that hits more severely highly indebted

[20] The maximum natural borrowing limit is around 150,000.

Figure A.8: A Convex Borrowing Interest Rate on the Amount of Debt

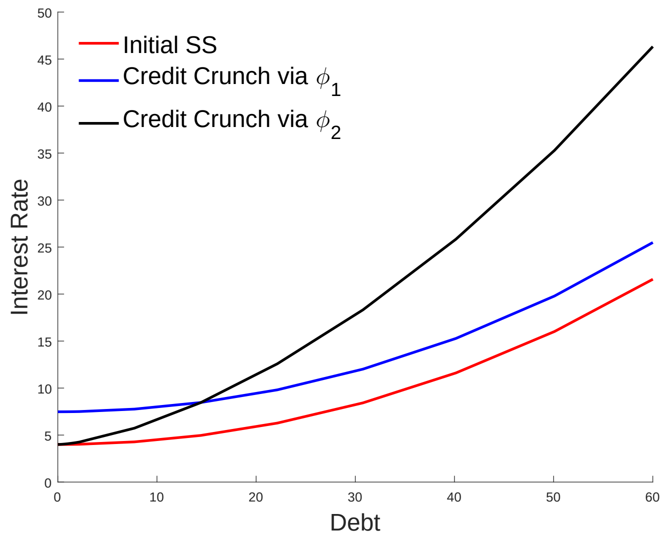

The functional form chosen to model a convex borrowing cost allows alternative definitions of a credit crunch.
Debt is in thousands of US dollars. Interest rate is reported in annual levels.
Initial SS corresponds to borrowing parameters $\phi_1 = 0$ and $\phi_2 = 1.9978$.
Credit Crunch via ϕ_1 corresponds to borrowing parameters $\phi_1 = 0.0083$ and $\phi_2 = 1.9978$.
Credit Crunch via ϕ_2 corresponds to borrowing parameters $\phi_1 = 0$ and $\phi_2 = 1.4967$.

borrowers. I separately recalibrate each of these borrowing parameters to obtain a reduction of 55.56% in the debt to income ratio in the new steady state. I find that such deleverage in the long run can be attained by an increase in ϕ_1 from 0 to 0.0083, while keeping constant the convexity of the borrowing cost function. In this setting, the annual borrowing rate of a \$20,000 debt would be 9.4%, whereas it would be 19.8% for a debt of \$50,000. Alternatively, the deleverage can be induced by a decrease in ϕ_2 to 1.4967 while keeping ϕ_1 equal to zero. Here, the annual borrowing rate of a \$20,000 debt would be 15.0%, whereas it would be 39.4% for a debt of \$50,000. As before, I study both the consequences of an immediate adjustment shock and of a gradual shock where the new borrowing parameters adjust linearly over six quarters.

Figures A.11 to A.12 show that the response to a credit crunch is still significantly milder in the soft constraint economy regardless if the credit crunch affected all the borrowers in a similar fashion (via ϕ_1) or hit disproportionately the heavily indebted (via ϕ_2). The aggregate consumption response appears only slightly lower for the former case. The mild response of consumption to a credit crunch operating through interest rates is then robust to the specification of the borrowing spread.

Figure A.9: Debt to Income and Wealth to Income Ratios. Models vs. Data

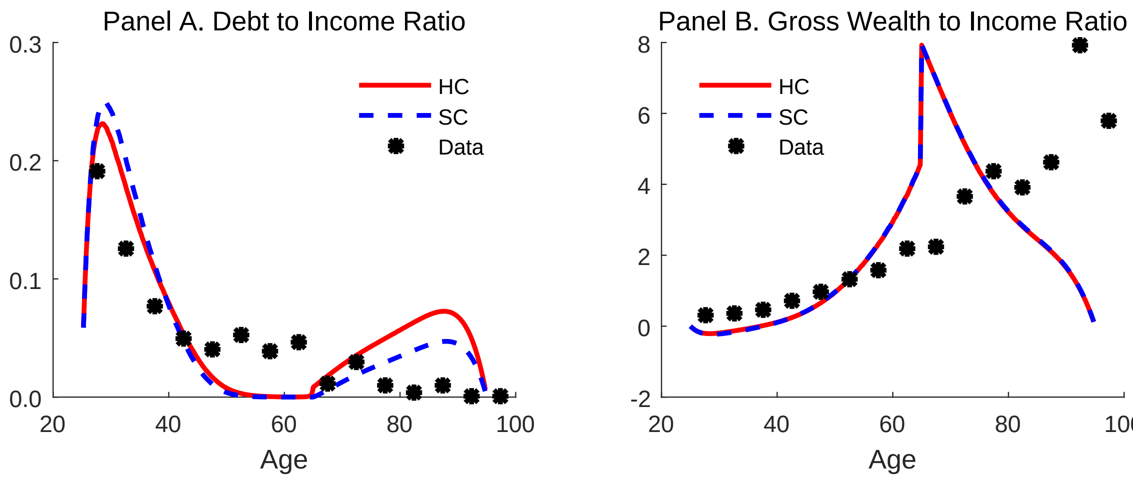

The soft constraint economy with a convex borrowing cost predicts very similar life cycle patterns of asset accumulation to the hard constraint economy ones.
Debt to income is the ratio of the absolute value of aggregate negative assets to aggregate annual income. Gross wealth to income is the ratio of aggregate positive assets to aggregate annual income.
HC: Hard Constraint Model. SC: Soft Constraint Model. Data: SCF 2007. Refer to the text for further details.

Figure A.10: Assets and Debt Distribution during Working Life

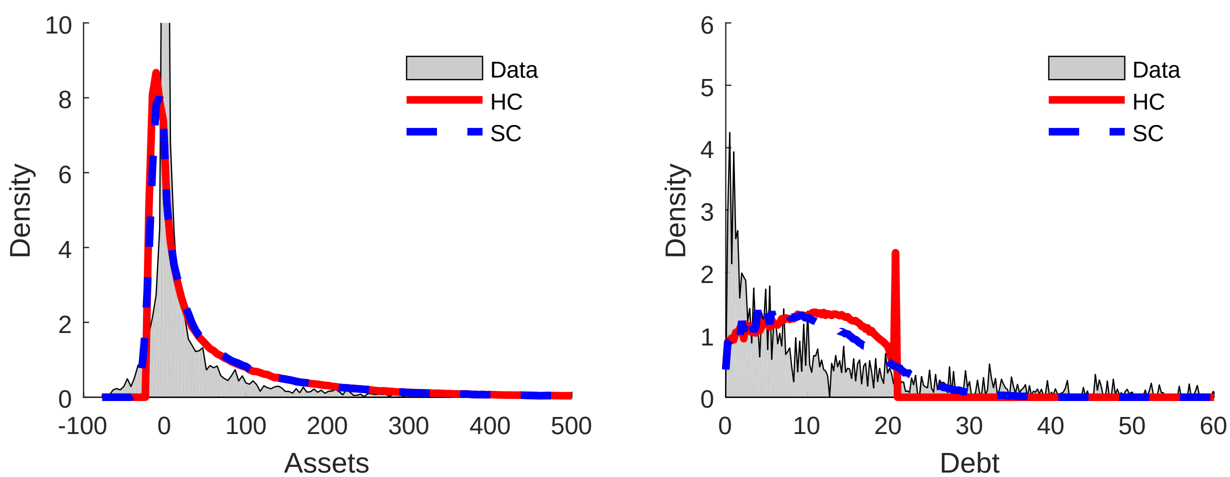

The soft constraint economy with a convex borrowing cost predicts very similar assets and debt distributions to the hard constraint economy ones.
Assets and debt are in thousands of US dollars. The axis have been truncated to facilitate visualization.
HC: Hard Constraint Model. SC: Soft Constraint Model. Data: SCF 2007. Refer to the text for further details.

Figure A.11: Unexpected Tightening of the Borrowing Conditions. Immediate Adjustment - Credit Crunch via ϕ_1

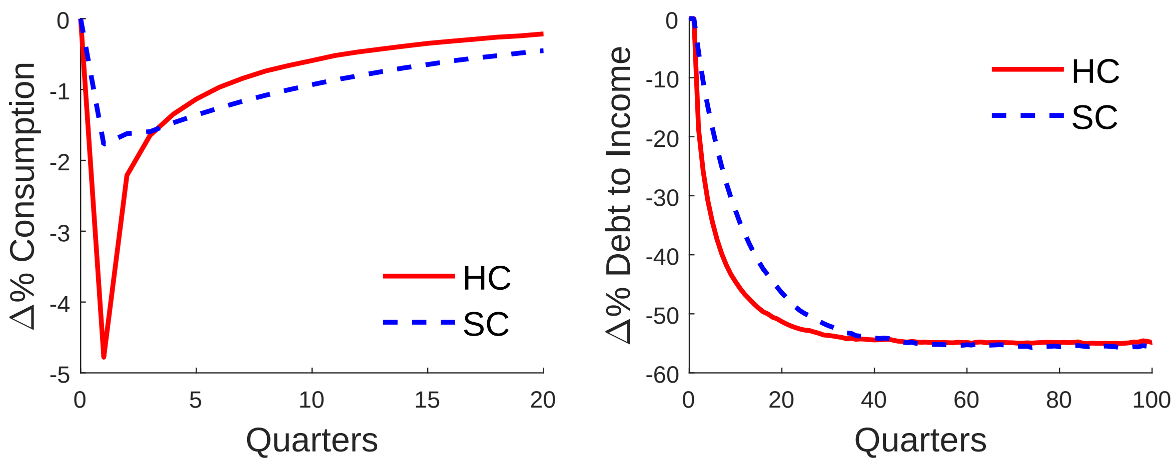

Consumption decreases more severely in hard constraint economy.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.

Figure A.12: Unexpected Tightening of the Borrowing Conditions. Adjustment in 6 quarters - Credit Crunch via ϕ_1

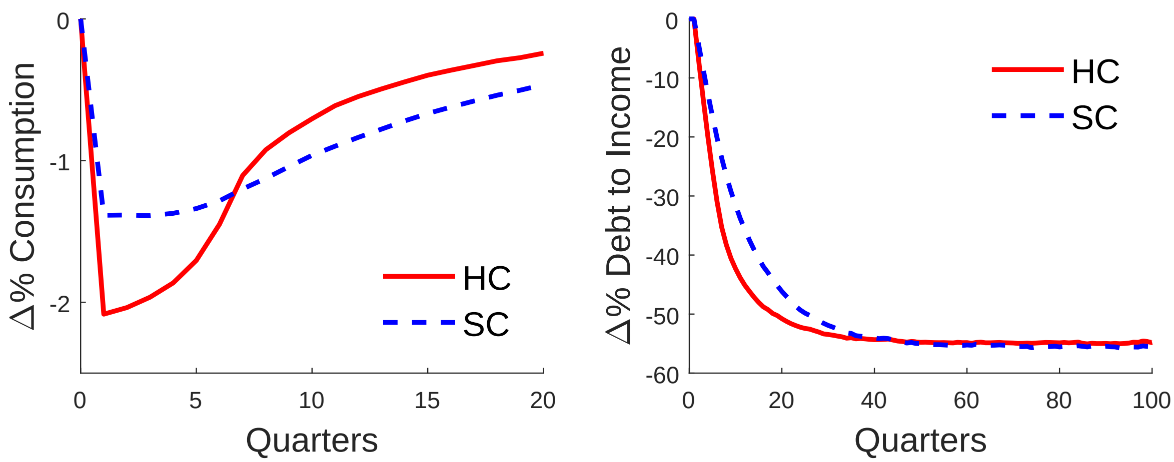

Consumption decreases more severely in hard constraint economy.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Borrowing parameters are assumed to adjust linearly over six quarters. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.

Figure A.13: Unexpected Tightening of the Borrowing Conditions. Immediate Adjustment - Credit Crunch via ϕ_2

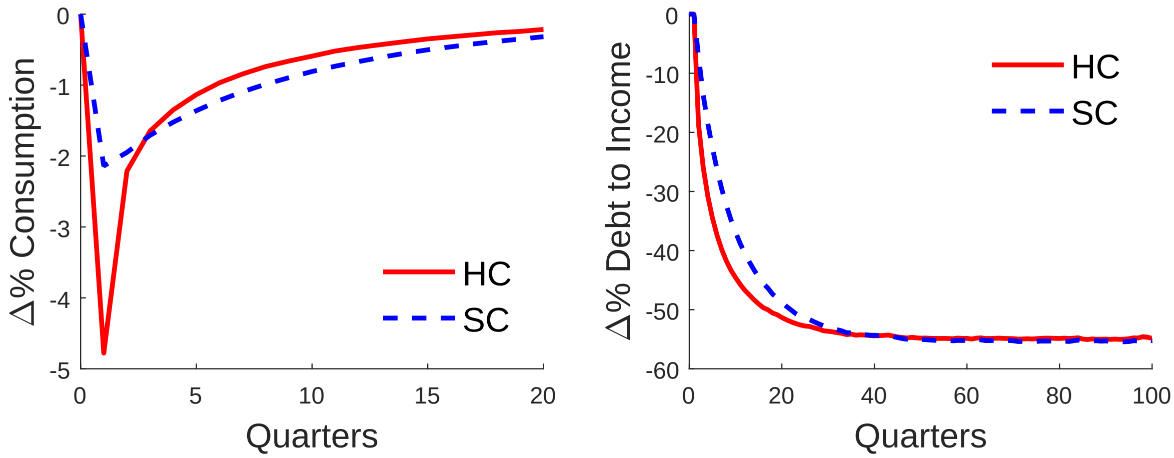

Consumption decreases more severely in hard constraint economy.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.

Figure A.14: Unexpected Tightening of the Borrowing Conditions. Adjustment in 6 quarters - Credit Crunch via ϕ_2

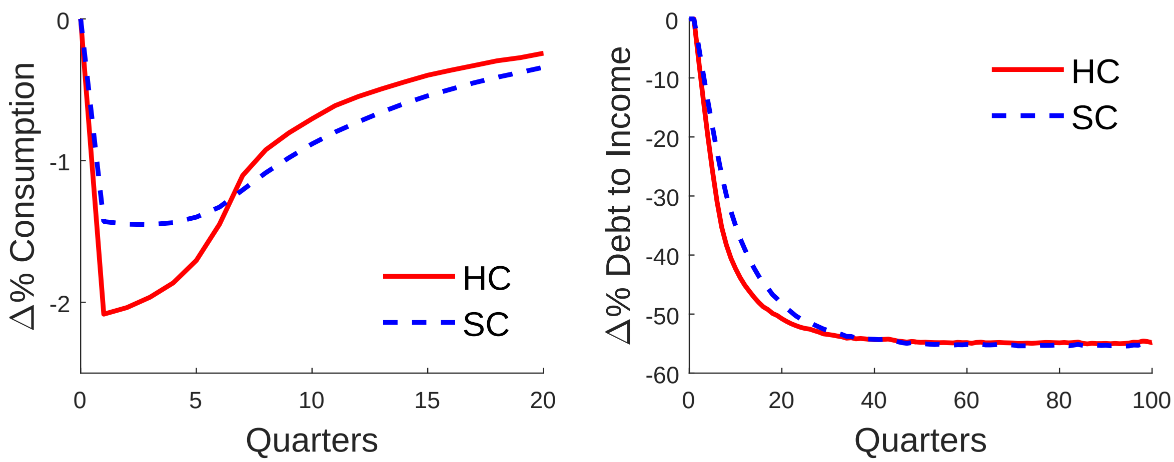

Consumption decreases more severely in hard constraint economy.
HC: Hard Constraint Model. SC: Soft Constraint Model.
Aggregate consumption and debt to income responses in each economy after an unexpected tightening of the credit conditions that decreases the debt to income ratio in the new steady state by 55.56%. Borrowing parameters are assumed to adjust linearly over six quarters. Percentages are with respect to the initial steady state levels. The x-axis measures the number of quarters after the credit crunch hits the economy.