Credit Limits and Consumption Behavior over the Life Cycle*

César Salinas†
Indiana University

This version: May 28, 2024

Abstract

This paper studies the role of time-varying credit limits through the lens of a life cycle incomplete markets model calibrated for the U.S. Changes in credit card limits are explained by observable household characteristics and the estimated unobservable variation is quite large. The quantitative exercise shows that even though young households are more indebted in an economy with stochastic borrowing limits, aggregate consumption is not greatly affected by transitory or persistent shocks of this type. However, in the presence of these shocks, households lose the ability to self-insure against other uninsurable idiosyncratic shocks, e.g., labor income shocks. A disaggregated analysis shows that the loss of self-insurance capacity is mainly explained by the effects that stochastic borrowing limits have on the wealth distribution, the precautionary savings channel households have to face unexpected risks.

*I would like to thank Bulent Guler, Laura Liu, Christian Matthes, Liu Mendoza, Todd B. Walker, seminar participants at Indiana University, and the Research Institute in Economics and Finance for their useful comments and suggestions. The author acknowledges the Indiana University Pervasive Technology Institute for providing HPC resources that have contributed to the research results reported within this paper. All remaining errors and omisions are my entirely responsibility.

†Department of Economics. Address: Wylie Hall, 100 South Woodlawn Avenue, Bloomington, IN 47405, USA. Email: salinasc@iu.edu.

1 Introduction

Most of the earlier literature on consumption and wealth dynamics over the life-cycle assumed that households can not borrow (zero borrowing limits), face the loosest possible debt limit (natural borrowing limits), or face a single and invariant borrowing limit. However, recent studies highlight the importance of time-varying changes on credit limits.[1] In this paper I quantify how credit card limits change with observable and unobservable household characteristics and assess the responsiveness of household decisions and consumption through the lens of a life cycle incomplete markets model.

Credit card borrowing is the main source of consumption smoothing, especially for individuals who have small liquid savings and/or are exposed to frequent changes in labor income. More than 70 percent of US households have a credit card with credit utilization rates being higher for young people.[2] Although roughly 70 percent of household debt in the last 20 years is attributable to mortgage debt (Brown, Stein and Zafar, 2015), there is a growing literature that documents an important role of households' access to unsecured revolving credit to face income and unemployment fluctuations in the short-term (Herkenhoff, Phillips and Cohen-Cole, 2016; Herkenhoff, 2019; Braxton, Herkenhoff and Phillips, 2020). The main hypothesis in this paper is that to the extend that households smooth shocks by borrowing and saving for precautionary reasons, changes in credit limits affects life-cycle consumption decisions.

Using public available data from the Survey of Consumer Finances (SCF), I first find that credit card limits increase with age and they are correlated with observable household characteristics such as income and measures of credit risk. After controlling for these variables, the remaining unobservable component shows an important and consistent volatility over the life-cycle. In fact, this almost invariant change in credit limits can be explained by two components. The first one, a transitory component whose shock variances increase slightly early in life and then remain constant before retirement. And the second, a persistent component whose shock variances decrease with age but their persistence rises over time.

To quantify the significance of credit card limits risk, I extend a standard life cycle model with incomplete markets (Storesletten, Telmer and Yaron, 2004*a*; Kaplan and Violante, 2010), where households are patient or impatient, have constant relative risk aversion (CRRA) utility, supply labor inelastically, and can insure themselves against risks by trading a non-state-contingent asset, but they are subject to changes in credit limits. In particular, households are subject to transitory

[1] Fulford (2015) investigates how variable credit limits incentivize households to hold both high interest debts and low interest savings at the same time. Guerrieri and Lorenzoni (2017) study how permanent tightening in borrowing limits affect consumers' debt, savings, interest rates and output. Jensen, Ravn and Santoro (2018) explore how exogenous changes in the loan-to-value ratios impact aggregate volatility and the comovement between debt and macroeconomic variables. Drechsel (2023) examines both theoretically and empirically how supply shocks are quantitatively more important when firms face earnings-based credit constraints than when they face collateral constraints.

[2] The value of the credit utilization rate, defined as the credit card debt to credit card limit ratio, is around 0.5 for credit card holders between 20 and 30 years old, and 0.2 for those between 60 to 70 years old (Fulford and Schuh, 2015).

and persistent shocks to earnings and credit limits while they work, and during retirement receive pension benefits, face the same borrowing limit that they had in the last work period, and face a survival probability. To focus the analysis solely on idiosyncratic risks faced during the working life, I assume perfect annuity markets so households are completely insured against survival risks and face no other uncertainty after retirement.

The benchmark model is calibrated to replicate key moments for wealth, unsecured debt and credit card limits in the U.S. The results show that young households are more indebted and use more than 50 percent of their available credit limit the first 10 years of the working life. Their debt to income ratio follows a similar evolution where debt represents 12 percent of the total income for households younger than 35 years old. With low savings and income accumulated by young households, they use debt as an additional mechanism to insure against idiosyncratic shocks.

I also show that compared to other specifications without stochastic borrowing limit shocks, young households are more indebted and utilize a larger fraction of their available credit limit in the benchmark economy. It turns out that this higher level of debt combined with a positive intertemporal savings motive for patient agents explains that the median consumption is high in this economy.

In addition to its impact on credit utilization and debt, stochastic borrowing limits affect the insurance mechanism households have to smooth consumption. I follow the formulation from Blundell, Pistaferri and Preston (2008) and Kaplan and Violante (2010) and compute the consumption insurance coefficients with respect to transitory and persistent shocks. The simulated insurance coefficients suggest that household consumption is more affected by income shocks than borrowing limit shocks. This result can be explained by the fact that not all households in the model use their entire credit limit. However, a disaggregated analysis of these coefficients indicates that young and less wealthy households are less insured against these shocks compared to older and wealthier households.

With respect to income shocks, the computed insurance coefficient of consumption for transitory shocks is 24% lower in the model with stochastic borrowing limits compared to other specifications without these shocks. The same coefficient is 50% lower for persistent income shocks. The presence of borrowing limit shocks not only directly affects consumption behavior, but also the distribution of assets that households accumulate throughout the life cycle. The lower capacity to insure against income shocks is mainly explained by the effect of stochastic borrowing limits on the distribution of assets, which impacts the precautionary savings channel that households have to face unexpected risks, and finally affects consumption behavior.

Related literature. This paper contributes to the existing literature that studies quantitatively how inequality in consumption and wealth change over the life-cycle (e.g., Storesletten et al., 2004*a*; Guvenen, 2007; Primiceri and Rens, 2009; Karahan and Ozkan, 2013). In this line, Kaplan (2012) extends the standard incomplete markets life-cycle model to include unemployment risk to better match the life-cycle profiles of the first and second moments of wages, hours worked and consumption. Arslan, Guler and Taskin (2020) extends a similar model with price search decisions to study how consumption and expenditure inequality differ with age and how price search can serve as an additional insurance mechanism to face income risk. I depart from these studies by introducing life-cycle changes in credit limits as an additional source of risk households face.

My paper is also related to the unsecured consumer credit literature (e.g., Gross and Souleles, 2002; Athreya and Simpson, 2006; Chatterjee, Corbae, Nakajima and Ríos-Rull, 2007; Livshits, MacGee and Tertilt, 2007; Telyukova, 2013; Herkenhoff, 2019; Dobbie and Song, 2020). In particular, this paper is related most directly with Athreya (2008); Crossley and Low (2011); Fulford (2015); Chatterjee, Gunawan and Kohn (2022); and Faccini, Lee, Luetticke, Ravn and Renkin (2024). Athreya (2008) includes the option to debt default to study the evolution of consumption and wealth over the life-cycle. Compared to the scenario with no default options, he finds that default policies in the U.S. create severe credit constraints for young people and induce them to borrow less. In fact, consumption inequality rises for the young but falls for the elderly because it is more convenient for them to default when they face risks later in life. Crossley and Low (2011) study the role of unemployment insurance when borrowing and self-insurance is costly throughout the lens of a tractable life-cycle model. The authors find that the value of these external sources of insurance depend on, among other things, the age at job loss, access to credit, and the return on savings.

Using credit cards data from Equifax and the Federal Reserve Bank of New York Consumer Credit Panel, Fulford (2015) shows that individuals gain and lose access to credit frequently and estimates of credit limit volatility are larger than most estimates of labor income volatility. By including these estimates into a model where infinitely lived households maximize their consumption facing labor income risk, the author shows that with stochastic credit limits, households optimally choose to hold debt and cash at the same time (the so-called *"credit card puzzle"*). Chatterjee et al. (2022) study the role of credit uncertainty via collateral constraints in an otherwise standard real business cycle (RBC) model. The authors estimate credit uncertainty as being the stochastic volatility in credit growth in the US, and incorporate it into an RBC model where firms finance their operations through working capital requirements. When the collateral constraint binds, credit uncertainty generates a simultaneous decline in output, consumption, investment, real wages and hours. Using administrative data from Denmark, Faccini et al. (2024) examine how consumer credit spreads affect consumption and wealth dynamics suggesting the marginal propensity to consume is countercyclical. Through the lens of a HANK model, the authors find that cyclical fluctuations in credit spreads are important to explain the heterogeneous effects of aggregate shocks across the wealth distribution.

In the spirit of these studies, this paper aims to take microdata on household credit more seriously by estimating how credit limits vary according to observable and unobservable household characteristics and incorporating these results into a life cycle model with heterogeneous households that optimally decide their level of consumption, debt and savings.

The paper is organized as follows. Section 2 presents the data and the econometric framework used to estimate how credit card limits change over the life cycle. Section 3 presents the model and the calibration strategy. I discuss the main results in section 4. In section 5 I asses the role of stochastic borrowing limits and conclude in section 6.

2 Credit Card Limits

In this section I show how credit card limits change over the life cycle and across observable household characteristics using publicly available data from the Survey of Consumer Finances (SCF) for the period 1989 to 2019.

2.1 Data

The SCF is a triennial survey that collects information about balance sheet, income, and other demographic characteristics of families in the U.S. but it is not a panel. In this sense, I will follow Devlin-Foltz and Sabelhaus (2016) and construct a synthetic panel grouping observations by income, education, marital status, and race to analyze how these group of households, or bins, are affected by changes in credit card borrowing limits over time.

To obtain these household bins I proceed as follows: for the period 1989-2019, I only look at households whose head's age is between 26 and 60 years old and have reported borrowing limits on their credit cards.[3] Then, I categorize households by income quartile (4 categories), head's education (4 categories), marital status (2 categories), and head's race (2 categories).[4] Finally, I collapse the main variables for each bin and follow how average household bins characteristics change over age and year. This criteria leave me with 2,930 household bins and 9,543 observations.

The resulting life-cycle profiles of real credit card limits are shown in figure 1. Panel (a) shows that, on average, credit card limits constantly increase during the working life, increasing by around 100 percent between the ages of 26 and 60. Panel (b) shows that credit limits variability increases slightly early in the working life and then slowly decreases before retirement.[5]

2.2 Estimation

To evaluate the role of changes in credit limits, it is important to understand what type of factors can account for these changes. Let b^i_{ht} denote the log of annual credit card limits of household i at age h in time t. Similar to the standard estimates of income processes, I assume credit limits depend on observable household characteristics, a persistent and a transitory shock (Carroll and Samwick, 1997; Floden and Lindé, 2001; Storesletten, Telmer and Yaron, 2004*b*). The log of the

[3]The remain sub-sample represents more than 70% of the total SCF 1989-2019 sample.

[4]The income quartile groups are: 0-24.9, 25-49.9, 50-74.9, 75-100. The education categories for the household head are: no high school diploma/GED, high school diploma or GED, some college, college degree. The marital status categories are: married or living with a partner, and neither married nor living with a partner. The household head race categories are: white non-Hispanic, and African American, Hispanic or other.

[5]Using data from the Consumer Credit Panel collected by the Federal Reserve Bank of New York, Fulford and Schuh (2017) find similar results as the ones reported in figure 1.

Figure 1: Credit card borrowing limits over the life-cycle

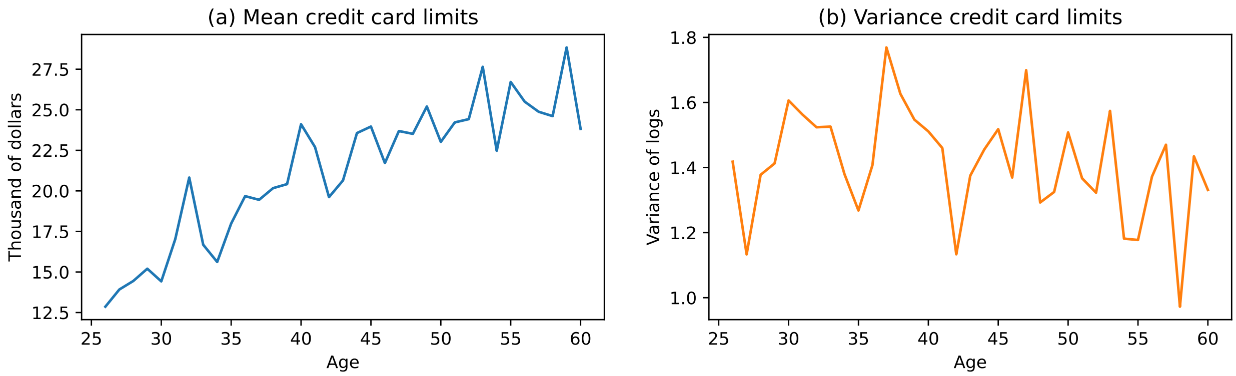

Notes: The figure on panel (a) shows the credit card borrowing limits average of the SCF household bins by age. The figure on panel (b) shows the variance of the log of credit card borrowing limits of the SCF household bins by age.

borrowing limits follows

$$\begin{aligned} b^i_{ht} &= f(\boldsymbol{x}^i_{ht};\boldsymbol{\theta}_{ht}) + u^i_{ht} + \phi_t n^i_h \\ u^i_{ht} &= \rho_{u,h-1} u^i_{h-1,t-1} + \pi_t v^i_h, \end{aligned} \tag{1}$$

where $\boldsymbol{x}^i_{ht}$ is the vector of observable household characteristics, u^i_{ht} and n^i_{ht} represent the persistent and transitory components, respectively. Both v^i_h and n^i_h are assumed to be identically and independently distributed over age and across households with means equal to zero and variances equal to $\sigma^2_{v,h}$ and $\sigma^2_{n,h}$, respectively. I allow the persistence parameter ($\rho_{u,h-1}$), and the variance of persistent and transitory shocks to vary with age. These age-profiles take into account potential age specific events throughout the life cycle, e.g., access to credit markets and credit history.

To estimate the process described in (1) I follow a two step procedure. I first estimate by OLS the observable component $f(\boldsymbol{x}^i_{h,t};\boldsymbol{\theta}_{h,t})$, and then estimate the remaining parameters removing the estimated observable component from the log of credit card borrowing limits.

As pointed out by Fulford and Schuh (2015), credit card limit changes are affected by both financial institutions and individuals decisions. To assess the effects that borrowing limits have on individuals decisions it is important to measure and isolate the variation that is given to individuals. Therefore, the selection of variables on which b^i_{ht} is projected depends on factors that can account for credit demand, i.e., variables that could affect individual decisions to change credit card limits.

Vector $\boldsymbol{x}^i_{h,t}$ includes a quadratic polynomial in age, year fixed effects, and discrete variables for race, education and working status. I also include some measures of credit risk, like number of credit cards, dummy for late debt payments, and dummy for high credit card utilization rates (refer to appendix A for the definition of variables).[6] The estimated coefficients of interest are reported on table 1.

Column (1) shows that as age increases, credit limits will also increase, but at a slightly slower

[6]The SCF does not have information about the household credit score, that is an important factor to determine borrowers' risk. Instead, I approximate this measure by including these three factors that creditors take into account when calculating credit scores.

Table 1: OLS estimation results for the log of credit card borrowing limits

	(1)	(2)	(3)	(4)	(5)
Age	0.0661***	0.0642***	0.0724***	0.0630***	0.0631***
	(0.0118)	(0.0117)	(0.0111)	(0.0093)	(0.0092)
Age^2	-0.0005***	-0.0005***	-0.0006***	-0.0005***	-0.0005***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Log income					0.0389***
					(0.0035)
Time fixed effects	No	Yes	Yes	Yes	Yes
Other control variables	No	No	Yes	Yes	Yes
Measures of credit risk	No	No	No	Yes	Yes
R-squared	0.027	0.047	0.163	0.416	0.424

Notes: Standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

rate. Columns (2) and (3) include time fixed effects and other household head characteristics. By including measures of credit risk and income in columns (4) and (5), the regressions are able to explain more of the variability of credit card limits. These results are expected since credit limits are typically set in proportion to the borrower's credit history, approximated by the household head age, credit risk, and income. To successfully measure how credit limits affect household decisions, our theoretical model should be able to take these heterogeneities into account.

To extract the risk which remains for the household bins after observable characteristics have been removed, I construct the residual credit limit $\tilde{b}^i_{ht} = b^i_{ht} - f(\boldsymbol{x}^i_{h,t}; \hat{\boldsymbol{\theta}}_{h,t})$. This variable is decomposed into a fixed effect, a persistent and a transitory component:

$$\begin{aligned} \tilde{b}^i_{ht} &= \alpha^i + u^i_{ht} + \phi_t n^i_h \\ u^i_{ht} &= \rho_{u,h-1} u^i_{h-1,t-1} + \pi_t v^i_h. \end{aligned} \tag{2}$$

According to the equation (2), α^i is an individual fixed effect capturing variation in initial conditions. The remaining components are allowed to change over age and year. n^i_h is a transitory component with variance σ^2_{nh} that captures temporary changes in credit limits and measurement error. u^i_{ht} is the persistent component that captures lasting changes in credit limits, and each individual is hit by a persistent shock of size v^i_h every period with lasting effect determined by the persistence parameter ρ_u. u^i_{1t} is drawn from a distribution with mean zero and variance equal to σ^2_{z1} and captures the initial variation in the persistent component. Similar to the previous studies, I also control for the change in residuals over time by including the parameters ϕ_t and π_t.

The variance-covariance structure of the specification in (2) is given by:

$$\begin{aligned} var(\tilde{b}^i_{ht}) &= \sigma^2_\alpha + var(u^i_{ht}) + \phi^2_t \sigma^2_{nh}, \\ cov(\tilde{b}^i_{ht}, \tilde{b}^i_{h+p,t+p}) &= \sigma^2_\alpha + \rho_h \rho_{h+1} \dots \rho_{h+p-1} var(u^i_{ht}), \\ var(u^i_{ht}) &= \rho^2_{u,h-1} var(u^i_{h-1,t-1}) + \pi^2_t \sigma^2_{vh}. \end{aligned} \tag{3}$$

Figure 2: Estimation results for age intervals

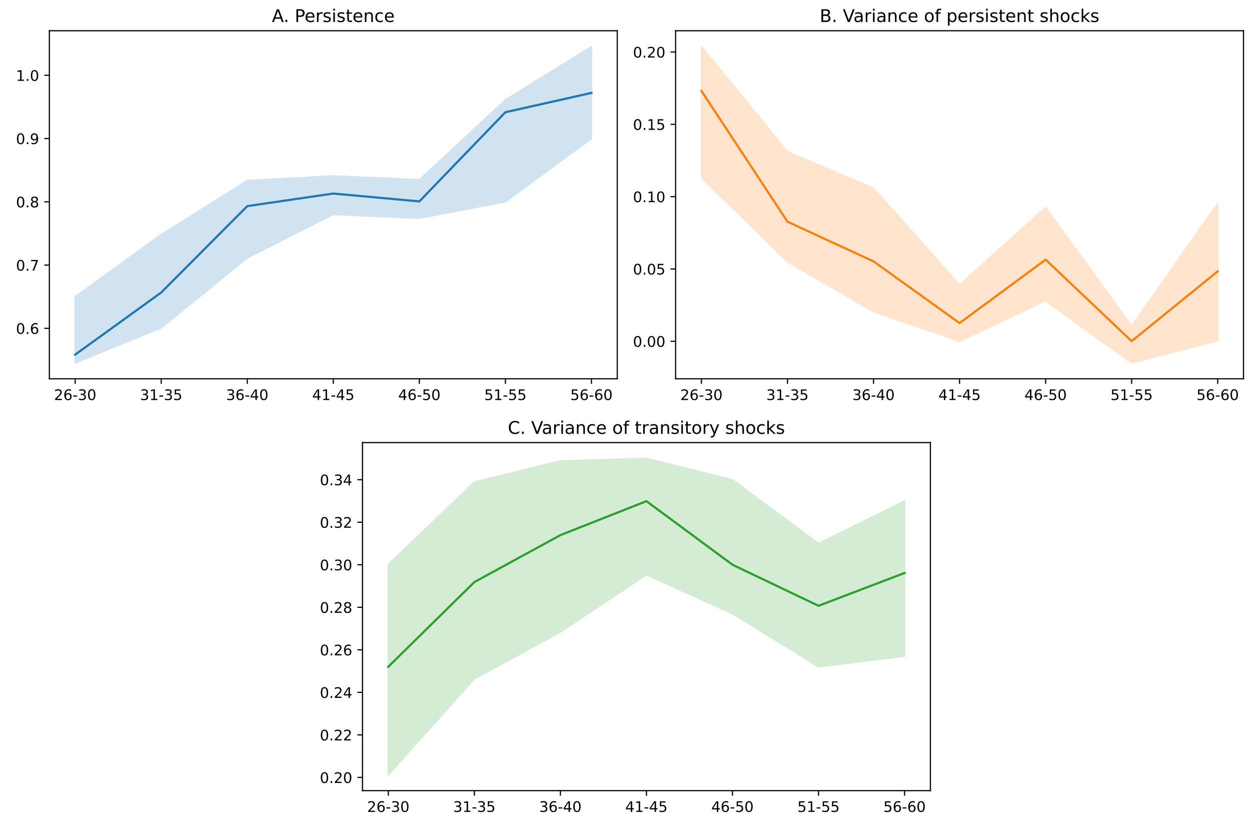

Notes: The solid line represents the point estimate. The shaded area represents the 95% bootstrap confidence interval.

I follow Karahan and Ozkan (2013) and identify the parameters of interest by using the normalizations: $\rho_{u1} = \rho_{u2}, \pi_1 = \phi_1 = 1, \pi_T = \phi_{T-1}$, and $\sigma^2_{vH} = \sigma^2_{v,H-1}$. An equally weighted minimum distance estimator between the moments of the empirical variance-covariance structure of residual credit limits and their theoretical counterparts in (3) is employed to estimate the vector of parameters $(\sigma^2_\alpha, \rho_{h-1}, \sigma^2_{n,h}, \sigma^2_{v,h}, \phi_t, \pi_t)$. In particular, I target all the variance and covariance terms over age $cov(\tilde{b}^i_h, \tilde{b}^i_{h+p})$ and time $cov(\tilde{b}^i_t, \tilde{b}^i_{t+q})$. This strategy generates 231 moments to target when at least 80 observations are used to calculate them.

2.3 Persistent and transitory components

I estimate the persistence and variance of shocks in two ways. First, I divide the working life into age intervals, where the persistence and variances are constant within an interval but they differ between intervals. Then, I assume the life cycle profiles follow a polynomial function of age and estimate its parameters. In this last way, I am able to reduce the number of parameters to be estimated from 128 to 36.[7]

[7] The unrestricted estimation would imply to find the variance of the fixed effect (1 parameter), the persistence coefficients (35 parameters), the variance of the persistent shock (35 parameters), the variance of the transitory shock

Figure 3: Estimation results for cubic specification

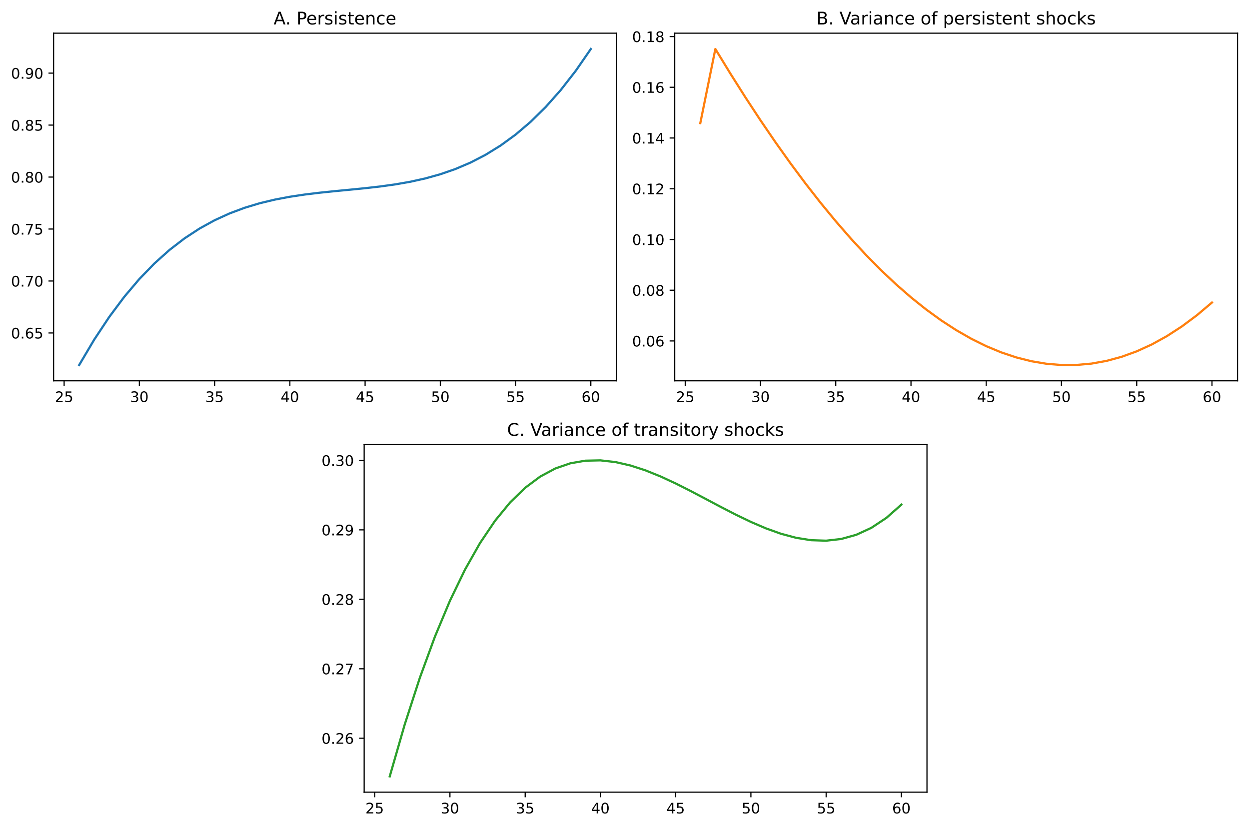

Figure 2 shows the estimated results for seven age intervals. The persistence for young workers is statistically smaller than that of old-age workers. At the same time, the variance of persistence shocks are larger early in life compared with those between age 50 and 60. These results are consistent with the idea that older individuals have a longer credit history, which can impact positively on their credit score and their access to better credit conditions, such as wide credit card limits. Lastly, the variance of transitory shocks are larger compared to the variance of persistent shocks, but they do not show a clear statistical pattern increasing slightly early in life and remain almost constant afterwards.

Overall, the age interval estimation results suggest that the persistence and the variance of persistent shocks have non-flat profiles over the life cycle and the variance of transitory shocks increases slightly at the beginning of working life. To account for these non-linearities, I assume that these profiles follow a cubic function on age. The point estimates for the parameters in the cubic specification as well as their bootstrap standard errors are reported in table B1 (see appendix B).

The first two top panels in figure 3 show the resulting profile for the persistent component.

(35 parameters), and the time loading factors (22 parameters). Assuming the age dependent coefficients follow a cubic specification, I only need to estimate 12 parameters for the persistence coefficients, the variance of the persistent shock, and the variance of the transitory shock.

Similar to the previous results, two facts are revealed from this figure. First, the degree of persistence of these shocks increases throughout the life cycle. Persistence starts at around 0.65 for people under 30 years old and increases up to 0.9 by around age 60. Second, the variance of persistent shocks is moderately high early in life and then decreases with age. This variance drops more than twice from 0.18 at around age 26 to 0.06 at age 50, and then slowly increases before retirement.

The bottom panel in figure 3 plots the variance of transitory shocks. There is an increase early in life from around 0.26 to 0.30 by age 40. After this age, the variance decreases slightly but remains stable around 0.29. Compared to the variance of persistence shocks, the variance of transitory shocks is larger and does not decrease considerably over the life cycle. Using data from Equifax and the Federal Reserve Bank of New York Consumer Credit Panel, Fulford and Schuh (2015) estimates that the variance of this type of transitory shocks is on average equal to 0.25 suggesting that temporary changes in credit card limits represent an important source of fluctuation for borrowers.

3 The Model

To study how households' decisions are affected by changes in credit limits I incorporate the previous estimates into an otherwise standard life cycle incomplete markets model with idiosyncratic labor income and borrowing limit shocks. Borrowing limits depend on observable household characteristics and exogenous stochastic components. I then simulate an artificial panel of households from the model to calculate how utilization rates, debt and consumption decisions are affected by time-varying credit limits.

Demographics. There is no aggregate uncertainty. The population consists of a continuum of households indexed by i. Agents live for H years and work for the first H^{work} years of their life, after which they retire. Retirement lasts for H^{ret} periods.

Preferences. Households have time-separable expected utility given by

$$\mathbb{E}\sum_{h=1}^{H}\beta^{i^{h-1}}\mathcal{S}_h u(C_h^i).$$

Altruism is assumed away and households supply labor inelastically over the working life. There are two types of households in this economy: patient and impatient households($\beta^p \geq \beta^i$). The type and share of these households is realized at the beginning of the working life. Unlike standard models with one type of households, this feature allows me to generate a realistic amount of debt for unsecured debt calculated with U.S. data.[8]

There is no chance of dying before retirement. After retirement, agents face a survival probability $\mathcal{S}_h < 1$ and all agents die by age H with certainty. There exist perfect annuity markets and

[8] Aguiar, Bils and Boar (2020) also allow for preference heterogeneity into an incomplete-markets life-cycle economy to account for different facts about the spending behavior of "hand-to-mouth" households found in the data, such as not having greater spending growth but facing greater spending volatility.

households are completely insured against survival risk. This assumption allows me to focus the analysis solely on household idiosyncratic risks during the working life.

Assets. Households can hold a non-state-contingent single asset. Positive holdings of the asset pays a gross return equal to R_+. Holdings can be negative but they are subject to a higher gross interest rate $R_- > R_+$ and they must satisfy the specified credit limit. The main departure of the model from the standard life cycle models is the presence of a stochastic credit limit.

If a household has debt greater than its new credit limit, it must pay at least the differential to meet the new credit limit. To give a better understanding, consider the case where a household is indebted and faces a decrease on its credit limit the next period ($A_h^i < B_h^i < 0$). In this case, the household must pay at least the differential on its debt ($B_h^i - A_h^i$) to hold the new level of credit limit next period. Of course, the household could choose to pay back more, but the differential will be the minimum payment. During the working life, the log of credit limits follows equation (1) and is given by

$$\begin{aligned} b_h^i &= \theta_0 + \theta_1 h + \theta_2 h^2 + \theta_y y_h^i + u_h^i + n_h^i \\ u_h^i &= \rho_{u,h-1} u_{h-1}^i + v_h^i, \end{aligned} \tag{4}$$

where $(\theta_0 + \theta_1 h + \theta_2 h^2)$ is the common deterministic profile that follows a quadratic polynomial on age, y_h^i is the log of labor income, and u_h^i and n_h^i represent the age-variant persistent and transitory borrowing limit components, respectively. During retirement, households do not face borrowing limits risk and their limit is constant and equal to what they had in the last year of work ($b_{H^{work}}^i = b_{h+H^{work}}^i$).

I also assume that there is no default option. Default is clearly an important element of the adjustment to a tighter credit regime but is beyond the scope of this paper.[9]

Labor income. During the working years, households receive labor income Y_h^i. I follow the standard literature that quantifies the evolution of income over the life cycle. The log of income follows

$$\begin{aligned} y_h^i &= \kappa_h + z_h^i + \varepsilon_h^i, \\ z_h^i &= \rho_z z_{h-1}^i + \eta_h^i, \end{aligned} \tag{5}$$

where κ_h is the non-stochastic experience profile, z_h^i and ε_h^i represent the persistent and transitory income components, respectively. Both η_h^i and ε_h^i are assumed to be identically and independently distributed over age and across households with means equal to zero and variances equal to σ_η^2 and σ_ε^2, respectively.

Social security transfers. Retired households receive after-tax social security transfers from the government, which are a function of the individual vector of gross earnings realizations $\mathcal{P}(\tilde{Y}_h^i)$.

[9] Consumers can default informally by simply stopping repayments and becoming delinquent, or consumers can default formally by declaring bankruptcy. I will abstract from these behaviors, which imply a more sophisticated but complex modelling approach, assuming that all households can at least pay the differential on their revolving debt.

To map net income earnings into gross income earnings I follow Castañeda, Díaz-Giménez and Ríos-Rull (2003) and Kaplan and Violante (2010) and use a nonlinear tax function given by

$$\tau(\tilde{Y}_h^i) = \tau_1 \left[\tilde{Y}_h^i - \left((\tilde{Y}_h^i)^{-\tau_2} + \tau_3 \right)^{-1/\tau_2} \right], \tag{6}$$

where gross labor income can be recovered by inverting the equation for net income $Y_h^i = \tilde{Y}_h^i - \tau(\tilde{Y}_h^i)$.

Budget constraints. The household's budget constraint follows

$$\begin{aligned} C_h^i + A_{h+1}^i &= R(A_{h-1})A_h^i + Y_h^i && \text{if } h \leq H^{work}, \\ C_h^i + \left(\frac{\mathcal{S}_h}{\mathcal{S}_{h+1}} \right) A_{h+1}^i &= R(A_{h-1})A_h^i + \mathcal{P}(\tilde{Y}_h^i) && \text{if } h > H^{work}, \\ A_{h+1}^i &\geq -B_h^i, \end{aligned}$$

where Y_h^i is the level of labor income after taxes and transfers, $\mathcal{P}(\tilde{Y}_h^i)$ is the function determining the pension income during retirement, A_{h+1}^i is the amount of the asset carried over by individual i from age h to $h+1$ and B_h^i is the level of borrowing limits that is known at age h.

3.1 Calibration and solution method

To understand the role of credit limit changes, I calibrate the model to reproduce key features of the unsecured debt in the U.S. economy.

The model period is one year. Each household enters the labor market at age 26, retires at age 60, and dies with certainty at age 95. Households start their working life with zero assets ($A_0^i = 0$). I choose a constant relative risk aversion (CRRA) specification for $u(C_h^i)$ with risk aversion parameter $\gamma = 2$, a standard value in the literature. I set β^p to match the aggregate wealth to income ratio of 3.1, and set β^i to match the aggregate unsecured revolving debt to income ratio of 0.05 for the U.S. economy. I also set the share of impatient agents in the simulated economy to match the share of wealth holdings at the top 20 percent of the net worth distribution according to the SCF that is approximately equal to 0.8. The choice of these three parameters is crucial to measure the households' decisions of asset holdings used to smooth income and credit limits shocks in the model.

The interest rate on savings is equal to 3 percent and I use the information on interest rates paid on credit cards reported in the SCF to set the average real borrowing interest rate equal to 10 percent. Credit card borrowing limits change with observable characteristics and unobservable stochastic variations. I use the estimated coefficients from column (5) in table 1 to account for the variation that can be explained by age and income. The age-variant estimates shown in figure 3 are used to account for the persistent and transitory borrowing limit components. The initial variance σ_{u1} is set to be equal to 0.1459 to match the dispersion of household credit card limits at age 26 as shown in table B1. Credit limits are then scaled proportionately by adjusting the value of θ_0 in equation (4) to match the aggregate credit card limits to income ratio of 27 percent according to the SCF.[10]

I follow standard parameter values for the earnings process used in the literature. The persistent income component follows a random walk ($\rho_z = 1$) with an initial variance σ_{z1} equal to 0.15 to match the dispersion of household earnings at age 26. The common deterministic profile κ_h, the variance of persistent shocks σ_η, and the variance of transitory shocks σ_ε are set to match the rise of earnings in levels and dispersion for households from age 25 to 60 in the PSID data (Blundell et al., 2008; Kaplan and Violante, 2010; Guvenen, 2009). Similar to the US Social Security system, pension benefits are a function of lifetime average individual gross earnings. More precisely, two bend points are defined at 0.18 and 1.10 times the cross-sectional average gross earnings. After retirement the household pension is equal to 90 percent of average past earnings up to the first bend point, 32 percent from the first bend point to the second bend point, and 15 percent beyond that. Benefits are then scaled proportionately to ensure a worker labor income average is entitled to a replacement rate of 45 percent (Mitchell and Phillips, 2006). The values for (τ_1, τ_2), as shown in equation (6), are taken from the estimates of Gouveia and Strauss (1994). The value of τ_3 is chosen to match the ratio of total tax receipts on labor income to total labor income of 25 percent. Table C1 shows the values of the externally and internally calibrated value parameters in the benchmark model.

Solution method. The model is solved using the endogenous grid method (Carroll, 2006) with 200 exponentially spaced grid points for assets. The grid for lifetime average earnings has 7 points. The decision rule is linearly interpolated between grid points. The earnings and credit limits persistent components are approximated using a discrete Markov chain with 9 equally spaced points. The earnings and credit limits transitory components are approximated using a discrete Markov chain with 7 equally spaced points. The results shown below are based on simulating an artificial panel of 50,0000 households from 26 to 95 years old (70 years).

Calibration results. Table 2 shows the model-generated moments against the data counterparts. Panel A shows that the benchmark model matches the targeted moments in the data. The features of the model also allow for realistic average utilization rates and debt variation to be reproduced, as shown in panel B.

The left panel of figure 4 shows the average credit card limits profile simulated in the model and the calculated from the SCF data proportional to aggregate income. In the model, aggregate income is calculated as the sum of aggregate gross labor income plus aggregate asset income. In the data, aggregate income was calculated as the sum of household wages and salaries income before deductions for taxes plus capital income. The credit limits simulated in the model closely follow their counterparts in the data.

[10]This is calculated as the ratio between the aggregate credit card limits to aggregate labor and capital income ratio in the SCF (1989-2019), where the top 5 percent of households in the wealth distribution are excluded. Income is defined as the sum of household annual income from wages and salaries before deductions for taxes and capital income.

Table 2: Moments in the benchmark model

	Data	Model
Panel A. Targeted moments		
Wealth to income ratio	3.10	3.10
Debt to income ratio	0.05	0.05
Capital share at top 20%	0.80	0.80
Credit limits to income ratio	0.27	0.27
Panel B. Non-targeted moments		
Utilization rate	0.26	0.25
Variance of log debt	16.48	16.02

Figure 4: Life-cycle profiles for mean and variance in credit limits: model versus data

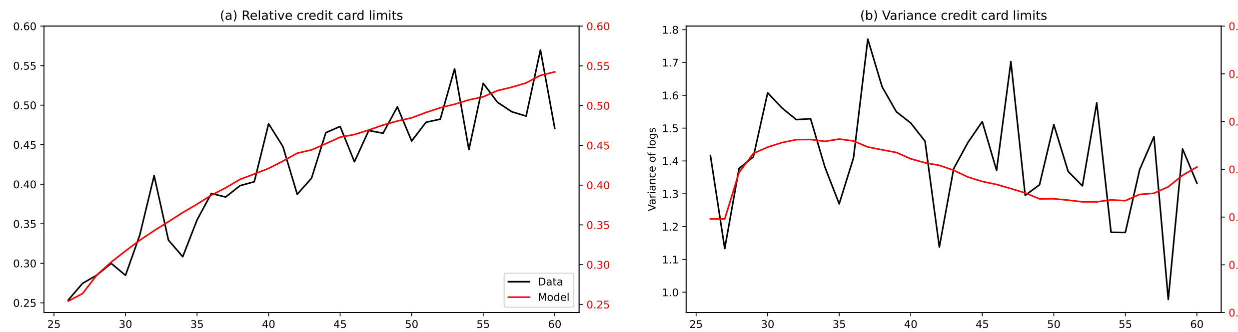

Notes: The average of credit card limits are relative to the income at age $h = 26$.

The right panel of figure 4 shows the variance of credit limits over the life cycle in the model and in the data. First, the model explains more than 30 percent of the variance in credit limits in the US. Second, the data exhibits large ups and downs that the minimum distance estimator manages to capture as an increase in variance early in life, a decrease after age 32-33, and a slightly increase just before retirement.

4 Results

As discussed in the introductory section, stochastic credit limits could alter households' decisions because they represent an additional source of uncertainty and they affect one of the insurance mechanisms households use to insure themselves against other adverse shocks, such as labor income. In this section, I quantitatively evaluate the importance of time-varying borrowing limits. The main results are calculated during the working life since households do not face idiosyncratic risk during retirement.

Figure 5: Consumption and wealth over the life cycle

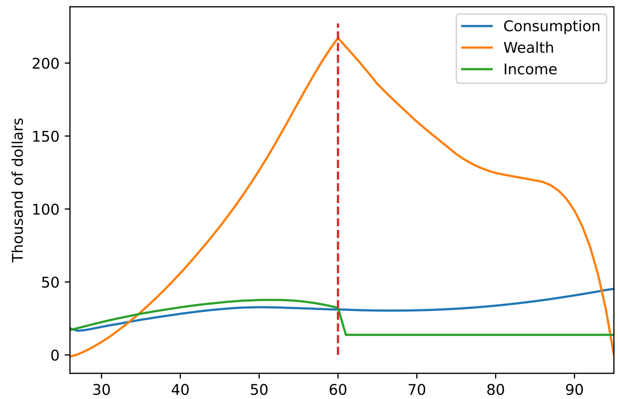

4.1 Consumption, wealth, utilization and debt

The life cycle profile for the mean of consumption, wealth and income are plotted in figure 5. Net labor income and social security benefits are exogenously fed into the model. Mean consumption grows before retirement mainly because of the precautionary saving motive, which is present in the model due to the existence of the income and borrowing limit shocks. The consumption trajectory continues to grow after retirement because of the positive intertemporal saving motive, i.e., $\beta^p R > 1$, for patient agents (see table C1). In order to generate the same amount of wealth and debt in the model as it is in the data, the discount factor for patient agents must be greater than one and at the same time the discount factor for impatient agents must be lower than one. Although the difference in discount rates traces different consumption paths, the increase in the consumption of patient agents after retirement dominates the evolution of aggregate consumption even though the population of impatient agents is greater in this economy.

Wealth dynamics follow the typical evolution path in standard models where households start to accumulate assets early in life for precautionary motives and reduce asset holdings after retirement. The economy, on average, is not indebted even though households have non-zero credit limits over their life cycle.

Figure 6 shows the life cycle profile for the mean of credit utilization; i.e., the debt to credit limit ratio, and the total debt to income ratio. Young households are more indebted on average and use approximately 60 percent of their available credit limit the first five years of their working life. This rate decreases almost linearly over the life cycle and is lower than 0.1 after age 50. The debt to income ratio follows a similar dynamic and household debt represents 15 percent of the available income between 26 to 30 years and represents less than 1 percent of the total income between 55 to 60 years.

This latter picture suggests that young households use debt as an additional mechanism to insure against idiosyncratic shocks. Once they start to accumulate assets and income they decrease their debt holdings as they can use precautionary savings to face adverse shocks during the working life.

Figure 6: Credit utilization and debt over the life cycle

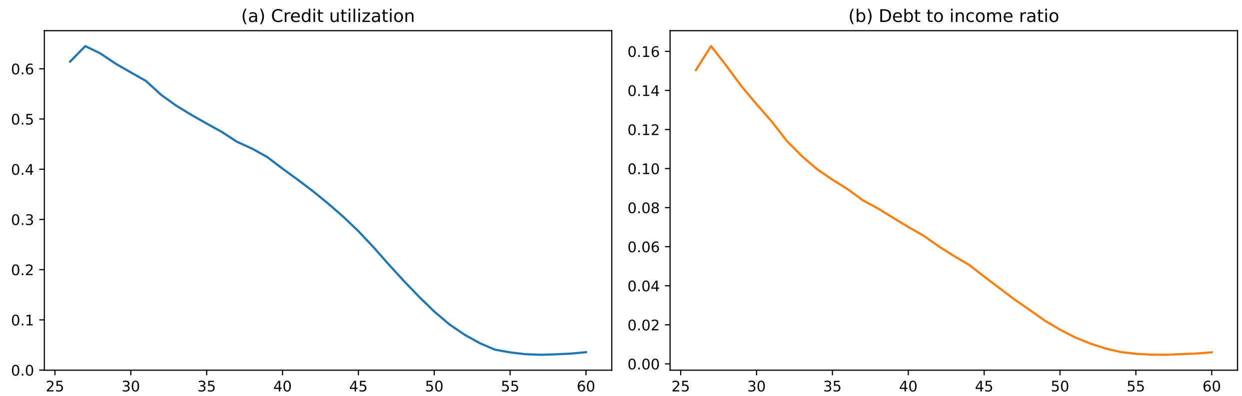

Table 3: Insurance coefficients

	Borrowing Limit Shock	Income Shock
Transitory	0.89	0.57
Persistent	0.85	0.11

4.2 Borrowing limits, income and self-insurance

Households can insure themselves against adverse shocks by accumulating precautionary savings or taking on debt. The extended framework in this paper is intended to assess the role of credit card limits as an additional source of uncertainty.

Following Blundell et al. (2008) and Kaplan and Violante (2010), I compute the insurance coefficients, i.e., the share of the variance of the shock ψ that does not translate into consumption growth, as follows:

$$\phi_\psi = 1 - \frac{cov(\Delta c_h^i, \psi_h^i)}{var(\psi_h^i)}, \tag{7}$$

where values close to one (zero) of ϕ_ψ denote a perfect (imperfect) insurance against the realizations of the shock ψ.

Using U.S. microlevel data, Blundell et al. (2008) estimate the insurance coefficients between income shocks and consumption are 0.95 and 0.36 for transitory and permanent shocks, respectively. Using a calibrated life-cycle model for the US economy, Kaplan and Violante (2010) calculate these coefficients in the ranges of 0.82-0.94 and 0.07-0.23, respectively. Lee and Sæverud (2023), motivated by results on subjective earnings expectations from Denmark, use an environment in which workers have partial information about income shocks and calculate consumption insurance coefficients for transitory income shocks of 0.79 and for permanent income shocks of 0.41. Other empirical studies have found that the level of insurance against transitory income shocks is in the range of 0.5-0.9 (e.g., Parker, Souleles, Johnson and McClelland (2013) using tax rebates, and Fagereng, Holm and Natvik (2021) using lottery winnings).

Table 3 shows the average insurance coefficients for the entire population in the benchmark economy. Both transitory and persistent borrowing limit shocks have a similar insurance coefficient implying that both shocks are important to explain consumption variation. Furthermore, the magnitude of these values close to 0.9 indicates that the consumption growth is not greatly affected by borrowing limit shocks. Even though the estimated stochastic variation in borrowing limits is high compared to standard estimates of income variation, not all households in this economy use their entire credit limit and go into debt as can be seen in figure 6.

In the case of income, the insurance coefficients are 0.57 (transitory shock) and 0.11 (persistent shock). The benchmark model is able to generate a degree of insurance to persistent and transitory income shocks that is similar and comparable to the empirical literature. However, the insurance coefficients are relatively lower than other similar quantitative studies suggesting a lower degree of insurance in this model. This is because the model incorporates a different channel of uncertainty, changes in credit conditions. When borrowing limits change stochastically, it is more difficult for households, especially those with little wealth or precautionary savings, to smooth consumption using debt.

4.3 Heterogeneous insurance coefficients

The magnitude of the effect of transitory and persistent shocks differs along different dimensions. Figure 7 shows the life cycle profile for the insurance coefficients. In general, young households are more affected by idiosyncratic shocks. The insurance coefficients for transitory borrowing limit shocks are below the aggregate average of 0.9 the first 15 years of the working life. After the age of 40, the size of the insurance coefficient is above the entire population average implying a higher degree of consumption insurance.

Surprisingly, persistent borrowing limit shocks have a similar degree of consumption insurance over the life cycle. Even though the variance of persistent shocks and their persistence follow opposite trajectories over time, households can partially insure themselves against these shocks at an early age and almost completely just before retirement.

The life cycle profile for income shocks follow a different trajectory. Both, transitory and persistent income shocks are important and crucial for consumption insurance for young households. Transitory income shocks can explain almost 50 percent of the variation in consumption in the first 10 years of working life, but over time households will accumulate enough assets to reverse this situation at later ages.

In contrast, insurance coefficients for persistent income shocks are lower compared to the ones for transitory income shocks and they are even negative over the first decade. According to equation (7), a negative value for this coefficient implies that consumption responds more than one-to-one to persistent income shocks. As discussed in Kaplan and Violante (2010), this result can be explained because of the interaction between transitory and persistent shocks in the model. When transitory and persistent shocks are realized, the marginal cost of not being able to smooth persistent income shocks rises compared to other shocks. The total response of consumption takes into account the direct effect of persistent income shocks and the decrease in precautionary wealth due to the

Figure 7: Insurance coefficients by age

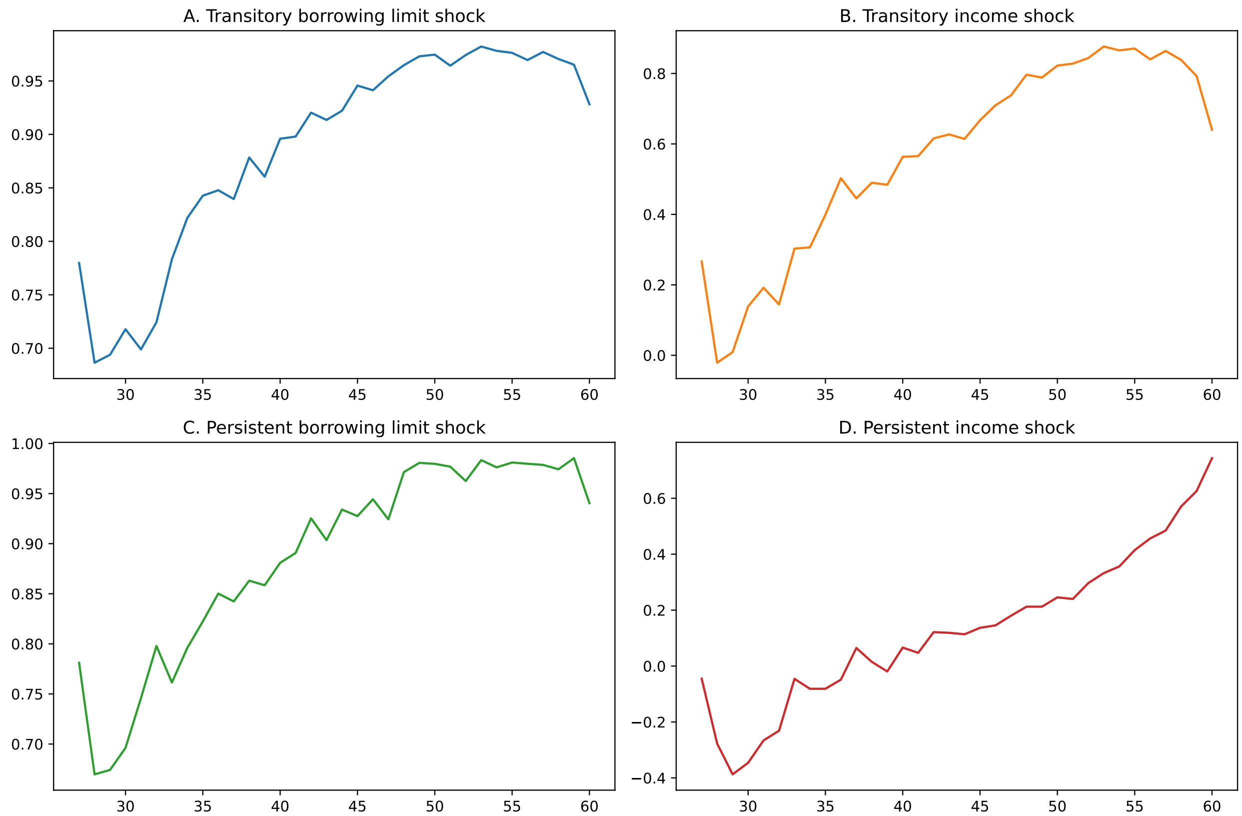

relative importance of other shocks for households. Starting at age 40, consumption responses to these shocks begin to decrease slowly but remain below the value of 0.7, indicating the importance of persistent income shocks throughout the entire working life.

Wealth is one of the self-insurance mechanism households have to face idiosyncratic shocks in this model, and we can confirm this by plotting the consumption insurance coefficients by wealth deciles. Households in the first four deciles are characterized by having negative wealth, while the upper deciles are not indebted and have positive wealth holdings. According to figure 8, insurance coefficients also follow a u-shaped curve for low levels of wealth.

The first four wealth deciles are the most affected by transitory and persistent shocks to borrowing limits. Consumption responds more to transitory borrowing limit shocks in the second and third deciles, in contrast to what happens in the first decile. This is because the latter type of households have more access to take debt and insure themselves after the shocks are realized in the economy. Once access to credit starts to decrease, households must adjust their consumption in response to the shocks they are facing. Something similar happens for persistent borrowing limit shocks. Wealth deciles above four can perfectly insured against borrowing limit shocks because they are not indebted and therefore they are not affected by borrowing limit shocks.

The second, third and fourth deciles are most affected by transitory income shocks. In these cases, the insurance coefficient is negative implying a more than one-to-one response of consumption

Figure 8: Insurance coefficients by wealth deciles

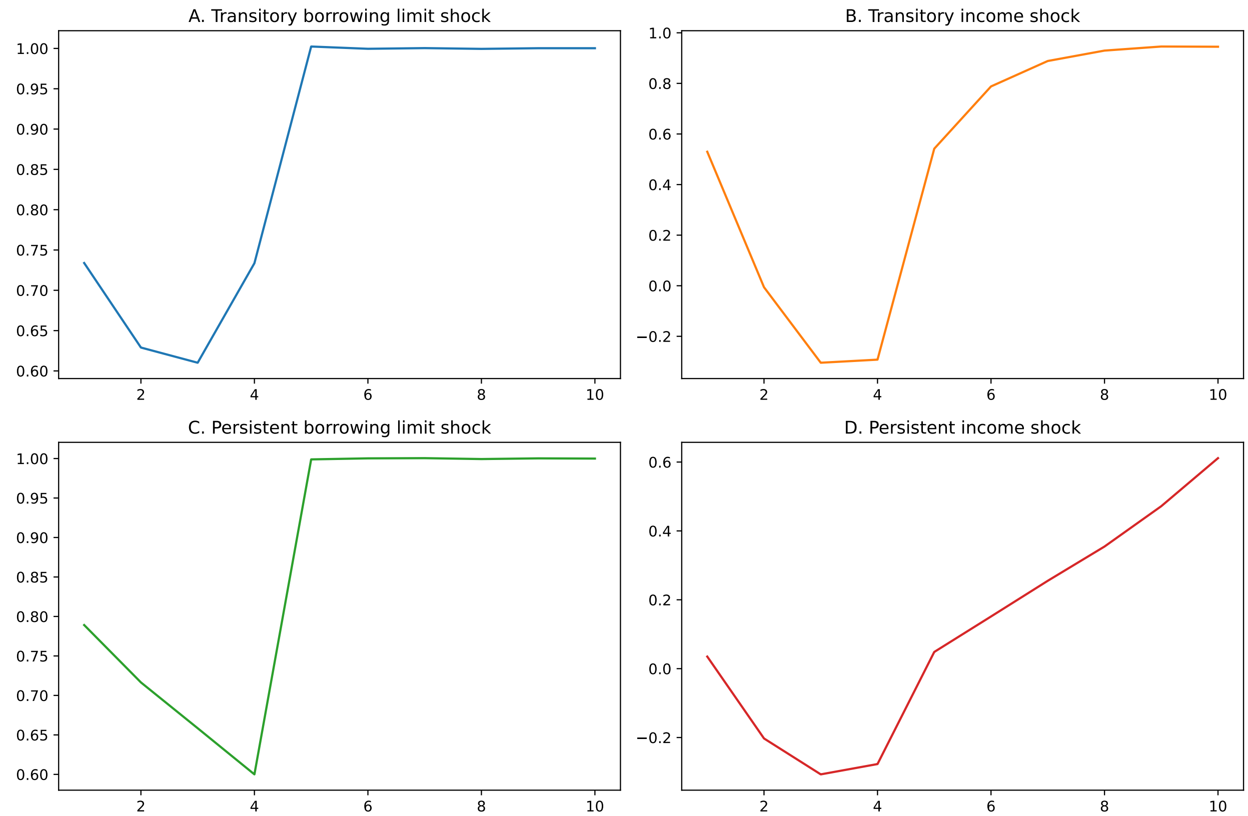

to these shocks. As explained above, households in the first decile of the wealth distribution face loose credit limits, but households in the next three deciles do not. For them, transitory shocks are relatively more important compared to borrowing limit shocks and consumption responds more to the former shocks. Finally, persistent income shocks are not easy to insure and the first five deciles in the wealth distribution are most affected by this type of shock.

5 How important are stochastic borrowing limits?

To study the relevance of stochastic borrowing limits, I solve the benchmark model by removing idiosyncratic variation in credit limits in three ways: (i) cancel transitory and persistent stochastic components, i.e., $u_h^i = n_h^i = 0$ in equation (4); (ii) cancel transitory and persistent stochastic components and eliminate the dependence of credit limits on labor income, i.e., $u_h^i = n_h^i = \theta_y = 0$ in equation (4); and (iii) assume a constant and invariant borrowing limit over the life cycle, i.e., $\theta_1 = \theta_2 = \theta_y = u_h^i = n_h^i = 0$ in equation (4). This last specification is similar to the one used in most of the previous literature that study consumption and wealth over the life cycle (Guvenen, 2007; Kaplan and Violante, 2010; Arslan et al., 2020).

When performing these three exercises, I recalibrate the discount factor for patient agents, the discount factor for impatient agents, the share of impatient agents and the intercept of the

Table 4: Moments for different model specifications

	Data	Benchmark	No Stoch. BL	No Stoch. BL + No Income	Invariant No Stoch. BL
Panel A. Targeted moments					
Wealth to income ratio	3.10	3.10	3.10	3.10	3.10
Debt to income ratio	0.05	0.05	0.05	0.05	0.05
Capital share at top 20%	0.80	0.80	0.80	0.80	0.80
Credit limits to income ratio	0.27	0.27	0.27	0.27	0.27
Panel B. Non-targeted moments					
Utilization rate	0.26	0.25	0.19	0.21	0.19
Variance of log debt	16.48	16.02	16.81	17.66	16.71

Table 5: Consumption and wealth distribution (age 26 to 60)

	Consumption		Wealth	
	Median	St. Dev.	Median	St. Dev.
Benchmark	22,790.25	17,285.78	7,483.74	175,494.45
No Stochastic BL	22,646.55	17,227.70	7,562.93	170,249.07
No Stochastic BL + No Income	22,681.25	17,586.94	5,501.32	178,009.15
Invariant No Stochastic BL	22,658.64	17,439.58	6,911.91	177,001.17

borrowing limit process to match the empirical moments shown in table 2.

5.1 Households' decisions

Table 4 shows that these alternative specifications can match the targeted moments in the data. This also applies to debt variation that is not targeted across different specifications. However, the average utilization rate generated in these models is below the reported one in the data, suggesting a relative gain of the benchmark model with respect to other specifications.

Table 5 shows the median and standard deviation for consumption and wealth distributions during the working life (from 26 to 60 years). In the benchmark model, the median consumption is the highest, which can be explained by two factors. First, the discount factors necessary to match the empirical moments are greater in the benchmark model and consumption grows faster in this case. Second, the benchmark model exhibits an additional source of uncertainty, stochastic borrowing limits. In the presence of this component, households accumulate more assets to self-insure against all the shocks they face and this accumulation will at some point increase their overall consumption (Dávila, Hong, Krusell and Ríos-Rull, 2012). This is confirmed by comparing the results in column 4 in table 5 and realizing that the median wealth distribution is high in the benchmark economy.

In order to match the credit limits to income ratio, the last two specifications (third and fourth rows) must have looser borrowing limits compared to the first two economies (first and second rows). With certain borrowing limits, households take the decision to borrow more or less depending solely on their income realizations, therefore the standard deviation for consumption and wealth are high in the last two specifications.

Figure 9: Credit utilization and debt

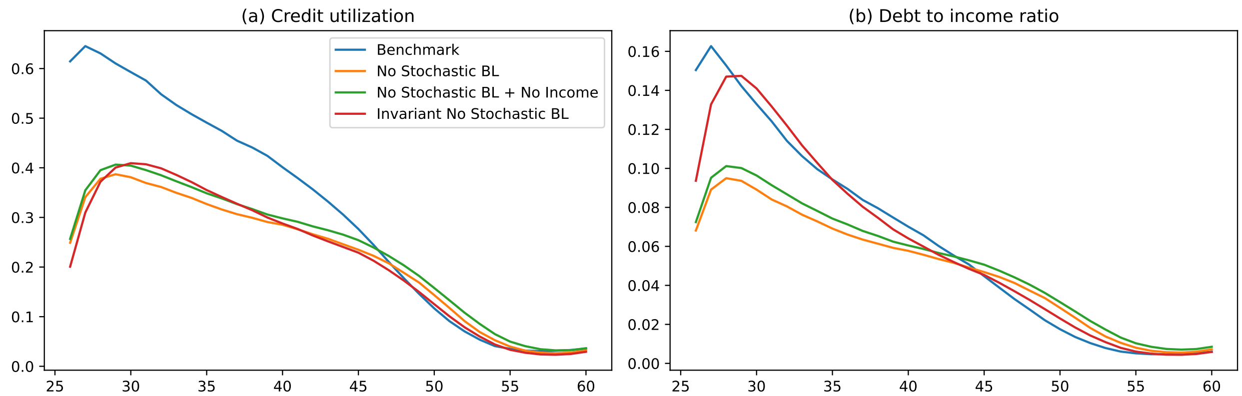

Although the variance of borrowing limit shocks is larger compared to income shocks, households prefer to take more debt specially when they are young. Figure 9 shows the credit utilization rate and the debt-to-income ratio across different model specifications.

Stochastic borrowing limits randomly create tight or loose credit conditions for households throughout the working life. Households take this into account and start to use a larger fraction of their available borrowing limit at early ages. However, since they also need to accumulate positive wealth to pay their debts and smooth consumption after retirement, the slope of credit utilization is steeper than economies where stochastic borrowing limits are absent.

The same intuition applies for the debt-to-income ratio where young households are more indebted and their debt as a fraction of their total income is higher. A particular case of interest is the economy with invariant and no stochastic borrowing limits. This economy shows flatter credit utilization rates but a high debt-to-income ratio early in life, similar to the benchmark model. This is because, in the three first models borrowing limits increase with age as they do in the data. However, in the last case, borrowing limits are constant and they are looser to match the credit limit to income ratio. Since young households are constrained, they can now take on more debt early in life.

Overall in net, stochastic borrowing limits create an incentive for households to take more debt when they are young and consume more by accumulating wealth later in life. The remaining question is how borrowing limit shocks interact with labor income shocks to affect consumption behavior over the life cycle.

5.2 Insurance coefficients

Table 6 shows the aggregate consumption insurance coefficients for the model specifications discussed above. Borrowing limit shocks are only included in the benchmark model, but income shocks are included in all model specifications.

There are two findings that we can highlight from this table. First, in the absence of idiosyncratic borrowing limit shocks, households can self-insure better against income shocks. Consumption

Table 6: Aggregate insurance coefficients

	Borrowing Limit Shock		Income Shock	
	Transitory	Persistent	Transitory	Persistent
Benchmark	0.89	0.85	0.57	0.11
No Stochastic BL			0.75	0.22
No Stochastic BL + No Income			0.74	0.21
Invariant No Stochastic BL			0.76	0.21

Figure 10: Insurance coefficients by age

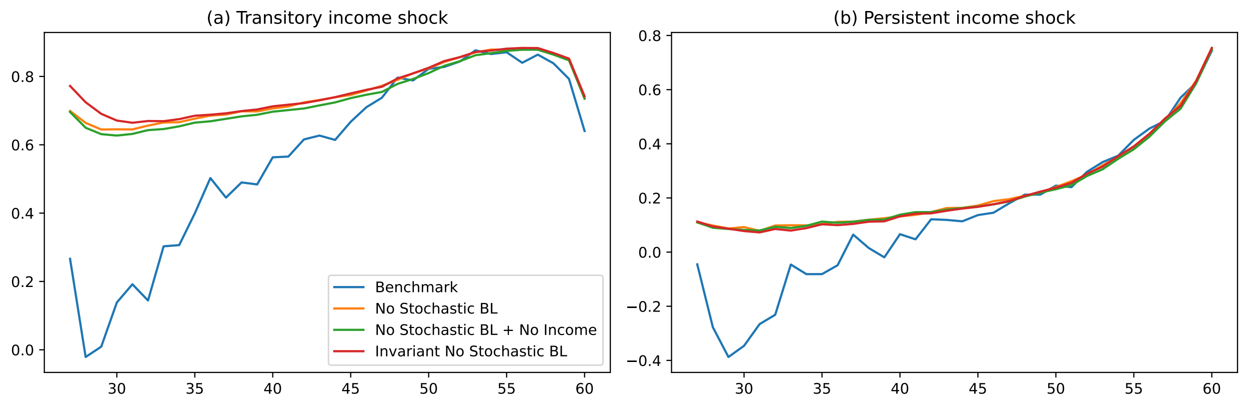

responds approximately 30% less to transitory income shocks (from 0.57 to 0.75) and 90% less to persistent income shocks (from 0.11 to 0.21) when there are no borrowing limit shocks.

Second, insurance with respect to persistent income shocks are aligned with findings from previous studies that report a coefficient in the range of 0.07-0.36. However, the insurance to transitory income shocks is lower (0.74-0.75) with respect to other studies (0.82-0.95). This is because in this model, impatient agents seek to consume more in the present at the expense of losing insurance power with respect to this type of shocks.

Different disaggregated analyzes show which types of agents are most exposed to income shocks. Figure 10 shows that without stochastic borrowing limit shocks, consumption of young households is less affected by transitory and persistent income shocks. In general, households can insure more than 50% of transitory income shocks but in the presence of stochastic borrowing limits and the desire to borrow early in life, households will insure less against transitory income shocks.

Panel (b) shows that it is more difficult to protect against persistent income shocks. In fact, households can insure less than 30% of these shocks before age 50. With stochastic borrowing limit shocks, households will insure less against these shocks at early ages but will eventually catch-up the same level of insurance before retirement.

Figure 11 shows the consumption insurance coefficients calculated for each decile in the wealth distribution. An important drop in consumption insurance is observed for the second, third and fourth deciles when borrowing limits are stochastic. As explained above, households that want to borrow will use their resources, precautionary savings and debt, to protect against borrowing limit

Figure 11: Insurance coefficients by wealth

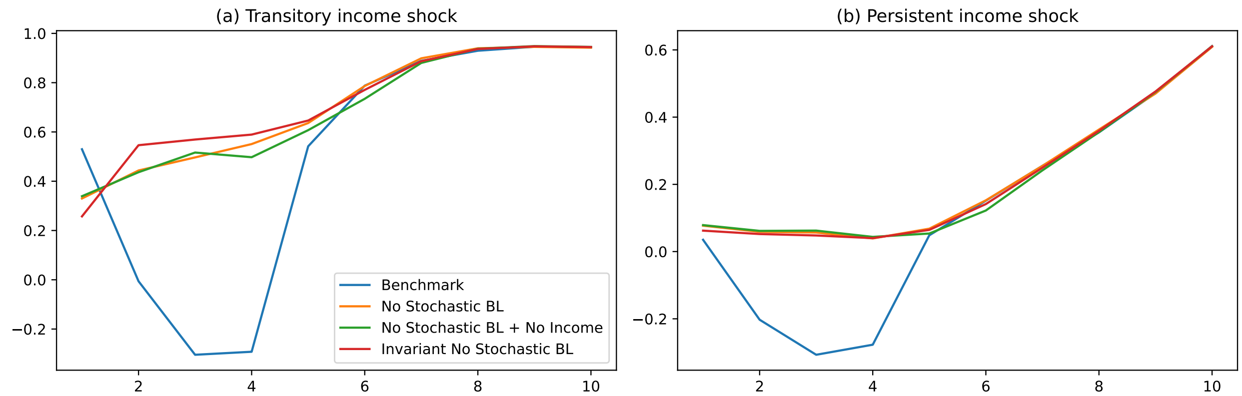

and income shocks.

Self-insure against transitory income shocks is more difficult to less wealthy households because even though they have access to debt, they have less precautionary resources to face unfavorable scenarios. Without stochastic borrowing limit shocks, the consumption insurance coefficient for the first five deciles in the wealth distribution is in the range of 0.3-0.6. Consumption insurance for upper deciles is similar across all the model specifications.

Consumption of less wealthy households is more affected by persistent income shocks. The insurance coefficients for the first five deciles in the wealth distribution is quite constant and lower than 0.15. The top five deciles of the distribution present similar insurance coefficients.

In sum, stochastic borrowing limit shocks affect the ability of household to smooth labor income shocks. Younger and low middle wealthy households are most affected by the presence of shocks to borrowing limits, while older households with higher levels of wealth show similar insurance coefficients with and without these shocks. This is because these households are naturally constrained in terms of precautionary savings (negative or low positive wealth) and access to credit (tight credit limits).

5.3 Decomposition of the effects of stochastic borrowing limits

What explains the lower degree of household insurance against income risk when borrowing limits are stochastic? To answer this question, it is important to note that the presence of risk in credit markets affects households' optimal consumption decisions in two ways. First, by adjusting consumption directly when borrowing limits and income shocks are realized (the *policy function* effect). Second, by inducing households to borrow more or less depending on their realized shock, modifying debt and saving decisions and indirectly affecting consumption. This last indirect channel could have important distributional effects on wealth and consumption because it impacts both the debt channel and the precautionary savings channel that households have to face adverse shocks (the *distributional* effect).

In general, the consumption insurance coefficient with stochastic borrowing limits depends on

Figure 12: Effects of stochastic borrowing limits

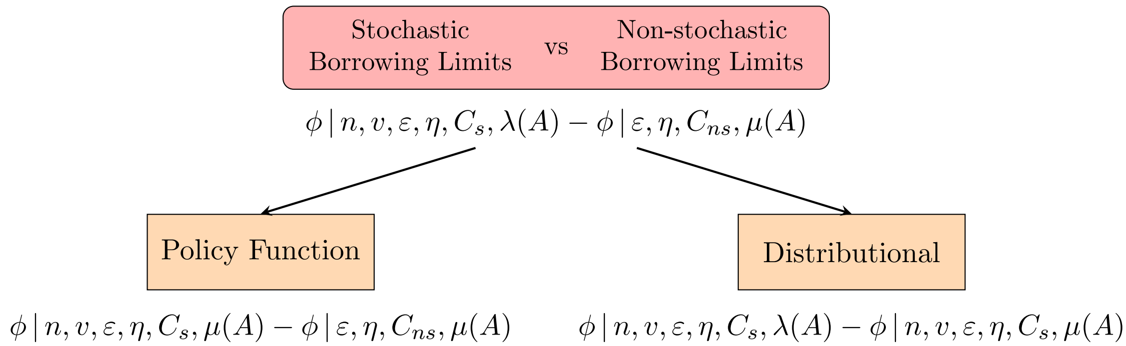

the realization of borrowing limit shocks and income shocks, the calculated consumption policy function, and the optimal choice of assets ($\phi \,|\, n, v, \varepsilon, \eta, C_s, \lambda(A)$). On the other hand, in the environment without borrowing limit shocks, the degree of consumption insurance depends on the realization of income shocks, the calculated consumption policy function, and the optimal choice of assets ($\phi \,|\, \varepsilon, \eta, C_{ns}, \mu(A)$). These two environments have their own generated consumption policy functions that capture the effect that state variables (stochastic and pre-determined) have on consumption.

To isolate the policy function effects of shocks to borrowing limits on consumption insurance, the choice of assets in the economy (asset distribution) should not respond to stochastic changes in borrowing limits. To calculate this counterfactual, I will resolve the benchmark model with the same calculated policy function and simulated idiosyncratic shocks, but use as input the asset distribution calculated from alternative specifications ($\phi \,|\, n, v, \varepsilon, \eta, C_s, \mu(A)$). The difference between this counterfactual and the insurance coefficient computed when stochastic borrowing limits are not stochastic, account for the policy function effect of borrowing limits shocks on consumption insurance ($\phi \,|\, n, v, \varepsilon, \eta, C_s, \mu(A) - \phi \,|\, \varepsilon, \eta, C_{ns}, \mu(A)$). If the difference is small, then the loss of consumption insurance against income shocks comes exclusively from the indirect effect of credit shocks through asset accumulation on consumption (see figure 12).

Following this logic, the distributional effect is calculated as the difference between the insurance coefficient in the benchmark model and the counterfactual insurance coefficient defined above ($\phi \,|\, n, v, \varepsilon, \eta, C_s, \lambda(A) - \phi \,|\, n, v, \varepsilon, \eta, C_s, \mu(A)$). This difference compares two economies facing the same risks, but assets in the benchmark economy are affected by both income and borrowing limit shocks $\lambda(A)$, while assets in the counterfactual economy are only affected by income shocks $\mu(A)$. When this difference is significantly negative, stochastic borrowing limits are impacting the distribution of wealth, creating incentives for households to take on more debt than they would if borrowing limits were not stochastic. Since agents necessarily have to pay their debt, self-insurance against income shocks is severely affected by changes in credit conditions.

Table 7 shows the decomposition of the effects of stochastic borrowing limits as a percentage of the degree of insurance in environments without credit risk. Columns 2 and 5 show the policy function effect on consumption insurance using as input the asset distribution from multiple speci-

Table 7: Decomposition of the effects of stochastic borrowing limits

	Transitory Income Shock			Persistent Income Shock		
	Pol. Func.	Distrib.	Total	Pol. Func.	Distrib.	Total
No Stochastic BL	18.48	-42.57	-24.10	24.44	-73.89	-49.45
No Stochastic BL + No Income	20.08	-43.02	-22.94	27.21	-75.68	-48.47
Invariant No Stochastic BL	17.03	-42.02	-24.99	28.84	-77.31	-48.47

Notes: The numbers above are proportional to the degree of consumption insurance in the model without stochastic borrowing limits.

fications. Compared to the original benchmark model, households can better insure against income shocks in the counterfactual economy. On average, consumption responds between 17% to 20% less to transitory income shocks and between 24% to 29% less to persistent income shocks. The difference takes into account the direct effect of stochastic borrowing limits on consumption decisions. In the presence of these shocks, self-insurance against risks is stronger and this is reflected in the benchmark policy function. Once the sources of risk in credit limits are attenuated, households would be perfectly insured against borrowing limit shocks and self-insurance resources would only be used to face income shocks.[11]

Columns 3 and 6 of table 7 show the distributional effects of stochastic borrowing limits by including the asset distribution from multiple specifications in the benchmark model. Compared to the counterfactual model, consumption insurance is lower in the benchmark model. On average, consumption responds between 42% to 43% more to transitory income shocks and between 74% to 77% more to persistent income shocks. The difference can be explained by the fact that households are taking on less debt in the absence of stochastic borrowing limits which allow them to better smooth consumption when income shocks are realized. Similar to the previous case, the distribution of assets in the counterfactual economy $(\phi \,|\, n, v, \varepsilon, \eta, C_s, \mu(A))$ is not affected by borrowing limit shocks, and the low level of insurance in the benchmark economy $(\phi \,|\, n, v, \varepsilon, \eta, C_s, \lambda(A))$ is explained by the high level of debt taken on by young and low middle wealth households. This finding can be confirmed with previous results showing high credit utilization and debt in the benchmark model (see figure 9).

Finally, the exercise above sheds light on the main beneficiaries in the absence of risks in borrowing limits. Figure D1 shows that young households could almost perfectly insure themselves against transitory income shocks. Similar to the previous discussion, the combination between small debt positions in the alternative specifications, and the relatively stronger response to idiosyncratic risks in the benchmark model account for the difference. A similar conclusion is reached in the case of persistent income shocks, but with the difference that young households are not better off than old households (see figure D2). Overall, the distributional effects of stochastic borrowing limits

[11]Even though the policy function in the benchmark model takes into account stochastic changes in borrowing limits, the level of assets held by households is not affected by these changes in the counterfactual specification. In fact, households are completely insured when the exercise of keeping the distribution of wealth constant and only changing the policy function is carried out (see table D1).

through asset accumulation mainly explain the loss of capacity to self-insure against income shocks.

6 Conclusion

In this paper, I study the role of time-varying borrowing limits on consumption by estimating how credit card limits vary over the life cycle and incorporating a realistic characterization of these changes into a calibrated model for the U.S. that reproduces key wealth and unsecured debt moments.

Credit card limits increase with household age and income, decrease with measures of credit risk, and the estimated unobservable variation is quite large. After decomposing this last variation into a transitory and persistent components that vary with age, I am able to find that the variance of transitory shocks increases slightly at early ages but remains constant for most of the life cycle. On the other hand, the variance of persistent shocks decreases throughout the life cycle while its persistence follows an opposite direction. These findings are consistent with better information that credit agencies have about borrowers, resulting in greater certainty about credit qualifications over time.

The quantitative exercise shows that households take on more debt and credit utilization is high when borrowing limits are stochastic. However, despite the estimated variability of borrowing limits, aggregate consumption is not greatly affected by transitory or persistent shocks of this type. The results also suggest heterogeneous consumption responses to uninsurable idiosyncratic shocks. Younger and less wealthy households are most affected by borrowing limits and income shocks. Moreover, in the presence of borrowing limit shocks, households lose the ability to self-insure against income shocks. The effects of these shocks on the distribution of asset and the precautionary channels households have to insure against risks mainly explain the low levels of consumption insurance.

The paper leaves open some avenues for future research. First, it would be interesting to extend the framework to include financial intermediation and endogenous borrowing limits. The model presented in this paper attempts to assess the degree of household response to exogenous variation in credit limits, but it is natural to think that part of this variation is also endogenous. Second, I have based the model framework by isolating one dimension of access to credit, changes in borrowing limits. The incorporation of other financial frictions, such as credit spreads or default options, could shed light on a better understanding of consumer insurance throughout the life cycle.

References

Aguiar, M. A., Bils, M. and Boar, C. (2020), 'Who Are the Hand-to-Mouth?'.

Arslan, Y., Guler, B. and Taskin, T. (2020), 'Price Search, Consumption Inequality and Expenditure Inequality over the Life-Cycle', *International Economic Review* **62**(1), 295–320.

Athreya, K. B. (2008), 'Default, Insurance, and Debt over the Life-Cycle', *Journal of Monetary Economics* **55**(4), 752–774.

Athreya, K. B. and Simpson, N. B. (2006), 'Unsecured Debt with Public Insurance: From Bad to Worse', *Journal of Monetary Economics* **53**(4), 797–825.

Blundell, R., Pistaferri, L. and Preston, I. (2008), 'Consumption Inequality and Partial Insurance', *American Economic Review* **98**(5), 1887–1921.

Braxton, J. C., Herkenhoff, K. F. and Phillips, G. M. (2020), 'Can the Unemployed Borrow? Implications for Public Insurance'.

Brown, M., Stein, S. and Zafar, B. (2015), 'The Impact of Housing Markets on Consumer Debt: Credit Report Evidence from 1999 to 2012', *Journal of Money, Credit and Banking* **47**(S1), 175–213.

Carroll, C. D. (2006), 'The method of endogenous gridpoints for solving dynamic stochastic optimization problems', *Economics Letters* **91**(3), 312–320.

Carroll, C. D. and Samwick, A. A. (1997), 'The nature of precautionary wealth', *Journal of Monetary Economics* **40**(1), 41–71.

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

Chatterjee, P., Gunawan, D. and Kohn, R. (2022), The Interaction Between Credit Constraints and Uncertainty Shocks, Technical report.

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

Crossley, T. F. and Low, H. (2011), 'Borrowing Constraints, the Cost of Precautionary Saving and Unemployment Insurance', *International Tax and Public Finance* **18**(6), 658–687.

Dávila, J., Hong, J. H., Krusell, P. and Ríos-Rull, J.-V. (2012), 'Constrained Efficiency in the Neoclassical Growth Model with Uninsurable Idiosyncratic Shocks', *Econometrica* **80**(6), 2431–2467.

Devlin-Foltz, S. and Sabelhaus, J. (2016), 'Heterogeneity in Economic Shocks and Household Spending in the US', *Fiscal Studies* **37**(1), 153–192.

Dobbie, W. and Song, J. (2020), 'Targeted Debt Relief and the Origins of Financial Distress: Experimental Evidence from Distressed Credit Card Borrowers', *American Economic Review* **110**(4), 984–1018.

Drechsel, T. (2023), 'Earnings-Based Borrowing Constraints and Macroeconomic Fluctuations', *American Economic Journal: Macroeconomics* **15**(2), 1–34.

Faccini, R., Lee, S., Luetticke, R., Ravn, M. O. and Renkin, T. (2024), 'Financial Frictions: Micro vs Macro Volatility'.

Fagereng, A., Holm, M. B. and Natvik, G. J. (2021), 'MPC Heterogeneity and Household Balance Sheets', *American Economic Journal: Macroeconomics* **13**(4), 1–54.

Floden, M. and Lindé, J. (2001), 'Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?', *Review of Economic Dynamics* **4**(2), 406–437.

Fulford, S. L. (2015), 'How Important is Variability in Consumer Credit Limits?', *Journal of Monetary Economics* **72**, 42–63.

Fulford, S. and Schuh, S. D. (2015), 'Consumer revolving credit and debt over the life cycle and business cycle'.

Fulford, S. and Schuh, S. D. (2017), 'Credit Card Utilization and Consumption over the Life Cycle and Business Cycle', *SSRN Electronic Journal* .

Gouveia, M. and Strauss, R. P. (1994), 'Effective federal individual income tax functions: An exploratory empirical analysis', *National Tax Journal* **47**(2), 317–339.

Gross, D. B. and Souleles, N. S. (2002), 'Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data', *The Quarterly Journal of Economics* **117**(1), 149–185.

Guerrieri, V. and Lorenzoni, G. (2017), 'Credit Crises, Precautionary Savings, and the Liquidity Trap', *The Quarterly Journal of Economics* **132**(3), 1427–1467.

Guvenen, F. (2007), 'Learning Your Earning: Are Labor Income Shocks Really Very Persistent?', *American Economic Review* **97**(3), 687–712.

Guvenen, F. (2009), 'An empirical investigation of labor income processes', *Review of Economic Dynamics* **12**(1), 58–79.

Herkenhoff, K. F. (2019), 'The Impact of Consumer Credit Access on Unemployment', *The Review of Economic Studies* **86**(6), 2605–2642.

Herkenhoff, K., Phillips, G. and Cohen-Cole, E. (2016), How Credit Constraints Impact Job Finding Rates, Sorting and Aggregate Output, Technical report, National Bureau of Economic Research.

Jensen, H., Ravn, S. H. and Santoro, E. (2018), 'Changing Credit Limits, Changing Business Cycles', *European Economic Review* **102**, 211–239.

Kaplan, G. (2012), 'Inequality and the Life Cycle', *Quantitative Economics* **3**(3), 471–525.

Kaplan, G. and Violante, G. L. (2010), 'How Much Consumption Insurance Beyond Self-Insurance?', *American Economic Journal: Macroeconomics* **2**(4), 53–87.

Karahan, F. and Ozkan, S. (2013), 'On the Persistence of Income Shocks over the Life Cycle: Evidence, Theory, and Implications', *Review of Economic Dynamics* **16**(3), 452–476.

Lee, E. and Sæverud, J. (2023), 'Earnings Shocks, Expectations, and Spending'.

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

Mitchell, O. S. and Phillips, J. W. (2006), 'Social Security Replacement Rates for Alternative Earnings Benchmarks.', *Benefits Quarterly* **22**(4).

Parker, J. A., Souleles, N. S., Johnson, D. S. and McClelland, R. (2013), 'Consumer Spending and the Economic Stimulus Payments of 2008', *American Economic Review* **103**(6), 2530–2553.

Primiceri, G. E. and Rens, T. v. (2009), 'Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality', *Journal of Monetary Economics* **56**(1), 20–39. Publisher: Elsevier BV.

Storesletten, K., Telmer, C. I. and Yaron, A. (2004*a*), 'Consumption and Risk Sharing over the Life Cycle', *Journal of Monetary Economics* **51**(3), 609–633.

Storesletten, K., Telmer, C. I. and Yaron, A. (2004*b*), 'Cyclical Dynamics in Idiosyncratic Labor Market Risk', *Journal of Political Economy* **112**(3), 695–717.

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

Appendix A Description of the data

A synthetic panel is created using data from the Survey of Consumer Finances from the period 1989-2019. I constructed household bins according to the income quartile group, household head's education category, household marital status, adn household head's race. According to this classification, observable variables are collapsed by their weighted average for all the bins defined for each age group and year. The final data used in the estimation results consists on an unbalanced panel data for 2,930 household bins and 9,543 observations.

Table A1 shows the definition of each variable used on the estimation of the credit limit process.

Table A1: Definition of variables

Variable	Definition and construction
Age category	Age group of the household reference person. Household head's age is between 26 and 60 years old.
Income quartile groups	The 4 categories are: 0-24.9, 25-49.9, 50-74.9, 75-100 according to the total amount of household income in 2019 dollars.
Education category	Education group of the household reference person. 4 categories: no high school diploma/GED, high school diploma or GED, some college and college degree.
Marital status category	Marital status of the household reference person. 2 categories: married or living with a partner and neither married nor living with a partner.
Race	Household bin's fraction of white non-Hispanic household heads. Calculated as the weighted average of the household reference person's race/ethnicity status whose values are equal to: 1 = white non-Hispanic; 0 = other race group.
Total credit cards borrowing limit	Average amount a household bin could borrow on credit cards. This is calculated as the weighted average of households' total credit cards borrowing limit for each bin on a given year. It only includes bank-type credit cards (Visa, Mastercard, Discover, Optima). Other store, gasoline, airlines and card rental cards are excluded.
Wage income	Average amount of the household bin's wage and salary income. Calculated as the weighted average of the household wage and salary income in millions of 2019 dollars.
Education	Household bin's average highest completed grade. Calculated as the weighted average of the household reference person's highest completed grade whose values are equal to: 1 = Less than 1st grade, 1st, 2nd, 3rd or 4th grade; 2 = 5th or 6th grade; 3 = 7th or 8th grade; 4 = 9th grade; 5 = 10th grade; 6 = 11th grade; 7 = 12th grade, no diploma; 8 = high school graduate - high school diploma or equivalent; 9 = some college but no degree; 10 = associate degree in college - occupation/vocation program; 11 = associate degree in college - academic program; 12 = bachelor's degree (e.g., BA, AB, BS); 13 = master's degree; 14 = doctorate or professional school degree.
Male	Household bin's fraction of male household heads. Calculated as the weighted average of the household reference person's gender whose values are equal to: 0 = female; 1 = male.

Table A1 (continued): Definitions of variables

Variable	Definition and construction
Work for someone else	Household bin's fraction of household heads working for someone else. Calculated as the weighted average of the household reference person's occupation category whose values are equal to: 1 = work for someone else; 0 = other groups.
Self-employed	Household bin's fraction of household heads who are self-employed or are in a partnership. Calculated as the weighted average of the household reference person's occupation category whose values are equal to: 1 = self-employed/partnership; 0 = other groups.
Number of credit cards	Average amount of the household bin's credit cards. Calculated as the weighted average of the total number of credit cards held by households.
Carry credit card balance	Household bin's fraction of households that carry a balance on credit cards. Calculated as the weighted average of the household balance on credit cards status whose values are equal to: 1 = if the household carry a balance on credit cards; 0 = if the household does not carry a balance on credit cards. Excludes charge accounts at stores.
Credit card balance	Average amount of the household bin's total value of credit card balances. Calculated as the weighted average of the total value of credit card balances held by households in millions of 2019 dollars. Balances do not include purchases made since the last account statement. From 1992 revolving debt at stores is also treated as credit card debt.
Interest rate on credit cards	Household bin's average interest rate paid on credit cards. Calculated as the weighted average of the interest rate the household pays on the credit card with the largest balance on a given year.
Bankruptcy in the past 5 years	Household bin's fraction of households have declared bankruptcy. Calculated as the weighted average of the variable indicating if the household has declared bankruptcy in the past 5 years whose values are equal to: 0 = no; 1 = yes.

Appendix B Other estimation results

Here I report the estimated parameters for the cubic specification.

Table B1: Estimation results for the cubic specification: age profiles

	γ_0	γ_1	γ_2	γ_3
σ^2_α	0.0218			
	(0.0002)			
ρ_h	0.5920	0.2864	-0.1460	0.0261
	(0.0008)	(0.0001)	(0.0000)	(0.0000)
$\sigma^2_{u,1}$	0.1459			
	(0.0073)			
$\sigma^2_{v,h}$	0.1955	-0.1051	0.0154	0.0014
	(0.0005)	(0.0001)	(0.0000)	(0.0000)
$\sigma^2_{n,h}$	0.2461	0.0880	-0.0449	0.0067
	(0.0000)	(0.0000)	(0.0000)	(0.0000)

Note: Bootstrap standard errors in parentheses. γ's are the coefficients of a cubic polynomial $x_h = \gamma_0 + \gamma_1 * (h/10) + \gamma_2 * (h/10)^2 + \gamma_3 * (h/10)^3$, for $x = \{\rho, \sigma^2_v, \sigma^2_n\}$.

Appendix C Calibrated parameters in the benchmark model

The table below shows the parameters externally and internally calibrated in the benchmark model.

Table C1: Benchmark model parameters

Parameter	Description	Value
Externally calibrated		
$R_+ - 1$	Saving interest rate	0.03
$R_- - 1$	Borrowing interest rate	0.1
γ	Relative risk aversion	2
ρ_z	Persistence of income process	1
σ_{z1}^2	Initial variance of persistent income	0.15
σ_η^2	Variance of persistent income shock	0.01
σ_ε^2	Variance of transitory income shock	0.05
θ_1	Age slope coefficient	0.0631
θ_2	Age squared slope coefficient	-0.0005
θ_3	Log income slope coefficient	0.0389
ρ_{uh}	Persistence of borrowing limit process	Figure 3
σ_{u1}^2	Initial variance of persistent borrowing limit	0.1459
σ_{vh}^2	Variance of persistent borrowing limit shock	Figure 3
σ_{nh}^2	Variance of transitory borrowing limit shock	Figure 3
τ_1	Tax function parameter	0.258
τ_2	Tax function parameter	0.768
τ_3	Tax function parameter	0.0006
Internally calibrated		
β^p	Patient agents discount factor	1.030
β^i	Impatient agents discount factor	0.901
λ^i	Share of impatient agents	0.680
θ_0	Intercept of the borrowing limit process	9.287

Appendix D Additional results of the indirect effects of stochastic credit limits

The table below shows the consumption insurance coefficients to borrowing limit shocks when the benchmark policy functions are used in combination with the distribution of assets generated by different specifications.

Table D1: Insurance coefficients to borrowing limit shocks with the benchmark policy function

Asset Distribution	Transitory BL shock	Persistent BL shock
Benchmark	0.8859	0.8468
No Stochastic BL	1.0011	1.0004
No Stochastic BL + No Income	0.9994	1.0009
Invariant No Stochastic BL	0.9993	1.0009

Figure D1 shows the life cycle profile of consumption insurance coefficients to transitory income shocks calculated in the benchmark model (blue benchmark line), in the alternative specification (calculated PF line), and with the benchmark policy functions using as input the asset distribution from alternative specifications (benchmark PF line).

Figure D1: Insurance coefficients to transitory income shocks by age

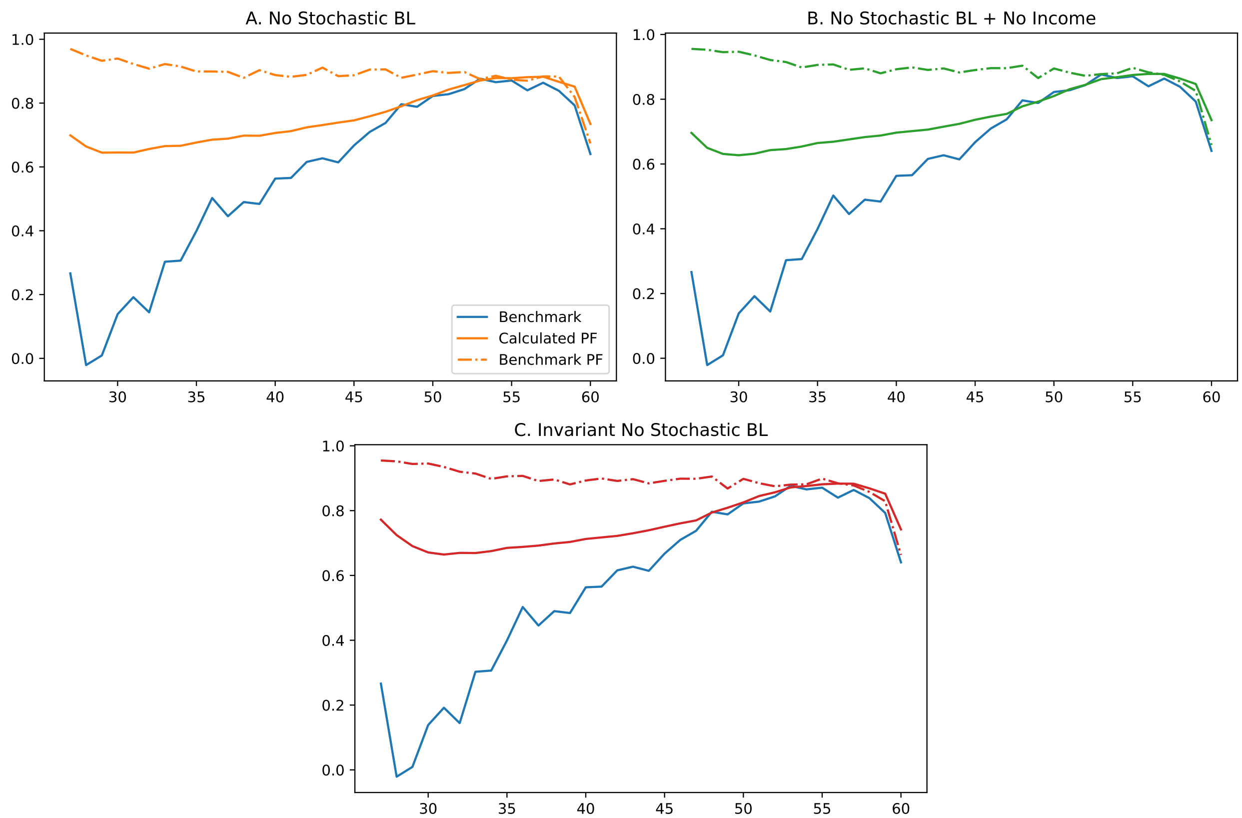

Figure D2 shows the life cycle profile of consumption insurance coefficients to persistent income shocks calculated in the benchmark model (blue benchmark line), in the alternative specification (calculated PF line), and with the benchmark policy functions using as input the asset distribution from alternative specifications (benchmark PF line).

Figure D2: Insurance coefficients to persistent income shocks by age