

Consumption insurance and credit shocks*

Stefan Wöhrmüller

This draft: January 2024

Abstract

This paper studies how credit shocks affect the pass-through of idiosyncratic productivity shocks to consumption. Using a heterogeneous-agent incomplete-markets model I simulate two different credit shock dynamics as observed in credit panel data, a permanent and a mean-reverting one, and measure consumption insurance along the entire transition path. I show that consumption insurance against idiosyncratic productivity shocks drops on impact for both kind of credit shocks, while they imply qualitative different consumption insurance paths in the medium run. Importantly, I find that these paths differ by current wealth holdings. Asset-poor households experience the largest decrease in consumption insurance, whereas asset-rich households actually have access to more consumption insurance subsequent to a credit shock. Finally, endogenous labor supply attenuates these dynamics.

Keywords: consumption insurance, credit shocks, incomplete markets

JEL codes: D31, D52, E21, E44

*I am especially thankful to Christian Stoltenberg for his comments, continuous guidance, and patience during this project. I am very grateful to Andrew Hannon for discussing my work and to Elisabeth Proehl and Daniel Schmidt for reviewing earlier drafts in more detail. I would like to thank Jeanne Commault as well as participants at various conferences for comments I have received both during and after the presentation. All errors are my own. *Email:* s.h.p.woehrmueller@uva.nl

1 Introduction


When future earnings are uncertain, households can self-insure by borrowing and lending to smooth consumption over time. For instance, in response to a negative uninsurable earnings shock, a household might borrow funds to keep consumption steady. Thus, the availability of credit is an important determinant of households' welfare and consumption smoothing patterns. However, the recent great financial crisis (GFC) in the United States exhibited a sudden decline in the capacity to borrow. Unsecured borrowing capacity, determined by the aggregate credit card limit-to-GDP ratio, decreased by around 32% from 25% to 17% and is still - years later - at this lower level. Aggregate credit card limits in absolute terms, on the other hand, decreased by about 30% during the GFC, but made up almost three quarters of the drop by 2020 (see left panel in Figure 1). Moreover, the *use* of unsecured credit - determined by the credit card debt in terms of its absolute values and relative to GDP - shows similar dynamics (see right panel in Figure 1).

This paper takes the stylized facts from Figure 1 and answers the following question through the lens of a heterogeneous-agent incomplete-markets model: How do permanent and mean-reverting credit shocks affect households' consumption smoothing patterns? In particular, I am interested in the effects of different credit limit dynamics on the amount of consumption insurance with respect to both persistent and transitory idiosyncratic labor productivity shocks.

My results can be briefly summarized as follows. The model suggests that both type of credit shocks reduce the economy-wide level of consumption insurance for both type of idiosyncratic shocks in the short run, while displaying qualitative different paths in the medium run. Importantly however, the aggregate consumption insurance response masks considerable heterogeneity along the wealth distribution, as poorer households' self-insurance worsens and richer households' self-insurance improves subsequent to a credit shock. At last, I show that endogenous labor supply attenuates these dynamics.

In particular, the heterogeneous-agent incomplete-markets environment I consider is similar to [Huggett \(1993\)](#) and [Guerrieri and Lorenzoni \(2017\)](#). Risk-averse households produce consumption goods using labor, are heterogeneous in terms of their labor productivity and seek insurance against both persistent and transitory idiosyncratic shocks to this productivity level. They can only trade one period risk-free bonds up to an exogenous (ad-hoc) borrowing limit and pay lump-sum taxes to a government that supplies a fixed amount of risk-free bonds. The credit shock is then modelled by gradually reducing the borrowing limit. Due to the absence of closed-form solutions in this particular class of models, solutions are obtained using numerical simulations with empirical-based

Figure 1: Credit card limits and balances in the GFC



Note. The left panel of this figure shows the aggregate credit card limit and the credit limit-to-GDP ratio in the United States. The right panel shows the aggregate credit card debt and the credit card debt-to-GDP ratio in the United States. Credit card data is obtained from the credit card panel of the New York Fed.

parameterization.

In my quantitative exercise, I distinguish between a permanent and a mean-reverting credit shock that resemble the empirical evidence as in Figure 1. First, the permanent credit shock is modelled to match the drop in credit limit-to-GDP ratio. Second, the mean-reverting credit shock is modelled to match the drop and mean-reversion of aggregate credit card debt as observed in the data. Having calibrated the transitional dynamics, I then compute consumption insurance coefficients - in the spirit of [Kaplan and Violante \(2010\)](#) - for both persistent and transitory idiosyncratic shocks along the transition path. These coefficients measure to what extent (log) consumption can be isolated from the idiosyncratic shocks.

Turning to the results in more detail, I first find that consumption insurance against persistent idiosyncratic shocks decreases by 1.5% for both credit shocks on impact. Consumption insurance against transitory idiosyncratic shocks decreases 0.75% on impact for the permanent credit shock and 1.5% for the mean-reverting credit shock. These dynamics can be explained by the deleveraging behavior of agents subsequent the credit shocks: households have to decrease their debt levels to move away from the tighter credit limit. Thereby, they accumulate assets and forgo consumption (see also [Guerrieri and Lorenzoni \(2017\)](#)).

Second, the main difference between the two credit shock specifications manifests itself in the medium run, once the credit limit stays at its lower level - in the permanent case - or starts to mean-revert to its initial level - in the mean-reverting case. On the one hand,

the coefficients in the permanent credit shock specification converge to the terminal value from below, indicating a smoothing out of the deleveraging phase. On the other hand, the coefficients in the mean-reverting specification overshoot their pre-shock value and then converge back to their initial value, as the deleverage phase reverses and agents start to decumulate assets when the credit limit loosens again.

Third, the self-insurance capacity of households depends critically on the position in the wealth distribution. The decrease in the insurance coefficient is larger - both in absolute and relative terms - for agents in the bottom 10% of the wealth distribution compared to households in the bottom 25%. The former are closer to the credit limit and hence, the deleveraging impact is stronger. On the other hand, consumption insurance coefficients for households in the top 10% and 25%, for instance, actually *increase*. Given the fixed supply of bonds in the economy, these households decumulate bond holdings which generates additional resources as a buffer.

Finally, another margin that households can use in the model to isolate consumption against idiosyncratic shocks is to change their labor supply. To analyze the impact of this margin, I repeat my quantitative exercise under exogenous labor supply of households. I find that the qualitative dynamics are unchanged, however, the decrease in consumption insurance is stronger. For instance, the consumption insurance against persistent idiosyncratic shocks decreases by about 5% for both credit shocks. This implies that the possibility to adjust their labor supply is an important mechanism for households to insure their consumption in response to credit shocks.

Related literature My paper relates to several works in the literature regarding credit shocks¹ and consumption insurance. I bring these two areas together by explicitly modelling different credit limit dynamics.

The theoretical and numerical part of my paper is most closely related to the work of [Guerrieri and Lorenzoni \(2017\)](#) and [Kaplan and Violante \(2010\)](#). [Guerrieri and Lorenzoni \(2017\)](#) study the effects of a credit crunch on aggregate spending and output. For their analysis, they use a heterogeneous-agent incomplete-markets model in which agents are subject to idiosyncratic shocks going back to [Bewley \(1977\)](#), [Huggett \(1993\)](#), and [Aiyagari \(1994\)](#). They document that the credit crunch induces a deleveraging process which depresses interest rates and generates an output drop; moreover, the precautionary savings motive amplifies this mechanism. I adapt their baseline model but focus on the effects on consumption insurance in the economy. Thereby, my main contribution is the explicit

¹ Other exemplary papers than the ones mentioned in the main text studying credit shocks are [Cúrdia and Woodford \(2010\)](#), [Eggertsson and Krugman \(2012\)](#), and [Jones, Midrigan and Philippon \(2022\)](#)

distinction between different credit limit dynamics; that is, between a transitory and a permanent credit shock.²

In a quantitative study [Kaplan and Violante \(2010\)](#) assess how much insurance agents can obtain with self-insurance in a standard incomplete markets (SIM) life-cycle model and compare it with empirical estimates by [Blundell, Pistaferri and Preston \(2008\)](#).³ They conclude that consumption insurance implied by canonical consumption-savings models is too little compared to the data; a conclusion that is shared by other papers in the literature ([Krueger and Perri, 2006](#); [Broer, 2013](#)).⁴ I follow their approach and simulate an artificial panel in the model to construct consumption insurance coefficients as in [Blundell et al. \(2008\)](#).

The rest of the paper proceeds as follows. Section 2 outlines the model and defines the equilibrium in transition. Section 3 presents data, calibration targets and parameterizations as well as a brief comparison of the initial and terminal steady state. Section 4 contains the main analysis of the transitional dynamics under different regimes and specifications. Section 5 concludes.

2 Model

This section describes the benchmark economy in which risk-averse households face uninsurable idiosyncratic productivity shocks as in [Guerrieri and Lorenzoni \(2017\)](#). The key is that households can only partially self-insure against these idiosyncratic shocks using a single one-period asset, for which borrowing is constrained by an ad-hoc limit, and adjusting hours worked. The next section brings the model to the data and outlines how to determine the path of the borrowing limit during the transition.

Households Time is discrete. The economy is populated by a continuum of infinitely-lived households indexed i of measure unity. Households receive a utility flow U from consuming $c_{it} > 0$ and leisure, l_{it} . The time endowment of households, which can be allocated between leisure and labor, n_{it} , is normalized to 1. I assume that the utility function

² [López-Salido, Stein and Zakrajšek \(2017\)](#) and [Nakajima and Ríos-Rull \(2019\)](#) also document and model mean reversion in credit conditions.

³ Other studies have since modelled additional channels of insurance to overcome the gap between model and data, such as, advance information ([Stoltenberg and Singh, 2020](#)), family labor supply ([Blundell, Pistaferri and Saporta-Eksten, 2016](#)), moving back to ones parents ([Kaplan, 2012](#)), or social insurance programs ([Hubbard, Skinner and Zeldes, 1995](#)), among others.

⁴ [Wu and Krueger \(2021\)](#), on the other hand, find that a Bewley model with two-earner households and endogenous labor supply successfully recovers the degree of consumption insurance as estimated in the empirical counterpart by [Blundell et al. \(2016\)](#).

$U(c_{it}, n_{it})$ is separable and isoelastic in consumption and labor; $U : \mathbb{R}_+ \times \mathbb{R}_{[0,1]} \rightarrow \mathbb{R}$ is strictly increasing in consumption, strictly decreasing in labor, strictly concave and satisfies the Inada conditions. The future is discounted at rate β :

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t U(c_{it}, n_{it}) \right], \quad (1)$$

where the expectation is taken over realizations of idiosyncratic labor productivity shocks. For now, there is no aggregate risk.

Each household i produces the consumption good with the linear technology

$$y_{it} = \theta_{it} n_{it} \quad (2)$$

by choosing labor supply n_{it} .⁵ θ_{it} is an idiosyncratic shock to the labor productivity which follows a first-order Markov chain over a finite state space, $\{\theta^1, \dots, \theta^M\}$.

The only asset traded in this economy is a one-period risk-free bond, b . Hence, household utility maximization is restricted by the following budget constraint:

$$\frac{1}{1+r_t} b_{it+1} + c_{it} \leq b_{it} + y_{it} - \tau_t,$$

where b_{it} denote bond holdings, r_t is the interest rate, and τ_t is the tax burden at time t .

Borrowing is allowed up to an exogenous limit:

$$b_{it+1} \geq -\phi_{t+1}, \quad (3)$$

with $\phi_{t+1} > 0$. I will study the effects of an unexpected one-time shock that reduces this limit.

Recursive formulation of the household problem The following formulation as well as the definition of equilibrium below refer to the problem during transition since this constitutes the main exercise of the present paper. The equivalent descriptions in a stationary environment are relegated to Appendix D.

A household is characterized by the pair $(b_{it} = b, \theta_{it} = \theta)$ — the individual state. Given a sequence of interest rates $\{r_t\}_{t=0}^{\infty}$, a sequence of government policies $\{\tau_t\}_{t=0}^{\infty}$, and a sequence of borrowing limits $\{\phi_t\}_{t=0}^{\infty}$, each household chooses $c_t(b, \theta)$, $b_{t+1}(b, \theta)$, and

⁵ Hence, productivity shocks equal wage shocks and I will use the two terms interchangeably.

$n_t(b, \theta)$ to solve

$$\begin{aligned}
V_t(b, \theta) = \max & \left[U(c_t(b, \theta), n_t(b, \theta)) + \beta \sum_{\theta_{t+1} \in \Theta} V_{t+1}(b_{t+1}(b, \theta), \theta_{t+1}) \Gamma_{\theta, \theta_{t+1}} \right] \\
& \text{subject to} \\
& \frac{1}{1+r_t} b_{t+1}(b, \theta) + c_t(b, \theta) \leq b + y_t - \tau_t \\
& b_{t+1} \geq -\phi_{t+1}, \quad n_t(b, \theta) \geq 0 \\
& y_t = \theta n_t(b, \theta).
\end{aligned} \tag{4}$$

The policy functions $c_t(b, \theta)$ and $n_t(b, \theta)$ are sufficient to determine the transition of bond holdings, as future bond holdings, $b_{t+1}(b, \theta)$, can be derived from the budget constraint.

Denote by Φ_t the distribution of agents over states at time t . This distribution is the *aggregate* state variable. Note that during the transition period, the value function and policies are also a function of time, since the interest rate, the tax schedule and the borrowing limit are time varying: (r_t, τ_t, ϕ_t) . Furthermore, the dynamics induced by the credit shock are deterministic, that is, the entire transition path of the borrowing limit is known. As a result, given an initial distribution Φ_0 , it is known how (r_t, τ_t) and Φ_t evolve over time. Therefore, in Equation (4) it is not necessary to make prices and policy functions dependent on the distribution; the time subscript is sufficient.

Government The government budget constraint is

$$B_t = \frac{1}{1+r_t} B_{t+1} + \tau_t,$$

where B_t denotes the aggregate supply of government bonds. It is assumed, that the government chooses the tax schedule, τ_t , to keep a balanced budget, while keeping bond supply constant at B .

Equilibrium Before defining the equilibrium, it is necessary to define an appropriate measurable space on which the distribution of agents Φ_t is defined. Let $A \equiv [\underline{b}, \bar{b}]$ be the set of possible values for b_{it} with some lower bound \underline{b} and some upper bound \bar{b} ; it holds that $\underline{b} \leq -\phi_t \forall t$. Define the state space $S \equiv A \times \Theta$ and let the σ -algebra Σ_s be defined as $B_A \otimes P(\Theta)$, where B_A is the Borel σ -algebra on A and $P(\Theta)$ is the power set of Θ . Finally, let $\mathcal{S} = (\mathcal{A} \times \Theta)$ denote a typical subset of Σ_s .

The equilibrium in transition is summarized in the following definition.

Definition 1. Given an initial distribution Φ_0 , and a sequence of borrowing limits $\{\phi_t\}_{t=0}^\infty$, a recursive competitive equilibrium is a sequence of value functions $\{V_t\}_{t=0}^\infty$, policy functions for households $\{c_t(b, \theta), n_t(b, \theta)\}_{t=0}^\infty$, future bond holdings $\{b_{t+1}(b, \theta)\}_{t=0}^\infty$, interest rates $\{r_t\}_{t=0}^\infty$, government policies $\{\tau_t\}_{t=0}^\infty$, and distributions $\{\Phi_t\}_{t=0}^\infty$, such that, for all t :

i) Given r_t and $\tilde{\tau}_t$, the policy functions $c_t(b, \theta)$ and $n_t(b, \theta)$ solve the household's problem (4) and $V_t(b, \theta)$ is the associated value function

ii) The government budget constraint is satisfied

$$\tau_t = \frac{r_t B}{1 + r_t}$$

iii) The asset market clears

$$\int_{(A \times \Theta)} b_{t+1}(b, \theta) d\Phi_t = B$$

iv) For all $\mathcal{S} \in \Sigma_s$, the joint distribution measure Φ_{t+1} satisfies

$$\Phi_{t+1}(\mathcal{S}) = \int_{(A \times \Theta)} Q_t((b, \theta), \mathcal{S}) d\Phi_t,$$

where Q_t is the transition function defined as

$$Q_t((b, \theta), \mathcal{S}) = \mathbf{1}_{\{b_{t+1}(b, \theta) \in \mathcal{A}\}} \sum_{\theta_{t+1} \in \Theta} \Gamma_{\theta, \theta_{t+1}}.$$

The goods-market clearing condition is redundant by Walras law and thus omitted.

The optimization problem of the household gives rise to two optimality conditions that are used to construct the sequences of policy functions $c_t(b, \theta)$ and $n_t(b, \theta)$. In the following I skip the formulation of individual state variables for readability. First, an Euler equation governing optimal intertemporal substitution between consumption today and tomorrow

$$\beta(1 + r_t) \mathbb{E} [U_c(c_{t+1}, n_{t+1})] \leq U_c(c_t, n_t); \quad (5)$$

Equation (5) holds with equality if Equation (3) is not binding.

Second, an optimality condition for intratemporal substitution between labor and consumption

$$U_n(c_t, n_t) \leq -\theta U_c(c_t, n_t); \quad (6)$$

Equation (6) holds with equality if $n_t > 0$.

3 Quantitative exercise

3.1 Data and calibration

In the following I will first describe the main datasets used in calibration. Thereafter, I will speak about steady-state calibration, briefly introduce the consumption insurance measure, and finally, how to calibrate the transitional dynamics of the borrowing limit. Eventually, the calibrated model can inform us about the size and shape of consumption insurance in the economy as the borrowing limit changes.

3.1.1 Data

Productivity data I compute a productivity measure using data from the Panel Study of Income Dynamics (PSID). The PSID is a longitudinal survey of a representative sample of U.S. individuals and other members of their households. The survey is biennial since 1997 and features low attrition as well as high response rates ([Andreski, Li, Samancioglu and Schoeni, 2014](#)). I follow [Flodén and Lindé \(2001\)](#) and define productivity as an "agent's hourly wage rate relative to all other agents" (p.416). In my case, agent refers to either "head", wife/"wife", or both as indicated in the PSID. An agent's hourly wage is defined as total yearly labor income divided by annual hours worked. Moreover, to be in line with the model I use post-tax wage rates as there are no distortionary taxes present that would affect the labor supply decision of agents. I combine post-tax wage rates within a household to compute productivity levels per household. A more detailed description of the data and the variables can be found in [Appendix A](#).

Credit data I employ two different credit data sources. Firstly, the credit limit data is from the 2020 Q2 Household Debt and Credit Report (HDCR) via the FRBNY Consumer Credit Panel (CCP).⁶ The CCP is a longitudinal database which comprises information on consumer debt and credit via Equifax credit reports. The sampling procedure makes use of the randomness in the last 4 digits of the social security numbers to create a nationally representative random sample of individuals who have a credit report. The information

⁶ The microdata can only be accessed by researchers of the Federal Reserve System, hence, I rely on aggregate data which is published with recent reports; in my case this report is the HCDR. See questions on Data Requests [here](#)

from credit reports on individual accounts is then further limited to accounts that have been updated within the last three months.

In particular, the credit limit in the HDCR refers to the limit on *bankcard accounts* (or credit card accounts) which are revolving accounts for banks, bankcard companies, national credit card companies, credit unions and savings & loan associations (Lee and van der Klaauw, 2010). Credit limits via home equity revolving accounts are considered separately and are not captured here.

The second data source is the consumer credit time series from the Federal Reserve Board Flow of Funds (FoF). Particularly, I am interested in the *evolution* of credit card debt as a fraction of GDP. Note that this is the same dataset which is used to calibrate the net supply of bonds, B ; that is, the datapoints in 2006 are used as a benchmark to calibrate the initial steady state, whereas I make use of the whole time series of consumer credit to calibrate my transitional dynamics.

Lastly, note that even though I use two different data sets in my calibration exercise below, consumer credit is one of the components in the FoF which can be directly compared to debt measures in the HDCR (CCP). The report states that the numbers obtained in the third quarter of 2009 for consumer credit from both datasets are very similar in magnitude. (\$2.6 trillion in the CCP vs \$2.5 trillion in the FoF; see HCDR).⁷

3.1.2 Steady-state calibration

The calibration strategy for the initial steady state aims to capture economic conditions, especially at the household level, prior to the GFC in 2006 and before. Table 1 summarizes the calibration.

Model functional forms and parameters

Preferences I assume preferences are additively separable and isoelastic in consumption and labor:

$$U(c_{it}, n_{it}) = \frac{c_{it}^{1-\gamma}}{1-\gamma} + \psi \frac{(1-n_{it})^{1-\eta}}{1-\eta},$$

where I already normalized time endowment for leisure and labor to 1, $l + n = 1$. The time period is a quarter.

The discount factor β is set to match an annual interest rate of 2.5% in the initial steady state. The elasticity of intertemporal substitution (EIS), $\frac{1}{\gamma}$, is set to a quarter. The param-

⁷ A small part of the remaining difference, for instance, can be explained by debt holdings from individuals without a social security number; these individuals are captured only in the FoF.

eters governing the curvature of utility from leisure and the utility weight of leisure, η and ψ , are set to match an average Frisch elasticity of 1 and average hours worked of 0.4, respectively.⁸ The latter is supported by evidence from [Nekarda and Ramey \(2010\)](#)⁹.

Labor productivity process An agent's stochastic labor productivity depends on two components and is given by

$$\log \theta_{it} = \kappa_{it} + \epsilon_{it} \quad (7)$$

$$\kappa_{it} = \rho_q \kappa_{it-1} + \zeta_{it} \quad (8)$$

The first component, κ , is an AR(1) process with persistence ρ_q and innovation variance $\sigma_{q,\zeta}^2$. The second component, ϵ , is a transitory i.i.d. shock variance $\sigma_{q,\epsilon}^2$. Both innovations are assumed to be normal with zero mean. Theoretical variance and covariance moments of the stochastic wage process in Equation (7) are matched to empirical moments constructed using data from the 2002-2006 PSID waves. A detailed description of the estimation procedure with annual data and subsequent conversion to the quarterly frequency can be found in Appendix A. The continuous wage process is then transformed onto a discretized grid using [Rouwenhorst \(1995\)](#)'s method. The mean of θ under the stationary distribution is normalized to 1.

Asset supply Parameters regarding the asset supply are jointly matched to reconstruct pre-crisis characteristics of households' balance sheets. Using data from the 2006 Federal Reserve Flow of Funds, I target a liquid asset to GDP ratio of 1.78. Particularly, liquid asset include Deposits, Treasury Securities, Agency- and GSE- backed securities, Municipal Securities, Corporate Foreign Bonds, Mutual Fund Shares, and Security Credit¹⁰. The net supply of Bonds, B , is defined to be equal to liquid assets ratio minus credit card debt holdings of households, all expressed in terms of annual output. Thus, given a credit card debt-to-GDP ratio, B is then adjusted to match the liquid asset to GDP ratio. The borrowing capacities in the model, hence, mirror unsecured revolving credit. In addition, ϕ is used to match the credit card debt-to-GDP ratio prior to the GFC.

⁸ A Frisch elasticity is motivated by the effect that the unit of observation is the household. Hence, the labor supply decision implicitly takes into account higher labor supply elasticities of females as well as labor supply decisions on the extensive margin. See also [Conesa, Kitao and Krueger \(2009\)](#) for an alike motivation.

⁹ This reference refers to an older version of their working paper, meanwhile titled "The Cyclical Behavior of the Price-Cost Markup". For the evidence considered here, see Figure 10 in the 2010 version

¹⁰ Data is taken from Table B.100 which can be found [here](#)

Table 1: Calibrated parameters in the steady state

<i>Parameter</i>		<i>Value</i>	<i>Target / Source</i>
Discount factor	β	0.980	Annual interest rate of 2.5%
Curvature of utility from leisure	η	1.50	Avg. frisch elasticity = 1
Coefficient on leisure in utility	ψ	22.36	Avg. hours of endowment worked = 0.4
Elast. of intertemp. substitution	$1/\gamma$	0.25	Guerrieri and Lorenzoni (2017)
Persistence productivity process	ρ_q	0.976	PSID
Var. of innovation to persistent shock	$\sigma_{q,\zeta}^2$	0.012	PSID
Var. of transitory productivity shocks	$\sigma_{q,\epsilon}^2$	0.049	PSID
Net bond supply (to annual GDP)	\mathbf{B}/Y_A	1.6	Flow of Funds liquid assets
Borrowing limit (to annual GDP)	ϕ/Y_A	0.372	Credit card debt-to-GDP

Note. This table summarizes the calibrated and baseline parameters of the benchmark model. Variables and parameters in bold indicate that these have been calibrated to match simulated moments. See main text for details.

3.2 Consumption insurance measure

To evaluate the extent of risk-sharing and consumption insurance I compute insurance coefficients as introduced by [Blundell et al. \(2008\)](#) and further examined by [Kaplan and Violante \(2010\)](#). To set the scene, let us quickly revise the notion of insurance in this framework. The insurance coefficient for a shock x_{it} to logged income or productivity is defined as

$$\varphi^x = 1 - \frac{\text{cov}(\Delta c_{it}, x_{it})}{\text{var}(x_{it})},$$

where Δc_{it} denotes the first difference in logged consumption, and the second moments are taken over the entire cross-section of households. The intuition is best captured by [Kaplan and Violante \(2010\)](#): the insurance coefficient is "the share of the variance of the x shock that does *not* translate into consumption growth" (p.57).

To compute this measure from the model perspective is straightforward, for the shock realizations are observed when simulating an artificial panel. The insurance coefficients from the model then read

$$\varphi^\zeta \equiv 1 - \frac{\text{cov}(\Delta c_{it}, \zeta_{it})}{\text{var}(\zeta_{it})} \quad \text{and} \quad \varphi^\epsilon \equiv 1 - \frac{\text{cov}(\Delta c_{it}, \epsilon_{it})}{\text{var}(\epsilon_{it})}. \quad (9)$$

φ^ζ is computed from simulated data using the quasi-difference operator $\tilde{\Delta}$, with $\tilde{\Delta}\kappa_t \equiv \kappa_t - \rho_q \kappa_{t-1}$.

Table 2 shows the consumption insurance coefficients from the model in both the initial and terminal steady state as well as under endogenous and exogenous labor supply of agents. The terminal steady state corresponds to an economy with a tighter credit limit, as is explained in more detail below. We see that the terminal steady state features lower insurance coefficients in all four cases. The changes, albeit small in relative magnitude, are bigger for i) the persistence productivity shock and ii) under exogenous labor supply. Moreover, comparing exogenous and endogenous labor supply in both steady states reveals that the labor supply decision has a stronger effect on insuring persistent productivity shocks. These observations connect to the existing literature to the effect that first, persistent shocks are harder to insure ([Blundell et al., 2008](#); [Kaplan and Violante, 2010](#)) and second, labor supply is used as a consumption smoothing device ([Low, 2005](#); [Pijoan-Mas, 2006](#); [Wu and Krueger, 2021](#)).

Table 2: Consumption insurance coefficients in the steady states of the model

	Endogenous labor supply		Exogenous labor supply	
	φ^{ζ}	φ^{ϵ}	φ^{ζ}	φ^{ϵ}
Initial steady state	0.6987	0.9620	0.4423	0.9554
Terminal steady state	0.6962	0.9609	0.4381	0.9531

Note. This table shows the steady state consumption insurance coefficients for the persistent productivity shock, φ^{ζ} , and the transitory productivity shock, φ^{ϵ} , as defined in Equation (9). The left two columns depict the benchmark case where agents choose their labor supply endogenously. In the right two columns, the labor supply is exogenous and set to 1. The terminal steady state refers to the economy with a reduced credit limit.

3.3 Credit shocks

The main quantitative exercise of this paper moves beyond steady state outcomes and studies the short-term/transitional effects of limited credit availability on consumption smoothing behavior of households. The dynamics of the credit limit should reflect the two different kind of credit dynamics in Figure 1. Hence, I employ two calibrations of credit shocks that reflect these dynamics: i) a Mean-reverting credit shock and ii) a permanent credit shock.

In both calibrations, the shock to the credit limit is not expected by the agents and hits the economy in the initial steady state ("MIT shock"). To abstract from households defaulting on their one-period bonds, I follow [Guerrieri and Lorenzoni \(2017\)](#) and assume that the credit shock lasts for six periods and follows a linear path.¹¹ In this way, households are able to incrementally adjust their portfolios without paying back unrealistically large amounts of debt within a single period. The complete path of the credit limit towards the terminal steady state is deterministic and fully known to households.

In the following, I will explain how I construct the path of this credit limit in both calibrations in the model and which data targets I match. Table 3 summarizes the calibrated parameters and targeted moments during transition.

i) Mean-reverting credit shock The mean-reverting credit shock captures the transitory notion of the aggregate credit card statistics as observed in Figure 1. Both aggregate credit card debt and credit card limits drop during the GFC and then revert. Hence, in this calibration, I match the initial drop in the credit card debt-to-GDP ratio as well as the mean-

¹¹ Note, however, that default or delayed payment can be a channel of insurance by itself. [Gelman, Kariv, Shapiro, Silverman and Tadelis \(2020\)](#) provide evidence that during the 2013 US government shutdown, many affected employees delayed recurring payments, for instance, for credit cards to smooth consumption. In a quantitative paper, [Hannon \(2022\)](#) shows how the option to enter mortgage delinquency can cushion the consumption drop after a housing crisis.

reversion back to its pre-GFC level. The top panel of Table 3 shows the corresponding parameter values and targets for this calibration.

In particular, the shock in the first 6 periods follows the linear specification

$$\phi_{t+1} = \bar{\phi} - \frac{(1-\chi)\bar{\phi}}{6} \cdot t \quad \text{for } t \leq 6,$$

where χ governs the size of the shock. Subsequent the shock I specify the evolution of the credit limit as

$$\phi_{t+1} = \varrho \bar{\phi} + (1-\varrho)\phi_t \quad \text{for } t > 6,$$

where ϱ governs the speed of mean-reversion of the borrowing limit and $\bar{\phi}$ denotes the long-run average. The parameters ϱ and χ are then *jointly* determined in order to match the two empirical targets.

The first empirical target is the drop in the credit card debt-to-GDP ratio. I target a credit-card debt-to-GDP ratio of 3.8%, the lowest point of this time series after 2006, the year that represented my initial steady state.

The second empirical target concerns the speed of the mean-reversion of aggregate credit card balances. Since I look at transitional dynamics starting from a stationary recursive equilibrium, I extract the stationary component in the time series and then target the autocorrelation in this component. In particular, I decompose the (inflation-adjusted) credit card balances into a trend and cyclical component using the Hodrick-Prescott filter. Thereafter, I estimate the autocorrelation of the cyclical component using a simple AR(1). The mean-reverting path of the credit limit in the model after the shock will be specified in such a way, that the autocorrelation of the model-implied aggregate debt balance along the transition will match precisely this AR(1) estimate from the data. Appendix B describes the filter in more detail and provides the AR(1) estimates.

ii) Permanent credit shock The permanent credit shock, on the other hand, captures the downward shift in credit card variables once they are scaled with GDP as observed in Figure 1. This exercise is similar to the one conducted in the benchmark case of [Guerrieri and Lorenzoni \(2017\)](#), as the shock boils down to a permanent drop in the credit limit.¹² The difference, however, is the fact that I can observe this limit from the data.

In this calibration, I match the percentage change in the credit limit-to-GDP ratio. I choose a borrowing limit in the terminal steady state, $\phi^{terminal}$, which implies a drop of

¹² [Guerrieri and Lorenzoni \(2017\)](#) also simulate a transitory credit shock in an economy with durable goods and a credit spread. This shock is modelled as a transitory increase in the credit spread, which affects agents over the whole wealth distribution and not mainly the ones near the credit limit.

30 percent, from 25% to 17.5%, as observed in Figure 1. The path of the credit limit is then as follows: In the first six periods, the borrowing limit decreases linearly to $\phi^{terminal}$ and stays there permanently. The bottom panel of Table 3 summarizes the corresponding parameter value and target for this calibration. Moreover, the red line in the left panel of Figure 2 shows the path of the borrowing limit (scaled by GDP).

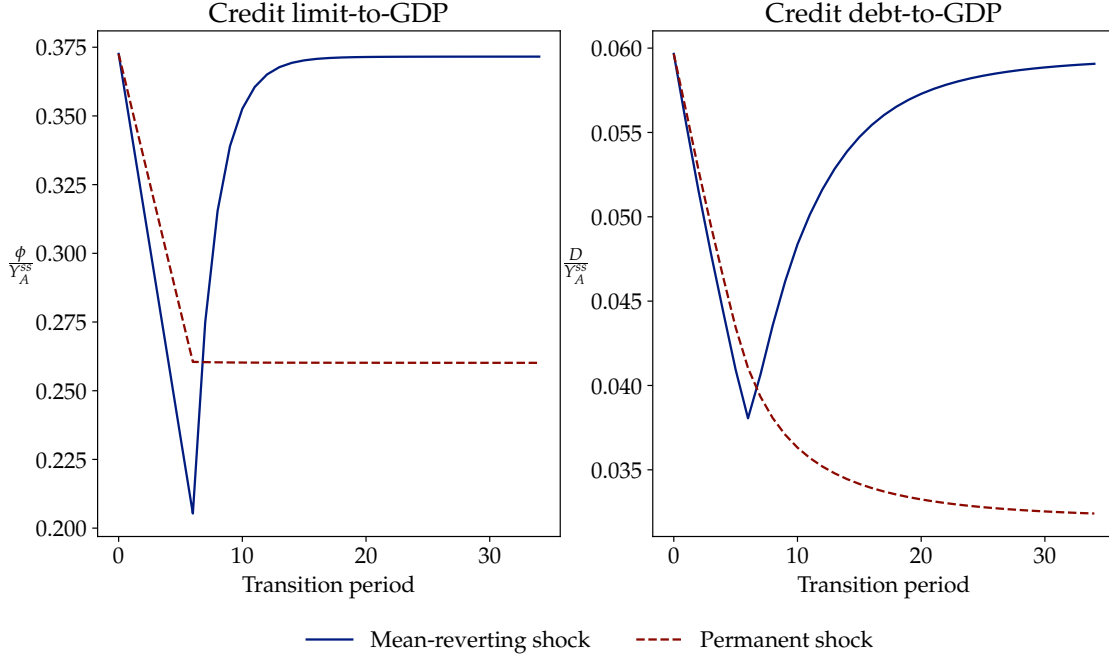
Table 3: Calibrated parameters and moments to match moments along the transition

i) Mean-reverting credit shock				
Parameter		Value	Target	
Credit-shock size	$1 - \chi$	0.372	Credit card debt-to-GDP low of 3.8%	
Mean-reversion	$1 - \varrho$	0.826	Autocorrelation in credit card debt of 0.841	
ii) Permanent credit shock				
Parameter		Value	Target	
Borrowing limit (to annual GDP)	$\phi^{terminal} / Y_A^{terminal}$	0.841	30% drop in the credit limit-to-GDP ratio	

Note. This table summarizes the calibrated parameters to match the transition of the credit shocks in the model to empirical moments. See main text for details.

Credit card debt and limit dynamics Overall, the mean-reversion parameter is similar in size to the matched autocorrelation of the credit debt-to-GDP ratio. The intuition for this results resembles the precautionary behavior subsequent credit shocks as documented, for instance, by Guerrieri and Lorenzoni (2017). When the shock hits the economy in the initial steady state, agents accumulate wealth to stay away from the borrowing limit. As the path for the credit limit is fully known to them, they can align deleveraging and leveraging phases with the dynamics of the limit. This perfect foresight behavior can also be seen from Figure 2 which shows the dynamics of the credit limit-to-GDP and the credit card debt-to-GDP ratios, respectively, under the different shock specifications. Even though the size of the transitory shock is larger than the permanent shock, in the former case the credit debt-to-GDP ratio is smaller after 6 periods as agents realize that the credit limit will be loosened from then on. In the latter case, however, agents keep deleveraging to stay away from the credit limit (Guerrieri and Lorenzoni, 2017).

Figure 2: Permanent and mean-reverting credit shocks



Note. The left panel of this figure shows the path of the credit limit under the transitory and permanent credit shock, respectively. The right panel shows the model implied credit debt-to-GDP ratio under both shock scenarios. Both panels show the first 35 periods of the transition.

4 Quantitative results


With a calibrated model in hand it is now possible to answer the main question of the paper: How do permanent and transitory credit shocks affect households' consumption smoothing patterns? To do so, I implement the transition as outlined above and simulate an artificial panel of households starting from the steady state over the entire transition. With their observed consumption choices and simulated shocks I compute the consumption insurance coefficients as in Equation (9).¹³

4.1 Transitional dynamics

Figure 3 shows the transition dynamics of the insurance coefficients for both type of credit shocks and productivity shocks, respectively. The left panel shows the dynamics after the permanent credit shock. We see that the consumption insurance coefficient drops

¹³ When simulating I take into account two issues. First, the households in the artificial panel only "enter" the transition once the simulation has converged and I draw from the stationary distribution in the initial steady state. Second, I always choose the same sequence of shocks in all my numerical exercises using a fixed seed for the random number generator.

Figure 3: Insurance coefficients over time



Note. The values depict percentage deviations of the initial steady state and are filtered using a 1-D gaussian filter with standard deviation (smoothing parameter) of 1.5.

moderately on impact for both productivity shocks. Consumption insurance for the persistent shock falls by 1.5% and 0.75% for the transitory shock. Moreover, the persistent coefficient immediately increases again after the initial shock, whereas the transitory coefficient drops further to a 1% reduction in period six, when the credit limit reached its trough, and only increases thereafter.

The right panel shows the dynamics after the mean-reverting credit shock. The dynamics in the first six periods are similar to the permanent shock. The persistent coefficient also drops to about 1.5% on impact and increases immediately thereafter. The transitory coefficient drops to 1.5% and also decreases further within the next 5 periods until 1.8%. Note, however, that one should keep in mind that the drop in the credit limit is larger for the mean-reverting shock (Figure 2).

The main difference between the two credit shock specifications manifests itself in the medium run; that is, period 7 to 25. While the coefficients in the permanent credit shock specification converge to the terminal steady state from below, the coefficients in the mean-reverting specification overshoot their pre-shock value and then converge back to the initial steady state (i.e. in this exercise also the terminal one) from above. What is the reason for this? The key to this observation is the asset accumulation and decumulation behavior of agents subsequent the shocks. Under the permanent credit shock, households

with debt have to start deleverage their portfolios, thereby reducing their nondurable consumption irrespective of their idiosyncratic shocks. With the credit limit at its new and tighter level, household slowly increase their precautionary savings (i.e. accumulate assets), and accept lower consumption insurance in the long run.

Under the mean-reverting credit shock, households realize that the economy will return to the initial steady state. Hence, after six periods when the shock dissipates and the credit limit moves back to its initial level, agents start decumulating assets and increase consumption (insurance). Given that the supply of assets is fixed, the interest rate, and thus the (precautionary) savings behavior of households, is a mirror image of the consumption insurance coefficients - decreasing under high accumulation of assets and vice versa. The dynamics of the interest rate and other aggregates can be seen in [Appendix C](#).

4.2 Exploring the mechanism


In what follows I will further explore two determinants of consumption smoothing behavior. First, the position in the wealth distribution and second, the possibility to adjust labor supply.

Heterogenous dynamics along the wealth distribution The aggregate responses of the consumption insurance coefficients masks heterogeneous responses along the wealth distribution. [Figure 4](#) illustrates this for the persistent productivity shock under a permanent credit limit reduction. The left panel shows the dynamics of the consumption insurance coefficients - in absolute terms - for the bottom 10% and 25% of the initial steady state wealth distribution. The decrease in the insurance coefficient is larger in both absolute and relative terms for the bottom 10%.

The right panel shows the dynamics of the consumption insurance coefficients for the top 10% and 25%. We see that the coefficients *increase*; again, to a larger extent for the top 10%. Why is this the case? Recall that the total amount of bonds available in the economy is unchanged during transition. Hence, the deleveraging episode of net borrowers induces net lenders at the top of the wealth distribution to sell some of their bond holdings. The additional resources facilitate consumption insurance and the value increases.

Comparing the bottom and top 10% (or 25%), however, the response of the former dominates in magnitude, giving rise to a reduction of consumption insurance in the aggregate. This is a common finding in the literature, that agents close to constraints (or close to kinks in their budget constraints) drive aggregate behavior after shocks ([Kaplan and Violante \(2014\)](#); [Guerrieri and Lorenzoni \(2017\)](#)).

Figure 4: Insurance coefficients of persistent idiosyncratic shocks over time by wealth groups



Note. This figure depicts the consumption insurance coefficients for the persistent idiosyncratic shock for different wealth groups. The wealth groups are determined from the initial steady state distribution.

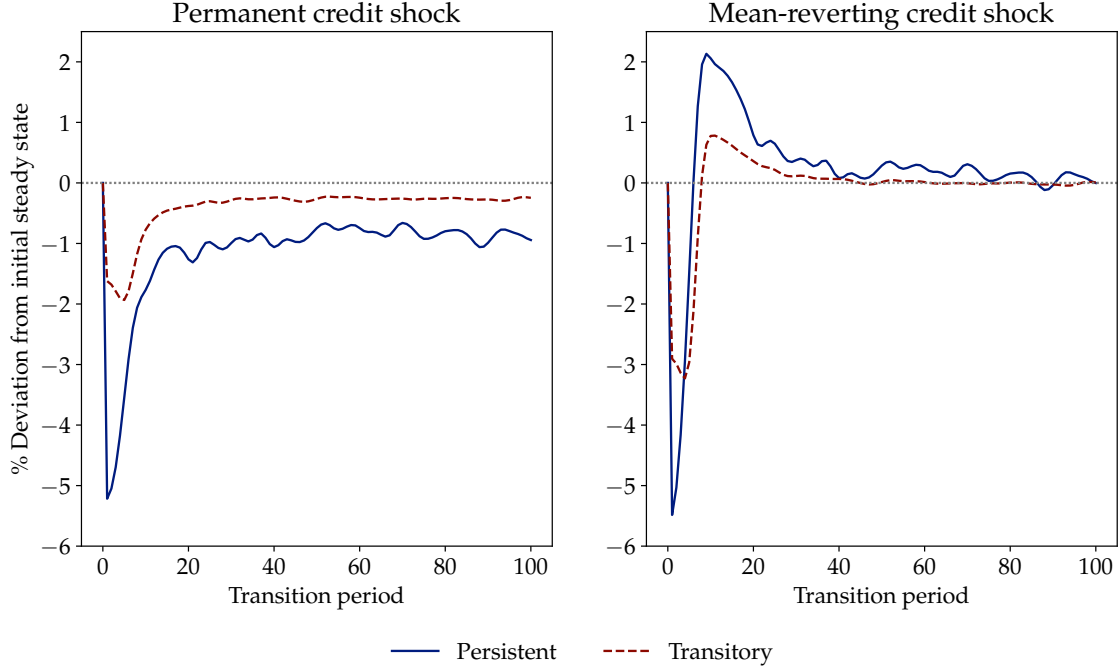
The quantitative relevance of endogenous labor supply In the benchmark economy households can also choose to adjust their labor supply in response to idiosyncratic shocks. How much does this margin contribute to the dynamics in the consumption insurance coefficients? To answer this question, I set the labor supply exogenously to 1 when conducting my quantitative experiment and recalibrate parameters to the same moments as in the benchmark case; both in the initial steady state and over the transition.¹⁴

Figure 5 shows the dynamics subsequent the credit shock when the labor supply of households is fixed. We see that removing the intensive margin of the labor supply decision "scales" the impact on the insurance coefficients. That is, the dynamics are qualitatively similar to the benchmark case but lower in magnitude. For instance, under the permanent credit shock, the consumption insurance with respect to persistent productivity shocks drops by 5% - 3.5 percentage points more than under the benchmark. Under the mean-reverting credit shock, the initial drop and the following overshooting of the persistent insurance coefficient are - in absolute terms - 4 and 1.5 percentage points higher, respectively.

Hence, these observations suggest that the possibility to adjust their labor supply is

¹⁴ While it is possible to match the empirical moments in the initial steady state and the permanent credit shock precisely, the model has difficulties to match the autocorrelation in the mean-reverting case.

Figure 5: Insurance coefficients over time with inelastic labor supply



Note. The values depict percentage deviations of the initial steady state under fixed labor supply and are filtered using a 1-D gaussian filter with standard deviation (smoothing parameter) of 1.5.

an important channel for households to insure their consumption in response to credit shocks. Note, however, that this is particularly relevant for households close to the borrowing limit. As [Guerrieri and Lorenzoni \(2017\)](#) note, the labor supply response is heterogeneous. Households close to the borrowing limit increase their labor supply, while households further away from it decrease it. Since wealth levels are correlated with productivity, overall labor supply increases whereas output decreases.

5 Conclusion

In this paper, I have studied the evolution of consumption insurance patterns in the economy after unexpected credit shocks. I focused on two different shock specifications as observed in the data: a mean-reverting credit shock following aggregate credit debt levels, and a permanent credit shock following credit debt-to-GDP ratios in the data. I have shown that the evolution of consumption insurance after both kind of shocks is similar in the short run, but differs in the medium run, as accumulation/deaccumulation episodes of assets differ. Moreover, I have shown that the quantitative effects are attenuated by endogenous labor supply, as households make use of the intensive margin to smooth con-

sumption. Lastly, I have shown that the dynamics differ considerably by present wealth holdings, as households at the lower end of the wealth distribution experiencing the brunt of the drop and households at the upper end actually increase their consumption insurance in response to credit shocks.

I want to mention at least two potential areas for future research related to this topic. First, this paper remains largely silent regarding other channels for insurance that are present in the data. While the sample selection for the empirical exercise tried to account for this, especially public insurance via taxes and transfers were still implicit in the model. A model that explicitly models the tax structure in the United States as, for instance, [Heathcote, Storesletten and Violante \(2017\)](#), would allow for a normative analysis whether a more progressive taxation would have been welfare improving; especially for agents at the lower end of the wealth distribution.

Second, my calibration only considers unsecured credit card debt. A richer model with secured debt, such as mortgages, could investigate the implication of decreasing collateral value for consumption insurance. For instance, modelling secured credit based on a second illiquid asset (housing) as in [Kaplan and Violante \(2014\)](#) or [Kaplan, Mitman and Violante \(2020\)](#) would endogenously give rise to a larger share of hand-to-mouth consumers consistent with the data ([Kaplan, Violante and Weidner, 2014](#)), potentially amplifying the aggregate responses. Moreover, in this model one would have to distinguish further between the type of credit that is affected by the shock.

References

- Aiyagari, S. Rao.** 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*, 109(3): 659–684.
- Altonji, Joseph G., and Lewis M. Segal.** 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *Journal of Business & Economic Statistics*, 14(3): 353–366.
- Andreski, Patricia, Geng Li, Mehmet Zahid Samancioglu, and Robert Schoeni.** 2014. "Estimates of Annual Consumption Expenditures and Its Major Components in the PSID in Comparison to the CE." *American Economic Review*, 104(5): 132–135.
- Auclert, Adrien, and Matthew Rognlie.** 2020. "Inequality and Aggregate Demand." 82.
- Bewley, Truman F.** 1977. "The Permanent Income Hypothesis: A Theoretical Formulation." *Journal of Economic Theory*, 16(2): 252–292.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston.** 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.

- Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten.** 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review*, 106(2): 387–435.
- Broer, Tobias.** 2013. "The Wrong Shape of Insurance? What Cross-Sectional Distributions Tell Us about Models of Consumption Smoothing." *American Economic Journal: Macroeconomics*, 5(4): 107–140.
- Carroll, Christopher D.** 2006. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics Letters*, 91(3): 312–320.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger.** 2009. "Taxing Capital? Not a Bad Idea After All!" *American Economic Review*, 99(1): 25–48.
- Cúrdia, Vasco, and Michael Woodford.** 2010. "Credit Spreads and Monetary Policy." *Journal of Money, Credit and Banking*, 42(s1): 3–35.
- Eggertsson, Gauti B., and Paul Krugman.** 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach*." *Quarterly Journal of Economics*, 127(3): 1469–1513.
- Flodén, Martin, and Jesper Lindé.** 2001. "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics*, 4(2): 406–437.
- French, Eric.** 2004. "The Labor Supply Response to (Mismeasured but) Predictable Wage Changes." *The Review of Economics and Statistics*, 86(2): 602–613.
- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis.** 2020. "How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown." *Journal of Public Economics*, 189: 103917.
- Guerrieri, Veronica, and Guido Lorenzoni.** 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *Quarterly Journal of Economics*, 132(3): 1427–1467.
- Hannon, Andrew.** 2022. "Falling Behind: Delinquency and Foreclosure in a Housing Crisis." *Manuscript*.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante.** 2010. "The Macroeconomic Implications of Rising Wage Inequality in the United States." *Journal of Political Economy*, 118(4): 681–722.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante.** 2017. "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics*, 132(4): 1693–1754.
- Hodrick, Robert J., and Edward C. Prescott.** 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, 29(1): 1–16.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes.** 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103(2): 360–399.

- Huggett, Mark.** 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control*, 17(5-6): 953–969.
- Jones, Callum, Virgiliu Midrigan, and Thomas Philippon.** 2022. "Household Leverage and the Recession." *Econometrica*, 90(5): 2471–2505.
- Kaplan, Greg.** 2012. "Moving Back Home: Insurance against Labor Market Risk." *Journal of Political Economy*, 120(3): 446–512.
- Kaplan, Greg, and Giovanni L Violante.** 2010. "How Much Consumption Insurance Beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2(4): 53–87.
- Kaplan, Greg, and Giovanni L. Violante.** 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica*, 82(4): 1199–1239.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner.** 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*, , (Spring): 77–138.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante.** 2020. "The Housing Boom and Bust: Model Meets Evidence." *Journal of Political Economy*, 128(9): 3285–3345.
- Krueger, Dirk, and Fabrizio Perri.** 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies*, 73(1): 163–169.
- Lee, Donghoon, and Wilbert van der Klaauw.** 2010. "An Introduction to the FRBNY Consumer Credit Panel." *Federal Reserve Bank of New York Staff Reports*, 479.
- Low, Hamish W.** 2005. "Self-insurance in a life-cycle model of labour supply and savings." *Review of Economic Dynamics*, 8(4): 945–975.
- López-Salido, David, Jeremy C. Stein, and Egon Zakrajšek.** 2017. "Credit-Market Sentiment and the Business Cycle." *The Quarterly Journal of Economics*, 132(3): 1373–1426.
- Nakajima, Makoto, and Jose-Victor Rios-Rull.** 2019. "Credit, Bankruptcy, and Aggregate Fluctuations." 61.
- Nekarda, Christopher J, and Valerie A. Ramey.** 2010. "The Cyclical Behavior of the Price-Cost Markup." *Mimeo*.
- Pijoan-Mas, Josep.** 2006. "Precautionary savings or working longer hours?" *Review of Economic Dynamics*, 9(2): 326–352.
- Rouwenhorst, Geert K., ed.** 1995. *Asset Pricing Implications of Equilibrium Business Cycle Models. Frontiers of Business Cycle Research*. Thomas F. Cooley, ed. Princeton ed., New Jersey: Princeton University Press.
- Stoltenberg, Christian A, and Swapnil Singh.** 2020. "Consumption insurance with advance information." *Quantitative Economics*, 11(2): 671–711.

Straub, Ludwig. 2019. "Consumption, Savings, and the Distribution of Permanent Income." *Manuscript*, 85.

Wu, Chunzan, and Dirk Krueger. 2021. "Consumption Insurance against Wage Risk: Family Labor Supply and Optimal Progressive Income Taxation." *American Economic Journal: Macroeconomics*, 13(1): 79–113.

Appendix

A Estimation of the productivity process

The full process to be estimated separately for head and spouse is

$$\begin{aligned}\log w_{m,t} &= \psi_m + \kappa_{m,t} + \epsilon_{m,t} + \nu_{m,t}, \\ \kappa_{m,t} &= \rho \kappa_{m,t-1} + \zeta_{m,t},\end{aligned}\tag{A.10}$$

where $w_{m,t}$ denotes the hourly wage of family member $m \in \{head, spouse\}$ at year t relative to the average hourly wage of all individuals at year t . This observed relative wage rate is the sum of an individual permanent component ψ_m , a persistent component κ , which is an AR(1) process with persistence ρ and innovation variance $\tilde{\sigma}_\zeta^2$, a temporary component ϵ with variance $\tilde{\sigma}_\epsilon^2$, and potential measurement error ζ with variance $\tilde{\sigma}_\nu^2$.¹⁵ To estimate the counterparts for the entire household, and to capture the productivity risk of a single consumption unit, I assume that both head and spouse have the same ρ and that their idiosyncratic shocks are independent. I denote these variables without a tilde. Moreover, since the measurement error cannot be separately identified from the transitory shock, I follow the literature and impose an external estimate for σ_ν^2 (Heathcote, Storesletten and Violante (2010); Straub (2019)). In particular, I set $\sigma_\nu^2 = 0.02$ as estimated by French (2004).

I now describe the strategy to estimate the relevant parameters ($\rho, \sigma_\zeta^2, \sigma_\epsilon^2$) from family wage residuals. The exposition largely follows Flodén and Lindé (2001) from whom I take the measure for productivity.

Data I use the Panel Study of Income Dynamics (PSID) data from 2003 to 2007 to estimate the parameters of interest. I choose this time horizon to be consistent with my overall calibration strategy to target the pre-crisis conditions. The PSID is a longitudinal survey of a representative sample of U.S. individuals and other members of their households. Since the PSID changed to a biennial frequency in 1997, I am effectively using the three survey waves from 2003, 2005, and 2006. In general, the questions on labor income and hours are retrospective. That is, survey questions from 2005 refer to the year 2004 etc. In addition to these standard questions, however, respondents are also asked to provide information about income and labor supply two years before. Collecting this information makes it possible to construct an *annual* panel of labor income and hours worked.

¹⁵ Variables with a tilde denote family member variables.

Until 2015, the PSID referred to the husband in a married couple as the Head of the household, irrespective of employment status and labor income. To capture the full labor productivity of a household, I include both Head and Wife/"Wife"¹⁶ (if present) when estimating the stochastic process. This is a better measure to use as it better describes the increase in resources a household has if either of its members increases their labor supply.

Sample selection I only look at households whose head or spouse follow the following criteria. The individual (i) is from the main Survey Research Center (SRC) sample, (ii) is between 25 and 59 years old, (iii) provides information on years of education (iv) has positive working hours with a maximum of 5840 (maximum possible value in the survey), (v) has an hourly wage rate more than half the minimum wage (in 2002 dollars), (vi) does not have unreasonable income swings between two years, and (vii) is observed in every year. Note that my sample selection applies these criteria individually to both head and spouse within a family. That is, if some of these criteria do not apply for the head, but do apply for the spouse, the latter is excluded in the sample, while the former is included. Starting from the individual heads and spouses in the SRC sample, Table A.4 shows how the selection process affects the total number of observations.

Table A.4: Sample selection in PSID

	# dropped	# remain
Initial sample (head and spouse)	...	31,161
Age between 25 and 59	4,924	26,237
No Education information	2,505	23,732
Hours worked ≤ 0 or > 5840	79	23,653
Hourly wage < 0.5 min wage (2002\$)	392	23,261
Wage fluctuations	876	22,385
Balanced panel	7850	14,535

Note. The initial sample already excludes the latino, the immigration as well as the SEO subsample. I look at PSID from 2002-2006, hence the final number of 14,535 amounts to 2,907 individuals observed over the entire time horizon. In total, I observe 10,865 families.

Variable definitions

Annual labor income The notion of labor income used for estimation includes wages and salaries, bonuses, overtime, tips, commissions, professional practice or trade, market

¹⁶ Until 2015, the PSID uses the term Wife for married females and "Wife" for a cohabiting female. Starting with the 2015 wave, it changed its terminology to *reference person* for the head, and *spouse* or *partner* for the Wife and "Wife", respectively. Going forward, I will use head and spouse.

gardening, farm income, and unincorporated business income. All variables refer to pre-tax values. Post-tax values are obtained by subtracting (labor income) taxes which are estimated with NBER's TAXSIM program.

Annual labor hours Total annual hours worked refers to self-reported hours worked in all jobs, including overtime. Due to missing data for the years 2003 and 2005, I cannot add time spent in unemployment or time spent away from work due to illness of the respondent or others to this variable. Results in [Flodén and Lindé \(2001\)](#) suggest, however, that the parameter estimates are largely unaffected by this omission.

Hourly earnings Hourly post-tax earnings (or post-tax hourly wage) are computed by dividing post-tax annual labor income by annual labor hours.

Table A.5 provides summary statistics of the hourly wage rate and yearly labor earnings of the final sample.

Table A.5: Summary statistics for the primary sample, PSID 2002-2006

Year	Mean age	Mean wage	Median wage	Var of log(wage)	Mean age	Mean wage	Median wage	Var of log(wage)
	Head				Spouse			
2002	40.93	20.28	16.19	0.35	29.58	15.83	13.64	0.27
2003	41.89	19.89	15.92	0.31	30.65	15.22	13.08	0.27
2004	42.89	21.44	17.00	0.33	31.40	16.76	14.48	0.27
2005	43.94	20.78	16.38	0.32	32.42	15.54	13.29	0.28
2006	44.93	22.45	17.58	0.34	33.19	17.48	14.80	0.28

Note. Wage variables are in 2002 dollars.

Estimation As [Flodén and Lindé \(2001\)](#), I refrain from modelling individual-specific intercepts to capture the permanent component. Instead, it is assumed that permanent wage differences can be captured by observed individual characteristics in 2002. These characteristics include age, completed education, occupation, sex, and race.¹⁷ That is, I estimate the following linear model using OLS:

$$\ln w_{m,2002} = \beta_0 + \beta_1 age_m + \beta_2 age_m^2 + \beta_3 sex_m + \beta_4 educ_m + \beta_5 o\vec{c}_m + v_{m,2002}, \quad (A.11)$$

where age denotes the individual's age, sex_m is a dummy variable for the individual's gender, $educ_m$ is years of completed education, and $o\vec{c}_m$ is a vector of 24 occupation dummies.

¹⁷ Education refers to years of completed schooling in 2002. The occupation variable captures 3-digit occupation codes from the 2000 BLS census of population and housing. In total, 24 categories are then constructed based on all major occupation profiles: https://www.bls.gov/oes/current/oes_stru.htm.

The estimation results are shown in Table A.6.

Table A.6: OLS estimation for permanent relative wage component

	Heads	Spouses
<i>age</i>	0.067*** (0.013)	0.006 (0.017)
$\frac{age^2}{100}$	-0.068*** (0.018)	-0.003 (0.021)
<i>sex</i>	0.152*** (0.039)	—
<i>educ</i>	0.074*** (0.007)	0.077***
<i>constant</i>	-2.625*** (0.365)	-1.636*** (0.570)
F-test	0.000	0.000
Adj R ²	0.235	0.284
N	1,919	988

Note. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Occupation dummies are included in the regression, but are not presented here. The sex variable is omitted in the specification for spouses as all observations are female.

The predicted wage $\log \hat{w}_{m,2002} \equiv \hat{\psi}_m$ is then used as an estimate for the permanent wage component for all years $t \in \{2002, 2003, 2004, 2005, 2006\}$ for head and spouse, respectively.

To isolate the residual within a household i , I sum the predicted values for all wage earners within a family and subtract it from the observed total family wage, that is,

$$w_{i,t}^{res} = \log w_{head,t} + \log w_{spouse,t} - (\hat{\psi}^{head} + \hat{\psi}^{spouse}).$$

The residual is essentially the stochastic part which captures all remaining productivity risk once observable permanent components have been removed.¹⁸

With the family wage residuals at hand, I can estimate the parameters ρ , σ_ζ^2 , and σ_ϵ^2

¹⁸ In particular, with the assumptions on the separate productivity processes of family members, we can write:

$$w_{i,t}^{res} = \underbrace{\kappa_{head,t} + \kappa_{spouse,t}}_{\kappa_{i,t}} + \underbrace{\epsilon_{head,t} + \epsilon_{spouse,t}}_{\epsilon_{i,t}} + \underbrace{v_{head,t} + v_{spouse,t}}_{v_{i,t}}$$

where $\kappa_{i,t} = \rho\kappa_{i,t-1} + \zeta_{i,t}$ and variances $(\sigma_\zeta^2, \sigma_\epsilon^2, \sigma_v^2)$.

from the following variance moment conditions for all t

$$\begin{aligned} \text{var}(w_{i,t}^{res}) - \text{var}(\kappa_{i,t}) - \text{var}(\epsilon_{i,t}) - \text{var}(v_{i,t}) &= 0 \\ \text{var}(w_{i,t}^{res}) - \frac{\sigma_{\zeta}^2}{1 - \rho^2} - \sigma_{\epsilon}^2 - \sigma_v^2 &= 0, \end{aligned} \quad (\text{A.12})$$

where the second line follows from stationarity of κ_t . And from the following covariance moment conditions for all $t > s$:

$$\begin{aligned} \text{cov}(w_{i,t}^{res}, w_{i,s}^{res}) - \text{cov}(\kappa_{i,t}, \kappa_{i,s}) &= 0 \\ \text{cov}(w_{i,t}^{res}, w_{i,s}^{res}) - \rho^{t-s} \frac{\sigma_{\zeta}^2}{1 - \rho^2} &= 0, \end{aligned} \quad (\text{A.13})$$

where the first line already excludes covariance terms with $\epsilon_{i,t}$ and $v_{i,t}$ due to the independence assumption.

In total I can construct 15 moments (five variance moments and ten covariance moments) to estimate three parameters; that is, the system is overidentified. Therefore, I use the generalized method of moments (GMM) to minimize the equally weighted distance¹⁹ between model and data determined by the theoretical moment conditions in Equation (A.12) and Equation (A.13), and their empirical counterparts.

Table A.7: GMM estimation results for the stochastic process

Parameter	Estimate	Standard Error
ρ	0.9368	0.0093
σ_{ζ}^2	0.0294	0.0043
σ_{ϵ}^2	0.0670	0.0044

Note. This table shows the estimated parameters of a process for log wage/productivity residuals at an annual frequency. Standard errors are block-bootstrapped with 500 iterations at the family/household level.

The GMM estimation results are shown in Table A.7.

Quarterly process A time period in the model is a quarter. Hence, I have to convert the estimates. To match the parameters from the quarterly AR(1) process to the annual moments, I use the expression for the variance and autocovariance of the *yearly average* of a quarterly AR(1) process. To illustrate this, take the following process z to represent a standard AR(1) process in quarterly terms, where t denotes the 4th quarter and $t - 3$ is

¹⁹ Equal weights are motivated by the results of Altonji and Segal (1996) who show that in small samples the equally weighted minimum distance estimator dominates the optimal distance estimator, especially when using higher than first-order moments.

the 1st quarter of a given year:

$$z_s = \rho_q z_{s-1} + \varepsilon_s \quad \forall s \in \{t, t-1, t-2, t-3\},$$

where I omit a constant as it does not affect either the variance or the covariance and ε is assumed to be identically and independently distributed with variance σ_ε^2 . Furthermore, I assume that the process is weakly stationary. Thus, the yearly average is given by $\bar{z} = \frac{1}{4}(z_t + z_{t-1} + z_{t-2} + z_{t-3})$. Given a zero mean and using that the unconditional variance of the quarterly AR(1) process is given by $\frac{\sigma_\varepsilon^2}{1-\rho_q^2}$, I can write the variance as

$$\begin{aligned} \text{Var}(\bar{z}) &= E[\bar{z}^2] = \frac{1}{4^2} E(z_t + z_{t-1} + z_{t-2} + z_{t-3})^2 \\ &= \frac{1}{4^2} (4 + 6\rho_q + 4\rho_q^2 + 2\rho_q^3) \frac{\sigma_\varepsilon^2}{1-\rho_q^2} \end{aligned} \quad (\text{A.14})$$

where I used that

$$\begin{aligned} E[z_t z_{t-h}] &= \rho_q^h E[z_{t-h}^2] \quad \forall h \in \{1, 2, 3\} \\ E[\varepsilon_t z_s] &= 0 \quad \text{for } s \leq t \end{aligned}$$

By analogy, the autocovariance between two yearly averages, \bar{z}_t and \bar{z}_{t+1} , is

$$\text{Cov}(\bar{z}_t \bar{z}_{t+1}) = \frac{1}{4^2} (\rho_q + 2\rho_q^2 + 3\rho_q^3 + 4\rho_q^4 + 3\rho_q^5 + 2\rho_q^6 + \rho_q^7) \frac{\sigma_\varepsilon^2}{1-\rho_q^2} \quad (\text{A.15})$$

The yearly estimates from Table A.7 are 0.9368 for the autocorrelation and 0.0294 for the conditional variance. Hence, I choose the quarterly parameters to match the yearly unconditional variance, $\frac{0.0294}{1-0.9368^2}$, for Equation (A.14) and the yearly autocovariance, $0.9368 \cdot \frac{0.0294}{1-0.9368^2}$, for Equation (A.15). This yields a quarterly autocorrelation of $\rho_q = 0.976$ and a quarterly variance of $\sigma_{q,\zeta}^2 = 0.012$. The last term to pin down is the quarterly variance of the transitory component, $\sigma_{q,\epsilon}^2$. To do so, I solve $\sigma_{q,\epsilon}^2 / \sigma_{q,\kappa}^2 = \sigma_\varepsilon^2 / \sigma_\kappa^2$. In other words, the relative contribution of the transitory component on a quarterly frequency is the same as on a yearly frequency.

B Hodrick-Prescott Filter on Credit Card Balances

Denote by y_t the aggregate credit card balance in the United States, the variable of interest. To isolate the short-run or cyclical component, I apply a standard Hodrick-Prescott




Figure B.6: Hodrick-Prescott Filter on Credit Card Balances

filter:

$$\min_{\{g_t\}} \left(\sum_{t=1}^T c_t^2 + \lambda \sum_{t=2}^{T-1} [(g_{t+1} - g_t) - (g_t - g_{t-1})]^2 \right),$$

where $c_t = y_t - g_t$ denotes the cyclical part, and g_t is the trend component. Figure B.6 shows the extracted trend component (Panel (a)) and the extracted cyclical component (Panel (b)). The episode prior to the great financial crisis is clearly visible when aggregate credit balances were substantially above the trend level.

Another common specification is to take the logarithm of y_t , as then $(g_{t+1} - g_t)$ is a growth rate and, hence, the second part can be interpreted as the change in trend growth (Hodrick and Prescott (1997)). I will look at both in my estimation.

The empirical moment I want to match with the model is the autocorrelation in the cyclical component. Hence, I estimate an AR(1) and target the coefficient on the lag. The results are shown in Table B.8. We see that the two coefficients are similar in magnitude and do not differ significantly. For my quantitative exercise I will target the coefficient in levels.

Table B.8: Autocorrelation in cyclical component


	Level	Log
Lag	0.841 (0.064)	0.829 (0.063)
Observations	67	67
R-squared	0.70	0.68

Note. This table lists the OLS estimates of the AR(1) process on the cyclical component of the aggregate credit card balances. The "Level/Log" specification refers to whether the Hodrick-Prescott filter was applied in levels or logs.

C Additional results

Transitional dynamics of aggregates Figure C.7 shows the responses of the interest rate, labor supply, and output, respectively, for both credit shock specifications. The permanent credit shock specification replicate qualitatively the results by [Guerrieri and Lorenzoni \(2017\)](#). The mean-reverting specification shows the opposing dynamics in the short and medim term, as discussed in the main text.

Figure C.7: Interest rate, labor supply, and output during the transition



Note. This figure shows the responses of the interest rate, labor supply, and output, respectively, for the permanent and mean-reverting credit shock specification. The interest rate is annualized, whereas aggregate labor supply and output are in percentage deviations from the initial steady state.

State-dependent credit limit I also experimented with a state-dependent form of the credit limit specified as $\phi_{t+1} = \tilde{\phi}_{t+1} Y_{t+1}$. The steady states are unaffected by this formula-

tion if I scale the state-dependent borrowing limit parameter accordingly: $\phi/Y = \tilde{\phi}$. I then modelled a permanent shock to $\tilde{\phi}_{t+1}$ that replicates the aggregate drop in the credit-limit-to-GDP ratio of 30% as observed in the data. However, as shown in Figure C.7, the drop in output under both credit shock specifications is fairly small with a maximum decline of 0.28% on impact. As a result, the transitional dynamics under the ad-hoc credit constraint from the main text and the state-dependent credit constraint introduced here are virtually identical.

D Steady state optimization and stationary recursive equilibrium

In this section, I first state the household problem in the steady states in recursive form. Second, I define the stationary recursive competitive equilibrium. The main difference to the transitional equilibrium definition in the main text is that the value function, policy functions, the interest rate and government policies are not indexed by Φ ; recall that the time subscript implicitly took care of this in Definition 1. The reason is that all conditions have to be satisfied only for the equilibrium measure Φ , which reproduces itself — it is stationary.

Recursive Problem

$$\begin{aligned}
V(b, \theta; \Phi) &= \max_{c, n, b'} \left[U(c, n) + \beta \mathbb{E}[V(b', \theta'; \Phi')] \right] \\
&\text{s.t.} \\
q(\Phi)b' + c &\leq b + y - \tau(\Phi) \\
b' &\geq -\phi \\
\Phi' &= H(\Phi),
\end{aligned} \tag{D.16}$$

where $H(\Phi)$ is the law of motion generated by the Markov process governed by Γ and the optimal policy functions of the household.

Denote the set of all probability measures on the measurable space (S, Σ_s) by \mathcal{M} . The stationary recursive equilibrium is summarized in the following definition.

Definition 2. A **stationary recursive competitive equilibrium** is a value function $V : S \rightarrow \mathbb{R}$, policy functions for the household $C : S \rightarrow \mathbb{R}$ and $N : S \rightarrow \mathbb{R}$, future bond holdings b' , an interest rate r , government policies $\tilde{\tau}$ and a measure $\Phi \in \mathcal{M}$ such that

i) V, C, N are measurable with respect to Σ_s , V satisfies the optimization problem of the household, Equation (D.16), and C, N are the associated policy functions, given r and τ .

ii) The government budget constraint is satisfied

$$\tau = \frac{rB}{1+r}$$

iii) The asset market clears

$$\int_{(A \times \Theta)} b'(b, \theta) d\Phi = B$$

iv) For all $\mathcal{S} \in \Sigma_s$

$$\Phi(\mathcal{S}) = \int_{(A \times \Theta)} Q((b, \theta), \mathcal{S}) d\Phi,$$

that is, the probability measure reproduces itself.

Again, the goods-market clearing condition is redundant by Walras law and thus omitted.

E Computational Details

E.1 Numerical Algorithm - Steady State

I describe now how I compute the policy functions and the invariant distribution for the initial and terminal steady states.

Policy Functions To obtain the policy functions, I employ the endogenous grid method by [Carroll \(2006\)](#).

1. I construct a grid on (b, θ) where $b \in G_b = \{b_1, \dots, -\phi, \dots, 0, \dots, b_N\}$ and $\theta \in \Theta = \{\theta^1, \dots, \theta^{13}\}$
2. I guess an initial policy function $\hat{C}_0(b_i, \theta_j) = \max\{c_{min}, r \cdot b_i\}$ where c_{min} is a pre-specified minimum consumption level.
3. Iterate over pairs $\{b'_i, \theta_j\}$ in this and the next step. Fix θ_j and iterate over all grid values of b'_i . For any pair $\{b'_i, \theta_j\}$ on the mesh $G_b \times \Theta$ where the borrowing limit is

not binding, construct

$$\tilde{C}(b'_i, \theta_j) = \left[\beta(1+r) \sum_{\theta' \in \Theta} \pi(\theta'|\theta_j) \hat{C}_0(b'_i, \theta')^{-\gamma} \right]^{-\frac{1}{\gamma}},$$

which is the Euler equation solved for consumption using my utility specification. $\pi(\theta'|\theta_j)$ denotes the probability of being type θ' tomorrow conditional on being θ_j today. Similarly, I solve for optimal labor supply and construct

$$\tilde{N}(b'_i, \theta_j) = \max \left\{ 0, 1 - \left[\frac{\theta_j \tilde{C}(b'_i, \theta_j)^{-\gamma}}{\psi} \right]^{-\frac{1}{\eta}} \right\}$$

where I use the consumption level from the consumption policy function above. If the borrowing limit is binding, I cannot use the Euler equation. Hence, I set the consumption level to c_{min} to calculate optimal labor supply n_{min} from the condition above. Using n_{min} , I compute the level of consumption, c^* , that solves

$$0 = -\phi - \frac{-\phi}{1+r} - c^* + \theta_j n_{min} - \tilde{\tau}.$$

This is the lowest consumption level that is generated by the consumption policy function. The consumption area between this level and the lowest consumption level where the constraint is not binding, is then generated by computing an evenly sized grid between these two points. Subsequently, labor supply for constrained households is computed identically as above, using the newly obtained consumption levels. Do this for all $\theta_j \in \Theta$.

These two approaches combined yield the policy function for consumption and labor supply, respectively, in any given iteration.

4. From the budget constraint, I solve for the value of assets today, $b^+(b'_i, \theta_j)$. For unconstrained agents, this is

$$b^+(b'_i, \theta_j) = \frac{b'_i}{1+r} + \tilde{C}(b'_i, \theta_j) - \theta_j \cdot \tilde{N}(b'_i, \theta_j) + \tau.$$

For constrained agents, I replace b'_i with $-\phi$. This implies that asset holdings of $b^+(b'_i, \theta_j)$ and an idiosyncratic shock of θ_j today would lead the agent to hold b'_i assets tomorrow. The function $b^+(b'_i, \theta_j)$ is not defined on the grid G_b and changes every iteration, i.e. endogenous grid. This is also important when I compute the invariant distribution.

5. Update the guess of the consumption policy function. To obtain a new guess $\hat{C}_1(b_i, \theta_j)$ I linearly inter-and extrapolate the values for $\{\tilde{C}(b_n^+, \theta_j), \tilde{C}(b_{n+1}^+, \theta_j)\}$ on the two most adjacent values $\{b_n^+, b_{n+1}^+\}$ that enclose the given grid point b_i . If some grid point values b_i are beyond b_N^+ , I extrapolate to obtain the new guess. After I obtained all new guesses, I impose the lower bound of c_{min} .
6. I declare convergence when

$$\max_{i,j} |\hat{C}_{n+1}(b_i, \theta_j) - \hat{C}_n(b_i, \theta_j)| < \varepsilon,$$

for some small ε and where $\hat{C}_n(b_i, \theta_j)$ denotes the consumption policy in the n 'th iteration. If convergence has not been achieved, repeat steps 3-5 using the latest guess of the consumption policy function and check convergence again.

Invariant Distribution

1. Assign weights for all possible bond holding values generated by the policy functions above that are proportional to the distance of the two most adjacent grid point values. For instance, let the current bond holdings be b^+ and let $\{b_{n-1}, b_n\}$ be the two most adjacent grid point values that enclose these bond holdings. The weight ζ^+ is then computed as

$$\zeta^+ = \frac{b^+ - b_{n-1}}{b_n - b_{n-1}}.$$

If current bond holdings are above (below) the highest (lowest) grid point value, set the weight for the highest (lowest) grid point equal to 1.

2. The initial guess for the initial distribution $\Phi_{(0)}(b|\theta)$ is the uniform distribution.
3. Fix θ_j and compute the distribution as follows:

$$\begin{aligned}\Phi_{(1)}(b_{n-1}|\theta_j) &= \sum_{\theta' \in \Theta} (1 - \zeta) \pi(\theta'|\theta_j) \Phi_{(0)}(b_{n-1}|\theta_j) \\ \Phi_{(1)}(b_n|\theta_j) &= \sum_{\theta' \in \Theta} \zeta \pi(\theta'|\theta_j) \Phi_{(0)}(b_n|\theta_j),\end{aligned}$$

where ζ denotes the particular weight for the bond holding that is enclosed by the grid points $\{b_{n-1}, b_n\}$. The weights adjust for the transition of off-grid values to values on the grid: when bond holdings are close to, say, b_n , the distribution $\Phi_{(1)}(b_n|\theta)$ gets a higher weight. Furthermore, the distribution $\Phi_{(k)}(b_n|\theta_j)$ is affected by the

probability of being type θ' tomorrow conditional on being θ_j today, denoted by $\pi(\theta'|\theta_j)$, and the current mass at particular bond holdings conditional on being type θ_j , $\Phi_{(k-1)}(b_n|\theta_j)$.

Repeat this for all $\theta_j \in \Theta$, $b_n \in G_b$ and sum the distributions to get $\Phi_{(k)}(b|\theta)$.

4. I check convergence by computing

$$|\Phi_{(k)}(b|\theta) - \Phi_{(k-1)}(b|\theta)| < \epsilon,$$

for some small ϵ and where $\Phi_{(k)}(b|\theta)$ denotes the distribution in the k 'th iteration. If convergence has not been achieved, repeat step 3 using the latest guess of the distribution and check convergence again.

E.2 Numerical Algorithm - Transition

I describe now how I compute the transition path for interest rates. Essentially, the algorithm iterates backward over the policy functions starting from the terminal steady state to obtain a sequence of policy functions. Subsequently, with the sequence of policy functions at hand, iterate the distributions forward starting from the stationary distribution in the initial steady state. This backward-forward iteration to obtain general equilibrium time paths of aggregate prices (and variables) has been employed by several recent articles, for instance, [Guerrieri and Lorenzoni \(2017\)](#) or [Auclert and Rognlie \(2020\)](#)

The economy at $t = 0$ is at the steady state with stationary distribution Φ over assets and productivity types. At the end of period $t = 0$, an unexpected credit crunch hits the economy that reduces the borrowing limit ϕ_t . I assume that the economy converges to the new terminal steady state after T periods, for T arbitrarily large but finite. The assumption on T allows us to solve the household problem by backward induction. To compute the equilibrium interest rate path, I follow these steps:

1. I set $T = 100$.
2. I specify the sequence of borrowing limits $\{\phi_t\}_{t=1}^T$ where the borrowing limit from the terminal steady state is obtained after 6 periods.
3. I compute the policy functions of the initial and terminal steady state using the algorithm in [Appendix E.1](#)
4. I guess an initial interest rate path of length T such that $r_t = r_T \quad \forall t > 1$, i.e. the sequence of interest rates equals the interest rate in the terminal steady state.

5. Since $\hat{C}_T(b, \theta)$ equals the consumption policy function from the terminal steady state, I can solve the household problem by backward induction and derive $\{\hat{C}_t(b, \theta)\}_{t=1}^{T-1}$ and $\{\hat{N}_t(b, \theta)\}_{t=1}^{T-1}$ using the endogenous grid method as outlined above. For every t , I use the corresponding interest rate and borrowing limit from the specified sequences. Again, bond holdings can be computed via the budget constraint.
6. Iterate the bond distribution forward starting from the initial steady state distribution at $t = 1$. Compute the aggregates, i.e. output, consumption, labor supply, and household bond demand, at time t , using the time t policy functions from the previous step.
7. For every iteration, I check bond market clearing for convergence:

$$\sqrt{\frac{\langle (B_{dem} - B), (B_{dem} - B) \rangle}{T}} < \epsilon, \quad (\text{E.17})$$

for a small ϵ and where B_{dem} denotes the aggregate bond demand vector of length T , B denotes the bond supply vector of length T , and $\langle \cdot, \cdot \rangle$ denotes the inner product.

8. If inequality Equation (E.17) is not satisfied for any iteration k , I update the interest rate path for next iteration, $r^{(k+1)}$, based on bond market clearing with the following linear updating rule

$$r^{(k+1)} = r^k - \epsilon(B_{dem}^{(k)} - B),$$

where B_{dem}^k denotes aggregate bond demand at iteration k and B denotes bond supply. ϵ are exponentially decaying weights, i.e. divergences in the bond markets at time periods closer to the credit crunch get a higher weight. Note that $r^{(k+1)}, r^k, \epsilon, B_{dem}^{(k)}$, and B are vectors of length T representing the whole path. Repeat steps 5-8 using the new interest rate path until convergence is achieved.