Aggregate Precautionary Savings Motives*

Pierre Mabille⁺

March 8, 2023

Abstract

This paper analyzes the effect of *aggregate* risk on households' precautionary savings, a new channel that complements the standard idiosyncratic precautionary motive. I build a general equilibrium model with incomplete markets, heterogeneous households, and aggregate risk to decompose the drivers of precautionary savings. The precautionary motive due to credit supply shocks is large, at odds with received wisdom about the low costs of aggregate fluctuations. It is larger for middle-class households, who are too rich to benefit from social programs but too poor to have enough liquid assets. Aggregate precautionary motives are important because they imply that aggregate shocks can have permanent effects even when they are temporary.

JEL classification: D14; D52; E21; E44; G11; G51.

Keywords: Household finance; precautionary savings; liquid assets; household debt; borrowing constraints.

^{*}First draft: February 2019. I thank Lorenzo Bretscher (discussant), Bernard Dumas, Vadim Elenev, Mark Gertler, Sasha Indarte, Virgiliu Midrigan, Thomas Philippon, Johannes Stroebel, Stijn Van Nieuwerburgh, Olivier Wang, and Stanley Zin for comments. This paper benefited from participants at the EFA annual meeting, the SED meeting, and the Becker-Friedman Institute Macro Finance Summer Session for Young Scholars. I gratefully acknowledge financial support from the Macro-Financial Modeling group of the Becker-Friedman Institute.

[†]Department of Finance, INSEAD, Boulevard de Constance, 77300 Fontainebleau, France. Email: pierre.mabille@insead.edu. Website: https://www.pierremabille.com.

1 Introduction

How does economy-wide–*aggregate*–risk affect households' precautionary savings? It is well established that households hedge against individual–*idiosyncratic*–risks (e.g., income, or health shocks) by accumulating precautionary savings above and beyond their needs for intertemporal consumption smoothing.¹ Such savings have important implications. They decisively shape households' balance sheets, and they contribute to lowering risk-free interest rates, which are key for most asset prices. Yet, little is known about the drivers of precautionary savings beyond idiosyncratic risk.

This paper evaluates the effects of aggregate risk on households' precautionary savings, and compares them with those resulting from the standard idiosyncratic precautionary motive. I build a general equilibrium model of precautionary savings with heterogeneous households and incomplete markets. I introduce two sources of real and financial aggregate risk to which the Great Recession and the Covid-19 crisis lent new urgency: fluctuations in (i) aggregate productivity and (ii) the tightness of households' borrowing constraints as a result of credit supply shocks. I use the model to decompose precautionary motives and quantify their real implications. The model addresses the empirical challenge of identifying the effects of aggregate shocks on households' savings, which are difficult to disentangle from those of idiosyncratic shocks in the data. To achieve identification, it departs from existing models with discrete aggregate shocks (e.g., Krusell and Smith (1998)), and incorporates continuous aggregate shocks that match the data over the full business cycle instead of a limited number of states of the world.

I obtain three new findings. First, aggregate risk significantly contributes to households' precautionary savings, in contrast with received wisdom about the low costs of aggregate fluctuations for households (e.g., Lucas (1987)). The contribution of credit supply shocks to higher savings and a lower risk-free rate is especially large. It represents 15% of the impact of idiosyncratic income shocks and it dwarfs the impact of aggregate productivity shocks, which is close to zero. Second, aggregate precautionary motives are the largest for "middle-class" households, which nuances the focus of economists on the top and the bottom of the wealth distribution. Such households are too rich to have enough public insurance from social safety nets, but too poor to have enough private insurance from liquid assets. Third, aggregate precautionary motives imply that aggregate shocks can have permanent real effects even when they are themselves temporary.

¹See, e.g., Zeldes (1989), Deaton (1991), Huggett (1993), Aiyagari (1994), Carroll (1997), Carroll and Samwick (1997), Gourinchas and Parker (2002), Kennickell and Lusardi (2004), Parker and Preston (2005), De Nardi, French and Jones (2010), Boar (2021).

The precautionary savings induced by credit supply shocks generate a low-debt environment where consumption is persistently depressed, which helps explain recent features of developed economies. Overall, credit supply and productivity shocks are critical to jointly match household balance sheet and macroeconomic moments, which has been a challenge for macro-finance models so far.

The model is populated by infinitely-lived, risk-averse households with heterogeneous income and wealth. Every period, households consume, elastically supply labor to competitive firms, and save in risk-free bonds or borrow subject to a credit limit, which depends on individual income and aggregate credit supply. The government raises progressive taxes and issues risk-free debt to finance progressive transfers and existing debt. The risk-free rate clears the market for savings, and the wage clears the labor market. Households face idiosyncratic productivity risk as in standard models, and two sources of aggregate risk: (i) aggregate productivity and (ii) their borrowing constraints are subject to continuous and mean-reverting shocks, which affect their income and their borrowing capacity.

Three ingredients are key for evaluating precautionary motives. First, markets are incomplete, which leads households to demand risk-free bonds to hedge against shocks. Incompleteness generates heterogeneity across households, which allows to separate the effects of idiosyncratic and aggregate shocks. Second, the model incorporates the general equilibrium feedback from households' savings to the risk-free rate, abstracting from which would lead to overstating precautionary motives. Aggregate risk increases precautionary savings, but less than with a fixed rate, since higher savings lower the equilibrium rate, which in turns makes saving less attractive. Third, households can hedge against risks using two forms of private and public insurance that are imperfect substitutes to savings: they can adjust their labor supply and receive government transfers in bad times. Without those, the role of savings would also be overstated.

I calibrate the model to match the level and cross-section of liquid assets and unsecured debt held by households in the United States. The supply of liquid assets is endogenous and consists of the bonds issued by the household sector and the government.² The stochastic process for aggregate productivity is calibrated externally using historical data on total factor productivity. The stochastic process for the tightness of households' borrowing constraints is calibrated internally to match the persistence and the volatility of the real risk-free rate in the data. The benefits of this approach are twofold. First, it

²This approach is the same as in Huggett (1993) and Aiyagari (1994) but focuses on households' precautionary savings instead of firm capital. Government debt in excess of households' savings is held by investors outside the model as in Guerrieri and Lorenzoni (2017)–a plausible assumption for the U.S. where debt held by the public is on average 40% of GDP.

circumvents the difficulty of identifying a long time series of credit supply shocks in the data. In comparison, the resulting values imply less variations than available estimates directly based on changes in credit card limits. Therefore, they provide a lower bound on aggregate precautionary savings due to credit supply shocks.³ Second, it guarantees that precautionary motives in the model are consistent with historical values of the risk-free rate, and therefore are correctly estimated.

My findings rely on a new decomposition of precautionary motives, which is based on the economy's departure from certainty equivalence with respect to the various sources of aggregate risk. Households have rational expectations and make optimal decisions knowing the stochastic processes that govern shocks. I solve for first- and second-order approximations of the stochastic dynamics of the economy with respect to aggregate shocks, around its deterministic steady state where these shocks are zero. The benefit of this approach is to provide an economics interpretation of the effect of aggregate risk. For each shock, the difference between second- and first-order terms captures the effect of volatility on household behavior. It leads to different amounts of savings for each source of risk, which would be otherwise difficult to identify in the data. This measure of precautionary savings assumes some "bounded rationality" by abstracting from the effects of the higher-order moments of aggregate shocks.⁴

The model highlights three precautionary motives. (1) The *standard motive* due to idiosyncratic income risk. It arises because of the prudence property of utility u'''(.) > 0, and because the combination of income shocks and borrowing constraints hampers consumption smoothing for some households. It has the largest impact. It increases liquid savings by 280% (from 31% to 119% of GDP, in annual terms) and decreases the risk-free rate by 6 percentage points (from 8.40% to 2.40%) compared to an economy without idiosyncratic risk where they are only determined by intertemporal substitution. (2) An *aggregate financial motive* due to credit supply shocks, whose impact further increases liquid savings by 45% (from 119% to 173% of GDP) and lowers the risk-free rate by 0.6 pp (from 2.40% to 1.80%). This impact is sizable and cannot be ignored when analyzing precautionary savings, especially since these estimates are a lower bound. (3) An *aggregate real motive* due to productivity shocks. While these shocks affect employment, interestingly, they have close to zero impact on liquid savings and the risk-free rate. Taken together,

³Because of data limitations, few papers directly estimate a stochastic process for households' borrowing constraints. For instance, Fulford (2015) estimates persistent and more volatile credit card limits, but assumes an exogenous fixed interest rate.

⁴Estimating these effects would require the households and the econometrician to have unrealistically high computing power. Abstracting from them is plausible given how households make aggregate forecasts in practice (e.g., Das, Kuhnen and Nagel (2019)), and it is quantitatively innocuous as the resulting market clearing errors are negligible.

these last two findings nuance received wisdom that aggregate fluctuations have low costs for household and thus would not lead to any precautionary behavior (e.g., Lucas (1987)).

What explains the large effect of credit supply shocks on precautionary savings? When the tightness of borrowing constraints varies, households' ability to smooth consumption is impaired. They insure against such changes by deleveraging and accumulating savings. This creates downward pressure on the risk-free rate, which in turn makes savings less attractive and determines their equilibrium quantity. Quantitatively, the magnitude of this effect depends on the volatility of credit supply shocks. In particular, it is larger in the post-Great Recession period on which this paper focuses.

To evaluate the real implications of the two aggregate precautionary motives, I then compute a variance decomposition of the business cycles in the model. Credit supply shocks explain more (about 60%) of the volatility of consumption and output than productivity shocks (about 40%). However, both shocks are critical to jointly match household balance sheet and macroeconomic moments.

First, they generate procyclical household debt. I use nonlinear impulse response functions to show that credit supply shocks increase households' liquid assets and decrease their debt through two effects. Through a first-order effect, a lower level of borrowing constraints forces constrained borrowers to deleverage, and those close to the constraint to increase precautionary savings to avoid hitting them because of income risk. Then, through a second-order effect, the volatility of borrowing constraints themselves makes them more likely to bind, which further increases savings. These effects lower output, which generates the positive correlation with household debt. With flexible labor supply, a lower risk-free rate creates an intertemporal substitution motive that increases households' consumption of leisure. It decreases hours worked by unconstrained households, which tend to be more productive. These households further decrease their hours because of a wealth effect as they consume more goods. Even if borrowing-constrained households work more to pay back their debt, the net effect is a decrease in aggregate employment, which translates into lower output.

Second, aggregate precautionary motives help explain the simultaneous recovery of consumption and the decrease in household debt (and the risk-free rate) after the Great Recession. It is a puzzle for standards models, in which an increase in future consumption should increase debt and the risk-free rate because of intertemporal substitution. I use this feature of the post-Great Recession period to test the empirical validity of the model. I apply a particle filter to estimate the sequence of structural productivity and credit supply shocks, which explain the observed paths for consumption and the risk-free

rate after 2005. This is the first paper to perform this exercise, which is computationally challenging, in a general equilibrium model with heterogeneous households, incomplete markets, and aggregate risk.

A persistent tightening of borrowing constraints across households by 15% explains the decrease in the risk-free rate from 2.5% to -1.5%. Its effect is exacerbated by a simultaneous short-lived 2% drop in aggregate productivity, which makes constraints more binding. Combined, both shocks generate a temporary decline in consumption early on, which recovers when productivity goes back to its prior level. The model also matches the decrease in household debt and employment, which are not targeted. It provides structural estimates of credit supply shocks, which are difficult to identify in the data. These estimates stay persistently low throughout the sample, indicating that borrowing constraints have remained tight long after the recession. Interestingly, they imply tighter borrowing constraints than standard measures of credit conditions such as bank lending standard, which suggests that structural estimates may be useful to capture the complete credit landscape faced by households.

Related literature. These results complement a longstanding literature that focuses on idiosyncratic income risk but abstracts from aggregate risk as a potential driver of households' precautionary savings (Zeldes (1989), Deaton (1991), Carroll (1997), Parker and Preston (2005)). Credit supply shocks, in particular, are large and frequent. Not only did households' borrowing constraints massively tighten after the Great Recession. They also vary in long time series, including the recent Covid-19 recession, and with monetary and macro-prudential policy.⁵

This is the first paper that decomposes the sources of precautionary savings in a general equilibrium model with heterogeneous households and aggregate risk. It is critical to endogenize the interest rate it as in Huggett (1993), Aiyagari (1994), and Heaton and Lucas (1996) as it is a key determinant of households' savings. My quantitative results reflect recent empirical evidence that aggregate risk lowers the risk-free rate (Hartzmark (2016), Pflueger, Siriwardane and Sunderam (2020)), which in turn affects the quantity of savings. In contrast, Ludvigson (1999) and Fulford (2015) analyze stylized models with stochastic borrowing constraints but with exogenous interest rates. They focus on consumption smoothing and the credit card puzzle (Bertaut, Haliassos and Reiter (2009)) rather than precautionary savings. Instead, this paper proposes a detailed model of precautionary

⁵See, for instance, Gross and Souleles (2002), Mian, Rao and Sufi (2013), Mian, Sufi and Verner (2017), Baker (2018), Cherry, Jiang, Matvos, Piskorski and Seru (forthcoming), Agarwal, Chomsisengphet, Mahoney and Stroebel (2018), and Acharya, Bergant, Crosignani, Eisfert and McCann (2022).

savings with various substitutable forms of insurance as in the data. As in Hubbard, Skinner and Zeldes (1995), social insurance programs lower precautionary savings, and so does flexible labor supply. The effect of stochastic borrowing constraints is consistent with Guiso, Jappelli and Terlizzese (1996), where expectations of future borrowing constraints increase households' demand for risk-free savings. Favilukis, Ludvigson and Van Nieuwerburgh (2017) study a richer model with time-varying borrowing constraints, but they focus on housing and only allow for two states of the business cycle where credit supply is either high or low. In corporate finance, Jermann and Quadrini (2012) also analyze time-varying borrowing constraints, but in a model with a representative firm and no idiosyncratic shocks, which abstracts from precautionary motives.

Finally, this paper estimates the time series of structural shocks that drive equilibrium prices and quantities using particle filtering in a model with heterogeneous households, incomplete markets, and aggregate risk. I use this novel approach to test the empirical validity of the model in recent data. It relies on continuous aggregate shocks instead of discrete shocks with a limited number of states as in existing models (e.g., Krusell and Smith (1998), Favilukis et al. (2017)), and on a second-order approximation of the dynamics of the model. Beyond this paper, this approach can help improve the realism of quantitative heterogeneous household models with aggregate risk.

Outline. The rest of the paper is organized as follows. Section 2 presents the model, and Section 3 the resulting decomposition of precautionary motives. Section 4 describes the calibration. Section 5 analyzes the main results on the impact of aggregate risks on precautionary savings and their implications for household balance sheets and the macroe-conomy. Section 6 tests the empirical validity of the model mechanism in the post-Great Recession period. Section 7 concludes.

2 Model

This section describes a general equilibrium model of precautionary savings with heterogeneous households, incomplete markets, and two sources of aggregate risk in a closed economy: shocks to aggregate productivity and to households' borrowing constraints. Households have rational expectations and time is discrete.

2.1 Households

Household choices. The economy is populated by a continuum of measure 1 of heterogeneous, risk-averse households. Households face idiosyncratic labor income risk. They consume c_t units of a single final good produced by competitive firms, and elastically supply n_t labor hours to these firms. Firms' profits are redistributed equally. There are progressive taxes on labor income and progressive transfers from the government that are conditional on income.

Households can save in liquid assets and borrow with unsecured debt by buying and selling one-period risk-free bonds b_{t+1} . Their balance sheets are summarized by their net bond holdings. When borrowing, they face stochastic borrowing constraints, which consist of an aggregate component that depends on credit supply and of an individual component that depends on household income.

Households choose consumption, labor supply, and their net bond holdings to maximize the expected discounted value of the utility flows from consumption net of the disutility of working:

$$\max_{\{c_{it}, n_{it}, b_{it+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{it}^{1-\gamma}}{1-\gamma} - \psi \frac{n_{it}^{1+\eta}}{1+\eta} \right],$$
(1)

subject to budget and borrowing constraints,

$$c_{it} + \frac{b_{it+1}}{1+r_t} + \tau_t \left(\theta_{it}, n_{it}\right) \le w_t \theta_{it} n_{it} + b_{it} + \pi_t + T\left(\theta_{it}\right)$$

$$b_{it+1} \ge -\overline{\phi}_t \phi\left(\theta_{it}\right).$$
(2)

Consumption *c*, net bond holdings b'/(1 + r), and tax payments τ , cannot exceed current savings or debt *b*, and income from labor earnings $w\theta n$, firm profits π , and government transfers *T*.

Borrowing constraints depend on aggregate and individual components, which interact multiplicatively. The individual component captures the effect of income on borrowing constraints through the function $\phi(.)$, which depends on productivity. The aggregate component captures the effect of credit supply.

To keep the model tractable while introducing aggregate risk, I assume that there is a single interest rate r_t such that households can borrow or save at the same rate. One benefit of this approach is to provide a lower bound on households' precautionary motives since a higher interest rate on borrowing would further lower debt and increase savings.

Idiosyncratic risk. Idiosyncratic productivity θ follows an AR(1) process (in logarithm). Its volatility is a decreasing function of aggregate productivity z_t , which makes idiosyncratic income risk countercyclical. Productivity is discretized as a finite Markov chain $\Theta(z_t) = \{\underline{\theta}(z_t), ..., \overline{\theta}(z_t)\}$ with transition matrix $\Pi_{\theta}(z_t)$ using the Rouwenhorst method:

$$\log \theta_{it} = \rho_{\theta} \log \theta_{it-1} + \sigma_{\theta} \left(z_t \right) \epsilon_{it}^{\theta}, \ \epsilon^{\theta} \sim \mathcal{N}(0, 1).$$
(3)

Aggregate risk. There are two sources of aggregate risk. Aggregate productivity z_t follows a standard AR(1) process:

$$\log z_t = \rho_z \log z_{t-1} + \epsilon_t^z. \tag{4}$$

Credit supply follows a mean-reverting AR(1) process with mean $\overline{\phi}$. A negative credit supply shock $\epsilon^{\phi} < 0$ induces a tightening of borrowing constraints for all households, irrespective of their individual income, while a positive shock $\epsilon^{\phi} \ge 0$ induces a relaxation:

$$\log \overline{\phi}_t - \log \overline{\phi} = \rho_\phi \left(\log \overline{\phi}_{t-1} - \log \overline{\phi} \right) + \epsilon_t^\phi \tag{5}$$

Aggregate productivity and credit supply shocks are correlated and follow a bivariate Normal distribution:

$$\begin{pmatrix} \epsilon^{\phi} \\ \epsilon^{z} \end{pmatrix} \stackrel{iid}{\sim} \mathcal{N} \left(\mathbf{0}, \begin{pmatrix} \sigma_{\phi}^{2} & \sigma_{\phi} \sigma_{z} \rho_{\phi z} \\ \sigma_{\phi} \sigma_{z} \rho_{\phi z} & \sigma_{z}^{2} \end{pmatrix} \right)$$
(6)

2.2 Firms

A continuum of competitive firms hires efficient units of labor from households every period, and combines them using a decreasing returns to scale production technology subject to aggregate productivity shocks. Firms choose total hours to solve a static profit maximization problem:

$$\max_{N_t} \pi_t = z_t N_t^{\alpha} - w_t N_t \tag{7}$$

Profits are redistributed to households equally, which provides a lower bound on precautionary motives as it relaxes constraints for poorer households relatively more. Firm shares are not tradable to focus on precautionary savings in risk-free assets.⁶

⁶The model focuses on households' savings and abstracts from firm capital, which is the subject of a rich literature in corporate finance (e.g., Jermann and Quadrini (2012)). This assumption can be relaxed with a larger state space.

In equilibrium, firms' profits and the wage bill are constant shares of output. Therefore, the firm sector transmits aggregate productivity shocks one to one to households' wages and profit shares:

$$\pi_t = (1 - \alpha)Y_t = (1 - \alpha)z_t N_t^{\alpha}$$

$$w_t N_t = \alpha Y_t = \alpha z_t N_t^{\alpha}$$
(8)

2.3 Government

The government raises a progressive tax on labor earnings and issues risk-free bonds to finance progressive transfers and outstanding debt:

$$\int T(\theta) d\lambda_t(\theta, b) + B_t \leq \int \tau_t(\theta, n(\theta, b)) d\lambda_t(\theta, b) + \frac{B_{t+1}}{1 + r_t}$$
(9)

Progressive income taxes are an affine function of labor earnings. Its slope depends on household productivity and its intercept adjusts such that the government budget constraint holds every period:

$$\tau_t \left(\theta_{it}, n_{it} \right) = \tau_{0t} + \tau_1 \left(\theta_{it} \right) w_t \theta_{it} n_{it} \tag{10}$$

2.4 Equilibrium

The model has heterogeneous households, incomplete markets, and aggregate risk. Therefore, the entire cross-sectional distribution of households over productivity and net bond holdings { $\lambda_t(\theta, b)$ } is an aggregate state variable. Households know the current risk-free rate r_t but need to forecast r_{t+1} next period to make intertemporal consumption and savings choices. The model is a closed economy where the supply of liquid assets comes from risk-free bonds issued by the household sector and the government. For a given supply of liquid assets, forecasting the risk-free rate is equivalent to forecasting the demand for liquid assets next period, which depends on the entire cross-sectional distribution. The distribution is time-varying because of aggregate shocks to productivity and borrowing constraints. Households must also forecast the current wage w_t , which is given by the intersection of firms' labor demand and households' labor supply, which is itself a function of the wage. The rational expectations equilibrium of the economy is a fixed point where households' forecasts for aggregate state variables coincide with their realized values (as, e.g., in Krusell and Smith (1998)).

Definition 1 (competitive equilibrium). Given a sequence of aggregate shocks to pro-

ductivity and borrowing constraints $\{z_t, \overline{\phi}_t\}$, a competitive equilibrium is a sequence of time-varying policy functions for household consumption, labor supply, and net bond holdings $\{c_t(\theta, b), n_t(\theta, b), b_{t+1}(\theta, b)\}$ and for firm labor demand N_t , and a sequence of risk-free rates and wages $\{r_t, w_t\}$, and government taxes $\{\tau_t\}$, such that:

(i) The optimality conditions for households' savings and labor supply choices hold:

$$c_{t}(\theta,b)^{-\gamma} = \beta(1+r_{t})\mathbb{E}_{t}\left[c_{t+1}(\theta,b)^{-\gamma}\right] + \mu_{t}(\theta,b)$$

$$\psi n_{t}(\theta,b)^{\eta} = (1-\tau_{1}(\theta))w_{t}\theta c_{t}(\theta,b)^{-\gamma}$$
(11)

where $\mu_t(\theta, b)$ denotes the multiplier on the borrowing constraint of household of type (θ, b) . The first intertemporal optimality condition states that for each household, the marginal cost of additional savings must equal the sum of the expected discounted gains of these savings next period when earning the risk-free rate and of the shadow price of relaxing the borrowing constraint. The second intratemporal optimality condition states that the marginal cost of an addition hour of work must equal the marginal utility benefits associated with the additional earnings net of taxes.

(ii) The optimality conditions for firms' labor demand choices hold:

$$\alpha z_t N_t^{\alpha - 1} = w_t \tag{12}$$

The marginal productivity of an additional work hours in efficiency units must equal the wage.

(iii) The government budget constraint holds:

$$\int T(\theta) d\lambda_t(\theta, b) + B_t = \int \tau_t(\theta, n(\theta, b)) d\lambda_t(\theta, b) + \frac{B_{t+1}}{1+r_t}$$
(13)

The government raises a progressive tax on labor earnings and issues risk-free bonds to finance progressive transfers and outstanding debt.

(iv) The markets for goods, labor, and savings clear:

$$\int c_t (\theta, b) d\lambda_t (\theta, b) = Y_t = z_t N_t^{\alpha}$$

$$\int \theta n_t (\theta, b) d\lambda_t (\theta, b) = N_t$$

$$\int b_{t+1} (\theta, b) d\lambda_t (\theta, b) = B_{t+1} = B$$
(14)

First, aggregate consumption must equal output. The price of goods is normalized to 1. Second, the wage adjusts such that aggregate work hours in efficiency units equal firms' demand for labor. Third, the risk-free rate adjusts such that the demand for liquid assets equals the total supply of liquid assets *B*, which is fixed and comes from the household sector and the government.⁷

(v) The cross-sectional distribution of households and aggregate shocks evolve according to their laws of motion. For idiosyncratic state variables, denote Θ × B the sigma-algebra associated with the Cartesian product of the discrete set of productivity and the compact set of net bond holdings, and (Θ, B) a subset of that sigma-algebra. The law of motion for the distribution is given by:

$$\lambda_{t+1}\left(\tilde{\Theta},\tilde{\mathcal{B}}\right) = \int_{\Theta\times\mathcal{B}} Q_{\overline{\phi}_{t},z_{t}}\left(\left(\theta,b\right),\left(\tilde{\Theta},\tilde{\mathcal{B}}\right)\right) d\lambda_{t}\left(\theta,b\right)$$

where $Q_{z_{t},\overline{\phi}_{t}}\left(\left(\theta,b\right),\left(\tilde{\Theta},\tilde{\mathcal{B}}\right)\right) = \mathbf{1}\left\{b_{t}'(\theta,b)\in\tilde{\mathcal{B}}\right\}\sum_{\theta'\in\tilde{\Theta}}\Pi_{\theta}\left(\theta'|\theta\right)$ (15)

The transition function $Q_{z_t,\overline{\phi_t}}$ depends on individual productivity and net bond holdings, and on aggregate productivity and credit credit supply. Aggregate shocks make the distribution time-varying.

Solution. The model has heterogeneous households, incomplete markets, and aggregate risk, so it is solved numerically. Given policy functions and the cross-sectional distribution, the wage can be solved for analytically using labor market clearing:

$$w_{t} = \alpha z_{t} N_{t}^{\alpha - 1} = \alpha z_{t} \left(\int \theta n\left(\theta, b\right) d\lambda_{t}\left(\theta, b\right) \right)^{\alpha - 1}$$
(16)

It is directly affected by aggregate productivity shocks z_t , and indirectly by credit supply shocks through through their effects on the distribution λ_t .

3 Decomposition of Precautionary Savings Motives

This section uses the model of Section 2 to decompose the contributions of idiosyncratic and aggregate risks to households' precautionary savings.

⁷Savers with positive net bond holdings hold both risk-free bonds issued by indebted households and the government. To close the model, these bonds are assumed to be perfect substitutes, which is a standard assumption (e.g., Guerrieri and Lorenzoni (2017)).

3.1 Model-Based Measure of Precautionary Motives

The decomposition of precautionary motives consists of three steps. First, variables and functions in the model (e.g., policy functions, the cross-sectional distribution) are approximated using projection methods to generate a discrete model with a finite number of parameters. Second, the stationary steady state of the discrete model *without* aggregate shocks is computed. The resulting solution of the model is global and nonlinear with respect to idiosyncratic state variables, and it exactly holds without aggregate shocks. Third, the solution of the discrete model is perturbed with respect to aggregate shocks around its stationary steady state where they are zero, to generate a new solution for the stochastic steady state of the model *with* aggregate shocks and rational expectations.

The model-based measure of precautionary motives relies on comparing the stochastic steady states of the model with first-order and second-order perturbations with respect to aggregate shocks. In the *first-order*, equilibrium variables depend linearly on the lagged values of shocks and states, and there certainty equivalence. Only the level of aggregate shocks affects household behavior, but not their volatility. Thus, there are only idiosyncratic precautionary motives, but no aggregate precautionary motives. In the *second-order*, variables depend nonlinearly on the lagged values of shocks and states, and they depend on the volatility of aggregate shocks. There are both idiosyncratic and aggregate precautionary motives, and the difference with the first-order solution provides a measure of the latter. Appendix A describes the algorithm and the numerical solution of the model in detail.

Step 1: Projection. Equilibrium conditions (i)-(v) are stacked in a multivariate vectorvalued function $\mathcal{F}(.)$ that represents the nonlinear system of equations defining the equilibrium:

$$\mathbb{E}_{t}\left[\mathcal{F}\left(\mathbf{y}_{t}, \mathbf{y}_{t+1}, \mathbf{x}_{t}, \mathbf{x}_{t+1}, \boldsymbol{\varepsilon}_{t+1}^{\phi}, \boldsymbol{\varepsilon}_{t+1}^{z}\right)\right] = 0$$
(17)

Variables are sorted into non-predetermined and predetermined variables. The vector of non-predetermined variables **y** contains projection coefficients for policy functions (labor supply n_t , net bond holdings b_{t+1} , and consumption c is obtained from the budget constraint), the risk-free rate r_t , the wage w_t , aggregate consumption C_t and employment N_t .⁸ The vector of predetermined variables **x** contains the histogram weights used to project the cross-sectional distribution λ_t , and aggregate shocks to productivity z_t and borrowing constraints $\overline{\phi_t}$.

⁸Policy functions are approximated using linear splines.

Step 2: Stationary steady state. Solving for the stationary steady state of the model without aggregate shocks is equivalent to solving the nonlinear system of equations defined by:

$$\mathcal{F}\left(\mathbf{y},\mathbf{y},\mathbf{x},\mathbf{x},0,0\right) = 0. \tag{18}$$

This is a more challenging problem than solving for the typical consumption-savings allocation because of flexible labor supply and endogenous risk-free rates, wages, and government taxes. To address it, I use a variant of the policy-time iteration method (e.g., Elenev, Landvoigt and Nieuwerburgh (2021)), which combines Broyden's numerical equation solver and automatic differentiation to compute exact derivatives (see Appendix A).

Step 3: Perturbations. The last step starts from the global and nonlinear solution for the stationary steady state of the model without aggregate shocks. Then, denote η the perturbation parameter that scales the quantity of aggregate risk in the economy around the stationary steady state of the model. The solution of the expectation difference equation 17, which defines the equilibrium with aggregate risk $\mathbb{E}_t [\mathcal{F}(.)] = 0$, has the following form:

$$\mathbf{x}_{t+1} = \mathbf{h} \left(\mathbf{x}_{t}, \eta \right) + \eta \begin{pmatrix} \mathbf{0} \\ \boldsymbol{\varepsilon}_{t+1}^{\phi} \\ \boldsymbol{\varepsilon}_{t+1}^{z} \end{pmatrix}$$

$$\mathbf{y}_{t} = \mathbf{g} \left(\mathbf{x}_{t}, \eta \right)$$
(19)

where $\mathbf{h}(.,\eta)$ and $\mathbf{g}(.,\eta)$ are nonlinear vector-valued functions, which relate future predetermined variables and non-predetermined variables to current predetermined variables and depend on aggregate risk. Variables are then written in deviations from the stationary steady state for simplicity.

First-order. The dynamics of the model with idiosyncratic precautionary motives but no aggregate motives is given by a first-order perturbation of the system of equations 19 with respect to aggregate shocks:

$$\widehat{\mathbf{x}_{t+1}} = \mathbf{h}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}} + \eta \begin{pmatrix} \mathbf{0} \\ \boldsymbol{\varepsilon}_{t+1}^{\phi} \\ \boldsymbol{\varepsilon}_{t+1}^{z} \end{pmatrix}$$

$$\widehat{\mathbf{y}_{t}} = \mathbf{g}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}}.$$
(20)

I solve for the vectors of coefficients $\mathbf{h}_{\mathbf{x}}(\mathbf{x}, 0)$ and $\mathbf{g}_{\mathbf{x}}(\mathbf{x}, 0)$, which linearly relate future predetermined variables and non-predetermined variables to current predetermined variables and the level of aggregate shocks, using the gensys algorithm (Sims (2001)). This step involves computing the Jacobian of the multivariate vector-valued function $\mathcal{F}(.)$.

Second-order. The dynamics of the model with both idiosyncratic and aggregate precautionary motives is given by a second-order perturbation of equations 19 with respect to aggregate shocks:

$$\widehat{\mathbf{x}_{t+1}} = \mathbf{h}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}} + \frac{1}{2} \mathbf{h}_{\mathbf{xx}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}}^{2} + \underbrace{\frac{1}{2} \mathbf{h}_{\eta\eta} \left(\mathbf{x}, 0 \right) \eta^{2}}_{\text{nonlinearity}} + \eta \begin{pmatrix} \mathbf{0} \\ \boldsymbol{\varepsilon}_{t+1}^{\phi} \\ \boldsymbol{\varepsilon}_{t+1}^{z} \end{pmatrix}$$

$$\widehat{\mathbf{y}_{t}} = \mathbf{g}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}} + \underbrace{\frac{1}{2} \mathbf{g}_{\mathbf{xx}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}}^{2}}_{\text{nonlinearity}} + \underbrace{\frac{1}{2} \mathbf{g}_{\eta\eta} \left(\mathbf{x}, 0 \right) \eta^{2}}_{\text{no certainty equivalence}} .$$

$$(21)$$

Solving for the new vectors of coefficients implies computing the Jacobian and the Hessian of the function $\mathcal{F}(.)$. Since typical second-order perturbation methods for representative agent models (e.g., Schmitt-Grohe and Uribe (2008)) cannot be applied due to the high dimension of the equation system, I apply a series of steps based on the gensys2 algorithm (Kim, Kim, Schaumburg and Sims (2008)) in order to reduce its dimension.

3.2 Interpretation

The difference between the second-order and the first-order perturbations of the model provide two new measures of interest.

Aggregate precautionary savings. First and foremost, the coefficient vectors $\mathbf{h}_{\eta\eta}(\mathbf{x}, 0)$ and $\mathbf{g}_{\eta\eta}(\mathbf{x}, 0)$ provide a measure of aggregate precautionary motives as the departure of the model from certainty equivalence with respect to aggregate shocks. With certainty equivalence, only the *level* of aggregate productivity and credit supply shocks affects household behavior. Without it, the *volatility* of aggregate shocks also modifies households' policy functions and aggregates, including risk-free rates and wages, in a way that reflect their precautionary behavior.

In the first-order, deviations from the stationary steady state of the model are on average zero, so that the stochastic steady state of a long-run simulation of this model coincides with the stationary steady state. In the second-order, the stochastic steady state of the model permanently differs from its stationary steady state as long as the volatility of aggregate shocks is positive and the second-order coefficients are nonzero. It can be interpreted as the average level of a long-run simulation of the model with aggregate shocks.

One limitation to this measure of precautionary savings is that it assumes partial "bounded rationality" from households by abstracting from the effects of higher-order moments of aggregate shocks, whose estimation requires unrealistically large computing power. However, this assumption is empirically realistic given how households make aggregate forecasts in practice (e.g., Das et al. (2019)), and I show below that it generates negligible market clearing errors.

The main results in the paper are based on this decomposition for the two sources of aggregate risk in the model (see Section 5). In particular, the aggregate precautionary motive due to credit supply shocks generates higher aggregate liquid savings, lower aggregate debt, and lower risk-free rates in the second-order perturbation of the model than in the first-order. This difference is measured by the coefficient vectors $\mathbf{h}_{\eta\eta}$ and $\mathbf{g}_{\eta\eta}$, and depends on the volatility of the shock. A decomposition of these coefficients then shows that the changes in aggregate variables result from individual changes in the policy functions of "middle-class" households who accumulate more precautionary savings to hedge against these shocks.

Typical models of household precautionary savings assume fixed borrowing constraints that can be subject to unexpected shocks leading to households suddenly and massively deleveraging (e.g., Guerrieri and Lorenzoni (2017)). In this economy, instead, households know the stochastic processes governing borrowing constraints but not their future realizations. This creates a precautionary motive in the Euler equation of savers:

$$c_t^{-\gamma} = \beta(1+r_t)\mathbb{E}_t \left[c_{t+1}^{-\gamma} \right] + \mu_t.$$
(22)

A high multiplier on current borrowing constraints implies higher savings and a lower risk-free rate. Iterating the equation forward, the average and higher-order moments of future multipliers on borrowing constraints $\{\mu_s\}_{s \ge t}$, including their volatility due to aggregate shocks, further increase savings and lower the rate.

Nonlinearity. Second, the coefficient vectors $\mathbf{h}_{xx}(\mathbf{x},0)$ and $\mathbf{g}_{xx}(\mathbf{x},0)$ capture the nonlinear dependence of future predetermined variables and non-predetermined variables on current predetermined variables. These terms help jointly match household balance sheet and macroeconomic moments (see Section 5) and improve the empirical fit of the model (see Section 6). One limitation is that it is not an exact measure of nonlinearity in the model as it abstracts from third-order terms and higher. Nevertheless, I show below that it provides an accurate approximation of the dynamics of the economy as the resulting

market clearing errors are negligible.

4 Calibration

The model is calibrated to match household balance sheet and macroeconomic moments in the stationary steady state of the economy. Table 1 summarizes the calibration, which is split between externally and internally calibrated parameters. To evaluate precautionary savings, I focus on matching the level and cross-sectional moments of liquid assets and unsecured debt held by U.S. households, as well as government taxes and transfers and the real risk-free rate, which are key for households' incentives to save. One period is a quarter. Average household income is normalized to 1.

Parameter	Explanation	Value	Target/source
Internal			
β	Discount factor	0.9925	Risk-free rate $= 1.80\%$ (FRB)
B	Liquid asset supply	6	Liquid assets/GDP = 1.78 (FRB)
$\overline{\phi}$	Borrowing constraints: aggregate average	2.6	Unsecured debt/GDP = 0.18 (FRB)
ρ_{ϕ}	Borrowing constraints: aggregate persistence	0.99	Risk-free rate persistence $= 0.65$ (FRB)
σ_{ϕ}	Borrowing constraint: aggregate volatility	0.025	Risk-free rate volatility = 1.9% (FRB)
η'	Curvature disutility of work	2	Frisch elasticity $= 1/2$
ψ	Disutility of work	11.5	Income normalization $Y = 1$
External			
γ	Risk aversion	5	See text
α	Labor share	2/3	Labor share of output $= 2/3$
$ au_1(heta)$	Tax progressivity by productivity	[0.05, 0.13, 0.17, 0.20, 0.28]	Tax distribution by income (CPS)
$T\left(heta ight)$	Government transfers by productivity	[1, 0.43, 0.24, 0.17, 0.13]	Transfer distribution by income (CPS)
$\phi\left(heta ight)$	Borrowing constraints: idiosyncratic	[1, 1.03, 1.06, 1.08, 2.33]	Debt distribution by income (SCF)
$ ho_{ heta}$	Idiosyncratic productivity persistence	0.977	Wage persistence
$\sigma_{ heta}$	Idiosyncratic productivity volatility	0.12	Wage volatility
$ ho_z$	Aggregate productivity persistence	0.86	TFP persistence
σ_z	Aggregate productivity volatility	0.0128	TFP volatility
$ ho_{\phi z}$	Productivity and borrowing constraint correlation	0.5	Debt-income correlation $= 0.9$ (FRB, BEA)

Table 1: Calibration: main parameters

Notes: One period is a quarter, targets are annualized. Taxes, transfers, and borrowing constraints depend on household idiosyncratic productivity $[\underline{\theta} = \theta_1, ..., \overline{\theta} = \theta_5]$. Sources: Current Population Survey, Survey of Consumer Finances, Federal Reserve Board, Bureau of Economic Analysis.

4.1 Internal Parameters

The following parameters are chosen to match household balance sheet and macroeconomic moments. **Discount factor.** The discount factor $\beta = 0.9925$ is chosen to match the average real riskfree rate. It is measured in the data as the average of annual Treasury Inflation Indexed long-term yields of 1.80% between 2000 and 2018 (Federal Reserve Board, H.15 Selected Interest Rates). A similar value of 1.86% is obtained for the one-year Treasury bill rate net of one-year survey expectations of inflation (GDP deflator from the Survey of Professional Forecasters) in a longer sample between 1973 and 2014.

Households with the lowest productivity level have a 20% lower discount factor, which makes them more impatient to consume. This assumption generates a large fraction of borrowing-constrained low-productivity households. Without discount factor heterogeneity, these households would be less constrained than slightly more productive households because they receive larger progressive transfers, which woud be at odds with the data.

Liquid assets. Households accumulate savings using liquid assets and borrow with unsecured debt by buying and selling one-period risk-free bonds. The demand for liquid assets arises from households' intertemporal consumption smoothing and precautionary motives. To close the model, the supply of liquid assets comes from the household sector and the rest of the economy consisting of the government.⁹ In the data, liquid assets are defined as the sum of all deposits and securities held directly by households, which are computed in the Flow of Funds (Federal Reserve Board, Z.1, table B.100) as the sum of inventory change (line 9), Treasury currency (16), checkable deposits and currency (19), time and savings deposits (20), money market fund shares (21), open market paper (24), and Treasury securities (25). In the model, the total supply of liquid assets *B* is chosen to match the resulting value of liquid assets to GDP of 1.78. I set *B* = 6 and obtain a value of 1.73.

Borrowing constraints: aggregate component. The aggregate component of borrowing constraints $\overline{\phi}$ captures the effect of credit supply. It is chosen to match the ratio of unsecured debt to GDP of 0.18 in the Flow of Funds. Unsecured debt is computed as total household liabilities minus mortgage debt (table B.100, line 34). I set $\overline{\phi} = 2.6$ and obtain a value of 0.23.

The persistence ρ_{ϕ} and the volatility σ_{ϕ} of the aggregate component of borrowing constraints are estimated by indirect inference to match the persistence and the volatility of

⁹See, e.g., Huggett (1993), Aiyagari (1994), and Guerrieri and Lorenzoni (2017) for the same approach in general equilibrium models of household savings in a closed economy. Government debt in excess of households' liquid savings is held by investors outside the model, which is a plausible assumption for the U.S. where debt held by the public is on average 40% of GDP.

the real risk-free rate in the data, which are respectively equal to 1.50% and 1.90% in annual terms (Federal Reserve Board). The benefits of this approach are twofold. First, it guarantees that precautionary motives in the model are consistent with the properties of the risk-free rate, and therefore are well estimated. Second, it circumvents the difficulty of identifying time series for credit supply shocks in the data. I obtain quarterly values of $\rho_{\phi} = 0.99$ and $\sigma_{\phi} = 0.025$, and I use sensitivity analyses to show that these estimates are well identified (Appendix Figure 7). In comparison, they imply less risk than available estimates directly based on changes in credit card limits, and therefore they provide a lower bound on aggregate precautionary savings in the model.¹⁰

The correlation between the aggregate component of borrowing constraints and productivity $\rho_{\phi z}$ is chosen to match the correlation between outstanding total consumer credit owned and securitized (Federal Reserv board, G.19 Consumer Credit) and linearly detrended personal income (Bureau of Economic Analysis), which is about 0.90. I obtain a quarterly value of $\rho_{\phi z} = 0.50$, because the model already endogenously generates a positive correlation between household debt and income (and the risk-free rate), though it is not large enough to quantitatively match the procyclical behavior of household debt.

Labor supply. The curvature of the disutility of work hours $\eta = 2$ generates a Frisch elasticity of labor supply of 1/2, in line with empirical estimates (Whalen and Reichling (2017)). The level of the disutility of work hours $\psi = 11.5$ is chosen to normalize average household income Y = 1.

4.2 External Parameters

The remaining parameters are externally calibrated.

Risk aversion. The coefficient of relative risk aversion γ is equal to 5, a standard value in finance. Higher values imply larger precautionary savings and a lower risk-free rate. Given the set of internally calibrated parameters, this value provides the best fit of the model with the data. It is lower than in Favilukis et al. (2017), who also analyze a general equilibrium model with stochastic borrowing constraints with a risk aversion of 8.

Borrowing constraints: idiosyncratic component. To calibrate the distribution of household debt by income given the aggregate component of borrowing constraints, I map the

¹⁰Because of data limitations, few papers directly estimate a stochastic process for households' borrowing constraints. Using recent data for credit card limits, Fulford (2015) estimates that credit supply shocks are close to permanent and about ten times more volatile as my estimates.

income distribution in the model by constructing the corresponding productivity groups in the data.¹¹ In equilibrium, 6.25% of households are in the lowest productivity group with θ_1 , 25% are with θ_2 , 37.5% with θ_3 , 25% with θ_4 , and 6.25% in the highest productivity group with θ_5 .

Unsecured credit is computed in the Survey of Consumer Finances as total household debt minus the total value of debt secured by primary residence (including mortgages and HELOC) and the total value of debt for other residential properties. This leaves other lines of credit, credit card balances, installment loans (including education and auto loans), and other debt. In the five income groups, the average (median) values for unsecured household debt are respectively \$130,190 (0), \$17,150 (450), \$25,340 (6,512), \$141,190 (9,431), and \$302,920 (0). While these values are equilibrium objects, the goal of the idiosyncratic component of borrowing constrains $\phi(.)$ is to capture empirical evidence that individual borrowing limits tend to increase with income. To reflect it, I smooth and normalize the distribution of $\phi(\theta_1), ..., \phi(\theta_5)$ by assuming that individual borrowing limits are uniformly distributed. The resulting values of [1, 1.03, 1.06, 1.08, 2.33] measure relative borrowing limits by income and reflect the SCF data. They allow the model to match the dispersion of household debt by income, which is key for households' response to borrowing constraint fluctuations.

Government taxes and transfers. Similarly, progressive taxes and transfers are chosen to replicate the distribution of taxes and transfers by household income in the data (Congressional Budget Office (2006), Exhibit 18). In the five income groups, the average total transfers to non-elderly households are respectively \$15,200, \$6,600, \$3,700, \$2,600, and \$2,000. Normalizing these values, I obtain the transfer function $T(\theta) = [1, 0.43, 0.24, 0.17, 0.13]$. Transfers represent 6.9% of average income. To match that share in the model, I apply a constant multiplicative factor to the transfers for all income groups. In the five income groups, the average taxes are respectively \$2,600, \$6,500, \$11,800, \$19,700, and \$68,100. The resulting slope of the tax function by productivity level is $\tau_1(\theta) = [0.05, 0.13, 0.17, 0.20, 0.28]$.

Idiosyncratic productivity risk. The persistence and volatility of the discretized productivity process are chosen to match the persistence and volatility of wages in **Kopecky** and **Suen** (2010), of 0.977 and 0.12. The larger volatility of idiosyncratic productivity in the model approximates exogenous unemployment risk, which is mot modeled explicitly

¹¹A similar exercise is done in Jappelli (1990).

for simplicity.¹²

I use the estimate of Storesletten, Telmer and Yaron (2004) to calibrate countercyclical income risk. In the data, the standard deviation of individual income increases by 0.09 for a 1.5% change in output from peak to trough. In the model, a negative one standard deviation shock to aggregate productivity *z* lowers steady state output by -0.5%. To match the data, I assume that such a shock increases the volatility of idiosyncratic income by 0.09/(1.5/0.3)=0.03.

Aggregate productivity risk. Estimates for the persistence and volatility of aggregate productivity are taken from Fernald (2014). They are respectively equal to $\rho_z = 0.86$ and $\sigma_z = 0.0128$.

4.3 Model Fit

The model matches key empirical moments of households' balance sheets. The upper panel of Table 2 reports targeted moments, and the lower panel reports untargeted moments. By virtue of the calibration, the model replicates well the ratios of aggregate liquid assets and unsecured debt to income. Interestingly, it also matches inequality in the wealth distribution captured by the ratio of average to median wealth. It generates a realistic cross-sectional distribution with a large mass of borrowing-constrained households and a decreasing fraction of households with larger asset levels. In particular, it matches the share of constrained households computed in The Pew Charitable Trusts (2015) as those reporting to be without savings, which is close to the 21% share of hand-to-mouth households in Kaplan and Violante (2014).

5 Main Results

This section presents the main results on the impact of aggregate risk on precautionary savings in three steps. First, I present average estimates of the impact of fluctuations in households' borrowing constraints and aggregate productivity. The impact of aggregate shocks can be permanent even if shocks themselves are temporary. Second, I decompose the impact of aggregate risk across households and highlight the role of "middle class" savers. Third, I analyze their real implications for the business cycle. Both fluctuations

¹²Adding an unemployment state when solving for households' labor supply in the system of equilibrium conditions is computationally challenging because the system would be non-differentiable at that point. In the model, zero labor supply can be interpreted as unemployment, but there is no involuntary unemployment.

Table 2: Household balance sheets

Statistic:	Data	Model
Aggregate liquid assets/aggregate income	1.78	1.73
Aggregate unsecured debt/aggregate income	0.18	0.23
Mean/median wealth	4.60	4.90
Share of constrained households	0.33	0.35

Notes: The upper panel reports targeted moments, the lower panel reports untargeted moments. One period is a quarter, targets are annualized. Source: Boar, Gorea and Midrigan (2021), The Pew Charitable Trusts (2015).

in borrowing constraints and aggregate productivity are key to jointly match household balance sheet and macroeconomic moments.

5.1 Impact of Aggregate Risk

Average estimates. As shown in Section 3, the difference between equilibrium coefficients in the second- versus the first-order solution of the model provide a measure of aggregate precautionary savings motives that relies on departure from certainty equivalence. Table 3 reports the corresponding averages of equilibrium variables. Since the impact of aggregate productivity risk is negligible, the table focuses on the volatility of households' borrowing constraints. The second column reports variables in the stationary steady state without aggregate shocks, and the third and fourth columns report deviations from these variables in the stochastic steady state of the model for two values of the volatility of households' borrowing constraints, which correspond to a long sample that ends before the Great Recession and to the post-Great Recession period.

Several key assumptions ensure that these estimates are a lower bound, as discussed previously: the risk-free rate is endogenous; households save and borrow at the same rate; households can increase their labor supply and receive government transfers in bad times; and the volatility of borrowing constraints matches is lower than available estimates to match the volatility of the risk-free rate.

Borrowing constraint risk. Aggregate precautionary motives imply that aggregate shocks can have a permanent effect even if they are themselves temporary. Small fluctuations in aggregate borrowing constraints lower the ratio of household debt to GDP and the risk-free rate. They are too small to affect other equilibrium variables. Larger fluctuations based on the volatility of borrowing constraints in the post-Great Recession period lower household debt to GDP by 45% and the risk-free rate by 25.4% relative to the station-

ary steady state. This generates a low-debt environment with a low risk-free rate similar to the post-Great Recession period. Due to households' precautionary behavior, aggregate fluctuations in borrowing constraints have permanent effects on household balance sheets and the macroeconomy even if aggregate shocks are themselves temporary. A large volatility $\sigma_{\phi} = 0.10$ decreases consumption, output, and profits by about 1.5%, and employment by 2.5%. The wage increases by 0.9% to prevent households from further decreasing their work hours. Even these large effects are still below the values that a model calibrated to match the volatility of credit card limits with $\sigma_{\phi} = 0.25$ would imply.¹³ Therefore, the contribution of aggregate fluctuations in borrowing constraints to precautionary savings is sizable.

The impact of fluctuations in borrowing constraints on precautionary behavior increases in risk aversion γ , in the aggregate persistence ρ_{ϕ} and volatility σ_{ϕ} of the constraints, in the dispersion of individual borrowing constraints by income $\phi(\theta)$, and in countercyclical income risk.

Variable	Stationary steady state	$\sigma_{\phi} = 0.025$ (1973-2005)	$\sigma_{\phi} = 0.10$ (post-2005)
Interest rate	2.397%	-0.1%	-25.4%
Wage	1.491	0%	+0.9%
Profits	0.333	0%	-1.5%
Employment	0.447	0%	-2.5%
Output	1.000	0%	-1.6%
Consumption	1.000	0%	-1.6%
Debt/GDP	0.229	-0.4%	-45%

Table 3: Average impact of aggregate risk

Notes: One period is a quarter, targets are annualized. Comparative statics analysis holding other parameters fixed. Columns 3 and 4 are percentage deviations from steady state values (column 2).

Productivity risk. Strikingly, the same comparative statics with respect to the volatility of aggregate productivity shocks shows that their contribution to precautionary behavior is negligible. This is true even for highly volatile shocks. In particular, increasing by a factor of five only slightly decreases the risk-free rate by 0.10%, while other variables remain unchanged. Aggregate fluctuations in borrowing constraints contribute much more to precautionary behavior than productivity risk. This is because borrowing constraints themselves are more volatile than aggregate productivity and because they have a higher

¹³See, e.g., Fulford (2015), who assumes permanent shocks to borrowing constraints, close to the persistence of $\rho_{\phi} = 0.99$ in the model.

impact for a given volatility since they decisively affect households' ability to insure by borrowing.

Decomposition. The model highlights three precautionary motives. (1) The standard motive due to idiosyncratic income risk. It arises because of the prudence property of the utility function, u''(.) > 0 (Kimball (1990)): the volatility of income increases future expected marginal utility because of Jensen's inequality, which implies a decrease in current consumption and an increase in savings in households' Euler equation. It also arises because the combination of income shocks and borrowing constraints hampers consumption smoothing for some households. That effect is stronger for households at or near their constraints. It has the largest impact. It increases liquid savings by 280% (from 31%) to 119% of GDP, in annual terms) and decreases the risk-free rate by 6 percentage points (from 8.40% to 2.40%) compared to an economy without idiosyncratic risk where they are only determined by intertemporal substitution. In such an economy, the risk-free rate is given by $1/\beta - 1$. (2) An aggregate financial motive due to credit supply shocks, whose impact further increases liquid savings by 45% (from 119% to 173% of GDP) and lowers the risk-free rate by 0.6 pp (from 2.40% to 1.80%). This impact is sizable and cannot be ignored when analyzing precautionary savings; it is likely even greater given that my estimates are a lower bound. Interestingly, it decreases long-run consumption by 1.6% on average, leading to substantial costs of aggregate fluctuations, which is a significant departure from existing settings that focuses on productivity shocks (e.g., Lucas (1987)). (3) An aggregate real motive due to productivity shocks. This motive acts in the same way as the first motive, since aggregate productivity shocks also result in labor income fluctuations. The difference with the first motive is that aggregate productivity changes across periods are smaller and less persistent than changes in idiosyncratic productivity, and they impact households uniformly. As a result, their impact on household behavior is negligible. They affect employment, but they have close to zero impact on liquid savings and the risk-free rate. Unlike in the case of (2), these results are consistent with existing low estimates of the cost of business cycles. Taken together, my results for (2) and (3) nuance received wisdom that aggregate fluctuations have low costs for household and thus would not lead to any precautionary behavior.

Mechanism. What explains the effect of fluctuations in borrowing constraints? Because of the departure from certainty equivalence, households anticipate future shocks to borrowing constraints based on their stochastic process and they insure every period against future binding borrowing constraints. The stochastic steady state of the economy shifts in

response to a higher precautionary motive: on average, households accumulate less debt, more liquid assets, and the risk-free rate is lower. For a low volatility of the credit shocks, the large negative financial response only mildly affects real variables, because prices are flexible and adjust to clear markets, leaving quantities relatively unaffected.

When the volatility of shocks to borrowing constraints is large, financial conditions affect real variables because of a composition effect that induces less productive households to work more and more productive ones to work less, and because of a wealth effect on labor supply. This gives rise to an economy with persistently low output, employment, debt, and risk-free rate. In terms of long-run aggregate consumption, aggregate volatility in borrowing constraints leads to much larger costs of business cycles than aggregate productivity shocks.

Credit shocks also affect the economy dynamically, from the time they hit households' borrowing constraints to the time they revert back to their steady state values. Nonlinearity in the second-order solution of the model amplify the economy's response to a tightening of borrowing constraints, by capturing the different responses of households depending on their net bond holdings. While nonlinearities are close to zero for aggregate productivity shocks, which affect all households identically, this is not the case for shocks to borrowing constraints. Therefore, aggregate prices and quantities do not respond linearly to aggregate shocks as they would in a representative agent economy–this property is called near-aggregation in Krusell and Smith (1998)–and the full cross-sectional distribution of households affects equilibrium variables.

5.2 Household Heterogeneity

How does aggregate risk affect precautionary behavior at the household level? Figure 1 decomposes the effect of aggregate precautionary motives across bond holding levels for a household with the median income in the post-Great Recession economy where $\sigma_{\phi} = 0.10$. The blue lines depict the household's policy functions in the stationary steady state without aggregate risk (first-order), and the orange line depicts them in the stochastic steady state with aggregate risk (second-order).

Aggregate fluctuations in borrowing constraints lead households to consume less, save and work more, and achieve higher precautionary savings, as revealed by the differences between the two lines. This effect is especially large for "middle class" households with some debt $b \leq 0$ but not the highest debt levels in the economy. In annual terms, such households have debt levels that are close or slightly higher than the average ratio of debt to GDP of 0.23.

Replicating the same comparison for productivity groups θ_1 to θ_5 shows that this result also holds for income. The effect of aggregate risk on precautionary behavior is larger for low income households with productivity θ_2 and θ_3 , while it is lower for poorer households with θ_1 and for richer households with θ_4 and θ_5 . "Middle class" households benefit less than the poorest households from the progressivity of government taxes and transfers, and they have less liquid assets to start with than richer households. Therefore, their precautionary motive dominates their impatience to consume. Interestingly, households with the highest income are the only ones to consume slightly more in the presence of aggregate risk. The reason is the lower equilibrium risk-free rate, which increases the incentive to consume more of the consumption good and leisure.



Figure 1: Impact of aggregate risk on household choices

Notes: Policy function of the median income household by net bond holdings, starting at the corresponding steady state credit limit $\overline{\phi}\phi(\theta)$ for: consumption, next period bond holdings, savings and labor supply. Order 1 (blue) vs order 2 (orange).

5.3 Real Implications

What are the real implications of these sizable aggregate precautionary savings motives? I conclude this section by showing that the combination of aggregate productivity and credit supply shocks is key to explain volatility in household balance sheet and macroe-conomic moments.

Business cycle volatility. Table 4 reports the result from a variance decomposition designed to quantify the contributions of aggregate productivity and borrowing constraint shocks to business cycle volatility. This computation uses the nonlinear laws of motion of the economy (see Appendix A.5 for details).

Credit supply shocks are responsible for a larger fraction of volatility than aggregate productivity, especially in household debt and the risk-free rate. The model assigns more than half of the volatility of output, consumption, and employment to credit supply shocks, except for wages, which scale one by one with aggregate productivity. This effect holds even in the absence of price rigidity, because changes in interest rates lead employment, hence output, to react negatively to fluctuations in borrowing constraints. This effect is amplified by the fact that the persistence and volatility of credit supply shocks are higher than for aggregate productivity.

Variable:	Borrowing constraints	Productivity
Interest rate	59%	41%
Employment	52%	48%
Wage	21%	79%
Profits	59%	41%
Output	59%	41%

Table 4: Contributions of aggregate risks to business cycle volatility

Notes: Variance decomposition: shares of the variance of variables in the first column accounted for by credit (second column) and aggregate productivity shocks (third column): risk-free rate, employment, wages, profits, output, taxes. Variance shares are computed by bootstrap, as the Monte-Carlo average of the variance decompositions of generalized forecast errors at a large forecasting horizon (H = 1000 periods). Computations use N = 500 simulations.

Mechanism. To explain these results, I compare the economy's impulse response functions (IRF) to aggregate productivity and borrowing constraint shocks, in the first- and in the second-order approximations of the model. The amplification of the responses of financial and real variables to credit shocks is key to replicate their short-run comovements, especially around the Great Recession. *Nonlinear impulse response functions.* Figure 2 plots the economy's response to a one standard deviation shock to borrowing constraints, under the linear dynamics with certainty equivalence and without aggregate risk (order 1), and the nonlinear dynamics with aggregate risk (order 2). Aggregates are computed using the time-varying path of individual policy functions and of histogram weights. Deviations are from the stationary steady state. As shown in the figure, variables stay longer at lower values following a tightening of households' borrowing constraints. In the first-order approximation, households respond to the levels of current and expected future shocks, but not to their volatility, because of certainty equivalence. Policy functions, the cross-sectional distribution and prices respond linearly to shocks, and are linear functions of their lagged values. In the second-order approximation, households anticipate aggregate shocks, whose volatility enter linearly. In addition, the economy evolves nonlinearly with respect to the level of shocks, with variables being linear and quadratic functions of their lagged values.

Amplification. Accounting for aggregate precautionary motives and nonlinearity significantly amplifies the response of aggregates to a credit shock.¹⁴ Amplification is larger for the risk-free rate and household debt. The initial response is amplified by a factor of 5 for debt to GDP, of 4 for the risk-free rate, of 1.5 for consumption, output and profit, and of 1.4 for the wage. While the risk-free rate decrease (in response to a one-time shock) is short-lived, other variables stay persistently low. The sharp decline in the rate causes consumption and employment to rebound (simultaneously, profits slightly increase and the wage slightly decreases). However, the rebound is short-lived, and the large persistence of the credit shock that induces borrowing constraints to stay persistently low, further decreases consumption and employment (the decrease in the risk-free rate) cannot offset the quantity restriction imposed by tighter borrowing constraints, which mechanically force constrained households to hold less debt.

Borrowing constraints. Borrowing constraints for all households are tightened, but lower income households are able to borrow less than richer ones, reflecting idiosyncratic differences in their ability to borrow. As a result, constrained households are forced to reduce their debt and increase their net bond holdings, thereby decreasing their consumption of goods and leisure. They trade-off working more to smooth consumption against the disutility of labor. Debt to GDP decreases, and stays persistently low, mostly because of the large persistence of the credit shock. The decrease in total consumption results from the composition of low-income constrained households decreasing their spending,

¹⁴These results contrast with recent findings for firms and aggregate productivity shocks (Winberry (2018)).



Figure 2: Nonlinear impulse responses to borrowing constraint shocks: household balance sheets

Notes: Impulse response functions to a one standard deviation credit shock. Credit constraints (upper left panel) as a fraction of annual steady state output (upper left panel) are for θ_1 (lowest line) to θ_5 households (highest line). Other panels plot IRF in the 1st versus the 2nd order approximation of the model. Initial period: deterministic steady state. One period is a quarter, variables are annualized.

and richer unconstrained households increasing theirs because they earn a lower return on their savings. The decrease in the risk-free rate allows to balance a larger savings demand from the former with a lower demand from the latter, and to clear the savings market. Aggregate employment decreases, causing a decline in output (see below). As implied by the economy's resource constraint, consumption falls with output.

Employment. Households also insure against shocks by adjusting their labor supply. Figure 3 plots the response of employment, which results from the composition of less productive, constrained households increasing their hours to smooth consumption when they are forced to deleverage, and of more productive, unconstrained households who consume more leisure as they decrease their savings (wealth effect). In addition, the sharp

decline in the risk-free rate creates an intertemporal substitution effect, which induces all households to consume more leisure in the current period.

The sign of the labor response depends on which effects dominates. In the model, output declines mainly because more productive agents work less, despite less productive agents working more. This result is due to stochastic borrowing constraints and departs from economies with fixed credit limits, where employment increases after a credit shock (e.g., Guerrieri and Lorenzoni (2017)). The reason is that the credit supply shock is mean-reverting. In economies with fixed credit limits and deterministic shocks that are perfectly foreseen by households, constrained households and those expecting to be constrained in the near future choose to work more. This effect is dampened in the model, because agents expect the credit shock to mean-revert. Even with flexible prices, small mean-reverting credit shocks can generate larger negative employment responses than unanticipated, large and permanent deleveraging episodes.

Without that effect, hours worked would increase after a tightening of credit supply, resulting in precautionary savings increasing less, and risk-free rates decreasing less. This result only holds in a setting with heterogeneous agents. With a representative agent, employment increases even after a credit supply shock, because this is the only way the agent can save more (e.g. Jones, Midrigan and Philippon (2022)).

Aggregate productivity shocks. Instead, when the economy hit by an aggregate productivity shock, borrowing constraints stay fixed. Because all households become less productive, output and profits drop. The wage decreases because of a lower marginal productivity of labor, and households supply more hours to compensate for the decrease in their income. The lower wage induces a higher labor demand from firms. Summing up these two effects, aggregate employment increases. To smooth their consumption when their incomes decline, lower income unconstrained households issue more debt, while richer households increase their savings. This increases wealth inequality. The risk-free rate increases to clear the savings market, with a larger demand for debt by low-income households and a larger demand for liquid assets by the rich.

6 Empirical Test of the Model

How plausible is the channel documented in the previous section? To test the empirical validity of the model, I investigate its ability to match the time series for aggregate consumption and the risk-free rate in the post-Great Recession period.





Notes: Impulse response functions to a one standard deviation credit shock for aggregate employment (left panel) and wages (right panel), in order 1 vs order 2. Initial period: deterministic steady state. One period is a quarter.

6.1 Post-Great Recession Puzzle

This test is a stringent requirement for two reasons. First, its goal is to match the full time series of the main variables in addition to their averages (see the calibration of Section 4). Second, the variations in consumption and the risk-free rate after the recession cannot be typically replicated in standard models, which gives rise to a quantitative puzzle that the model should address. After the recession, aggregate consumption recovered but real risk-free rates remained low. In models with productivity risk and no fluctuations in borrowing constraints, a growing and stable consumption path should imply a higher risk-free rate, as households seek to borrow to smooth consumption over time (intertemporal substitution motive). Conversely, a lower risk-free rate should induce households to consume more early on instead of saving and increasing their future consumption.

The following results demonstrate that the model is successful at matching the post-Great Recession data and addressing this puzzle thanks to the combination of aggregate risk in productivity and borrowing constraints risks.

6.2 Model Fit

I apply a particle filter to estimate the path of structural shocks to productivity and borrowing constraints that generate the observed paths for consumption and the real risk-free rate after the recession. Aggregate precautionary behavior due to departure from certainty equivalence and the model nonlinearity allow to match the data even in times of high volatility. Appendix A.6.2 details the estimation procedure.

The time series for the real risk-free rate is measured as the 5-Year Treasury Inflation-Indexed Securities Constant Maturity rate (Federal Reserve Board). The series for aggregate consumption is measured using Real Personal Consumption Expenditures (Bureau of Economic Analysis), for which I compute the deviation from its initial value in the sample.¹⁵

Figure 4 reports the results for the model fit. The model is successful at matching both time series in the post-recession sample. It exactly matches the dynamics of aggregate consumption and it closely matches the risk-free rate, except in the middle of the sample where it is slightly higher.

As shown in Figure 5, the model also closely matches the dynamics of the ratio of debt to GDP and employment, which were not targeted by the estimation. In the data, house-hold debt is measured as Total Revolving Credit Owned and Securitized (Federal Reserve Board). Employment is measured as Civilian Employment-Population Ratio (Bureau of Labor Statistics).¹⁶ The model replicates the hump-shaped dynamics of household debt to GDP, which starts with the run-up to the crisis until 2008, and then the decrease and eventually the increase in credit around 2015. The model overstates the decrease in household debt and employment in the last part of the sample, for which the large persistence of borrowing constraint shocks and the wealth effect on labor supply may be responsible.

6.3 Structural Aggregate Shocks

Figure 6 plots the time series of estimates for the structural aggregate shocks. These estimates suggest that: (i) Productivity only fell during the Great Recession itself (2008-2009) and then quickly reverted to its pre-recession level; (ii) The aggregate component of households' borrowing constraints kept decreasing until the middle of the decade, and stayed persistently low throughout the post-recession period.

¹⁵Because consumption is non-stationary, I detrend the series with a Hodrick-Prescott filter, and subtract the resulting initial value to normalize the detrended deviation to zero in the first period of the sample.

¹⁶The model has a continuum of measure 1 of households, so N is the ratio of employed to the entire population.





Notes: Risk-free rate (annualized, percentage) and consumption deviation from 2006Q3 value, predicted by particle filtering in the nonlinear version of the model (solid line) vs the data (dashed line). N = 20,000 particles simulated. The risk-free rate (left axis, blue) is measured as the 5-Year Treasury Inflation-Indexed Security, Constant Maturity rate, not seasonally adjusted (source: Board of Governors of the Federal Reserve System). Consumption (right axis, orange) is measured as Real Personal Consumption Expenditures, Billions of Chained 2009 Dollars, quarterly, seasonally adjusted (source: U.S. Bureau of Economic Analysis). Quarterly sample, 2006Q3-2017Q2. One period is a quarter.

With these two sources of aggregate risk, the model is able to simultaneously replicate the increase in consumption and the decrease in the risk-free rate. The decrease in the risk-free from 2.5% to -1.5% (in annual terms) results from a large tightening of the aggregate component of households' borrowing constraints. Compared to its initial value, aggregate credit supply decreases by more than 15%, and remains consistently low throughout the sample. This tightening prevents households from using debt to smooth consumption fluctuations. The resulting precautionary motive is exacerbated by the short-lived drop in aggregate productivity, which induces constrained households to deleverage and save quickly. To clear the savings market, the risk-free rate decreases and stays persistently low as long as borrowing constraints remain tight.

Credit shocks are a slow-moving variable, which gives rise to low-frequency changes in debt and in the risk-free rate. In contrast, changes in aggregate productivity are more frequent and they track changes in aggregate consumption. A 2% decline in aggregate productivity hits the economy at the beginning of the recession and helps match the



Figure 5: Out-of-sample predictions: household debt and employment

Notes: Debt/GDP (left panel) and employment (right panel) implied by risk-free rate and consumption data, recovered by particle filtering. Model (solid line) versus data (dashed line). Results for N = 20,000 particles simulated. Variables are in log-deviations from their 2006Q3 values. Quarterly sample, 2006Q3-2017Q2. One period is a quarter.

decrease in aggregate consumption. Productivity then reverts to its pre-recession level within less than two years, and continues to increase by more than 2% by the end of the sample. Finally, the model estimates for structural productivity shocks strongly correlate with empirical measures of total factor productivity, which provides further external validation (see Appendix Figure 9). The same is true for the aggregate component of households' borrowing constraints, which closely track lending standards in the data (see Appendix Figure 10). Interestingly, however, model estimates imply tighter borrowing constraints than bank lending standard alone, which suggests that the latter may not capture the complete landscape of credit conditions faced by households.



Figure 6: Estimates of structural borrowing constraint and productivity shocks

Notes: Structural credit (left axis, blue) and aggregate productivity shocks (right axis, orange) estimated by particle filtering. Variables are in log-deviations from their 2006Q3 values. Quarterly sample, 2006Q3-2017Q2. One period is a quarter.

7 Conclusion

This paper uses a structural model of household savings to estimate the effect of a new *ag-gregate* precautionary motive, which arises from economy-wide fluctuations in productivity and the tightness of borrowing constraints. This motive is quantitatively important, at odds with received wisdom about the low costs of aggregate fluctuations for households. It leads to a large decrease in the risk-free rate and an increase in savings, especially for middle-class households, which nuances the recent focus of economists on the top and the bottom of the wealth distribution. This motive is critical to jointly explain household balance sheet and macroeconomic moments. In particular, the combination of persistently tight borrowing constraints and rising productivity since the Great Recession can explain why the recovery of consumption despite low levels of household debt and the risk-free rate-a puzzle for standard models with fixed borrowing constraints. These estimates imply tighter borrowing constraints than standard measures of credit conditions such as bank lending standard, which suggests that model-based measures may be useful to capture the complete credit landscape faced by households.

References

- Acharya, Viral V., Katharina Bergant, Matteo Crosignani, Tim Eisfert, and Fergal Mc-Cann, "The Anatomy of the Transmission of Macroprudential Policies," *The Journal of Finance*, 2022, 77 (5), 2533–2575.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?," *Quarterly Journal of Economics*, September 2018, 133 (1).
- Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, August 1994, *109* (3), 659–684.
- **Baker, Scott R.**, "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data," *The Journal of Political Economy*, August 2018, *126* (4).
- **Bertaut, Carol C., Michael Haliassos, and Michael Reiter**, "Credit Card Debt Puzzles and Debt Revolvers for Self Control," *Review of Finance*, 2009, *13*, 657–692.
- **Boar, Corina**, "Dynastic Precautionary Savings," *The Review of Economic Studies*, 2021, *88* (6), 2735–2765.
- _ , Denis Gorea, and Virgiliu Midrigan, "Liquidity Constraints in the U.S. Housing Market," The Review of Economic Studies, 2021.
- **Carroll, Christopher D.**, "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 1997, 112 (1), 1–55.
- ____, "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, June 2006, *91*, 312–320.
- and Andrew A. Samwick, "The nature of precautionary wealth," *Journal of Monetary Economics*, 1997, 40 (1), 41–71.
- Cherry, Susan, Erica Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru, "Government and Private Household Debt Relief during COVID-19," *Brookings Papers on Economic Activity*, forthcoming.
- **Congressional Budget Office**, "The Distribution of Federal Spending and Taxes in 2006," Technical Report 2006.

- **Das, Sreyoshi, Camelia M. Kuhnen, and Stefan Nagel**, "Socioeconomic Status and Macroeconomic Expectations," *The Review of Financial Studies*, 2019, 33 (1), 395–432.
- **De Nardi, Mariacristina, Eric French, and John B. Jones**, "Why Do the Elderly Save? The Role of Medical Expenses.," *Journal of Political Economy*, 2010, *118* (1), 39–75.
- **Deaton, Angus**, "Savings and Liquidity Constraints," *Econometrica*, September 1991, 59 (5), 1221–1248.
- Elenev, Vadim, Tim Landvoigt, and Stijn Van Nieuwerburgh, "A Macroeconomic Model with Financially Constrained Producers and Intermediaries," *Econometrica*, 2021, *89* (3), 1361–1418.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh, "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 2017, 125 (1), 140–223.
- **Fernald, John G.**, "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Technical Report, Federal Reserve Bank of San Francisco 2014.
- **Fulford, Scott L.**, "How important is variability in consumer credit limits?," *Journal of Monetary Economics*, 2015, 72, 42–63.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption Over the Life Cycle," *Econometrica*, 2002, 70 (1), 47–89.
- **Gross, David B. and Nicholas S. Souleles**, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *The Quarterly Journal of Economics*, February 2002, 117 (1), 149–185.
- **Guerrieri, Veronica and Guido Lorenzoni**, "Credit Crises, Precautionary Savings, and the Liquidity Trap," *The Quarterly Journal of Economics*, 2017, 132 (3), 1427–1467.
- **Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese**, "Income Risk, Borrowing Constraints, and Portfolio Choice," *The American Economic Review*, March 1996, *86* (1), 158–172.
- Hartzmark, Samuel M., "Economic Uncertainty and Interest Rates," *The Review of Asset Pricing Studies*, February 2016, 6 (2), 179–220.

- Heaton, John and Deborah J. Lucas, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *The Journal of Political Economy*, June 1996, 104 (3), 443–487.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance," *The Journal of Political Economy*, April 1995, *103* (2), 360–399.
- Huggett, Mark, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 1993, 17 (5-6), 953–969.
- Jappelli, Tullio, "Who is Credit Constrained in the U.S. Economy?," *The Quarterly Journal of Economics*, February 1990, 105 (1), 219–234.
- Jermann, Urban and Vincenzo Quadrini, "Macroeconomic Effects of Financial Shocks," *American Economic Review*, February 2012, 102 (1), 238–71.
- Jones, Callum, Virgiliu Midrigan, and Thomas Philippon, "Household Leverage and the Recession," *Econometrica*, 2022, 90 (5), 2471–2505.
- **Kaplan, Greg and Giovanni L. Violante**, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, July 2014, *8*2 (4), 1199–1239.
- **Kennickell, Arthur and Annamaria Lusardi**, "Disentangling the Importance of the Precautionary Saving Motive," Technical Report 10888, National Bureau of Economic Research 2004.
- Kim, Jinill, Sunghyun Kim, Ernst Schaumburg, and Christopher A. Sims, "Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models," *Journal of Economic Dynamics and Control*, 2008, 32, 3397–3414.
- **Kimball, Miles S.**, "Precautionary Savings in the Small and in the Large," *Econometrica*, 1990, *58*, 53–73.
- Klein, Paul, "Using the generalized Schur form to solve a multivariate linear rational expectations model," *Journal of Economic Dynamics and Control*, 2000, 24, 1405–1423.
- **Kopecky, Karen and Richard Suen**, "Finite State Markov-chain Approximations to Highly Persistent Processes," *Review of Economic Dynamics*, July 2010, *13* (3), 701–714.
- Krusell, Per and Anthony Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, *106* (5), 867–896.

- Lanne, Markku and Henri Nyberg, "Generalized Forecast Error Variance Decomposition for Linear and Nonlinear Multivariate Models," Oxford Bulletin of Economics and Statistics, 2016, 78 (4), 595–603.
- Lucas, Robert E., Models of Business Cycles, Wiley-Blackwell, 1987.
- Ludvigson, Sydney, "Consumption and credit: a model of time-varying liquidity constraints," *The Review of Economics and Statistics*, August 1999, *81* (3), 434–447.
- Mian, Atif, Amir Sufi, and Emil Verner, "Household Debt and Business Cycle Worldwide," *The Quarterly Journal of Economics*, November 2017, 132 (4), 1755–1817.
- _, Kamalesh Rao, and Amir Sufi, "Household Balance Sheets, Consumption, and the Economic Slump," The Quarterly Journal of Economics, 2013, 128 (4), 1687–1726.
- Parker, Jonathan A. and Bruce Preston, "Precautionary Savings and Consumption Fluctuations," *American Economic Review*, September 2005, 95 (4), 1119–1143.
- Pflueger, Carolin, Emil Siriwardane, and Adi Sunderam, "Financial Market Risk Perceptions and the Macroeconomy," *Quarterly Journal of Economics*, 2020, 135 (3), 1443–1491.
- **Reiter, Michael**, "Solving Heterogeneous-Agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control*, 2009, 33 (3), 649–665.
- Schmitt-Grohe, Stephanie and Martin Uribe, "Solving dynamic general equilibrium models using a second-order approximation to the policy function," *Journal of Economic Dynamics and Control*, 2008, 28, 755–775.
- Sims, Christopher A., "Solving linear rational expectations models," *Computational Economics*, 2001, 20 (1-2), 1–20.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *The Journal of Political Economy*, 2004, 112 (3), 695–717.
- The Pew Charitable Trusts, "Americans' Financial Security," Brief March 2015.
- **Whalen, Charles and Felix Reichling**, "Estimates of the Frisch Elasticity of Labor Supply: A Review," *Eastern Economic Journal*, 2017, 43, 37–42.
- Winberry, Thomas, "A Method for Solving and Estimating Heterogeneous Agent Macro Models," *Quantitative Economics*, November 2018, 9 (3).
- Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, May 1989, 104 (2), 275–298.

Appendix

A Decomposition of Precautionary Savings

A.1 Algorithm

1. (a) Variables are index by time *t* to denote their dependence on aggregate states $(\overline{\phi_t}, z_t, \lambda_t)$. The distribution of households over $\Theta \times \mathcal{B}$ is approximated as a histogram by a finite number of mass points on the Cartesian product of $\Theta = \{\theta_i\}_{i=1}^{N_{\theta}}$ and a fine grid $\{b_j\}_{j=1}^{N_{b}^{f}}$. $\Phi_t(\theta_i, b_j)$ denotes the fraction of households with productivity θ_i and net bond holdings b_j . Its evolution is implied by policy functions according to:

$$\Phi_{t+1}\left(\theta_{i'}, b_{j'}\right) = \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{b}^{f}} \Pi_{\theta}\left(\theta_{i'}|\theta_{i}\right) \omega_{i,j,j',t} \times \Phi_{t}\left(\theta_{i}, b_{j}\right)$$
where
$$\omega_{i,j,j',t} = \begin{cases} \frac{b' - b_{j'-1}}{b_{j'} - b_{j'-1}} & \text{if } b_{t}'\left(\theta_{i}, b_{j}\right) \in \left[b_{j'-1}, b_{j}\right] \\ \frac{b_{j'+1} - b'}{b_{j'+1} - b_{j'}} & \text{if } b_{t}'\left(\theta_{i}, b_{j}\right) \in \left[b_{j'}, b_{j'+1}\right] \\ 0 & \text{otherwise,} \end{cases}$$
(23)

where $b_{j'-1}, b_{j'}, a_{j'+1}$ are asset points on the fine grid that bracket the value of next period assets implied by the policy function. ω depend on t because policy functions depend on the aggregate state, i.e. $b'_t(\theta_i, b_j) = b'(\theta_i, b_j; \overline{\phi_t}, z_t, \lambda_t)$. For instance, if credit shocks $\overline{\phi_t}$ are low, tightening borrowing constraints, this distorts and shifts upwards the function b'(.) because households are forced to save more, which through its impact on ω results in less mass on low asset values.

(b) Household saving and labor supply policy functions are interpolated using linear splines with respectively N_b and N_n knots. Households' saving function b'(.) is characterized by a critical level of assets χ_{θ} at which their borrowing constraints start binding, which depends on productivity. For every $\theta \in \Theta$, let $b_{\theta,j} = \chi_{\theta} + x_j$, with $0 = x_1 < ... < x_{N^b}$ denote the splines' knots for b' at which households' Euler equations hold with equality. For $b \leq \xi_{\theta}$, savings $b'(\theta, b) = -\overline{\phi}\phi(\theta) h(\theta)$ are determined by the borrowing limit ($\overline{\phi}_t = \overline{\phi}$ in the deterministic steady state). It defines the collocation nodes at which we force households' optimality conditions to hold to solve for policy functions. For a given aggregate state ($\overline{\phi}, z, \Phi$), the saving function is finitely represented

by $N_{\theta} \times (N_b + 1)$ coefficients giving the value of savings at the knots and the threshold below which households are constrained. So is the labor supply function, with $N_{\theta} \times N_n$ values at the knots for labor (which may differ from the knots for savings). The consumption function at the saving knots is backed out from the budget constraint:

$$c_{t}\left(\theta, b_{\theta,j}\right) = b_{\theta,j} + \left(1 - \tau_{1}\left(\theta\right)\right) w_{t} \theta n_{t}\left(\theta, b_{\theta,j}\right) + T\left(\theta\right) + \pi_{t} - \tau_{0t} - \frac{b_{t}'\left(\theta, b_{\theta,j}\right)}{1 + r_{t}}$$
(24)

(c) Equilibrium conditions for the discrete model are listed below. The first set of equations and the following two involve predetermined variables: the histogram weights (because weights should sum to 1, we keep only track of the number of weights minus 1), the credit and aggregate productivity shocks. The next sets of equations involve jump variables: the asset price, aggregate labor demand, the wage, profits, aggregate output, aggregate consumption, and the (discretized versions of) policy functions for labor and savings (including values of coefficients at knot points and borrowing constraint thresholds). The inclusion of some variables among jump variables, whose dynamics we want to solve for, is not strictly speaking necessary (it is the case for aggregate labor demand, the wage, profits, aggregate output and aggregate consumption). Their equation counterparts are definitional, and their values can be backed out from the other jump variables without including them explicitly in the equilibrium system of equation. However, including them makes the system dynamics better behaved numerically, because it provides more information to the code when taking derivatives with automatic differentiation.

In words, these equations are: the laws of motion for the distribution, credit and aggregate productivity; the market clearing conditions for assets and labor; the definitions of aggregate output,¹⁷ consumption, the wage and profits; the intratemporal optimality condition for households' labor supply, and the intertemporal optimality condition for savings/consumption (Euler equation). In the Euler equations, the *t*-conditional expectation is about t + 1 values of aggregate shocks (next period borrowing constraints and wage influence current

¹⁷Given the goods market clearing condition implied by the remaining equilibrium conditions and Walras law, aggregate output should equal aggregate consumption. During simulations, I recompute aggregate output fully nonlinearly using the policy functions and distributions implied by the perturbed solution, as $Y_t = z_t K^{1-\alpha} \left[\sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{b}^{f}} \theta_i n_t (\theta_i, b_j) \Phi_t (\theta_i, b_j) \right]^{\alpha}$. I check that the deviation from goods market clearing is close to 0.

decisions), and is taken with respect to their values at *t*.

$$\begin{split} \Phi_{t+1}\left(\theta_{i'}, b_{j'}\right) &- \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{\theta}^{f}} \Pi_{\theta}\left(\theta_{i'}|\theta_{i}\right) \left(\omega_{i,j,j',t} \Phi_{t}\left(\theta_{i}, b_{j}\right)\right) = 0, \quad i' \in [|1, N_{\theta}|], j' \in [|1, N_{\theta}^{f}|] \\ \log \overline{\phi}_{t+1} &- \log \overline{\phi} - \rho_{\phi}(\log \overline{\phi}_{t} - \log \overline{\phi}) - \epsilon_{t+1}^{q} = 0 \\ \log z_{t+1} - \rho_{z} \log z_{t} - \epsilon_{t+1}^{z} = 0 \\ \begin{pmatrix} \epsilon^{\phi} \\ \epsilon^{z} \end{pmatrix}^{iid} \mathcal{N}\left(\mathbf{0}, \begin{pmatrix} \sigma_{\phi}^{2} & \sigma_{\phi}\sigma_{z}\rho_{\phi z} \\ \sigma_{\phi}\sigma_{z}\rho_{\phi z} & \sigma_{z}^{2} \end{pmatrix}\right) \\ B - \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{\theta}^{f}} b_{j}\Phi_{t+1}\left(\theta_{i}, b_{j}\right) = 0 \\ N_{t} - \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{\theta}^{f}} \theta_{i}n_{t}\left(\theta_{i}, b_{j}\right) \Phi_{t}\left(\theta_{i}, b_{j}\right) \\ Y_{t} - \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{\theta}^{f}} c_{t}\left(\theta_{i}, b_{j}\right) \Phi_{t}\left(\theta_{i}, b_{j}\right) \\ \mathcal{C}_{t} - \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{\theta}^{f}} c_{t}\left(\theta_{i}, b_{j}\right) \Phi_{t}\left(\theta_{i}, b_{j}\right) \\ w_{t} = \alpha z_{t}\left(\frac{1}{N_{t}}\right)^{1-\alpha} \\ \pi_{t} = (1 - \alpha) z_{t} K^{1-\alpha} N_{t}^{\alpha} \\ (1 - \tau_{1}\left(\theta_{i}\right)) w_{t} \theta_{i} c_{t}\left(\theta_{i}, b_{j}\right)^{-\gamma} - \psi n_{t}\left(\theta_{i}, b_{j}\right)^{\eta} = 0, \quad i \in [|1, N_{\theta}|], j \in [|1, N_{\theta$$

Denote as \mathbf{y}_t the $6 + N_\theta \times (N_n + N_b + 1)$ vector of current jump (control) variables. Denote as \mathbf{x}_t the $N_\theta \times N_b^f - 1 + 2$ vector of current state (predetermined) variables. Equilibrium conditions are stacked in a multivariate, vector-valued function $\mathcal{F}(.)$ that represents the nonlinear system of equations that defines the equilibrium:

$$\mathbb{E}_{t}\left[\mathcal{F}\left(\mathbf{y}_{t}, \mathbf{y}_{t+1}, \mathbf{x}_{t}, \mathbf{x}_{t+1}, \boldsymbol{\epsilon}_{t+1}^{q}, \boldsymbol{\epsilon}_{t+1}^{z}\right)\right] = 0$$
(26)

2. Solving for the deterministic steady state of the economy (without aggregate shocks) amounts to finding **y**, **x** that solve the following system of equation, which has as many unknowns as equations:

$$\mathcal{F}(\mathbf{y}, \mathbf{y}, \mathbf{x}, \mathbf{x}, 0, 0) = 0 \tag{27}$$

In theory, it could be solved directly using a nonlinear equation solver. In practice, there is no guarantee that numerical equation solvers will converge when we use projection methods to approximate policy functions. In addition to solving the households' consumption problem, the difficulty comes from having endogenous labor supply, endogenous government taxes, and solving for two equilibrium prices (wage and interest rate). I also solve for the value of the disutility of labor ψ that

normalizes steady state output *Y* to 1.

Therefore, to make the problem more stable, I use the following variant of policy time iteration. First, given a guess for x and y,¹⁸ compute government taxes for all agents. Given taxes and the guess, solve for households' labor supply policy. Given that policy, solve then for households' savings policy. Using the policy functions, compute the implied stationary distribution (using an eigenvector method), and the new taxes. The process is repeated until policy functions converge. I use Broyden's method every time a numerical solver is needed, and automatic differentiation to compute exact derivatives. Since the convergence of the numerical solver is not guaranteed under any initial guess and parameter combination, I calibrate the steady state of the model with a homotopy method. That is, I slowly change parameters until the target is reached, starting from a combination under which the model steady state is easily computed. If needed, I modify the state space boundaries over that process.

- (a) Start with a guess for the risk-free rate and labor demand $(p^{(0)}, N^{(0)}, \psi^{(0)})$, for policy function values $(\mathbf{b}^{\prime(0)}(.), \mathbf{n}^{(0)}(.))$, and the cross-sectional distribution $\mathbf{\Phi}^{(0)}(.)$ (it is only needed to compute the first iterate of government taxes). It is easier to solve for the risk-free rate and labor demand demand, and back out the interest rate 1/p 1 and the wage (from the firm's optimal labor choice) than solving directly for the latter. Thus having $(p^{(0)}, N^{(0)})$ is equivalent to having $(r^{(0)}, w^{(0)})$.
- (b) Given those, use the endogenous grid method of Carroll (2006) to iterate backwards on the household's optimality conditions (the Euler and the labor intratemporal equations), and obtain a new guess for policy functions that will be supplied to the nonlinear policy solver solving the household's problem, (b'⁽¹⁾(.), n⁽¹⁾(.)). This requires computing endogenous government taxes (fixed every period because we are at the steady state), which is why we need a guess for the cross-sectional distribution.
- (c) The guess for prices is supplied to a second nonlinear solver wrapped around the policy solver, which solves for the prices clearing the savings and the labor market, and for the disutility of labor normalizing steady state output to 1. Within the price solver, I ensure that prices and labor disutility are positive

¹⁸A good guess is obtained by using the endogenous grid method of Carroll (2006) to iterate backwards on the household's optimality conditions, starting from any feasible guess.

 $(p^{(n)}, N^{(n)}, \psi^{(n)} > 0)$, and the stability condition $\beta/p^{(n)} \leq 1$ holds at every iteration *n*. The following steps occur within the price solver, and their iterates start at n = 1.

- (d) Given the exogenous law of motion for idiosyncratic income and the policy functions, compute the associated stationary distribution of households $\Phi^{(1)}(.)$ (I use the eigenvector method). Also compute the wage and profits from the firm's optimality condition: $w^{(1)} = \alpha \left(\frac{1}{N^{(1)}}\right)^{1-\alpha}$, and $\pi^{(1)} = (1-\alpha) \left(N^{(1)}\right)^{\alpha}$. Then, given prices, policy functions and the distribution, compute endogenous government taxes $\tau_0^{(1)}$.
- (e) Given prices, profits, taxes, and savings policies $\mathbf{b}'^{(1)}(.)$, solve the household's labor supply equation (using $\mathbf{n}^{(1)}(.)$ as a guess), and denote $\mathbf{n}^{(2)}(.)$ the new labor supply policy. It should always be non-negative. Here I use a nonlinear equation solver with Broyden's method, and supplies it with the Jacobian of the system of intratemporal equations. Here and later, derivatives are computed exactly with automatic differentiation, implemented with Julia's ForwardDiff package.
- (f) Back out the associated consumption function from the budget constraint. If it has a non-positive entry at a point in the state space, adjust $\mathbf{n}^{(2)}(.)$ at that point such that the household consumes $c_{min} = 0.001$. This step helps with convergence of the solver when solving for savings in the next step.
- (g) Given prices, profits, taxes and the new labor policy $\mathbf{n}^{(2)}(.)$, solve the house-hold's Euler equation (using $\mathbf{b}'^{(1)}(.)$ as a guess), and denote $\mathbf{b}'^{(2)}(.)$ the new savings policy. Use the same solver as for labor.
- (h) This completes one iterate in the loop solving for policy functions given prices. If the new policy functions $(\mathbf{n}^{(2)}(.)\mathbf{b}'^{(2)}(.))$ are close enough to the previous ones $(\mathbf{n}^{(2)}(.)\mathbf{b}'^{(2)}(.))$, then stop and we have solved the household's problem given prices $(p^{(0)}, N^{(0)}, \psi^{(0)})$. Otherwise, iterate on steps (d)-(g). That is, given $(p^{(0)}, N^{(0)}, \psi^{(0)})$ (hence the same wages and profits), compute new government taxes $\tau_0^{(n+1)}$. Then solve for new policy functions $(\mathbf{n}^{(n+1)}(.)\mathbf{b}'^{(n+1)}(.))$, compare them to the previous ones $(\mathbf{n}^{(n)}(.)\mathbf{b}'^{(n)}(.))$, and stop when they are close enough. This completes the solution of the household's problem given prices.
- (i) Using the law of motion of the exogenous income shock and the optimal sav-

ings function, compute the stationary distribution $\Phi^{(2)}$. Use it with policy functions to compute aggregate values for savings, labor supply and output. The price solver then chooses new values for prices and disutility of labor, $(p^{(1)}, N^{(1)}, \psi^{(1)})$, to solve the following three equations:

$$B - \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{b}^{f}} b_{j} \Phi_{t+1}^{(2)} (\theta_{i}, b_{j}) = 0$$

$$N^{(1)} - \sum_{i=1}^{N_{\theta}} \sum_{j=1}^{N_{b}^{f}} \theta_{i} n (\theta_{i}, b_{j}) \Phi^{(2)} (\theta_{i}, b_{j})$$

$$Y^{(1)} - 1 = 0 \Leftrightarrow (N^{(1)})^{\alpha} - 1 = 0$$
(28)

- (j) Then go back to step (a) with the new prices, and iterate until convergence, i.e. policy functions and the stationary distribution have converged, and the three equations are satisfied. We then obtain prices, policy functions and a distribution that solve the model in the deterministic steady state.
- 3. Do a first- and a second-order perturbation of the discrete model around its steady state. The solutions to the equilibrium expectational difference equation $\mathbb{E}_t [\mathcal{F}(.)] = 0$ are of the following form (Schmitt-Grohe and Uribe (2008)):

$$\mathbf{x}_{t+1} = \mathbf{h} \left(\mathbf{x}_{t}, \eta \right) + \eta \begin{pmatrix} \mathbf{0} \\ \epsilon_{t+1}^{q} \\ \epsilon_{t+1}^{z} \end{pmatrix}$$

$$\mathbf{y}_{t} = \mathbf{g} \left(\mathbf{x}_{t}, \eta \right)$$
(29)

where η is the perturbation parameter (there is only one such parameter) scaling the amount of aggregate uncertainty in the economy. The goal is to solve for approximations of the functions **h**, **g**.

(a) For the first-order approximation of the model, several methods can be used. I check existence and uniqueness, and verify that I obtain identical results using Sims' gensys (Sims (2001)) and Klein's methods (Klein (2000)), commonly used in the macro literature. I briefly describe the input and the output of Klein's method because it has a clear interpretation in terms of jump and predetermined variables. We solve for a first-order approximation of **g**, **h**. Writing variables in deviations from their steady state values (denoted as \hat{x}, \hat{y}) and linearizing equilibrium conditions around 0 (where variables equal their steady state values), we obtain

$$\mathcal{F}_{\mathbf{y}_{t}}\widehat{\mathbf{y}_{t}} + \mathcal{F}_{\mathbf{y}_{t+1}}\mathbb{E}_{t}\left[\widehat{\mathbf{y}_{t+1}}\right] + \mathcal{F}_{\mathbf{x}_{t}}\widehat{\mathbf{x}_{t}} + \mathcal{F}_{\mathbf{x}_{t+1}}\mathbb{E}_{t}\left[\widehat{\mathbf{x}_{t+1}}\right] + \mathcal{F}_{\boldsymbol{\varepsilon}_{t+1}^{q}}\mathbb{E}_{t}\left[\widehat{\boldsymbol{\varepsilon}_{t+1}}^{q}\right] + \mathcal{F}_{\boldsymbol{\varepsilon}_{t+1}^{z}}\mathbb{E}_{t}\left[\widehat{\boldsymbol{\varepsilon}_{t+1}}^{z}\right] = 0$$
(30)

where the derivatives of \mathcal{F} are evaluated at the steady state. They are submatrices of the Jacobian of \mathcal{F} , computed exactly with automatic differentiation. \hat{y}, \hat{x} terms are vectors, so their (matrix) products with the derivative matrices of \mathcal{F} are vectors. The Jacobian is a matrix of dimension

$$\left\{ \begin{bmatrix} N_{\theta} \times N_{b}^{f} - 1 + 2 \end{bmatrix} + \begin{bmatrix} 6 + N_{\theta} \times (N_{n} + N_{b} + 1) \end{bmatrix} \right\} \\ \times \left\{ 2 \times \begin{bmatrix} N_{\theta} \times N_{b}^{f} - 1 + 2 \end{bmatrix} + 2 \times \begin{bmatrix} 6 + N_{\theta} \times (N_{n} + N_{b} + 1) \end{bmatrix} + 2 \right\}$$

First-order approximations of the solution have the following form:

$$\widehat{\mathbf{x}_{t+1}} = \mathbf{h}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}} + \eta \begin{pmatrix} \mathbf{0} \\ \epsilon_{t+1}^{q} \\ \epsilon_{t+1}^{z} \\ \epsilon_{t+1}^{z} \end{pmatrix}$$
(31)
$$\widehat{\mathbf{y}_{t}} = \mathbf{g}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}}$$

(b) For the second-order approximation of the model, I do a second-order approximation of equilibrium conditions around the steady state. It involves the Hessian of *F*, a large three-dimensional array computed by automatic differentiation, of dimension:

$$\left\{ \begin{bmatrix} N_{\theta} \times N_{b}^{f} - 1 + 2 \end{bmatrix} + \begin{bmatrix} 6 + N_{\theta} \times (N_{n} + N_{b} + 1) \end{bmatrix} \right\} \\ \times \left\{ 2 \times \begin{bmatrix} N_{\theta} \times N_{b}^{f} - 1 + 2 \end{bmatrix} + 2 \times \begin{bmatrix} 6 + N_{\theta} \times (N_{n} + N_{b} + 1) \end{bmatrix} + 2 \right\}^{2}$$

The second-order approximation of the solution has the form:

$$\widehat{\mathbf{x}_{t+1}} = \mathbf{h}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}} + \frac{1}{2} \mathbf{h}_{\mathbf{x}\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}}^{2} + \frac{1}{2} \mathbf{h}_{\eta\eta} \left(\mathbf{x}, 0 \right) \eta^{2} + \eta \begin{pmatrix} \mathbf{0} \\ \boldsymbol{\varepsilon}_{t+1}^{q} \\ \boldsymbol{\varepsilon}_{t+1}^{z} \end{pmatrix}$$
(32)
$$\widehat{\mathbf{y}_{t+1}} = \mathbf{g}_{\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}} + \frac{1}{2} \mathbf{g}_{\mathbf{x}\mathbf{x}} \left(\mathbf{x}, 0 \right) \widehat{\mathbf{x}_{t}}^{2} + \frac{1}{2} \mathbf{g}_{\eta\eta} \left(\mathbf{x}, 0 \right) \eta^{2}$$

where the terms equal to zero (in h_{η} , g_{η} , $h_{x\eta}$, $h_{\eta x}$, $g_{x\eta}$, $g_{\eta x}$) were canceled. $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ terms are vectors, $\mathbf{g}_{\mathbf{x}}$, $\mathbf{h}_{\mathbf{x}}$ terms are matrices, $\mathbf{h}_{\mathbf{xx}}$, $\mathbf{g}_{\mathbf{xx}}$ are 3-dimensional arrays, and $\mathbf{h}_{\eta\eta}$, $\mathbf{g}_{\eta\eta}$ are vectors. Thus products of $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ vectors with first-order deriva-

tive matrices are matrix products, those with second-order arrays are tensor products, and those with η are simple constant times vectors products. I use Kim et al. (2008)'s gensys2 method to solve for the unknown coefficients. Schmitt-Grohe and Uribe (2008) propose instead to solve for the second-order coefficients in a linear system of equations involving the Jacobian and the Hessian of \mathcal{F} , and the first-order coefficients. While most papers with representative agent models use this method, it is not tractable in a setting with heterogeneous agents where the cross-sectional distribution is discretized as a histogram, since it involves constructing and inverting a matrix whose dimensions increases exponentially with the number of state variables. gensys2 allows to reduce the dimensionality of the system of equation to solve by applying a sequence of linear operations to the original system (Schur and singular value decompositions).

A.2 Stochastic Steady State

To compute the deviations of the stochastic steady state from the deterministic one, I compute a fixed point of the pruned laws of motion of the economy.¹⁹ The impulse response functions (IRF) to credit and aggregate productivity shocks are computed by feeding the laws of motion with nonzero innovations in the first period and iterating on them. I verify that market-clearing errors are close to zero over the simulated paths (Appendix A.4).

Pruning essentially computes first-order projections of second-order terms, based on a first-order expansion of the conditional expectation of the system's deviation from steady state, according to the following steps.

First, gensys2 solves a linearly transformed system, where original variables $(\hat{\mathbf{x}} \ \hat{\mathbf{y}})'$ that solve $\mathbb{E}_t [\mathcal{F}(.)] = 0$ are replaced by $Z' (\hat{\mathbf{x}} \ \hat{\mathbf{y}})'$, where Z is a square, non-singular matrix. To simplify notation, denote the transformed variables as $(\hat{\mathbf{x}} \ \hat{\mathbf{y}})'$ too. The second-order solution to the transformed system has the form (see the paper for details):

$$\widehat{\mathbf{x}_{t+1}} = F_1 \widehat{\mathbf{x}_t} + F_2 \eta \mathbf{f} \mathbf{f}_{t+1} + F_3 \eta^2 + \frac{1}{2} F_{11} \widehat{\mathbf{x}_t}^2 + F_{12} \widehat{\mathbf{x}_t} \mathbf{f} \mathbf{f}_{t+1} \eta + \frac{1}{2} F_{22} \eta^2 \mathbf{f} \mathbf{f}_{t+1}^2$$

$$\widehat{\mathbf{y}_t} = \frac{1}{2} M_{11} \widehat{\mathbf{x}_t}^2 + M_2 \eta^2$$
(33)

The presence of cross-derivative terms in the transformed solution does not contradict

¹⁹I use pruned laws of motion to alleviate the well-known problem that iterating on second-order laws of motion gives rise to higher-order terms that do not increase the accuracy of the approximation and are likely to lead to explosive paths. Pruning essentially computes first-order projections of second-order terms, based on a first-order expansion of the conditional expectation of the system's deviation from steady state.

their absence in the original solution, since they can be canceled by *Z*. Then, it implies that for s > 0:

$$\mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s}} \right] = F_{1} \mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s-1}} \right] + F_{3} \eta^{2} + \frac{1}{2} F_{11} \mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s-1}}^{2} \right] + \frac{1}{2} F_{22} \eta^{2} \Omega_{s}$$

$$= F_{1} \mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s-1}} \right] + F_{3} \eta^{2} + \frac{1}{2} F_{11} \left(\mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s-1}} \right]^{2} + \Sigma_{s-1} \right) + \frac{1}{2} F_{22} \eta^{2} \Omega_{s}$$

$$\mathbb{E}_{t} \left[\widehat{\mathbf{y}_{t+s}} \right] = \frac{1}{2} M_{11} \mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s}}^{2} \right] + M_{2} \eta^{2}$$

$$= \frac{1}{2} M_{11} \left(\mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s}} \right]^{2} + \Sigma_{s} \right) + M_{2} \eta^{2}$$

$$\Sigma_{s+1} = \eta^{2} F_{2} \Omega_{t} F_{2} + F_{1} \Sigma_{s} F_{1}$$
(34)

where Ω_s is the *t*-conditional variance-covariance matrix of \mathbf{ffl}_{t+s} , and Σ_s is the *t*-conditional variance-covariance matrix of $\widehat{\mathbf{x}_{t+s}}$, defined recursively by a discrete Lyapunov equation (from the law of motion of $\widehat{\mathbf{x}_{t+1}}$).

Then, projecting $\mathbb{E}_t \left[\widehat{\mathbf{x}_{t+s-1}} \right]$ terms on their first-order counterparts, denoted $\mathbb{E}_t^1 \left[\widehat{\mathbf{x}_{t+s-1}} \right]$, we obtained the pruned law of motion of the transformed solution:

$$\mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s}} \right] = F_{1} \mathbb{E}_{t} \left[\widehat{\mathbf{x}_{t+s-1}} \right] + F_{3} \eta^{2} + \frac{1}{2} F_{11} \left(\mathbb{E}_{t}^{1} \left[\widehat{\mathbf{x}_{t+s-1}} \right]^{2} + \Sigma_{s-1} \right) + \frac{1}{2} F_{22} \eta^{2} \Omega_{s} \\
\mathbb{E}_{t} \left[\widehat{\mathbf{y}_{t+s}} \right] = \frac{1}{2} M_{11} \left(\mathbb{E}_{t}^{1} \left[\widehat{\mathbf{x}_{t+s}} \right]^{2} + \Sigma_{s} \right) + M_{2} \eta^{2} \\
\mathbb{E}_{t}^{1} \left[\widehat{\mathbf{x}_{t+s}} \right] = F_{1} \mathbb{E}_{t}^{1} \left[\widehat{\mathbf{x}_{t+s-1}} \right] \\
\Sigma_{s+1} = \eta^{2} F_{2} \Omega_{t} F_{2} + F_{1} \Sigma_{s} F_{1}$$
(35)

To compute the steady state of the second-order solution to the original system, we first compute the steady state of the transformed system using its laws of motion. In particular, we solve for the steady state value of expected deviations of transformed variables from their steady state (set $\eta = 1$):

$$\mathbb{E}\left[\widehat{\mathbf{x}}\right] = \left(I - F_{1}\right)^{-1} \left(F_{3} + \frac{1}{2}F_{22}\Omega + \frac{1}{2}F_{11}\Sigma\right)$$

$$\mathbb{E}\left[\widehat{\mathbf{y}}\right] = \frac{1}{2}M_{11}\Sigma + M_{2}$$
(36)
where $\Sigma = F_{2}\Omega_{t}F_{2} + F_{1}\Sigma F_{1}$

Finally, we back out the steady state values of original variables as $Z'^{-1} \left(\mathbb{E} \left[\widehat{\mathbf{x}} \right] \ \mathbb{E} \left[\widehat{\mathbf{y}} \right] \right)'$.

A.3 Nonlinear Impulse Response Functions

To compute the economy's impulse response functions, we use the pruned version of the law of motion for transformed variables (for $\eta = 1$), for $t \ge 0$:

$$\widehat{\mathbf{x}_{t+1}} = F_1 \widehat{\mathbf{x}}_t + F_2 \mathbf{ffl}_{t+1} + F_3 + \frac{1}{2} F_{11} \widehat{\mathbf{x}_t^1}^2 + F_{12} \widehat{\mathbf{x}_t^1} \mathbf{ffl}_{t+1} + \frac{1}{2} F_{22} \mathbf{ffl}_{t+1}^2$$

$$\widehat{\mathbf{y}_t} = \frac{1}{2} M_{11} \widehat{\mathbf{x}_t^1}^2 + M_2$$

$$\widehat{\mathbf{x}_{t+1}^1} = F_1 \widehat{\mathbf{x}_t^1} + F_2 \mathbf{ffl}_{t+1}$$
(37)

We then back out the path of original variables as $\left\{ Z'^{-1} \begin{pmatrix} \widehat{\mathbf{x}}_t & \widehat{\mathbf{y}}_t \end{pmatrix}' \right\}_t$.

A.4 Market Clearing Errors

I measure the accuracy of the first- and second-order approximations by computing the residuals of equilibrium conditions, in particular market clearing conditions for savings, consumption and labor. They are small in the first-order approximation of the model, and further decrease towards zero in the second-order approximation, proving the good fit of the model (Table 5).

Table 5: Solution accuracy

Market:	Savings	Good	Labor
order 1	0.01% (0.03%)	0.04% (0.04%)	0.01% (0.01%)
order 2	0.00% (0.02%)	0.00% (0.00%)	0.00% (0.00%)

Notes: Average market clearing errors for IRF (sup norm in parentheses), computed as percentage differences normalized by the steady state value of the variable or by the initial value of the series.

A.5 Variance decomposition

A.5.1 First Order

The vector $Y = \begin{pmatrix} x & y \end{pmatrix}$ of equilibrium objects contains the predetermined and the jump variables. It is in deviation from steady state, but it doesn't matter for this exercise because we can just add the steady state vector, which will cancel out when taking variances. The output from gensys is a law of motion for *Y*, consisting of an AR(1) matrix Φ and an impact matrix *Z*:

$$(I - \Phi L) Y_{t+1} = Z \epsilon_{t+1} \tag{38}$$

where $\epsilon_{t+1} = (\epsilon_{t+1}^{\phi} \ \epsilon_{t+1}^{z})'$ is the vector of the two shocks, with covariance matrix $\tilde{\Sigma}_{\epsilon} = \begin{pmatrix} 1 & \rho_{\phi,z} \\ \rho_{\phi,z} & 1 \end{pmatrix}$, and where the rows of *Z* corresponding to ϵ_{t+1}^{ϕ} and ϵ_{t+1}^{z} are $\begin{pmatrix} \sigma_{\phi} & 0 \\ 0 & \sigma_{z} \end{pmatrix}$. Thus $\operatorname{Var} \left(\begin{pmatrix} \sigma_{\phi} & 0 \\ 0 & \sigma_{z} \end{pmatrix} \tilde{\Sigma}_{\epsilon} \right) = \begin{pmatrix} \sigma_{\phi}^{2} & \rho_{\phi,z}\sigma_{\phi}\sigma_{z} \\ \rho_{\phi,z}\sigma_{\phi}\sigma_{z} & \sigma_{z}^{2} \end{pmatrix} = \Sigma_{\epsilon}$. First, we transform the shocks with covariance $\tilde{\Sigma}_{\epsilon}$ so that they are orthogonal, i.e. their covariance matrix is the identity matrix. This is done by Cholesky factorization. The new orthogonal shocks are defined as $\nu_{t} = Q\epsilon_{t}$, with *Q* such that $\mathbb{E} [\nu_{t}\nu_{t}'] = I$. Denoting

S = Q^{-1} , $\epsilon_t = Sv_t$ and $SS' = \tilde{\Sigma}_{\epsilon}$. S is a lower triangular matrix given by the Cholesky factorization of $\tilde{\Sigma}_{\epsilon}$.

Then, we transform the economy's law of motion from an AR(1) to an MA(∞) representation, using the fact that the eigenvalues of Φ are within the unit circle (we denote *L* the lag operator). We also substitute for $\epsilon_{t+1} = S\nu_{t+1}$.

$$(I - \Phi L) Y_{t+1} = Z\epsilon_{t+1}$$

$$\Rightarrow Y_{t+1} = (I - \Phi L)^{-1} ZS\nu_{t+1}$$

$$Y_{t+1} = \sum_{k=0}^{\infty} \Phi^k L^k ZS\nu_{t+1}$$

$$Y_{t+1} = \sum_{k=0}^{\infty} \Phi^k ZS\nu_{t+1-k}$$

$$\Rightarrow Y_{t+h} = \sum_{k=0}^{\infty} \tilde{\Phi}^{(k)} \nu_{t+h-k}$$
(39)

for any forecasting horizon h > 0, and where $\tilde{\Phi}^{(k)} = \Phi^k ZS$ is a matrix of dimension (number of variables, number of shocks). Here we consider *N* variables and 2 shocks. Then, forecast errors at horizon h > 0 are:

$$e_{t+h} = Y_{t+h} - \mathbb{E}_t [Y_{t+h}] = \tilde{\Phi}^{(0)} \nu_{t+h} + \tilde{\Phi}^{(1)} \nu_{t+h-1} + \tilde{\Phi}^{(2)} \nu_{t+h-2} + \dots + \tilde{\Phi}^{(h-1)} = \sum_{i=1}^h \tilde{\Phi}^{(h-i)} \nu_{t+i}$$
(40)

For variable Y_j , $j \in \{1, ..., N\}$,

$$e_{j,t+h} = \sum_{i=1}^{h} \tilde{\Phi}_{j,.}^{(h-i)} \nu_{t+i}$$

= $\sum_{i=1}^{h} \left(\tilde{\Phi}_{j,1}^{(h-i)} \nu_{1,t+i} + \tilde{\Phi}_{j,2}^{(h-i)} \nu_{2,t+i} \right)$ (41)

So the total forecast error variance at horizon h > 0 for variable Y_j is, using the fact that ν 's are mutually independent, identically distributed and serially uncorrelated:

$$\operatorname{Var}\left(e_{j,t+h}\right) = \sum_{i=1}^{h} \left(\left(\tilde{\Phi}_{j,1}^{(h-i)}\right)^{2} + \left(\tilde{\Phi}_{j,2}^{(h-i)}\right)^{2} \right)$$
(42)

Finally, the share of the forecast error variance of variable Y_j at horizon h > 0 accounted for by ν^1 and ν^2 (transformed versions of the original shocks ϵ^{ψ} and ϵ^z) are respectively:

$$\frac{\sum_{i=1}^{h} \left(\tilde{\Phi}_{j,1}^{(h-i)}\right)^{2}}{\sum_{i=1}^{h} \left(\left(\tilde{\Phi}_{j,1}^{(h-i)}\right)^{2} + \left(\tilde{\Phi}_{j,2}^{(h-i)}\right)^{2}\right)} \text{ and } \frac{\sum_{i=1}^{h} \left(\tilde{\Phi}_{j,2}^{(h-i)}\right)^{2}}{\sum_{i=1}^{h} \left(\left(\tilde{\Phi}_{j,1}^{(h-i)}\right)^{2} + \left(\tilde{\Phi}_{j,2}^{(h-i)}\right)^{2}\right)}$$
(43)

Results are sensitive to whether the matrix obtained from the Cholesky factorization is lower or upper triangular. A lower triangular *S* implies that v^2 has no effect on v^1 . Note that because of the factorization, the v shocks are not clearly interpretable as credit and aggregate productivity shocks.

A.5.2 Second Order

I use a generalized forecast error variance decomposition for nonlinear models (Lanne and Nyberg (2016)). The starting point is the nonlinear (quadratic) model given by gensys2, which can be written as

$$Y_{t+1} = f(Y_t, \epsilon_{t+1}) \tag{44}$$

where *G* is a nonlinear function of the equilibrium vector and of innovations. As above, the interpretation of shocks is clearer when $\rho_{\phi,z} = 0$.

The generalized impulse-response function (GIRF) at horizon i > 0 (i.e. at date t + i) of variable Y_j , with respect to a credit shock (or aggregate productivity shock) of magnitude $\delta_{\phi,t+1}$ (or $\delta_{z,t+1}$) hitting at date t + 1, conditional on history of states $\omega_t = y_t$, is defined as:

$$GI_{j}(i, \delta_{\phi,t+1}, \omega_{t}) = \mathbb{E}_{t} \left[Y_{j,t+i} | \epsilon_{t+1}^{\phi} = \delta_{\phi,t+1}, \omega_{t} \right] - \mathbb{E}_{t} \left[Y_{j,t+i} | \omega_{t} \right]$$

and $GI_{j}(i, \delta_{z,t+1}, \omega_{t}) = \mathbb{E}_{t} \left[Y_{j,t+i} | \epsilon_{t+1}^{z} = \delta_{z,t+1}, \omega_{t} \right] - \mathbb{E}_{t} \left[Y_{j,t+i} | \omega_{t} \right]$ (45)

Then, the generalized forecast error variance decomposition (GFEVD) of variable Y_j at horizon h > 0, is between the fraction of variance explained by credit shocks, and that

explained by aggregate productivity shocks, respectively:

$$GFEVD_{j}(h, \delta_{\phi,t}) = \frac{\sum_{i=0}^{h} GI_{j}(i, \delta_{\phi,t+1}, \omega_{t})^{2}}{\sum_{i=0}^{h} GI_{j}(i, \delta_{\phi,t+1}, \omega_{t})^{2} + \sum_{i=0}^{h} GI_{j}(i, \delta_{z,t+1}, \omega_{t})^{2}}$$

$$GFEVD_{j}(h, \delta_{z,t}) = \frac{\sum_{i=0}^{h} GI_{j}(i, \delta_{z,t+1}, \omega_{t})^{2}}{\sum_{i=0}^{h} GI_{j}(i, \delta_{\phi,t+1}, \omega_{t})^{2} + \sum_{i=0}^{h} GI_{j}(i, \delta_{z,t+1}, \omega_{t})^{2}}$$
(46)

Because GIRF are nonlinear, GFEVD depend on the sign and size of the innovations δ . I therefore compute average GFEVD using bootstrap. First, because the solution of the model is based on perturbations around the steady state, we can get rid of the history dependence in ω . Then, I simulate a history of credit and aggregate productivity innovations of length T = 1000, $\left\{ \epsilon_t^{\phi}, \epsilon_t^z \right\}_{t=0}^T = \left\{ \delta_{\phi,t}, \delta_{z,t} \right\}_{t=0}^T$ using $\begin{pmatrix} \epsilon_t^{\phi} \\ \epsilon^z \end{pmatrix} \stackrel{iid}{\sim} \mathcal{N}(0, I_2)$ (with gensys2 the innovation variances σ_{ϕ}^2 and σ_z^2 are incorporated in the GIRF matrices). For each innovation $\delta_{\phi,t}$, I compute the associated $GFEVD_j(h, \delta_{\phi,t})$ for variable Y_j at horizon h. Finally, the average GFEVD is obtained by averaging over individual $GFEVD_j(h, \delta_{\phi,t})$'s by using the probability associated to each $\delta_{\phi,t}$ by the standard normal p.d.f. (Because $\mathcal{N}(0,1)$ is symmetric, we should get something like an average of the GFEVD for a shock $\delta = -1$ and a shock $\delta = +1$.) Computations are parallelized over the N dimension. It takes about 17 hours to run the case N = 500, H = 1000 using 28 cores.

A.6 Estimation of structural shocks

A.6.1 First Order: Kalman filter

A linear state space representation of the model is obtained from gensys. Using the above notation, the transition and the measurement equations are respectively:

$$Y_{t+1} = \Phi Y_t + Z \epsilon_{t+1}, \quad \epsilon_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, Q)$$

$$Y_{t+1}^{obs} = H' Y_{t+1} + v_t, \quad v_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, R)$$
(47)

 Φ and *Z* are readily obtained from gensys and $Q = I_2$ (variance-covariance terms are in *Z* by design). *H* is a selection matrix filled everywhere with zeros, and with ones for the entries corresponding to the observable variables in Y_{t+1} (risk-free rate and consumption). There is no noise in the measurement equation, i.e. $R = 0_{2\times 2}$: the risk-free rate and consumption are perfectly observed.

Using standard notation, denote $Y_{t|t-1} = \mathbb{E} [Y_t | Y^{obs,t-1}]$ (best linear predictor of Y_t given the history of observables Y^{obs} until t-1), $Y_{t|t-1}^{obs} = \mathbb{E} [Y_t^{obs} | Y^{obs,t-1}]$, and $Y_{t|t} = \mathbb{E} [Y_t | Y^{obs,t}]$.

Also denote $\Sigma_{t|t-1} = \mathbb{E}\left[\left(Y_t - Y_{t|t-1}\right)\left(Y_t - Y_{t|t-1}\right)'|Y^{obs,t-1}\right]$ (predicting error variancecovariance matrix of Y_t given the history of observables until t-1),

$$\begin{split} \Omega_{t|t-1} &= \mathbb{E}\left[\left(Y_t^{obs} - Y_{t|t-1}^{obs}\right)\left(Y_t^{obs} - Y_{t|t-1}^{obs}\right)'|Y^{obs,t-1}\right],\\ \Sigma_{t|t} &= \mathbb{E}\left[\left(Y_t - Y_{t|t}\right)\left(Y_t - Y_{t|t}\right)'|Y^{obs,t}\right]. \end{split}$$

The goal of the Kalman filter here is to back out the sequences of forecasted observable variables and underlying states $\{Y_{t|t-1}^{obs}, Y_{t|t}\}$ implied by the model, given a sequence of observable variables $\{Y_t^{obs}\}$ taken from the data. The algorithm proceeds as follows:

1. At t = 1, initial conditions $Y_{1|0}$, $\Sigma_{1|0}$ are set equal to their (deterministic) steady state values. That is, $Y_{1|0} = 0$ (the initial system of equations was written in log deviations from steady state), and $\Sigma_{1|0}$ is the solution to the Riccati equation $\Sigma_{1|0} = \Phi\Sigma_{1|0}\Phi' + ZI_2Z'$, which is solved by iterating on a symmetric, positive definite guess $\Sigma_{1|0}^{(0)}$ (using the stability of the system). I verify that the solution $\Sigma_{1|0}^{(\infty)} = \Sigma_{1|0}$ is symmetric and positive definite too. Following steps are for $t \ge 1$.

2. Given
$$\Sigma_{t|t-1}$$
, Y_t^{obs} , $Y_{t|t-1}^{obs}$, compute $\Omega_{t|t-1} = H' \Sigma_{t|t-1} H + R = H' \Sigma_{t|t-1} H$.

3. Compute
$$\operatorname{Cov}_{t-1}(Y_t^{obs}, Y_t) = \mathbb{E}\left[\left(Y_t^{obs} - Y_{t|t-1}^{obs}\right)\left(Y_t - Y_{t|t-1}\right)' | Y^{obs,t-1}\right] = H' \Sigma_{t|t-1}.$$

4. Compute the Kalman gain
$$K_t = \Sigma_{t|t-1} H \left(H' \Sigma_{t|t-1} H + R \right)^{-1} = \Sigma_{t|t-1} H \Omega_{t|t-1}^{-1}$$
.

- 5. Compute $Y_{t|t} = Y_{t|t-1} + K_t \left(Y_t^{obs} H'Y_{t|t-1} \right)$ ("nowcast" of the state).
- 6. Compute $\Sigma_{t|t} = \Sigma_{t|t-1} K_t H' \Sigma_{t|t-1}$ (variance-covariance matrix associated with the "nowcast" error).
- 7. Compute $\Sigma_{t+1|t} = \Phi \Sigma_{t|t} \Phi' + ZQZ' = \Phi \Sigma_{t|t} \Phi' + ZZ'$ (next period forecast error variance-covariance matrix).
- 8. Finally, compute $Y_{t+1|t} = \Phi Y_{t|t}$ and $Y_{t+1|t}^{obs} = H'Y_{t+1|t}$ (next period implied state, and next period forecasted observables).

A.6.2 Second Order: Particle Filter

A nonlinear state space representation of the model is obtained from gensys2. Using the above notation, the transition and the measurement equations are respectively:

$$Y_{t+1} = f(Y_t, W_{t+1}), \quad W_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, Q)$$

$$Y_{t+1}^{obs} = H'Y_{t+1} + v_t, \quad v_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, R)$$
(48)

f is the quadratic mapping (from gensys2) used to compute impulse responses in the second-order solution of the model (see above). $Q = I_2$ (variance-covariance terms are in the matrices part of *f* by design), and *H* is a selection matrix filled everywhere with zeros, and with ones for the entries corresponding to the observable variables in Y_{t+1} (risk-free rate and consumption). I assume that there is some but very little noise in the measurement equation, i.e. $R = 10^{-6} \times I_2$: the risk-free rate and consumption are close to perfectly observed. This is because the joint density of measurement errors is needed in the algorithm, so *R* cannot be zero.

Particles are i.i.d. draws $\{Y_{t-1}^i, W_{t-1}^i\}_{i=1}^N$ from the joint density $p(W_{t-1}, Y_{t-1}|Y_{t-1}^{obs})$. Proposed particles are i.i.d. draws $\{Y_{t|t-1}^i, W_{t|t-1}^i\}_{i=1}^N$ from the joint density $p(W_t, Y_{t-1}|Y_{t-1}^{obs})$. There are N of each of them. Here, the structural innovations W are independent of the vector of predetermined and jump variables Y. Therefore, drawing from the proposed joint density boils down to drawing from the innovations' density, and then applying the nonlinear mapping f to the previous proposed Y and the new innovations w, to get the new proposed particle Y. As before, the sequence of observable variables $\{Y_t^{obs}\}_{t=0}^T$ is taken from the data, with $Y_0^{obs} = 0$. That is, I assume w.l.o.g. that the beginning of the sample represents the deterministic steady state (hence log-deviations are zero). The algorithm proceeds as follows.

- 1. At t = 1, set the initial condition $Y_{0|0}^i = Y_0^i = W_0^i = 0$ for all i = 1, ...N, i.e. the log-deviation from the deterministic steady state is assumed to be zero at t = 0.
- 2. Generate *N* i.i.d. draws of proposed particles $\{Y_{t|t-1}^{i}, W_{t|t-1}^{i}\}_{i=1}^{N}$ from $p(W_{t}, Y_{t-1}|Y_{t-1}^{obs})$. That is, draw $w_{t|t-1}^{i}$ innovations from $\mathcal{N}(0, I_{2})$ and obtain the associated $Y_{t|t-1}^{i}$ from *f*.
- 3. Evaluate the conditional density $p\left(Y_t^{obs}|w_{t|t-1}^i, Y_{t-1}^{obs}, Y_{t|t-1}^i\right)$ using the measurement equation and the distribution of measurement errors v. That is,

$$p\left(Y_{t}^{obs}|w_{t|t-1}^{i}, Y_{t-1}^{obs}, Y_{t|t-1}^{i}\right) = \phi\left(Y_{t}^{obs} - H'Y_{t}|w_{t|t-1}^{i}, Y_{t-1}^{obs}, Y_{t|t-1}^{i}\right)$$

where ϕ is the (conditional) density of the multivariate standard normal distribution.

- 4. Evaluate the relative weights $q_t^i = \frac{p\left(Y_t^{obs}|w_{t|t-1}^i, Y_{t-1}^{obs}, Y_{t|t-1}^i\right)}{\sum_{j=1}^N p\left(Y_t^{obs}|w_{t|t-1}^j, Y_{t-1}^{obs}, Y_{t|t-1}^j\right)}$, normalized to be probabilities.
- 5. Re-sample, with replacement, *N* values $\{Y_{t|t-1}^{i}, W_{t|t-1}^{i}\}_{i=1}^{N}$ from the sample we had so far, now using the $\{q_{t}^{i}\}_{i=1}^{N}$ as probabilities. These new values are the particles, denoted $\{Y_{t}^{i}, W_{t}^{i}\}_{i=1}^{N}$.
- 6. Go back to step 2 for t + 1, generate new innovations and use the new swarm of particles $\{Y_t^i, W_t^i\}_{i=1}^N$ to generate a new swarm of proposed particles $\{Y_{t+1|t}^i, W_{t+1|t}^i\}_{i=1}^N$. Then iterate until reaching the end of the sample t = T.

Thus we obtain a sequence of swarms of particles $\left\{ \left\{ Y_t^i, W_t^i \right\}_{i=1}^N \right\}_{t=0}^T$, which represent empirical conditional densities at every point in time for the state Y, which are implied by the model, given the sequence of observables $\left\{ Y_t^{obs} \right\}_{t=0}^T$ from the data. In the main text, I plot the sample averages of these empirical conditional densities at t = 0, ...T. This paper is to my knowledge the first paper to apply nonlinear filtering to the perturbation-based solution of a heterogeneous agents model with aggregate shocks. Computations are parallelized over the *N* dimension. It takes about 12 hours to run the case N = 20,000, T = 44 using 28 cores.

B Calibration

B.1 Computation

On a 3.4 GHz Intel Core i5-7500 desktop with 8 GB of RAM, it takes 55s to solve for the model steady state, 7s and 821s to compute the Jacobian and the Hessian using automatic differentiation (in Julia), 8s and 170s to call gensys and gensys2 (in MATLAB). Overall, the model is solved in 15-20min.

B.2 Identification

In Figure 7, I separately plot surfaces for the risk-free rate autocorrelation and volatility, as functions of the credit shock autocorrelation and volatility, to show that the latter are

Parameter	Explanation	Value
$N_{ heta}$	Nb. idiosyncratic income states	5
N_h^f	Length grid for distribution	60
N_b	Length grid for savings	20
N_n	Length grid for labor supply	20
\overline{b}	Max. grid	90
x_1	Min. <i>x</i> added to χ	0.001
C _{min}	Min. consumption	0.001
-	Nb. iterations endogenous grid for initial guess	150
-	Solver tolerance for policy functions	10^{-6}
-	Solver tolerance for prices and ψ	10^{-6}

Table 6: Computation parameters

well-identified.



Figure 7: Identification of credit shock volatility and persistence with risk-rate data

Notes: Risk-free rate autocorrelation (upper panel) and annual % volatility (lower panel), as functions of the credit shock autocorrelation and volatility, estimated in a simulation of the linearized model with T = 10,000 periods. In each graph, the black dot represents the model calibration for the credit shock process. It is identified as it lies in non-flat areas of the (ρ , σ) surfaces.

C Additional Impulse Response Functions



Figure 8: Response to aggregate productivity shock

Notes: Impulse response function to a one standard deviation TFP shock: order 1 vs 2. The upper left panel plots the response of borrowing constraints to output for all income types (θ_1 for the lowest line, θ_5 for the highest), here zero. Initial period: deterministic steady state.

D Data





Notes: Total Factor Productivity at Constant National Prices for United States. Source: Penn World Table 9.0. Shaded area represents NBER recession.

Figure 10: Lending standards, unsecured household credit

Net percent 100 80 Jan. survey 60 40 20 0 Non-credit card (auto and other) Credit card Auto -20 -40 Other consumer -60 -80 -100 1991 1993 1995 1997 1999 2001 2003 2005 2007 2009 2011 2013 2015 2017 Note: For data starting in 2011:Q2, changes in standards for auto loans and consumer loans excluding credit card and auto loans are reported separately. In 2011:Q2 only, new and used auto loans are reported separately and equally weighted to calculate the auto loans series.

Net Percent of Domestic Respondents Tightening Standards for Consumer Loans



Net Percent of Domestic Respondents Reporting Increased Willingness to Make Consumer Installment Loans

Source: Ferederal Reserve Board, April 2017 Senior Loan Officer Opinion Survey on Bank Lending Practices. Quarterly frequency.