

Aggregate Precautionary Savings Motives*

Pierre Mabille[†]

March 5, 2024

Abstract

This paper analyzes the effect of *aggregate* risk on households' precautionary savings, a new channel that complements the standard idiosyncratic precautionary motive. I calibrate a general equilibrium model with incomplete markets, heterogeneous households, and aggregate risk to decompose the drivers of precautionary savings. Precautionary motives due to borrowing constraint shocks are large, nuancing received wisdom about the low costs of business cycles. They are larger for middle-class households, who are too rich to benefit from social programs but too poor to have enough liquid assets. In aggregate, they are key to help popular models explain the post-Great Recession credit landscape.

JEL classification: D14; E21; G11; G51.

Keywords: precautionary savings; borrowing constraints; heterogeneous household models; household finance.

*I thank Bernard Dumas, Mark Gertler, Michael Haliassos, Virgiliu Midrigan, Thomas Philippon, Stijn Van Nieuwerburgh, Olivier Wang, and Stanley Zin for comments. This paper benefited from participants at the WFA Meeting, the FIRS Conference, the AFA Annual Meeting, the EFA Annual Meeting, the SED Meeting, the CEPR Paris Symposium, the EUROFIDAI-ESSEC Paris December Finance Meeting, and the Becker-Friedman Institute Macro Finance Summer Session. 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 insure 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 and their consequences.

This paper fills this gap by evaluating the effects of aggregate risk on households’ precautionary savings and comparing them with the standard precautionary motive. I build a general equilibrium model of precautionary savings with heterogeneous households and incomplete markets. In addition to idiosyncratic risk, 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 aggregate productivity and in the tightness of households’ borrowing constraints, which limit their ability to smooth consumption. For instance, credit card limits fell by 25% in 2008 and 20% of borrowers had an account closed, including those with an excellent credit score (Federal Reserve Bank of New York, Consumer Credit Panel). In 2020, the net percentage of banks tightening credit standards on all consumer loans increased from 14% in the first quarter to 72% in the third quarter (Federal Reserve Board, Senior Loan Officer Opinion Survey on Bank Lending Practices).

I use a new representation of 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, I depart from existing heterogeneous agent models with a small number of aggregate states (based, e.g., on [Krusell and Smith \(1998\)](#)), and perturb the model with respect to *continuous* aggregate shocks. This setting is, to the best of my knowledge, the first to provide a dynamic decom-

¹See, e.g., [Zeldes \(1989\)](#), [Deaton \(1991\)](#), [Huggett \(1993\)](#), [Aiyagari \(1994\)](#), [Bertaut and Haliassos \(1997\)](#), [Carroll and Samwick \(1997\)](#), [Gourinchas and Parker \(2002\)](#), [Parker and Preston \(2005\)](#), [De Nardi, French, and Jones \(2010\)](#).

position of precautionary motives over the business cycle, thanks to tractable nonlinear impulse response functions to aggregate shocks.

I obtain three new findings, which highlight the importance of aggregate precautionary motives for households' balance sheets and the macroeconomy. First, aggregate risk significantly contributes to precautionary savings, which nuances received wisdom about the low cost of business cycles (e.g., Lucas (1987)). The contribution of borrowing constraint fluctuations to higher savings and a lower risk-free rate is especially large. It represents 60% of the impact of idiosyncratic income risk in terms of average income and it dwarfs the impact of aggregate productivity risk, which is close to zero. Second, aggregate precautionary motives are more important for "middle-class" households, in contrast with the focus of economists on the top and 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 their own liquid assets. Third, aggregate precautionary motives imply that shocks can have permanent effects even when they are themselves temporary. The resulting low-debt environment is key to help widely used heterogeneous household models explain periods of "creditless recovery" such as the post-Great Recession era.

The model is populated by infinitely-lived, risk-averse households with heterogeneous income and wealth. Every period, households consume and elastically supply labor to competitive firms. They save in risk-free bonds or borrow subject to a *stochastic* borrowing constraint, which depends on individual income and an aggregate component that is common across households and can be interpreted as capturing credit supply. The government raises progressive taxes and issues debt to finance progressive transfers and existing debt. The real risk-free rate and the wage endogenously clear the markets for savings and labor. Households face countercyclical idiosyncratic income risk and continuous aggregate shocks to productivity and their borrowing constraints, from which models with fixed or deterministic constraints typically abstract.

Three ingredients are key for evaluating precautionary motives. First, markets are incomplete, which generates heterogeneity across households and 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, which is itself a key determinant of savings. Abstracting from this feedback would lead to overstating precautionary savings: aggregate risk increases 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 also adjust their labor supply and receive government transfers, which are imperfect substitutes to savings; without these margins, the role of savings would also be overstated.

I calibrate the model using indirect inference to match the level and cross-section of liquid savings and unsecured debt that U.S. households use to smooth consumption, and the dynamics of the risk-free rate in the post-Great Recession period. My findings rely on a dynamic decomposition of precautionary motives, which is based on the economy's departure from certainty equivalence with respect to the various sources of aggregate risk. I solve for first- and second-order perturbations with respect to aggregate shocks. For each shock, the difference between second- and first-order terms captures the effect of volatility at each point in time. This leads to different amounts of savings for each source of risk, which would be difficult to identify in the data or in a standard model simulation.²

This setting leads to three contributions. First, I use the model to quantify the various precautionary motives, an open question for empirical analyses identifying one source of risk at a time. The *idiosyncratic motive* due to income risk arises because of the prudence property of utility, and because the combination of income shocks and static borrowing constraints limits consumption smoothing. It has the largest impact on average, and increases savings by 88 percentage points from 31% to 119% of average income compared to a riskless economy where they are only determined by intertemporal substitution. The *aggregate financial motive* due to stochastic borrowing constraints further increases savings by 54 pp from 119% to 173% of average income, which represents about 60% of the effect of idiosyncratic risk. This sizable impact cannot be ignored when analyzing households' balance sheets. Importantly, the assumptions of an endogenous risk-free rate, flexible la-

²For instance, a candidate option may be to simulate the model by alternatively turning on the various sources of risk to isolate their effects. However, this approach would capture the effects of both the level and the volatility of the shocks. Therefore, it would not separately identify precautionary motives, which arise because of the volatility of the shocks that is anticipated by households, including when the level effect of the shocks itself is small.

bor supply, and government transfers guarantee that households have alternative ways of smoothing consumption, and therefore that these estimates are a *lower bound* on precautionary savings. Interestingly, the *aggregate real motive* due to productivity shocks has almost zero impact on precautionary savings. While these results are consistent with low costs of business cycles alone for households, they highlight their large costs when associated with credit cycles. Finally, two caveats apply. First, this decomposition relies on the differentiability of equilibrium conditions. Keeping it tractable requires abstracting from frictional portfolio choices with stocks or housing, which are likely to amplify the effect of risk because of adjustment costs and, hence, precautionary savings. Second, this decomposition assumes some bounded rationality by abstracting from higher than second-order effects of aggregate shocks.³

Second, I analyze the heterogeneous effects of borrowing constraint fluctuations in the cross-section of households using nonlinear policy functions. When borrowing constraints contract, households' ability to smooth consumption deteriorates because of two effects. In the first order, a lower level of borrowing constraints forces constrained borrowers to deleverage, and those close to the constraint to increase savings to avoid becoming constrained because of future income shocks. In the second order, the volatility of borrowing constraints themselves makes them more likely to bind, which further increases savings and is absent from models without aggregate precautionary motives. These effects are larger for "middle-class" households with moderate amounts of debt. They do not have enough savings to ignore the risk of becoming unable to borrow, but their income is too high to benefit from government transfers.

Third, I highlight the real effects of aggregate precautionary motives. Using a variance decomposition, I find that aggregate shocks to borrowing constraints explain almost 60% of the volatility of consumption. The model uncovers a new recessionary effect of borrowing constraints that is driven by the differential labor supply responses of households, for which I find empirical support. The lower risk-free rate creates an intertemporal substitution motive that tends to decrease hours worked by unconstrained households that are

³Estimating these effects would require households to have unrealistically high computing power. Abstracting from them is plausible given how households make forecasts in practice (e.g., [Das, Kuhnen, and Nagel \(2019\)](#)). Using market clearing errors, I show that this assumption is innocuous (Appendix Table A1).

more productive. Even if constrained households work more to pay back their debt, the net effect is a decrease in consumption because they are less productive.

My findings have important empirical implications. Aggregate precautionary motives are key to help widely used heterogeneous household models explain the post-Great Recession data. The period since 2009 is characterized by the coincidence of low household debt and interest rates and, at the same time, a quick consumption recovery. These central facts are a puzzle for models with fixed borrowing constraints, in which higher future consumption should increase current debt and the risk-free rate through intertemporal substitution. Using a particle filter in the model with stochastic borrowing constraints, I estimate the sequences of aggregate shocks to productivity and borrowing constraints that explain the observed paths for consumption and the risk-free rate. I validate the calibration by showing that these estimates align with measures of borrowing constraints in the data, without being subject to survey error. This is the first paper to perform this challenging exercise in a general equilibrium model with heterogeneous households, incomplete markets, and aggregate risk, which has broad applications for future work in household finance and asset pricing.

These dynamic estimates uncover two patterns, which refine popular narratives of credit conditions that focus on the impact of shocks but abstract from their duration. First, a *persistent* tightening of borrowing constraints of 18% below their average level throughout the period can explain the decrease in household debt. Second, its effect is exacerbated by a *V-shaped* recession in aggregate productivity of 2%. These estimates are consistent with the onset and magnitude of the 25% decrease in credit card limits in the data (Federal Reserve Bank of New York, Consumer Credit Panel) and the increase in the share of lenders who report tightening lending standards (Federal Reserve Board, Senior Loan Officer Opinion Survey of Bank Lending Practices). While such survey measures can be hard to interpret, the model estimates the pass-through of lending standards to actual borrowing constraints. Interestingly, it indicates that the credit tightening was both slightly weaker and more persistent than surveys suggest. Such estimates can be a useful tool for policymakers to understand the credit landscape faced by borrowers as they potentially call for different responses.

Related literature. This work contributes to a longstanding literature that focuses on households' idiosyncratic risk but abstracts from aggregate risk as a driver of precautionary savings. The model accounts for shocks to households' borrowing constraints that are large and frequent, both at the individual and aggregate levels. Not only were constraints massively tightened after the Great Recession. They also vary in long time series, including the recent Covid-19 recession, and with monetary and macro-prudential policy. For instance, credit card limits have changed on average for more than a third of accounts every quarter since 1999, decreased by 36% conditional on changing, and 7% of households have completely lost credit access (Federal Reserve Bank of New York, Consumer Credit Panel).⁴ Yet, little is known so far about how households respond to this risk.

This is the first paper that decomposes idiosyncratic and aggregate precautionary savings motives in a general equilibrium model with heterogeneous households and incomplete markets. Bertaut and Haliassos (1997) study precautionary savings against idiosyncratic risk in income, return, and mortality, and Haliassos and Hassapis (2002) discuss their interactions with borrowing constraints. My estimates for the financial precautionary motive are consistent with Guiso, Jappelli, and Terlizzese (1996), in which expectations of future borrowing constraints increase households' savings.

Ludvigson (1999) and Bertaut, Haliassos, and Reiter (2009) analyze models with stochastic borrowing constraints but without aggregate risk and with exogenous interest rates. While they focus on consumption and the credit card puzzle, I analyze precautionary savings in a general equilibrium model with substitutable forms of private and public insurance. As in the data, flexible labor supply and social insurance programs lower savings (Hubbard, Skinner, and Zeldes (1995)). As in the models of Huggett (1993), Aiyagari (1994), and Heaton and Lucas (1996), endogenizing the interest rate is critical as it is a key determinant of savings. While previous work has shown that changes in the level of borrowing constraints affect consumption, housing, and employment (Favilukis, Ludvigson, and Van Nieuwerburgh (2017), Guerrieri and Lorenzoni (2017), Jones, Midrigan, and

⁴See also Appendix Figure A4, which provides further evidence of business cycle changes in lending standards for consumer loans. On changes in borrowing constraints, see, e.g., Ludvigson (1999), Gross and Souleles (2002), Mian, Rao, and Sufi (2013), Fulford (2015), Mian, Sufi, and Verner (2017), Baker (2018), Agarwal, Chomsisengphet, Mahoney, and Stroebe (2018), Cherry, Jiang, Matvos, Piskorski, and Seru (2022), and Acharya, Bergant, Crosignani, Eifert, and McCann (2022).

Philippon (2022)), my result highlight that their *volatility* is key for precautionary savings.

Finally, the new model representation in this paper allows to structurally estimate the time series of aggregate shocks that drive equilibrium quantities and prices by applying a particle filter in a model with heterogeneous households, incomplete markets, and aggregate risk. This approach can help better take the dynamic implications of these models to the data. It is related to recent work in macroeconomics, which has analyzed the first-order effect of aggregate shocks but abstracted from the second-order effect of risk, which is key in household finance applications.⁵

Outline. The rest of the paper is organized as follows. Section 2 presents the decomposition of precautionary motives in the model. Section 3 describes the calibration. Section 4 analyzes the main results on the financial and real effect of aggregate precautionary savings motives. Section 5 highlights their empirical implications for household credit since the Great Recession, and Section 6 concludes.

2 Model Decomposition of Precautionary Motives

This section describes a general equilibrium, closed economy model of precautionary savings with heterogeneous households, incomplete markets, and two sources of aggregate risk: shocks to aggregate productivity and to the tightness of households' borrowing constraints. It presents a novel representation of the model, which provides a dynamic decomposition of the contributions of aggregate risks to precautionary savings over the business cycle.

2.1 Environment

The economy is populated by a continuum of measure 1 of heterogeneous, risk-averse households with rational expectations. Markets are incomplete. Time is discrete.

⁵See, e.g., Reiter (2009) and Ahn, Kaplan, Moll, Winberry, and Wolf (2017).

Preferences. Households have time- and state-separable preferences, which are additively separable in consumption and labor. They have a constant relative risk aversion (CRRA) utility function over a nondurable consumption good c_t produced by competitive firms, and increasing and convex utility costs over hours worked n_t that are elastically supplied to firms:

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

Household balance sheets. Households can save in liquid assets and borrow with unsecured debt by buying and selling one-period bonds b_{t+1} at the real risk-free rate r_t . Their balance sheets are summarized by their net savings. When borrowing, they face stochastic borrowing constraints, which consist of an idiosyncratic component $\phi(\theta_t)$ that varies with individual household earnings θ_t and an aggregate component $\bar{\phi}_t$ that is common across households and varies over the business cycle, which can be interpreted as reflecting credit supply:

$$b_{it+1} \geq -\phi(\theta_{it}) \bar{\phi}_t. \quad (2)$$

The stochastic borrowing constraint captures changes in households' borrowing capacity over time as a result of both changes in their individual characteristics (income, or credit score) and in aggregate credit conditions that are common across households (credit supply, or lending standards). The multiplicative interaction of the individual and aggregate components of borrowing constraints reflects their complementarity in the data. The borrowing capacity $|\phi(\theta_t) \bar{\phi}_t|$ of households with a low $\phi(\theta_t)$, who face tighter credit constraints, tends to contract more when aggregate credit conditions $\bar{\phi}_t$ deteriorate. Conversely, households who face loose constraints tend to enjoy larger increases in their borrowing capacity when aggregate credit conditions improve.

Households face progressive taxes on labor income $\tau_t(\theta, n)$ and progressive transfers $T(\theta)$ from the government that are conditional on earnings. Firms' profits π_t are redistributed equally. This assumption provides a lower bound on estimated precautionary savings, since it relaxes the constraints of poor households who face more risk relatively

more. Households' budget constraints imply that their consumption, net savings, and tax payments cannot exceed their current savings or debt, income from labor earnings and firms' profits, and government transfers:

$$c_{it} + \frac{b_{it+1}}{1+r_t} + \tau_t(\theta_{it}, n_{it}) \leq w_t \theta_{it} n_{it} + b_{it} + \pi_t + T(\theta_{it}). \quad (3)$$

Finally, to keep the decomposition of precautionary savings tractable while introducing aggregate risk, I assume that there is a single interest rate r_t at which households can borrow or save. The benefit of this assumption is to also provide a lower bound on estimated precautionary motives, since a higher interest rate on borrowing would further decrease debt and increase savings.

Choices. Households choose sequences for consumption, labor supply, and net savings to maximize the expected discounted value of the utility flows from consumption net of the disutility of working, subject to their budget and stochastic borrowing constraints.

Types of risk. Households face idiosyncratic and aggregate risk, which affect their labor income, their financial income from risk-free savings and risky firms' profits, and their ability to borrow.

Countercyclical idiosyncratic income risk. The logarithm of individual productivity θ follows a persistent AR(1) process with autocorrelation ρ_θ :

$$\log \theta_{it} = \rho_\theta \log \theta_{it-1} + \sigma_\theta(z_t) \epsilon_{it}^\theta, \quad \epsilon_{it}^\theta \sim \mathcal{N}(0, 1). \quad (4)$$

The volatility of individual productivity σ_θ is a decreasing function of aggregate productivity z_t . Idiosyncratic income risk countercyclical, such that individual income volatility increases in bad times and decreases in good times.

Aggregate productivity risk. The logarithm of aggregate productivity z_t captures the state of the business cycle. It follows a persistent AR(1) process with autocorrelation ρ_z and volatility σ_z :

$$\log z_t = \rho_z \log z_{t-1} + \epsilon_t^z. \quad (5)$$

Aggregate borrowing constraint risk. The logarithm of the aggregate component of borrowing constraint follows a persistent, mean-reverting AR(1) process with mean $\bar{\phi}$, autocorrelation ρ_ϕ , and volatility σ_ϕ . A negative shock $\epsilon^\phi < 0$ induces a tightening of borrowing constraints that is common to all households, irrespective of their individual income, while a positive shock $\epsilon^\phi \geq 0$ relaxes them:

$$\log \bar{\phi}_t - \log \bar{\phi} = \rho_\phi (\log \bar{\phi}_{t-1} - \log \bar{\phi}) + \epsilon_t^\phi \quad (6)$$

Aggregate shocks to productivity and borrowing constraints follow a bivariate Normal distribution, which captures the correlation between the business cycle and aggregate credit conditions in the data:

$$\begin{pmatrix} \epsilon^\phi \\ \epsilon^z \end{pmatrix} \overset{iid}{\sim} \mathcal{N} \left(\mathbf{0}, \begin{pmatrix} \sigma_\phi^2 & \sigma_\phi \sigma_z \rho_{\phi z} \\ \sigma_\phi \sigma_z \rho_{\phi z} & \sigma_z^2 \end{pmatrix} \right) \quad (7)$$

Firms. A continuum of competitive firms hires efficient units of labor $\theta_t n_t$ from households every period, and combines them using a decreasing returns to scale production function that is subject to aggregate productivity shocks. Firms choose total employment in efficiency units N_t to solve a static profit maximization problem:

$$\max_{N_t} \pi_t = z_t N_t^\alpha - w_t N_t \quad (8)$$

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

$$\begin{aligned} \pi_t &= (1 - \alpha) Y_t = (1 - \alpha) z_t N_t^\alpha \\ w_t N_t &= \alpha Y_t = \alpha z_t N_t^\alpha \end{aligned} \quad (9)$$

To focus on the precautionary savings in risk-free assets that *households* use to smooth consumption at business cycle frequency, I assume that firms' shares are not tradable and abstract from risky capital that *firms* tend to accumulate for precautionary reasons over

longer horizons, which is the subject of a separate literature in corporate finance.⁶

Government. The government raises progressive taxes on labor earnings and issues risk-free debt to finance progressive transfers and outstanding debt. Its budget constraint is the following:

$$\int T(\theta) d\lambda_t(\theta, b) + B_t \leq \int \tau_t(\theta, n(\theta, b)) d\lambda_t(\theta, b) + \frac{B_{t+1}}{1+r_t} \quad (10)$$

Progressive taxes are an affine function of labor earnings:

$$\tau_t(\theta_{it}, n_{it}) = \tau_{0t} + \tau_1(\theta_{it}) w_t \theta_{it} n_{it}. \quad (11)$$

The slope τ_1 of the tax schedule depends on individual productivity, and the intercept τ_0 adjusts such that the government budget constraint holds every period.

2.2 Equilibrium

The model has heterogeneous households, incomplete markets, and idiosyncratic and aggregate risk. Therefore, the cross-sectional distribution of households over their productivity θ and net savings b , $\lambda_t(\theta, b)$, is an aggregate state variable. The distribution is time-varying because of aggregate shocks to productivity and borrowing constraints. Households take the current risk-free rate r_t as given but must forecast the next-period rate r_{t+1} to make intertemporal savings choices. The model is a closed economy in which the supply of liquid assets comes from the 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, which depends on the entire future cross-sectional distribution. Households must also forecast the current wage w_t that

⁶See, e.g., [Duchin, Gilbert, Harford, and Hrdlicka \(2017\)](#). Although this assumption could be relaxed at the cost of expanding the state space, it is unclear that it would change my estimates. *Households'* precautionary savings typically exclude risky assets, which do not allow to smooth consumption over the business cycle as much as risk-free savings because their price is uncertain. In a model with capital where households also use stocks for precautionary savings, aggregate risk may lead to higher capital accumulation. An increase in borrowing constraint risk would then be expansionary, which would be at odds with the post-Great Recession data.

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 states coincide with their realized values.

Definition. Given a sequence of aggregate shocks to productivity and borrowing constraints $\{z_t, \bar{\phi}_t\}$, a *competitive equilibrium* is a sequence of time-varying policy functions for households' consumption, labor supply, and net savings $\{c_t(\theta, b), n_t(\theta, b), b_{t+1}(\theta, b)\}$ for firms' labor demand N_t , risk-free rates and wages $\{r_t, w_t\}$, and government taxes $\{\tau_t\}$, such that the following conditions hold.

(i) Households optimally choose savings and labor supply:

$$\begin{aligned} c_t(\theta, b)^{-\gamma} &= \beta(1 + r_t)\mathbb{E}_t \left[c_{t+1}(\theta, b)^{-\gamma} \right] + \mu_t(\theta, b) \\ \psi n_t(\theta, b)^\eta &= (1 - \tau_1(\theta)) w_t \theta c_t(\theta, b)^{-\gamma} \end{aligned} \quad (12)$$

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

(ii) Firms optimally choose labor demand, so the marginal productivity of an additional work hours in efficiency units equals the wage:

$$\alpha z_t N_t^{\alpha-1} = w_t \quad (13)$$

(iii) The government budget constraint holds:

$$\int T(\theta) d\lambda_t(\theta, b) + B_t = \int \tau_t(\theta, n(\theta, b)) d\lambda_t(\theta, b) + \frac{B_{t+1}}{1+r_t} \quad (14)$$

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

$$\begin{aligned}
\int c_t(\theta, b) d\lambda_t(\theta, b) &= Y_t = z_t N_t^\alpha \\
\int \theta n_t(\theta, b) d\lambda_t(\theta, b) &= N_t \\
\int_+ b_{t+1}(\theta, b) d\lambda_t(\theta, b) &= \int_- b_{t+1}(\theta, b) d\lambda_t(\theta, b) + B_{t+1}
\end{aligned} \tag{15}$$

First, aggregate consumption must equal output. Consumption is the numeraire and its price is normalized to 1. Second, the wage adjusts such that total employment in efficiency units equals firms' demand for labor. Third, the risk-free rate adjusts such that the demand for savings, $\int_+ b_{t+1} d\lambda_t$, equals the total supply of liquid assets, which comes from the risk-free bonds issued by the household sector and the government, $\int_- b_{t+1} d\lambda_t$ and B_{t+1} .⁷

(v) The cross-sectional distribution of households and the aggregate shocks evolve according to their respective laws of motion. Denote $\Theta \times \mathcal{B}$ the sigma-algebra associated with the Cartesian product of the discrete set of individual productivity and the compact set of individual net savings, and $(\tilde{\Theta}, \tilde{\mathcal{B}})$ a subset of that sigma-algebra. The law of motion for the cross-sectional distribution of households is given by:

$$\begin{aligned}
\lambda_{t+1}(\tilde{\Theta}, \tilde{\mathcal{B}}) &= \int_{\Theta \times \mathcal{B}} Q_{\bar{\phi}_t, z_t}((\theta, b), (\tilde{\Theta}, \tilde{\mathcal{B}})) d\lambda_t(\theta, b) \\
\text{where } Q_{z_t, \bar{\phi}_t}((\theta, b), (\tilde{\Theta}, \tilde{\mathcal{B}})) &= \mathbf{1}\{b'_t(\theta, b) \in \tilde{\mathcal{B}}\} \sum_{\theta' \in \tilde{\Theta}} \Pi_\theta(\theta' | \theta)
\end{aligned} \tag{16}$$

The transition function $Q_{z_t, \bar{\phi}_t}$ depends on individual productivity and net savings, and on the aggregate components of productivity and borrowing constraints.

Solution. The model is solved numerically using a new representation described in the next subsection. Given policy functions and the distribution, the wage can be solved for analytically using labor market clearing:

$$w_t = \alpha z_t N_t^{\alpha-1} = \alpha z_t \left(\int \theta n(\theta, b) d\lambda_t(\theta, b) \right)^{\alpha-1} \tag{17}$$

⁷These two types of bonds are perfect substitutes, which is a standard assumption that allows to close the model (see, e.g., [Guerrieri and Lorenzoni \(2017\)](#)).

The wage directly depends on productivity shocks z_t , and indirectly on borrowing constraint shocks $\bar{\phi}_t$ through their effects on the distribution λ_t .

2.3 Decomposition

The decomposition of precautionary motives is based on a new representation of the model using projection and perturbation methods. I outline the main steps, which lend themselves well to economic interpretation. Appendix A describes the solution algorithm in detail.

Steps. First, variables and functions (policy functions, the cross-sectional distribution) are approximated using projections 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. Importantly, this solution of the model is exact without aggregate shocks; it is global and nonlinear with respect to idiosyncratic state variables. Third, the solution of the discrete model is perturbed with respect to aggregate shocks around the stationary steady state where they are zero, to generate the final solution for the stochastic steady state of the model *with* aggregate shocks.

The measure of precautionary motives relies on comparing two stochastic steady states of the model, with first-order and then second-order perturbations with respect to aggregate shocks. In the *first order*, certainty equivalence holds with respect to aggregate risk: only the level of shocks affects household behavior, but not their volatility. 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. The model departs from certainty equivalence as variables depend on the volatility of aggregate shocks. There are both precautionary motives, and the difference with the first-order solution is a measure of aggregate motives.

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

$$\mathbb{E}_t \left[\mathcal{F} \left(\mathbf{y}_t, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t+1}, \epsilon_{t+1}^\phi, \epsilon_{t+1}^z \right) \right] = 0. \quad (18)$$

Variables are sorted into non-predetermined and predetermined variables. The vector of non-predetermined variables \mathbf{y} contains projection coefficients for policy functions, which are approximated using linear splines, prices and aggregate quantities. The vector of predetermined variables \mathbf{x} contains the histogram weights used to project the cross-sectional distribution λ_t , and aggregate shocks to productivity z_t and borrowing constraints $\bar{\phi}_t$.

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

$$\mathcal{F}(\mathbf{y}, \mathbf{y}, \mathbf{x}, \mathbf{x}, 0, 0) = 0. \quad (19)$$

Two model ingredients make this more challenging than the standard consumption-savings problem: flexible labor supply and the endogeneity of the risk-free rate and government taxes. To solve the problem, I use a variant of the policy-time iteration method, which combines Broyden's numerical equation solver and automatic differentiation to compute exact derivatives.

Step 3: Perturbations. This step starts from the global and nonlinear solution for the stationary steady state of the model without aggregate shocks. Denote η the perturbation parameter that scales the quantity of aggregate risk in the economy. The solution of the expectation difference equation 18 defines the equilibrium with aggregate risk $\mathbb{E}_t[\mathcal{F}(\cdot)] = 0$. Define $\mathbf{h}(\mathbf{x}, \eta)$ and $\mathbf{g}(\cdot, \mathbf{x}, \eta)$ as nonlinear vector-valued functions that relate future predetermined variables and non-predetermined variables to current predetermined variables. Without loss of generality, the final solution of the model has the form:

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{h}(\mathbf{x}_t, \eta) + \eta \begin{pmatrix} \mathbf{0} \\ \epsilon_{t+1}^\phi \\ \epsilon_{t+1}^z \end{pmatrix} \\ \mathbf{y}_t &= \mathbf{g}(\mathbf{x}_t, \eta). \end{aligned} \quad (20)$$

These equations generalize the dynamic representation of representative agent models as in, e.g., [Schmitt-Grohe and Uribe \(2008\)](#), to a that setting with heterogeneous agents, incomplete markets, and aggregate risk.

First-order. Variables are written in deviations from the stationary steady state. 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 20 with respect to aggregate shocks:

$$\begin{aligned}\widehat{\mathbf{x}}_{t+1} &= \mathbf{h}_x(\mathbf{x}, 0) \widehat{\mathbf{x}}_t + \eta \begin{pmatrix} \mathbf{0} \\ \epsilon_{t+1}^\phi \\ \epsilon_{t+1}^z \end{pmatrix} \\ \widehat{\mathbf{y}}_t &= \mathbf{g}_x(\mathbf{x}, 0) \widehat{\mathbf{x}}_t.\end{aligned}\tag{21}$$

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

Second-order. The dynamics of the model with both idiosyncratic and *aggregate precautionary motives* is given by a second-order perturbation of equations 20 with respect to aggregate shocks:

$$\begin{aligned}\widehat{\mathbf{x}}_{t+1} &= \mathbf{h}_x(\mathbf{x}, 0) \widehat{\mathbf{x}}_t + \underbrace{\frac{1}{2} \mathbf{h}_{xx}(\mathbf{x}, 0) \widehat{\mathbf{x}}_t^2}_{\text{nonlinearity}} + \underbrace{\frac{1}{2} \mathbf{h}_{\eta\eta}(\mathbf{x}, 0) \eta^2}_{\text{agg. precautionary motive}} + \eta \begin{pmatrix} \mathbf{0} \\ \epsilon_{t+1}^\phi \\ \epsilon_{t+1}^z \end{pmatrix} \\ \widehat{\mathbf{y}}_t &= \mathbf{g}_x(\mathbf{x}, 0) \widehat{\mathbf{x}}_t + \underbrace{\frac{1}{2} \mathbf{g}_{xx}(\mathbf{x}, 0) \widehat{\mathbf{x}}_t^2}_{\text{nonlinearity}} + \underbrace{\frac{1}{2} \mathbf{g}_{\eta\eta}(\mathbf{x}, 0) \eta^2}_{\text{agg. precautionary motive}}.\end{aligned}\tag{22}$$

Solving for the new vectors of coefficients implies computing the Jacobian and the Hessian of the function $\mathcal{F}(\cdot)$. I solve for the vectors of coefficients $\mathbf{h}_{\eta\eta}(\mathbf{x}, 0)$ and $\mathbf{g}_{\eta\eta}(\mathbf{x}, 0)$, which relate future predetermined and non-predetermined variables to the *volatility* of aggregate shocks. Since perturbation methods for representative agent models cannot be applied due to the high dimension of the equation system, I apply a series of steps to reduce its dimension, which build on an extension of the gensys2 algorithm (Kim, Kim, Schaumburg, and Sims (2008)).

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

Nonlinearity. The coefficient vectors $\mathbf{h}_{\mathbf{x}\mathbf{x}}(\mathbf{x}, 0)$ and $\mathbf{g}_{\mathbf{x}\mathbf{x}}(\mathbf{x}, 0)$ measure the nonlinear dependence of future predetermined variables and non-predetermined variables on current predetermined variables. These terms *amplify* households' responses to a credit contraction compared to a linear approximation, as shown by the impulse response functions below (see Section 4). Therefore, they are crucial to help the model match the dynamics of household debt and consumption in the post-Great Recession data (see Section 5). They are not an exact measure of nonlinearity as my approach abstracts from higher than second-order terms. However, I show that the second-order approximation accurately describes the dynamics of the economy by computing the resulting market clearing errors, which are close to zero (Appendix Table A1).

Aggregate precautionary motive. The coefficient vectors $\mathbf{h}_{\eta\eta}(\mathbf{x}, 0)$ and $\mathbf{g}_{\eta\eta}(\mathbf{x}, 0)$ measure aggregate precautionary motives as the departure of the model from certainty equivalence with respect to aggregate shocks. With certainty equivalence, only the level of shocks to aggregate productivity and borrowing constraints would affect household behavior. Without it, the volatility of shocks affect policy functions, the cross-sectional distribution, and prices because of precautionary motives. Therefore, the stochastic steady state permanently differs from the stationary steady state. It can be interpreted as the time series average of a long-run simulation of the model with aggregate shocks.

Aggregate precautionary motives arise because households make optimal decisions given the stochastic processes governing aggregate shocks. For instance, their anticipation of future potential changes in borrowing constraints generates a precautionary motive in the Euler equation, $c_t^{-\gamma} = \beta(1 + r_t)\mathbb{E}_t \left[c_{t+1}^{-\gamma} \right] + \mu_t$. A high multiplier μ_t on current borrowing constraints implies higher savings. Iterating the equation forward, the average and higher-order moments of future multipliers on borrowing constraints $\{\mu_s\}_{s \geq t}$, especially their volatility due to aggregate shocks, further increase savings. In contrast, standard models typically assume that households have perfect foresight over fixed borrowing constraints. Fully unexpected shocks to the constraints tend to lead to more extreme responses than in the data as households suddenly and massively deleverage in

response to a credit contraction.⁸

3 Calibration

This section describes how the dynamic decomposition of precautionary motives in the model of Section 2 is mapped to U.S. post-Great Recession business cycle data. The model is calibrated to match the level and cross-sectional moments of liquid savings and unsecured debt that households use to smooth consumption over the business cycle. Importantly, the model also matches the dynamics of the risk-free rate and government taxes and transfers to accurately account for households' incentives to save.

Table 1 summarizes the calibration, which is split between externally, internally calibrated parameters, and sources of aggregate risk. Given external parameters, the fit of the model is measured as the sup norm between the vector of targeted moments and their empirical counterparts. One period is a quarter and moments are annualized. Average income is normalized to 1. Table 2 reports the fit of the model.

3.1 External Parameters

Idiosyncratic labor income shocks. The persistence and the volatility of the productivity process θ are chosen to match the persistence and volatility of wages of 0.92 and 0.24 in Kopecky and Suen (2010). The AR(1) process is discretized as a five state Markov chain as in Rouwenhorst (1995), which better matches the dynamics of income when the process is persistent than the standard method from Tauchen (1986). In steady state, 6.25% of households are in the lowest productivity group, 25% in the second lowest, 37.5% in the middle, 25% in the second highest, and 6.25% in the highest group.

Labor income shocks also generate changes in borrowing constraints at the individual level through the vector $[\phi(\theta)]_\theta$, which reflects changes in ability to borrow in the data that are due to individual factors (credit score, or employment status).

⁸See, e.g., Jones, Midrigan, and Philippon (2022), who discuss the model implications of households' slow deleveraging in the data.

Table 1: Calibration: Main parameters

Parameter	Explanation	Value	Target/source
<i>External</i>			
α	Labor share	2/3	Labor share of output = 2/3
$\tau_1(\theta)$	Tax progressivity by productivity	see text	Tax distribution by income (CPS)
$T(\theta)$	Government transfers by productivity	see text	Transfer distribution by income (CPS)
$\phi(\theta)$	Borrowing constraints: idiosyncratic	see text	Debt distribution by income (SCF)
ρ_θ	Idiosyncratic productivity persistence	0.92	Wage persistence
σ_θ	Idiosyncratic productivity volatility	0.24	Wage volatility
<i>Internal</i>			
γ	Risk aversion	5	See text
β	Discount factor	0.93	Risk-free rate = 1.80% (FRB)
B	Liquid savings supply	6	Liquid savings/income = 1.78 (FRB)
η	Curvature disutility of work	2	Frisch elasticity = 1/2
ψ	Disutility of work	11.5	Income normalization $Y = 1$
<i>Aggregate risk</i>			
ρ_z	Productivity persistence	0.55	TFP persistence
σ_z	Productivity volatility	0.026	TFP volatility
$\bar{\phi}$	Borrowing constraints average	2.6	Unsecured debt/income = 0.18 (FRB)
ρ_ϕ	Borrowing constraints persistence	0.96	Risk-free rate persistence = 0.65 (FRB)
σ_ϕ	Borrowing constraint volatility	0.10	Risk-free rate volatility = 1.9% (FRB)
$\rho_{\phi z}$	Productivity and borrowing constraint correlation	0.5	Debt-income correlation = 0.9 (FRB, BEA)

Notes: One model period is a quarter, parameters and targets are annualized. Sources: Current Population Survey (CPS), Survey of Consumer Finances (SCF), Federal Reserve Board (FRB), Bureau of Economic Analysis (BEA).

Cross-sectional distribution of borrowing constraints. The vector $[\phi(\theta)]_\theta$ for the idiosyncratic component of borrowing constraints is chosen to match, by construction, the distribution of unsecured household debt across the five income group in the data for the bottom 80% of households.⁹ Unsecured debt 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. Normalizing $\phi(\theta_1) = 1$ in the lowest income group, the resulting vector of relative borrowing limits is equal to $[1, 1.03, 1.06, 1.08, 2.33]$. These values capture the increase of individual borrowing limits with income. They replicate the dispersion of household debt by income that is key for households' response to borrowing constraint shocks.

⁹There is no mechanism in the model to generate high wealth inequality at the top (e.g., return heterogeneity, "superstar" income levels).

Cross-sectional distribution of progressive transfers and taxes. The vector of progressive transfers $[T(\theta)]_\theta$ and the vector of income tax slopes $[\tau_1(\theta)]_\theta$ are similarly chosen to replicate the distribution of transfers and taxes on working-age households across the five income groups in the data ([Congressional Budget Office \(2006\)](#), Exhibit 18). Normalizing $T(\theta_1) = 1$ in the lowest income group, the resulting vector of relative transfers is $T(\theta) = [1, 0.43, 0.24, 0.17, 0.13]$. Average transfers are decreasing in household income, and they are steeply decreasing at low income levels. In level, transfers represent 6.9% of average income. I apply a constant multiplicative factor to the vector of transfers to match this share in the model. The vector for tax slopes by income is $\tau_1(\theta) = [0.05, 0.13, 0.17, 0.20, 0.28]$. Income taxes are increasing with income, especially at high income levels.

3.2 Internal Parameters

The following parameters are chosen to match household balance sheets and the dynamics of the risk-free rate.

Risk aversion. The coefficient of relative risk aversion γ is chosen to provide the best overall fit of the model with the data for a given set of internally calibrated parameters. I obtain $\gamma = 5$, which has conservative implications for savings given that it is lower than standard values in finance. Higher values imply higher savings and a lower risk-free rate. For instance, [Favilukis, Ludvigson, and Van Nieuwerburgh \(2017\)](#) study a model where stochastic borrowing constraints that be either high or low and choose $\gamma = 8$.

Liquid savings. The demand for savings arises from households' intertemporal consumption smoothing and precautionary motives. Households accumulate liquid savings and borrow with unsecured debt by buying and selling one-period risk-free bonds. The supply of savings is endogenous and comes from risk-free bonds issued by the household sector, $\int_- b_{t+1} d\lambda_t$, and the government, $B_{t+1} = B$, as in standard closed economy

models.¹⁰ The demand for liquid savings consist of the positive part of the distribution of bonds, $\int_+ b_{t+1} d\lambda_t$. In the data, liquid savings 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 net supply of liquid savings $B = \int b_{t+1} d\lambda_t = \int_+ b_{t+1} d\lambda_t - \int_- b_{t+1} d\lambda_t$ is chosen to match the value of liquid assets to income, $\int_+ b_{t+1} d\lambda_t / Y_t$, of 1.78. I obtain $B = 6$, which delivers a close value of 1.73.

Discount factor. The average discount factor $\beta = 0.93$ is chosen to match the average real risk-free rate. The risk-free rate is measured in the data as the average of annual Treasury Inflation Indexed 5-year yield of 1.80% between 2000 and 2018 (Federal Reserve Board, H.15 Selected Interest Rates).

Households have heterogeneous discount factors, which helps the model generate the same amount of wealth inequality as in the data, as in [Krusell and Smith \(1998\)](#) and [Favilukis, Ludvigson, and Van Nieuwerburgh \(2017\)](#). The discount factor for households with the lowest of all five productivity levels, who account for 6.25% of the population, is set to be 20% than the average β . This value is chosen to generate the same fraction of borrowing-constrained households in this group as in the data by making them more impatient to consume. Without this heterogeneity, the transfers to these households in the data are so large that they would be less constrained than households in higher income group, which would be counterfactual.

Labor supply. The curvature of the disutility of work hours $\eta = 2$ generates a Frisch elasticity of labor supply of 0.5 that aligns with empirical estimates for micro data ([Whalen and Reichling \(2017\)](#)). This is a key moment to match as households can also adjust their labor supply to insure against aggregate risk, which typically leads them to save less.

¹⁰See, e.g., [Huggett \(1993\)](#) and [Aiyagari \(1994\)](#). Government debt in excess of households' savings is held by investors outside the model. This is a plausible assumption for the U.S. where debt held by the public is on average 40% of average income. For tractability, all closed economy models of household savings follow this approach.

The level of the disutility of work hours $\psi = 11.5$ is chosen to normalize average quarterly household income to $Y = 1$.

3.3 Aggregate Risk

Productivity shocks. The persistence and the volatility of aggregate productivity z are calibrated externally using historical data on total factor productivity (Fernald (2014)). They are set to $\rho_z = 0.55$ and $\sigma_z = 0.026$.

Countercyclical income risk, i.e., the dependence of the volatility σ_θ on aggregate productivity z , is calibrated by extrapolating estimates from Storesletten, Telmer, and Yaron (2004). 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 $0.09/(1.5/0.5)=0.03$, which represents 25% of the average volatility.¹¹

Borrowing constraint shocks. The average aggregate component of borrowing constraints that is common across households, $\bar{\phi}$, is chosen to match the ratio of average unsecured debt to income of 0.18 in the Flow of Funds. It can be interpreted as capturing average credit supply. Unsecured debt is computed as total household liabilities minus mortgage debt (table B.100, line 34). The resulting parameter $\bar{\phi} = 2.6$ delivers a close value for unsecured debt to income value of 0.23.

To guarantee that my estimates for precautionary motives are consistent with empirical values for the risk-free rate, which is a key determinant of household savings, I internally calibrate the process for the tightness of borrowing constraints to match the dynamics of the real risk-free rate in the data. Importantly, in Section 5, I empirically validate these estimates by showing that they imply borrowing constraints that are con-

¹¹Countercyclical income risk does not affect the deterministic steady state with zero aggregate productivity shocks, but it slightly amplifies the economy's response to shocks in both the first- and the second-order solutions of the model. It affects the first-order solution of the model, despite only changing moments of higher order, because it increases the dispersion of individual income, hence the probability that households hit their borrowing constraints and are prevented from smoothing consumption.

sistent with changes in credit limits and lending standards in the post-Great Recession period. The persistence ρ_ϕ and the volatility σ_ϕ of the aggregate component of borrowing constraints are chosen to match the persistence and the volatility of the real risk-free rate that are respectively equal to 0.65 and 1.90%, in annual terms. I obtain quarterly values of $\rho_\phi = 0.69$ and $\sigma_\phi = 0.10$. Using sensitivity analyses, I show that these estimates are well identified (Appendix Figure A1). Overall, they imply slightly less variation than some estimates directly based on changes in credit card limits, which can be as large as income changes (see, e.g., Fulford (2015)). Therefore, my estimates of precautionary motives should be interpreted as a *lower bound*.

This modeling approach, as in other work on stochastic borrowing constraints, tractably captures the fact that borrowing limits on some outstanding loans and on all new vary over time. A potential concern is that it can overstate the volatility in borrowing if changes in credit limits do not immediately lead to changes in household debt. However, several considerations mitigate this concern. First, in practice, credit card lenders are allowed to change credit limits, which accounted for about a third of consumer credit in 2008. For instance, credit limits have changed on average for more than a third of accounts every quarter since 1999, decreased by 36% conditional on changing, and 7% of households have completely lost credit access (Federal Reserve Bank of New York, Consumer Credit Panel). While lenders cannot ask for repayment when limits change on outstanding loans, they can apply extra charges after 45 days, which have a similar effect (one period is a quarter in the model). Second, credit limits on outstanding student and auto loans are fixed but they can be changed when these loans become delinquent.

Finally, the correlation between borrowing constraints and productivity $\rho_{\phi z}$ is chosen to match the correlation between outstanding total consumer credit owned and securitized (Federal Reserve 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, though it is not large enough to quantitatively match the procyclicality of household debt.

3.4 Model Fit

Table 2 describe the fit of the model, which matches key empirical moments of households' balance sheets. The upper panel reports targeted moments, and the lower panel reports untargeted moments that capture the tightness of borrowing constraints. The model replicates well the ratios of aggregate liquid assets to income of 1.78 and unsecured debt to income of 0.18. Importantly, it exactly matches the average and the dynamics of the risk-free rate, which is a key determinant of household savings. This guarantees that precautionary motives in the model are correctly estimates. By construction, the model also replicates the distribution of unsecured debt for the five productivity groups in the model, from the lowest to the highest income. Borrowing capacity increases with households' income group, though only slightly for the first four groups, before almost doubling for the highest group. Finally, the model correctly accounts for two alternative margins that households can use to smooth consumption in addition to savings. First, the model matches the distribution of government transfers and taxes across the five income groups. Transfers, hence the insurance value of social safety nets, are strongly decreasing at the bottom of the income distribution, and then decrease almost linearly for higher income groups. Therefore, they are an important substitute to precautionary savings for the poorest households. Second, the model matches the micro Frisch elasticity of labor supply of 0.5, which measures the change in hours worked in response to wage changes over the business cycle. Households can increase their labor supply instead or in addition to saving when their borrowing capacity deteriorates.

The model also generates realistic non-targeted moments. Crucially, it matches the share of 33% of borrowing-constrained households in the data, which is computed as the fraction of households that report being without savings ([The Pew Charitable Trusts \(2015\)](#)). This value is also close to the 21% share of hand-to-mouth households, which is another measure of the tightness of borrowing constraints ([Kaplan and Violante \(2014\)](#)). The model generates substantial inequality in the distribution of liquid savings across households, with an average to median ratio of savings equal to 4.90. This ratio captures respectively the low and high savings at the bottom and at the top of the distribu-

tion. It is about three quarters of its value of 6.57 in the data. Finally, the model matches well the distribution of liquid assets to income at the bottom of the distribution, as illustrated by the 10th and 25th percentiles that are close to zero. The borrowing constraints of these households are either binding or close to. The precautionary motive due to borrowing constraint fluctuations is the largest for slightly richer, “middle-class” households in higher percentiles of the distribution, who save much more in liquid assets relative to their income. Compared to the data, the median household tends to save more, but this difference, as well as the difference for the ratio of average to median savings, reflects the fact that the model is calibrated for the bottom 80% of households and does not have a mechanism to generate high wealth inequality at the top of the distribution (return heterogeneity, “superstar” income levels).

Table 2: Aggregate and distribution moments

	Data	Model
<i>Targeted</i>		
Aggregate liquid savings/income	1.78	1.73
Aggregate unsecured debt/income	0.18	0.23
Unsecured debt distribution by income	[1, 1.03, 1.06, 1.08, 2.33]	[1, 1.03, 1.06, 1.08, 2.33]
Government transfer distribution by income	[1, 0.43, 0.24, 0.17, 0.13]	[1, 0.43, 0.24, 0.17, 0.13]
Government taxes distribution by income	[0.05, 0.13, 0.17, 0.20, 0.28]	[0.05, 0.13, 0.17, 0.20, 0.28]
Frisch elasticity of labor supply	0.5	0.5
Risk-free rate average	1.80%	1.80%
Risk-free rate volatility	1.90%	1.90%
Risk-free rate persistence	0.65	0.65
<i>Non-targeted: borrowing constraint tightness</i>		
Share borrowing-constrained households	0.33	0.35
Mean/median savings	6.57	4.90
P10 liquid savings/income	0	0
P25 liquid savings/income	0.01	0
P50 liquid savings/income	0.15	0.30

Notes: Upper panel: targeted moments. Lower panel: non-targeted moments describing the tightness of households’ borrowing constraints. One period is a quarter, targets are annualized. Sources: FRB, CPS, SCF, [Boar, Gorea, and Midrigan \(2021\)](#), [Whalen and Reichling \(2017\)](#), [The Pew Charitable Trusts \(2015\)](#).

4 The Financial and Real Effects of Aggregate Precautionary Motives

This section presents the main results on the dynamic effects 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 changes in borrowing constraints is especially large, contrary to received wisdom about the low cost of business cycles. Second, I decompose the impact of aggregate risk in the cross-section of households and highlight the understudied role of "middle class" net savers. Third, I analyze the real implications of borrowing constraint risk for consumption and highlight a new recessionary mechanism.

4.1 Average Effect: The Role of Borrowing Constraint Risk

Estimates. As shown previously, the difference between equilibrium coefficients in the second-order and first-order perturbations of the model provides a measure of aggregate precautionary savings motives that relies on the model departure from certainty equivalence with respect to aggregate risk. Table 3 reports the corresponding averages for equilibrium variables to highlight the roles of aggregate shocks to borrowing constraints and productivity. The first column reports average values in the baseline steady state of the economy. The second column reports the contribution of aggregate borrowing constraint volatility σ_ϕ , and the third column reports the contribution of aggregate productivity volatility σ_z , in percentage deviations from the stationary steady state without aggregate shocks.

Borrowing constraint risk. Aggregate fluctuations in borrowing constraints increase households' savings to income and lower debt to income by about 45%. In equilibrium, they lead to a 25.4% lower risk-free rate than in a model with fixed borrowing constraints. Even when aggregate shocks themselves are temporary, aggregate volatility generates the low-debt and low-rate environment of the post-Great Recession period. Furthermore, borrowing constraint risk has a recessionary effect, which Subsection 4.3 explores in more

Table 3: Average effect of aggregate precautionary motives

Variable	(1) Baseline model	(2) Borrowing constraint volatility σ_ϕ	(3) Productivity volatility σ_z
Savings/income	1.73	+45.4%	< 0.1%
Debt/income	0.23	-45.1%	< 0.1%
Interest rate	1.80%	-25.4%	-0.1%
Wage	1.49	+0.9%	< 0.1%
Profits	0.33	-1.5%	< 0.1%
Employment	0.44	-2.5%	< 0.1%
Output	1.00	-1.6%	< 0.1%
Consumption	1.00	-1.6%	< 0.1%

Notes: Columns 2 and 3 report equilibrium differences, in percentage deviations, between the baseline model with aggregate risk (column 1) and counterfactual models without aggregate risk. Markets clear for each model, other parameters are fixed at their baseline values. One period is a quarter, variables are annualized.

detail. The volatility σ_ϕ decreases consumption, output, and profits by about 1.5%, and employment by 2.5%. In equilibrium, the wage increases by 0.9% to prevent households from further decreasing their work hours. Overall, the contribution of aggregate fluctuations in borrowing constraints to precautionary savings is sizable. This impact is increasing in risk aversion γ , in the persistence ρ_ϕ and volatility σ_ϕ of borrowing constraints, and in inequality in borrowing constraints between households $\phi(\theta)$.

Several assumptions guarantee that these estimates are a *lower bound* on aggregate precautionary motives. First, the endogenous risk-free rate ensures that savings are not overstated as they would be with a fixed rate. The model captures the equilibrium feedback loop through which higher risk increases savings, which decreases the equilibrium rate, which makes saving less attractive. It also matches the level and the dynamics of the risk-free rate, which largely determine savings. Second, households save and borrow at the same rate. With a higher borrowing rate, households would have a stronger incentive to save. Third, households can increase their labor supply and receive government transfers instead of saving when their borrowing constraints become more binding. So savings are not overstated for the poorest households.

Productivity risk. Importantly, the same counterfactual with respect to the volatility of

aggregate productivity shocks σ_z demonstrates that their contribution to precautionary behavior is negligible. Aggregate productivity risk only slightly decreases the risk-free rate by 0.10%, while other variables remain unchanged. These results are consistent with low costs of business cycles in the literature though they highlight that these costs are much higher when changes in credit conditions are accounted for. The role of borrowing constraint risk is explained both by the higher volatility of the constraints themselves and their much higher impact on households' ability to smooth consumption.

Decomposition. The estimates in Table 3 generate a decomposition of the various precautionary motives, which is obtained by solving the model with and without the various sources of idiosyncratic and aggregate risk. The model separately identifies three motives.

First, the standard *idiosyncratic motive* in heterogeneous household models has the largest impact on savings. The combination of income risk, prudence $u''' > 0$, and static borrowing constraints (see, e.g., [Kimball \(1990\)](#)) increases liquid savings by 88 pp from 31% to 119% of average income (in annual terms) compared to an economy without risk where they are only determined by intertemporal substitution. In equilibrium, it decreases the risk-free rate by 5.4 pp from 7.80%, when it is given by $1/\beta - 1$, to 2.40%. Individual income volatility increases future expected marginal utility because of Jensen's inequality, which increases savings in households' Euler equation. Finally, the combination of income shocks and borrowing constraints hampers consumption smoothing for households that are at or near their borrowing constraints.

Second, the *aggregate financial motive* due to stochastic borrowing constraints further increases savings by 54 pp from 119% to 173% of average income, which represents about 60% of the effect of idiosyncratic risk. This sizable impact cannot be ignored when analyzing households' balance sheets. In equilibrium, it further lowers the risk-free rate by 0.6 pp from 2.40% to 1.80%, which represents more than 10% of the effect of idiosyncratic risk. Interestingly, the aggregate financial motive leads to substantial costs of business cycles as it decreases consumption by 1.6% on average. This result is an important departure from existing settings that focuses on real business cycles and find close to zero effect.

Third, the impact of the *aggregate real motive* due to productivity shocks is negligible. This motive acts in the same way as the idiosyncratic motive as aggregate productivity shocks also result in individual labor income changes. However, these shocks are both much smaller and less persistent than idiosyncratic income shocks. Moreover, they have a uniform impact across the distribution of households.

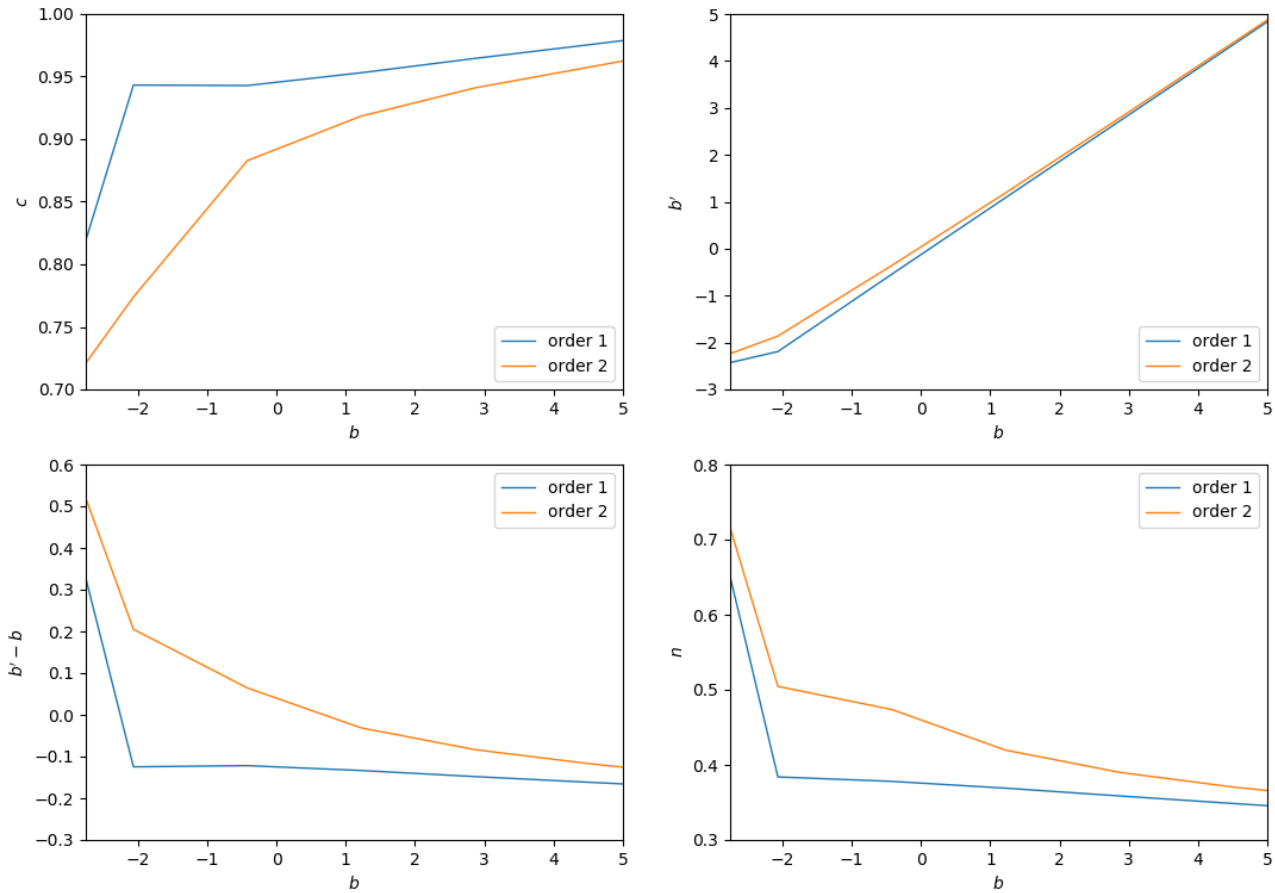
4.2 Heterogeneous Effects: The Role of “Middle-Class” Households

Given these large average effects, what are the effects of aggregate precautionary motives in the cross section of households? Figure 1 reports nonlinear policy functions to decompose their effects across various levels of net savings b for a household with the median income. The blue lines depict 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).

Borrowing constraint risk leads households to consume less c , save $b' - b$ and work n more, as shown by the differences between the orange and the blue 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 annual ratio of debt to income of 0.23, between 10% and 50% of average income. Compared to these households, precautionary motives are lower for the very poorest group by at least a factor of two.

Replicating the same comparison for productivity groups θ_1 to θ_5 instead of net savings 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 both poorer households with θ_1 and richer households with θ_4 and θ_5 . Middle-class households θ_2 and θ_3 , benefit less than the poorest from the progressivity of government taxes and transfers, and they have less liquid assets than the richest. The highest-income households are the only ones to consume slightly more. The lower risk-free rate increases their incentive to consume the consumption good and leisure.

Figure 1: Heterogeneous effects of aggregate precautionary motives on households' choices



Notes: Policy functions describing how households with the median income optimally choose consumption (upper left panel), the stock and flow of net savings (resp. upper right and lower left), and hours worked (lower right) as a function of their current savings. The blue lines (order 1) depict policy functions without aggregate precautionary motives and the orange lines (order 2) show their heterogeneous effects for various levels of net savings.

Mechanism. What explains the large and heterogeneous effects of stochastic borrowing constraints? Because of the departure from certainty equivalence, households anticipate the potential volatility of borrowing constraints and they save to insure 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.

The strongly nonlinear effect of shocks to borrowing constraints also explains their heterogeneous impact across households. The second-order solution of the model shows

that the economy's response to a tightening of borrowing constraints is amplified because households tend to respond more when they are close to their constraints. In contrast, nonlinearities are negligible for aggregate productivity shocks, which affect all households identically. This nonlinear effect is also absent from models with the near-aggregation property, in which prices and quantities respond linearly to aggregate shocks as they would in a representative agent economy (Krusell and Smith (1998)).

4.3 Real Effects: The Recessional Impact of Borrowing Constraints

What are the real effects of these motives? This subsection shows that, contrary to aggregate productivity shocks, changes in households' borrowing constraint have large impacts on consumption and employment, which lead to higher estimates of the cost of business cycles than previously thought. The reason is a new recessionary mechanism of borrowing constraints, for which I provide empirical support in Section 5 below.

Impact on business cycle volatility. Table 4 reports the result from a variance decomposition that quantifies the contributions of aggregate shocks to productivity and borrowing constraints to business cycle volatility. This computation is a new exercise in a heterogeneous agent model with incomplete markets and aggregate risk. It relies on the solution approach in Section 2 and uses the nonlinear laws of motion of the economy (see Appendix B.2 for details).

The results show that borrowing constraint fluctuations are responsible for a large fraction of business cycle volatility that dominates the role of aggregate productivity. More than half of the volatility in output, consumption, and employment is due to borrowing constraints. In contrast, aggregate productivity mostly explains changes in wages, which scale one for one with z . Interestingly, fluctuations in borrowing constraints have real effects even without price rigidity because they lead to a decrease in total employment when credit contracts.

Recessionary mechanism. To explain these effects, I compare the economy's response to aggregate productivity and borrowing constraint shocks, both in the model without

Table 4: Contributions of aggregate shocks to business cycle volatility

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

Notes: Variance decomposition: shares of the volatility of business cycle moments that are accounted for by borrowing constraints (second column) and productivity shocks (third column). 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). The computations use $N = 500$ simulations.

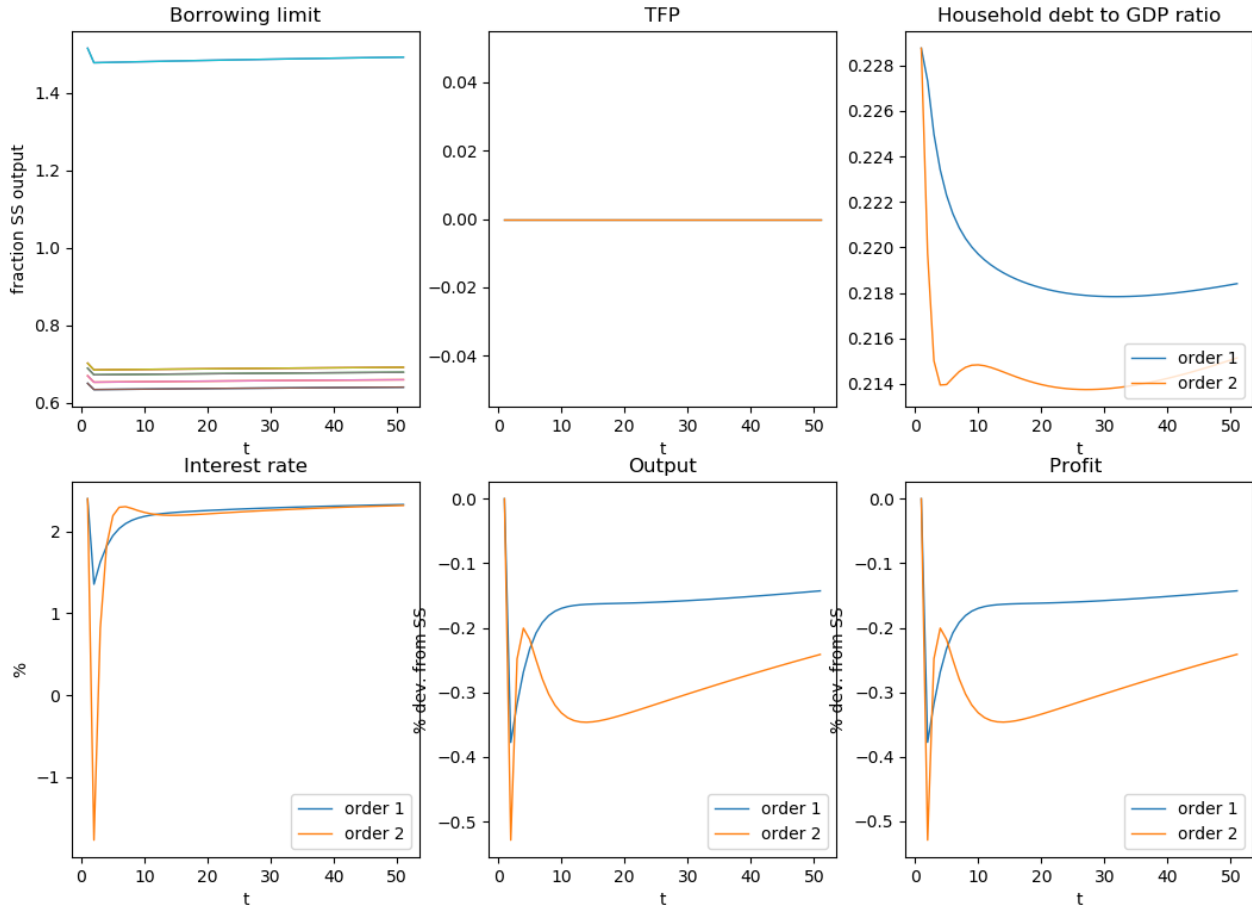
and with aggregate precautionary motives.

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 precautionary motives (order 1), and the nonlinear dynamics with precautionary motives (order 2). Aggregates are computed using the time-varying path of individual policy functions and histogram weights. Deviations are from the steady state. These results show that all variables stay persistently low following a tightening of households’ borrowing constraints with precautionary motives.

Borrowing constraints are tightened for all households, but lower income groups are able to borrow less than higher income groups, reflecting individual differences in their ability to borrow. Constrained households are forced to reduce their debt and increase their net savings. This leads them to decrease their consumption of goods and leisure. They trade off working more to smooth consumption against the disutility of labor. Debt to income decreases persistently, as well as output, which is equal to consumption in equilibrium. The decrease in consumption results from the composition of low-income constrained households decreasing their spending, and richer unconstrained households increasing theirs because they earn a lower return on their savings. The decrease in the risk-free rate equates the larger savings demand from the former with the lower demand from the latter.

Amplification. The amplification of shocks to borrowing constraints is large for house-

Figure 2: Household responses to aggregate borrowing constraint shocks



Notes: Nonlinear impulse response functions to a one standard deviation shock to borrowing constraints ϵ^ϕ . Borrowing constraints (upper left panel) are shown as a fraction of annual steady state output for low to high income households (resp. lowest to highest line). The other five panels compare impulse response functions in the first-order (linear, in blue) and second-order perturbations (nonlinear, in orange) of the model. Initial period: steady state. One period is a quarter, variables are annualized.

hold debt and the risk-free rate. The initial impact is amplified by a factor of 5 for debt to income, 4 for the risk-free rate, 1.5 for consumption, output and profit, and 1.4 for the wage. While the decrease in the risk-free rate (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. The large persistence ρ_ϕ , which induces borrowing constraints to stay persistently low, further decreases consumption, employment, and debt to income. The decrease in the risk-free rate cannot offset the quantity restriction imposed by tighter constraints.

Employment. Appendix Figure A3 plots the response of employment, which households can use to insure against shocks to their borrowing capacity. It shows a recessionary effect of borrowing constraints, for which I find empirical support in Section 5.

There is a composition effect between, first, less productive and constrained households who work more to smooth consumption when they are forced to deleverage and, second, more productive and unconstrained households who consume more leisure as they decrease their savings because of a wealth effect. The sign of the overall output and employment response depends on which effects dominates. In the model, output declines because more productive agents work less, despite less productive agents working more. This result is due to stochastic borrowing constraints and is absent from models with fixed borrowing constraints. In these models with deterministic shocks and perfect foresight, households near the constraint choose to work more because they anticipate the full trajectory of credit shocks, at odds with the data. In contrast, in a model with stochastic borrowing constraints, households anticipate that shocks will mean-revert, which leads less of them to increase their work hours.

This effect is reinforced by the intertemporal substitution effect that is due to the decrease in the risk-free rate, which induces all households to consume more leisure in the current period. Overall, this leads to a decrease in total employment. This generates an economy with persistently low household debt and rates, but also depressed output, which helps the model match the post-Great Recession data.

5 Empirical Implications for Household Credit

This section highlights the empirical implications of aggregate precautionary motives and provides empirical support for the model calibration. Using novel structural estimates of shocks to borrowing constraints from the model, I show that precautionary motives are key to explain the simultaneous household deleveraging and consumption recovery, which is a central feature of the post-Great Recession period. Structural estimates of borrowing constraints are consistent with the data, though they uncover a slightly milder and more persistent credit contraction than suggested by lending surveys. Such estimates

help better understand the credit landscape faced by borrowers.

5.1 Explaining the Post-Great Recession Puzzle

Empirical test. I empirically validate the model calibration by examining its ability to match the *dynamics* of household balance sheets and consumption in the post-Great Recession period. From 2009 to 2019, household debt and the risk-free rate remained persistently low while consumption quickly recovered from the recession. This “creditless recovery” is common to multiple credit crunches over time and across countries (see, e.g., [Claessens, Kose, and Terrones \(2009\)](#)). Yet, models with fixed borrowing constraints typically fail to explain such changes. Indeed, higher future consumption should increase current debt and the risk-free rate through an intertemporal substitution motive. Conversely, a lower rate should induce households to consume more early on, instead of saving and increasing their future consumption. Therefore, addressing this puzzle is a stringent test of the model. In addition, this exercise requires matching the entire time series for the main variables, in addition to their averages in the calibration of Section 3. The results show that the combination of stochastic borrowing constraints and aggregate productivity risk is key to help the model explain the data.

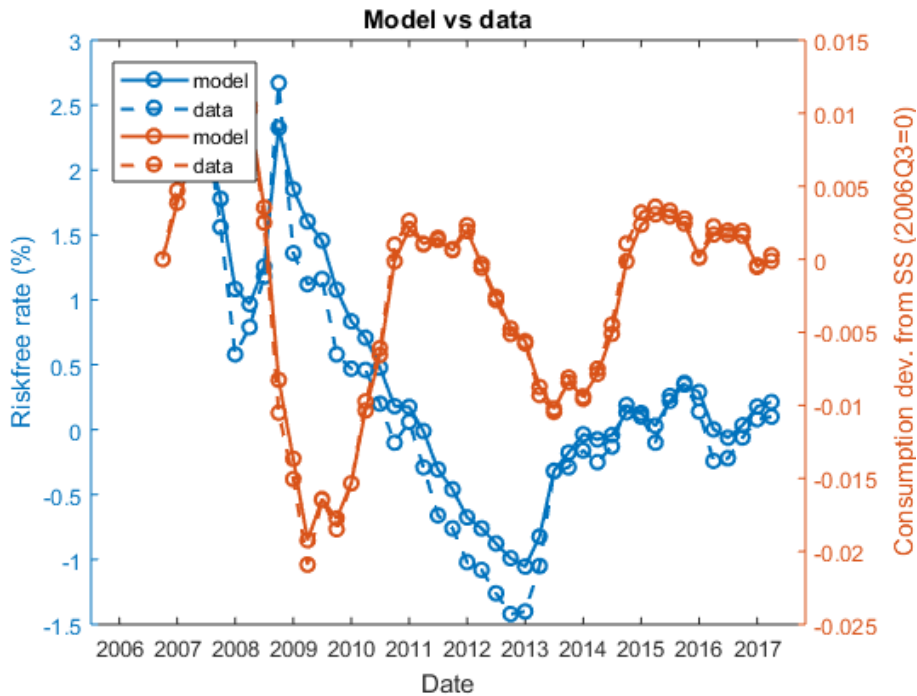
Model fit. Using a particle filter, I estimate the trajectories of structural shocks to productivity and borrowing constraints that generate the observed dynamics for consumption and the real risk-free rate. Aggregate precautionary motives and the nonlinear dynamics of the model are key to match the data in times of high volatility. These are, to the best of my knowledge, the first dynamic estimates of structural shocks in a heterogeneous agent model with incomplete markets and aggregate risk. This is a computationally challenging exercise, which relies on the model representation of Section 2. Appendix B.3.2 details the estimation procedure.

In the data, the time series for the real risk-free rate is measured as the 5-year treasury inflation-indexed securities constant maturity rate (H.15 Selected Interest Rates, Federal Reserve Board). 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 to make it comparable with the model.

Figure 3 reports the fit of the model after estimating the shocks. The model successfully addresses the post-Great Recession puzzle. It exactly matches the dynamics of aggregate consumption and it almost exactly matches the risk-free rate, except around 2013 in the vicinity of the Taper Tantrum when it is slightly higher.

Figure 3: Model fit and solution of post-Great Recession puzzle




Notes: Risk-free rate (in annual percentage terms; on left axis, in blue) and consumption deviation from 2006Q3 value (on right axis, in orange), predicted by particle filtering in the nonlinear version of the model with precautionary savings (solid lines) vs. the data (dashed lines). $N = 20,000$ particles simulated. Sources: FRB, BEA (quarterly samples). One period is a quarter.

Out-of-sample fit. I further validate the calibration by examining its ability to replicate the dynamics of debt and employment. Figure 4 shows that it also closely matches these dynamics, which were not targeted by the estimation. In the data, household 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) since the model has a continuum of measure 1 of households and therefore aggregate employment N is the ratio of the employed to the entire population.

First, the model replicates the hump-shaped dynamics of household debt to income, which starts with the run-up to the crisis until 2008, then decreases and eventually increases again around 2015. The model overstates the decrease in household debt in the last part of the sample, for which the large persistence of borrowing constraint shocks may be responsible. Second, importantly, the model matches the decrease in employment in the data that follows the credit contraction and is partly responsible for the decrease in consumption. This result provides empirical support for the recessionary effect of borrowing constraints highlighted in Section 4. Similarly to household debt, the decrease in employment is slightly overstated by the model, by about 10%, at the end of the sample.

Figure 4: Out-of-sample model fit for household debt and employment



Notes: Debt/income (left panel) and employment (right panel) implied by the nonlinear estimation of the model with a particle filter. The solid lines depict model results and the dashed lines the corresponding data. Variables are in log-deviations from their 2006Q3 values. One period is a quarter.

5.2 The Dynamics of Borrowing Constraints

Model estimates. I conclude by showing that the model implies realistic dynamics for both borrowing constraints and productivity, providing empirical support for their calibration in Section 3. I highlight what policymakers can learn from these estimates in

addition to lending surveys.


Figure 5 plots the dynamic estimates for aggregate shocks to households' borrowing constraints and productivity in the post-Great Recession period. They uncover two central patterns. First, there was a *persistent* tightening in borrowing constraints throughout the period, which were far from having recovered at the end of the sample. Constraints fell by almost 20% from 2007 to 2014. Then, they rebounded by 6 pp, before slowing down their recovery as they only increased by 1 pp from 2014 to 2018. Second, there was a *V-shaped* recession in productivity. Productivity only fell by slightly more than 2% during the NBER-dated recession itself (2008-2009), but then it quickly reverted to its previous level in less than two years.

These two ingredients allow the model to simultaneously explain the consumption recovery and the persistent decrease in the risk-free rate and household debt. The decrease in household debt to income by almost 30% and the 3 pp decrease in the risk-free rate (in annual terms) result from the large aggregate precautionary motive that arises from the tightening of households' borrowing constraints. The precautionary motive, which is due to slow-moving borrowing constraints, is exacerbated by the short-lived drop in aggregate productivity, which induces constrained households to deleverage and save quickly. The risk-free rate decreases to clear the savings market and stays persistently low as constraints remain tight.

Comparison with data. Importantly, the model generates realistic estimates of aggregate shocks, which closely track estimates from survey data. First, the 18% tightening of households' borrowing constraints in the model aligns with the 25% decrease in credit card limits in the data (Consumer Credit Panel, Federal Reserve Bank of New York). Credit limits steadily declined from 2008 to 2011 before only slowly going back to their previous values, which they reached only in 2020. During the same period, 20% of borrowers had a credit card account closed, which completely prevented them from accessing credit. These changes have important effects on borrowing constraints as more than three quarters of U.S. households have a credit card.

Second, the onset of the tightening of borrowing constraints in the model exactly

Figure 5: Dynamic estimates of borrowing constraints and productivity



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.

matches the data, including at quarterly frequency. The results are the same as in the Consumer Credit Panel, in which credit card limits first decreased in the first quarter of 2008. I also use survey data from the Senior Loan Officer Opinion Survey on Bank Lending Practices (Federal Reserve Board) to confirm this result, which are directly reported in Appendix Figure A4. There was a sharp increase in the net percentage of lenders that reported tightening their lending standards in the first quarter of 2008 as well. Though this measure is hard to translate into actual changes in borrowing constraints, the model provides the first structural estimate of this pass-through, which is easier to interpret. In the survey, the share of lenders tightening credit conditions increased by 70 percentage points for both credit cards and auto loans from 2008 to 2011. In comparison, the model estimates suggest that the actual tightening of constraints was milder but more persistent.

Third, the model estimates for aggregate productivity shocks closely track empirical measures of total factor productivity in the data, both qualitatively and quantitatively, providing further external validation. TFP variations in the data are reported in Appendix Figure A5. They are measured from the Penn World Tables at constant national prices for the U.S. The widely used measure of Fernald (2014) delivers similar results. Aggregate

productivity fell between 1% and 2% compared to its previous trend between 2007 and 2010. Then, it quickly reverted back to its previous trend. The model exactly matches this pattern without targeting it.

6 Conclusion

Little is known about the drivers of households' precautionary savings beyond idiosyncratic risk to their income, health, or family situation. Perhaps especially surprisingly, while economy-wide changes in borrowing constraints are the subject of a rich recent empirical literature, it is still unclear whether and how households adjust to this risk.

Using a model with heterogeneous households and incomplete markets, this paper presents a new dynamic decomposition of precautionary motives that highlights the role of such aggregate risk. The aggregate precautionary motive that arises from fluctuations in borrowing constraints is especially large, which refines received wisdom about the low costs of business cycles. This motive is the strongest and leads to relatively higher savings for middle-class households with average levels of debt, in contrast with the recent focus of economists on the top and bottom of the wealth distribution.

Aggregate precautionary motives are a key ingredient to help widely used models capture periods of "creditless recoveries" such as the post-Great Recession era. Interestingly, while the dynamic estimates of borrowing constraints from the model track measures of credit conditions in the data based on lending surveys, they imply a slightly milder and more persistent credit contraction, which paints a more nuanced picture of the recent credit landscape. These estimates are also useful to translate popular measures of credit conditions (such as the share of lenders tightening standards) into direct changes in borrowing constraints, which can be more easily interpreted by policymakers and call for different responses.

More generally, my results highlight that environments with various degrees of aggregate risk have widely different implications for household savings. Studying these differences across countries may be a useful direction for work on international household finance (e.g., [Badarinza, Campbell, and Ramadorai \(2016\)](#) and [Haliassos, Jansson,](#)

and Karabulut (2017)).

References

- Acharya, V. V., K. Bergant, M. Crosignani, T. Eifert, and F. McCann. 2022. The anatomy of the transmission of macroprudential policies. *The Journal of Finance* 77:2533–75.
- Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel. 2018. Do banks pass through credit expansions to consumers who want to borrow? *Quarterly Journal of Economics* 133.
- Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf. 2017. When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics Annual* .
- Aiyagari, S. R. 1994. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics* 109:659–84.
- Badarınza, C., J. Y. Campbell, and T. Ramadorai. 2016. International comparative household finance. *Annual Review of Economics* 8:111–44.
- Baker, S. R. 2018. Debt and the response to household income shocks: Validation and application of linked financial account data. *The Journal of Political Economy* 126.
- Bertaut, C. C., and M. Haliassos. 1997. Precautionary portfolio behavior from a life-cycle perspective. *Journal of Economic Dynamics and Control* 21:1511–42.
- Bertaut, C. C., M. Haliassos, and M. Reiter. 2009. Credit card debt puzzles and debt revolvers for self control. *Review of Finance* 13:657–92.
- Boar, C., D. Gorea, and V. Midrigan. 2021. Liquidity constraints in the u.s. housing market. *The Review of Economic Studies* .
- Carroll, C. D. 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters* 91:312–20.

- Carroll, C. D., and A. A. Samwick. 1997. The nature of precautionary wealth. *Journal of Monetary Economics* 40:41–71.
- Cherry, S., E. Jiang, G. Matvos, T. Piskorski, and A. Seru. 2022. Government and private household debt relief during covid-19. *Brookings Papers on Economic Activity* .
- Claessens, S., M. A. Kose, and M. E. Terrones. 2009. What happens during recessions, crunches and busts? *Economic Policy* 24:653–700.
- Congressional Budget Office. 2006. The distribution of federal spending and taxes in 2006. Working Paper.
- Das, S., C. M. Kuhnen, and S. Nagel. 2019. Socioeconomic status and macroeconomic expectations. *The Review of Financial Studies* 33:395–432.
- De Nardi, M., E. French, and J. B. Jones. 2010. Why do the elderly save? the role of medical expenses. *Journal of Political Economy* 118:39–75.
- Deaton, A. 1991. Savings and liquidity constraints. *Econometrica* 59:1221–48.
- Duchin, R., T. Gilbert, J. Harford, and C. Hrdlicka. 2017. Precautionary savings with risky assets: When cash is not cash. *The Journal of Finance* 72:793–852.
- Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh. 2017. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy* 125:140–223.
- Fernald, J. G. 2014. A quarterly, utilization-adjusted series on total factor productivity. Working Paper, Federal Reserve Bank of San Francisco.
- Fulford, S. L. 2015. How important is variability in consumer credit limits? *Journal of Monetary Economics* 72:42–63.
- Gourinchas, P.-O., and J. A. Parker. 2002. Consumption over the life cycle. *Econometrica* 70:47–89.

- Gross, D. B., and N. S. Souleles. 2002. Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. *The Quarterly Journal of Economics* 117:149–85.
- Guerrieri, V., and G. Lorenzoni. 2017. Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics* 132:1427–67.
- Guiso, L., T. Jappelli, and D. Terlizzese. 1996. Income risk, borrowing constraints, and portfolio choice. *The American Economic Review* 86:158–72.
- Haliassos, M., and C. Hassapis. 2002. *Borrowing constraints, portfolio choice, and precautionary motives*. . Computational Methods in Decision-Making, Economics and Finance. Springer US.
- Haliassos, M., T. Jansson, and Y. Karabulut. 2017. Incompatible european partners? cultural predispositions and household financial behavior. *Management Science* 63:3780–808.
- Heaton, J., and D. J. Lucas. 1996. Evaluating the effects of incomplete markets on risk sharing and asset pricing. *The Journal of Political Economy* 104:443–87.
- Hubbard, R. G., J. Skinner, and S. P. Zeldes. 1995. Precautionary saving and social insurance. *The Journal of Political Economy* 103:360–99.
- Huggett, M. 1993. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control* 17:953–69.
- Jones, C., V. Midrigan, and T. Philippon. 2022. Household leverage and the recession. *Econometrica* 90:2471–505.
- Kaplan, G., and G. L. Violante. 2014. A model of the consumption response to fiscal stimulus payments. *Econometrica* 82:1199–239.
- Kim, J., S. Kim, E. Schaumburg, and C. A. Sims. 2008. Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models. *Journal of Economic Dynamics and Control* 32:3397–414.

- Kimball, M. S. 1990. Precautionary savings in the small and in the large. *Econometrica* 58:53–73.
- Klein, P. 2000. Using the generalized schur form to solve a multivariate linear rational expectations model. *Journal of Economic Dynamics and Control* 24:1405–23.
- Kopecky, K., and R. Suen. 2010. Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics* 13:701–14.
- Krusell, P., and A. Smith. 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106:867–96.
- Lanne, M., and H. Nyberg. 2016. Generalized forecast error variance decomposition for linear and nonlinear multivariate models. *Oxford Bulletin of Economics and Statistics* 78:595–603.
- Lucas, R. E. 1987. *Models of business cycles*. Wiley-Blackwell.
- Ludvigson, S. 1999. Consumption and credit: a model of time-varying liquidity constraints. *The Review of Economics and Statistics* 81:434–47.
- Mian, A., K. Rao, and A. Sufi. 2013. Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics* 128:1687–726.
- Mian, A., A. Sufi, and E. Verner. 2017. Household debt and business cycle worldwide. *The Quarterly Journal of Economics* 132:1755–817.
- Parker, J. A., and B. Preston. 2005. Precautionary savings and consumption fluctuations. *American Economic Review* 95:1119–43.
- Reiter, M. 2009. Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control* 33:649–65.
- Rouwenhorst, K. G. 1995. *Asset pricing implications of equilibrium business cycle models*, 294–330. Princeton University Press. ISBN 9780691043234.

- Schmitt-Grohe, S., and M. Uribe. 2008. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control* 28:755–75.
- Sims, C. A. 2001. Solving linear rational expectations models. *Computational Economics* 20:1–20.
- Storesletten, K., C. I. Telmer, and A. Yaron. 2004. Cyclical dynamics in idiosyncratic labor market risk. *The Journal of Political Economy* 112:695–717.
- Tauchen, G. 1986. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters* 20:177–81.
- The Pew Charitable Trusts. 2015. Americans' financial security. Brief.
- Whalen, C., and F. Reichling. 2017. Estimates of the frisch elasticity of labor supply: A review. *Eastern Economic Journal* 43:37–42.
- Zeldes, S. P. 1989. Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics* 104:275–98.

Appendix

A Dynamic Decomposition of Precautionary Motives

A.1 Detailed Solution Approach

Step 1: Projection. Variables are indexed by time t to denote their dependence on aggregate states $(\bar{\phi}_t, z_t, \lambda_t)$. 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:

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

where $b_{j'-1}, b_{j'}, b_{j'+1}$ 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. $b'_t(\theta_i, b_j) = b'(\theta_i, b_j; \bar{\phi}_t, z_t, \lambda_t)$. For instance, if credit shocks $\bar{\phi}_t$ are low, tightening borrowing constraints, this distorts and shifts upwards the function $b'(\cdot)$ because households are forced to save more, which through its impact on ω results in less mass on low asset values.

1. Household saving and labor supply policy functions are interpolated using linear splines with respectively N_b and N_n knots. Households' saving function $b'(\cdot)$ 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 $0 = x_1 < \dots < x_{N_b}$ denote the splines' knots for b' at which households' Euler equations hold with equality. For $b \leq \zeta_\theta$, savings $b'(\theta, b) = -\bar{\phi}\phi(\theta)h(\theta)$ are determined by the borrowing limit ($\bar{\phi}_t = \bar{\phi}$ 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 $(\bar{\phi}, z, \Phi)$, 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:

$$c_t(\theta, b_{\theta,j}) = b_{\theta,j} + (1 - \tau_1(\theta)) w_t \theta n_t(\theta, b_{\theta,j}) + T(\theta) + \pi_t - \tau_{0t} - \frac{b'_t(\theta, b_{\theta,j})}{1 + r_t} \quad (24)$$

2. 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, 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 bor-

rowing constraints and wage influence current decisions), and is taken with respect to their values at t :

$$\begin{aligned}
& \Phi_{t+1}(\theta_{i'}, b_{j'}) - \sum_{i=1}^{N_\theta} \sum_{j=1}^{N_b^f} \Pi_\theta(\theta_{i'} | \theta_i) (\omega_{i,j,j',t} \Phi_t(\theta_i, b_j)) = 0, \quad i' \in [1, N_\theta], j' \in [1, N_b^f] \\
& \log \bar{\phi}_{t+1} - \log \bar{\phi} - \rho_\phi (\log \bar{\phi}_t - \log \bar{\phi}) - \epsilon_{t+1}^\phi = 0 \\
& \log z_{t+1} - \rho_z \log z_t - \epsilon_{t+1}^z = 0 \\
& \begin{pmatrix} \epsilon^\phi \\ \epsilon^z \end{pmatrix} \overset{iid}{\sim} \mathcal{N} \left(\mathbf{0}, \begin{pmatrix} \sigma_\phi^2 & \sigma_\phi \sigma_z \rho_{\phi z} \\ \sigma_\phi \sigma_z \rho_{\phi z} & \sigma_z^2 \end{pmatrix} \right) \\
& B - \sum_{i=1}^{N_\theta} \sum_{j=1}^{N_b^f} b_j \Phi_{t+1}(\theta_i, b_j) = 0 \\
& N_t - \sum_{i=1}^{N_\theta} \sum_{j=1}^{N_b^f} \theta_i n_t(\theta_i, b_j) \Phi_t(\theta_i, b_j) \\
& Y_t - \sum_{i=1}^{N_\theta} \sum_{j=1}^{N_b^f} c_t(\theta_i, b_j) \Phi_t(\theta_i, b_j) \\
& C_t - \sum_{i=1}^{N_\theta} \sum_{j=1}^{N_b^f} c_t(\theta_i, b_j) \Phi_t(\theta_i, b_j) \\
& w_t = \alpha z_t \left(\frac{1}{N_t} \right)^{1-\alpha} \\
& \pi_t = (1 - \alpha) z_t K^{1-\alpha} N_t^\alpha \\
& (1 - \tau_1(\theta_i)) w_t \theta_i c_t(\theta_i, b_j)^{-\gamma} - \psi n_t(\theta_i, b_j)^\eta = 0, \quad i \in [1, N_\theta], j \in [1, N_b] \\
& c_t(\theta_i, b_j)^{-\gamma} - \beta(1 + r_t) \mathbb{E}_t \left\{ \sum_{i'=1}^{N_\theta} c_{t+1}(\theta_{i'}, b'(\theta_i, b_j))^{-\gamma} \right\} = 0, \quad i \in [1, N_\theta], j \in [1, N_b]
\end{aligned} \tag{25}$$

Step 2: Stationary steady state. 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 $\mathcal{F}(\cdot)$ that represents the nonlinear system of equations that defines the equilibrium:

$$\mathbb{E}_t [\mathcal{F}(\mathbf{y}_t, \mathbf{y}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t+1}, \epsilon_{t+1}^q, \epsilon_{t+1}^z)] = 0 \tag{26}$$

Solving for the deterministic steady state of the economy (without aggregate shocks) amounts to finding \mathbf{y}, \mathbf{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 pro-

jection 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 \mathbf{x} and \mathbf{y} ,¹² 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.

1. Start with a guess for the risk-free rate and labor demand $(p^{(0)}, N^{(0)}, \psi^{(0)})$, for policy function values $(\mathbf{b}'^{(0)}(\cdot), \mathbf{n}^{(0)}(\cdot))$, and the cross-sectional distribution $\Phi^{(0)}(\cdot)$ (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 $(r^{(0)}, w^{(0)})$.
2. Given those, use the endogenous grid method 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, $(\mathbf{b}'^{(1)}(\cdot), \mathbf{n}^{(1)}(\cdot))$. This requires comput-

¹²A good guess is obtained by using the endogenous grid method (Carroll (2006)) to iterate backwards on the household's optimality conditions, starting from any feasible guess.

ing 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.

3. 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 ($p^{(n)}, N^{(n)}, \psi^{(n)} > 0$), 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 $n = 1$.
4. Given the exogenous law of motion for idiosyncratic income and the policy functions, compute the associated stationary distribution of households $\Phi^{(1)}(\cdot)$ (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 $\tau_0^{(1)}$.
5. Given prices, profits, taxes, and savings policies $\mathbf{b}'^{(1)}(\cdot)$, solve the household's labor supply equation (using $\mathbf{n}^{(1)}(\cdot)$ as a guess), and denote $\mathbf{n}^{(2)}(\cdot)$ 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.
6. 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 $\mathbf{n}^{(2)}(\cdot)$ 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.
7. Given prices, profits, taxes and the new labor policy $\mathbf{n}^{(2)}(\cdot)$, solve the household's Euler equation (using $\mathbf{b}'^{(1)}(\cdot)$ as a guess), and denote $\mathbf{b}'^{(2)}(\cdot)$ the new savings policy. Use the same solver as for labor.

8. This completes one iterate in the loop solving for policy functions given prices. If the new policy functions $\left(\mathbf{n}^{(2)}(\cdot) \mathbf{b}'^{(2)}(\cdot)\right)$ are close enough to the previous ones $\left(\mathbf{n}^{(1)}(\cdot) \mathbf{b}'^{(1)}(\cdot)\right)$, then stop and we have solved the household's problem given prices $\left(p^{(0)}, N^{(0)}, \psi^{(0)}\right)$. Otherwise, iterate on steps (d)-(g). That is, given $\left(p^{(0)}, N^{(0)}, \psi^{(0)}\right)$ (hence the same wages and profits), compute new government taxes $\tau_0^{(n+1)}$. Then solve for new policy functions $\left(\mathbf{n}^{(n+1)}(\cdot) \mathbf{b}'^{(n+1)}(\cdot)\right)$, compare them to the previous ones $\left(\mathbf{n}^{(n)}(\cdot) \mathbf{b}'^{(n)}(\cdot)\right)$, and stop when they are close enough. This completes the solution of the household's problem given prices.
9. Using the law of motion of the exogenous income shock and the optimal savings 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, $\left(p^{(1)}, N^{(1)}, \psi^{(1)}\right)$, to solve the following three equations:

$$\begin{aligned}
B - \sum_{i=1}^{N_\theta} \sum_{j=1}^{N_b^f} b_j \Phi_{t+1}^{(2)}(\theta_i, b_j) &= 0 \\
N^{(1)} - \sum_{i=1}^{N_\theta} \sum_{j=1}^{N_b^f} \theta_i n(\theta_i, b_j) \Phi^{(2)}(\theta_i, b_j) & \\
Y^{(1)} - 1 = 0 \Leftrightarrow \left(N^{(1)}\right)^\alpha - 1 = 0 &
\end{aligned} \tag{28}$$

10. 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.

Step 3: Perturbations. Implement a first- and a second-order perturbation of the discrete model around its steady state. The solutions to the equilibrium expectational differ-

ence equation $\mathbb{E}_t [\mathcal{F}(\cdot)] = 0$ are of the following form (Schmitt-Grohe and Uribe (2008)):

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{h}(\mathbf{x}_t, \eta) + \eta \begin{pmatrix} \mathbf{0} \\ \epsilon_{t+1}^q \\ \epsilon_{t+1}^z \end{pmatrix} \\ \mathbf{y}_t &= \mathbf{g}(\mathbf{x}_t, \eta) \end{aligned} \quad (29)$$

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 \mathbf{h}, \mathbf{g} .

1. 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)), which commonly used in macroeconomics. I briefly describe the input widely 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 \mathbf{g}, \mathbf{h} . Writing variables in deviations from their steady state values (denoted as \widehat{x}, \widehat{y}) and linearizing equilibrium conditions around 0 (where variables equal their steady state values), we obtain

$$\mathcal{F}_{\mathbf{y}_t} \widehat{\mathbf{y}}_t + \mathcal{F}_{\mathbf{y}_{t+1}} \mathbb{E}_t [\widehat{\mathbf{y}}_{t+1}] + \mathcal{F}_{\mathbf{x}_t} \widehat{\mathbf{x}}_t + \mathcal{F}_{\mathbf{x}_{t+1}} \mathbb{E}_t [\widehat{\mathbf{x}}_{t+1}] + \mathcal{F}_{\epsilon_{t+1}^q} \mathbb{E}_t [\widehat{\epsilon}_{t+1}^q] + \mathcal{F}_{\epsilon_{t+1}^z} \mathbb{E}_t [\widehat{\epsilon}_{t+1}^z] = 0 \quad (30)$$

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

$$\begin{aligned} & \left\{ \left[N_\theta \times N_b^f - 1 + 2 \right] + \left[6 + N_\theta \times (N_n + N_b + 1) \right] \right\} \\ & \times \left\{ 2 \times \left[N_\theta \times N_b^f - 1 + 2 \right] + 2 \times \left[6 + N_\theta \times (N_n + N_b + 1) \right] + 2 \right\} \end{aligned}$$

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

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

2. 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 \mathcal{F} , a large three-dimensional array computed by automatic differentiation, of dimension:

$$\begin{aligned}& \left\{ \left[N_\theta \times N_b^f - 1 + 2 \right] + \left[6 + N_\theta \times (N_n + N_b + 1) \right] \right\} \\ & \times \left\{ 2 \times \left[N_\theta \times N_b^f - 1 + 2 \right] + 2 \times \left[6 + N_\theta \times (N_n + N_b + 1) \right] + 2 \right\}^2\end{aligned}$$

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

$$\begin{aligned}\widehat{\mathbf{x}}_{t+1} &= \mathbf{h}_x(\mathbf{x}, 0) \widehat{\mathbf{x}}_t + \frac{1}{2} \mathbf{h}_{xx}(\mathbf{x}, 0) \widehat{\mathbf{x}}_t^2 + \frac{1}{2} \mathbf{h}_{\eta\eta}(\mathbf{x}, 0) \eta^2 + \eta \begin{pmatrix} \mathbf{0} \\ \epsilon_{t+1}^q \\ \epsilon_{t+1}^z \end{pmatrix} \\ \widehat{\mathbf{y}}_{t+1} &= \mathbf{g}_x(\mathbf{x}, 0) \widehat{\mathbf{x}}_t + \frac{1}{2} \mathbf{g}_{xx}(\mathbf{x}, 0) \widehat{\mathbf{x}}_t^2 + \frac{1}{2} \mathbf{g}_{\eta\eta}(\mathbf{x}, 0) \eta^2\end{aligned}\tag{32}$$

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

cretized 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.¹³ 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 Table A1).

Pruning 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 $\mathbb{E}_t [\mathcal{F}(\cdot)] = 0$ 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):

$$\begin{aligned}\widehat{\mathbf{x}}_{t+1} &= F_1\hat{\mathbf{x}}_t + F_2\eta\mathbf{ffl}_{t+1} + F_3\eta^2 + \frac{1}{2}F_{11}\hat{\mathbf{x}}_t^2 + F_{12}\hat{\mathbf{x}}_t\mathbf{ffl}_{t+1}\eta + \frac{1}{2}F_{22}\eta^2\mathbf{ffl}_{t+1}^2 \\ \hat{\mathbf{y}}_t &= \frac{1}{2}M_{11}\hat{\mathbf{x}}_t^2 + M_2\eta^2\end{aligned}\tag{33}$$

The presence of cross-derivative terms in the transformed solution does not contradict their absence in the original solution, since they can be canceled by Z . Then, it implies

¹³I 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.

that for $s > 0$:

$$\begin{aligned}
\mathbb{E}_t [\widehat{\mathbf{x}}_{t+s}] &= F_1 \mathbb{E}_t [\widehat{\mathbf{x}}_{t+s-1}] + F_3 \eta^2 + \frac{1}{2} F_{11} \mathbb{E}_t [\widehat{\mathbf{x}}_{t+s-1}^2] + \frac{1}{2} F_{22} \eta^2 \Omega_s \\
&= F_1 \mathbb{E}_t [\widehat{\mathbf{x}}_{t+s-1}] + F_3 \eta^2 + \frac{1}{2} F_{11} \left(\mathbb{E}_t [\widehat{\mathbf{x}}_{t+s-1}^2] + \Sigma_{s-1} \right) + \frac{1}{2} F_{22} \eta^2 \Omega_s \\
\mathbb{E}_t [\widehat{\mathbf{y}}_{t+s}] &= \frac{1}{2} M_{11} \mathbb{E}_t [\widehat{\mathbf{x}}_{t+s}^2] + M_2 \eta^2 \\
&= \frac{1}{2} M_{11} \left(\mathbb{E}_t [\widehat{\mathbf{x}}_{t+s}^2] + \Sigma_s \right) + M_2 \eta^2 \\
\Sigma_{s+1} &= \eta^2 F_2 \Omega_t F_2 + F_1 \Sigma_s F_1
\end{aligned} \tag{34}$$

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

$$\begin{aligned}
\mathbb{E}_t [\widehat{\mathbf{x}}_{t+s}] &= F_1 \mathbb{E}_t [\widehat{\mathbf{x}}_{t+s-1}] + F_3 \eta^2 + \frac{1}{2} F_{11} \left(\mathbb{E}_t^1 [\widehat{\mathbf{x}}_{t+s-1}]^2 + \Sigma_{s-1} \right) + \frac{1}{2} F_{22} \eta^2 \Omega_s \\
\mathbb{E}_t [\widehat{\mathbf{y}}_{t+s}] &= \frac{1}{2} M_{11} \left(\mathbb{E}_t^1 [\widehat{\mathbf{x}}_{t+s}^2] + \Sigma_s \right) + M_2 \eta^2 \\
\mathbb{E}_t^1 [\widehat{\mathbf{x}}_{t+s}] &= F_1 \mathbb{E}_t^1 [\widehat{\mathbf{x}}_{t+s-1}] \\
\Sigma_{s+1} &= \eta^2 F_2 \Omega_t F_2 + F_1 \Sigma_s F_1
\end{aligned} \tag{35}$$

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 $\eta = 1$):

$$\begin{aligned}
\mathbb{E} [\widehat{\mathbf{x}}] &= (I - F_1)^{-1} \left(F_3 + \frac{1}{2} F_{22} \Omega + \frac{1}{2} F_{11} \Sigma \right) \\
\mathbb{E} [\widehat{\mathbf{y}}] &= \frac{1}{2} M_{11} \Sigma + M_2 \\
\text{where } \Sigma &= F_2 \Omega_t F_2 + F_1 \Sigma F_1
\end{aligned} \tag{36}$$

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

B Applications

B.1 Nonlinear Impulse Response Functions to Aggregate Shocks

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

$$\begin{aligned}\widehat{\mathbf{x}}_{t+1} &= F_1 \widehat{\mathbf{x}}_t + F_2 \mathbf{ffl}_{t+1} + F_3 + \frac{1}{2} F_{11} \widehat{\mathbf{x}}_t^2 + F_{12} \widehat{\mathbf{x}}_t \mathbf{ffl}_{t+1} + \frac{1}{2} F_{22} \mathbf{ffl}_{t+1}^2 \\ \widehat{\mathbf{y}}_t &= \frac{1}{2} M_{11} \widehat{\mathbf{x}}_t^2 + M_2 \\ \widehat{\mathbf{x}}_{t+1}^1 &= F_1 \widehat{\mathbf{x}}_t^1 + F_2 \mathbf{ffl}_{t+1}\end{aligned}\tag{37}$$

I then back out the path of original variables as $\left\{ Z'^{-1} \left(\widehat{\mathbf{x}}_t \quad \widehat{\mathbf{y}}_t \right)' \right\}_t$.

B.2 Variance Decomposition: Contributions of Aggregate Risks

B.2.1 First Order: No Precautionary Motives

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 $\epsilon_{t+1} = \begin{pmatrix} \epsilon_{t+1}^\phi & \epsilon_{t+1}^z \end{pmatrix}'$ 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 $\text{Var} \left(\begin{pmatrix} \sigma_\phi & 0 \\ 0 & \sigma_z \end{pmatrix} \tilde{\Sigma}_\epsilon \right) = \begin{pmatrix} \sigma_\phi^2 & \rho_{\phi,z} \sigma_\phi \sigma_z \\ \rho_{\phi,z} \sigma_\phi \sigma_z & \sigma_z^2 \end{pmatrix} = \Sigma_\epsilon$.

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 $v_t = Q \epsilon_t$, with Q such that $\mathbb{E}[v_t v_t'] = I$. Denoting $S = Q^{-1}$, $\epsilon_t = S v_t$ 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 $\epsilon_{t+1} = Sv_{t+1}$.

$$\begin{aligned}
(I - \Phi L) Y_{t+1} &= Z\epsilon_{t+1} \\
\Rightarrow Y_{t+1} &= (I - \Phi L)^{-1} ZSv_{t+1} \\
Y_{t+1} &= \sum_{k=0}^{\infty} \Phi^k L^k ZSv_{t+1} \\
Y_{t+1} &= \sum_{k=0}^{\infty} \Phi^k ZSv_{t+1-k} \\
\Rightarrow Y_{t+h} &= \sum_{k=0}^{\infty} \tilde{\Phi}^{(k)} v_{t+h-k}
\end{aligned} \tag{39}$$

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:

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

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

$$\begin{aligned}
e_{j,t+h} &= \sum_{i=1}^h \tilde{\Phi}_{j,\cdot}^{(h-i)} v_{t+i} \\
&= \sum_{i=1}^h \left(\tilde{\Phi}_{j,1}^{(h-i)} v_{1,t+i} + \tilde{\Phi}_{j,2}^{(h-i)} v_{2,t+i} \right)
\end{aligned} \tag{41}$$

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

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

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:

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

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

B.2.2 Second Order: The Role of Precautionary Motives

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

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

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:

$$\begin{aligned} GI_j(i, \delta_{\phi,t+1}, \omega_t) &= \mathbb{E}_t \left[Y_{j,t+i} | \epsilon_{t+1}^\phi = \delta_{\phi,t+1}, \omega_t \right] - \mathbb{E}_t \left[Y_{j,t+i} | \omega_t \right] \\ \text{and } GI_j(i, \delta_{z,t+1}, \omega_t) &= \mathbb{E}_t \left[Y_{j,t+i} | \epsilon_{t+1}^z = \delta_{z,t+1}, \omega_t \right] - \mathbb{E}_t \left[Y_{j,t+i} | \omega_t \right] \end{aligned} \quad (45)$$

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:

$$\begin{aligned} GFEVD_j(h, \delta_{\phi,t}) &= \frac{\sum_{i=0}^h GI_j(i, \delta_{\phi,t+1}, \omega_t)^2}{\sum_{i=0}^h GI_j(i, \delta_{\phi,t+1}, \omega_t)^2 + \sum_{i=0}^h GI_j(i, \delta_{z,t+1}, \omega_t)^2} \\ GFEVD_j(h, \delta_{z,t}) &= \frac{\sum_{i=0}^h GI_j(i, \delta_{z,t+1}, \omega_t)^2}{\sum_{i=0}^h GI_j(i, \delta_{\phi,t+1}, \omega_t)^2 + \sum_{i=0}^h GI_j(i, \delta_{z,t+1}, \omega_t)^2} \end{aligned} \quad (46)$$

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 $T = 1000$, $\{\epsilon_t^\phi, \epsilon_t^z\}_{t=0}^T = \{\delta_{\phi,t}, \delta_{z,t}\}_{t=0}^T$ using $\begin{pmatrix} \epsilon_t^\phi \\ \epsilon_t^z \end{pmatrix} \stackrel{iid}{\sim} \mathcal{N}(0, I_2)$ (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.

B.3 Dynamic Estimation of Aggregate Shocks

B.3.1 Kalman Filter: No Precautionary Motives

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

$$\begin{aligned} Y_{t+1} &= \Phi Y_t + Z \epsilon_{t+1}, \quad \epsilon_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, Q) \\ Y_{t+1}^{obs} &= H' Y_{t+1} + v_t, \quad v_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, R) \end{aligned} \quad (47)$$

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 consump-

tion). 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 variance-covariance matrix of Y_t given the history of observables until $t-1$),

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

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

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 $\Sigma_{1|0} = \Phi \Sigma_{1|0} \Phi' + Z I_2 Z'$, 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 $t \geq 1$.
2. Given $\Sigma_{t|t-1}, Y_{t|t-1}^{obs}, Y_{t|t-1}^{obs}$, compute $\Omega_{t|t-1} = H' \Sigma_{t|t-1} H + R = H' \Sigma_{t|t-1} H$.
3. Compute $\text{Cov}_{t-1} (Y_t^{obs}, Y_t) = \mathbb{E} \left[\left(Y_t^{obs} - Y_{t|t-1}^{obs} \right) \left(Y_t - Y_{t|t-1} \right)' | Y^{obs,t-1} \right] = H' \Sigma_{t|t-1}$.
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).

B.3.2 Particle Filter: The Role of Precautionary Motives

The second-order, nonlinear state space representation of the model is obtained from `gensys2`. Using the above notation, the transition and the measurement equations are respectively:

$$\begin{aligned} Y_{t+1} &= f(Y_t, W_{t+1}), \quad W_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, Q) \\ Y_{t+1}^{obs} &= H'Y_{t+1} + v_t, \quad v_{t+1} \stackrel{iid}{\sim} \mathcal{N}(0, R) \end{aligned} \tag{48}$$

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 $i = 1, \dots, N$, 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 $\left\{ Y_{t|t-1}^i, W_{t|t-1}^i \right\}_{i=1}^N$ 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 $p(Y_t^{obs} | w_{t|t-1}^i, Y_{t-1}^{obs}, Y_{t|t-1}^i)$ using the measurement equation and the distribution of measurement errors v . That is,

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

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

4. Evaluate the relative weights $q_t^i = \frac{p(Y_t^{obs} | w_{t|t-1}^i, Y_{t-1}^{obs}, Y_{t|t-1}^i)}{\sum_{j=1}^N p(Y_t^{obs} | w_{t|t-1}^j, Y_{t-1}^{obs}, Y_{t|t-1}^j)}$, normalized to be probabilities.
5. Re-sample, with replacement, N values $\left\{ Y_{t|t-1}^i, W_{t|t-1}^i \right\}_{i=1}^N$ 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 $\left\{ Y_{t+1|t}^i, W_{t+1|t}^i \right\}_{i=1}^N$. 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 $\{Y_t^{obs}\}_{t=0}^T$ from the data. In the main text, I plot the sample averages of these empirical conditional densities at $t = 0, \dots, T$. 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 par-

allelized over the N dimension. It takes about 12 hours to run the case $N = 20,000, T = 44$ using 28 cores.

C Additional Model Results

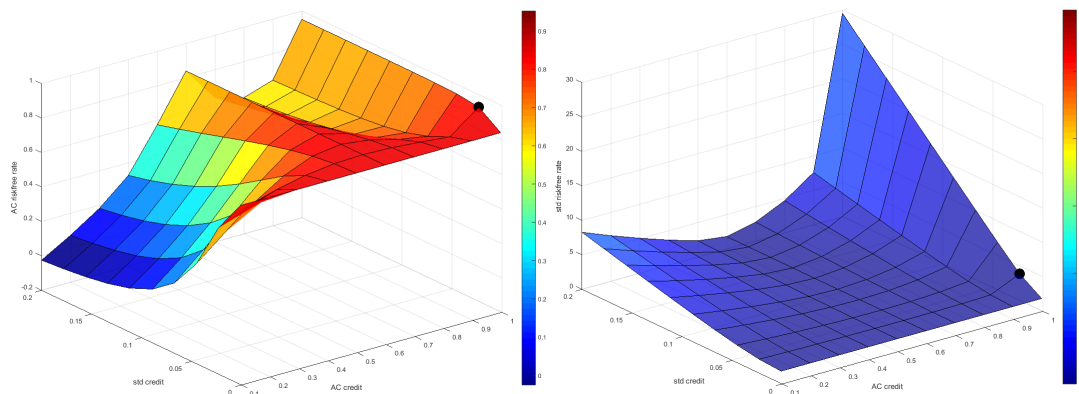
C.1 Solution Accuracy

Table A1: Market clearing errors

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: Market clearing errors for impulse response functions, average (main) and maximum values (in parentheses). Errors are computed as percentage differences normalized by the steady state value of the variable or by the initial value of the series. They are reported for the first-order (first row) and the second-order perturbations of the model (second row).

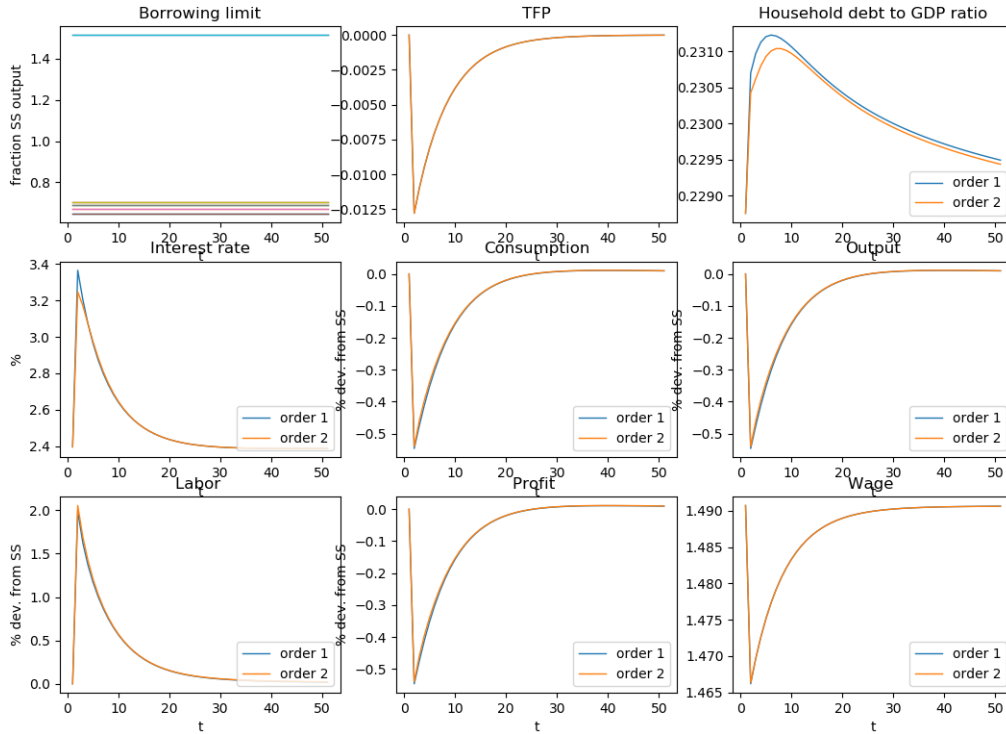
Figure A1: Effect of the volatility and persistence of borrowing constraints on equilibrium risk-free rate



Notes: These two graphs show that the volatility and persistence of borrowing constraints in the calibration are well identified by matching these same moments for the risk-free rate. They depict the persistence (left panel) and the volatility (in annual percentage terms, right panel) of the risk-free rate as two-dimensional functions of the persistence and volatility of borrowing constraints. These equilibrium values are estimated in a simulation of the model with $T = 10,000$ periods. On each graph, the black dot represents the baseline calibration. In both cases it lies in non-flat areas of the (ρ, σ) surfaces.

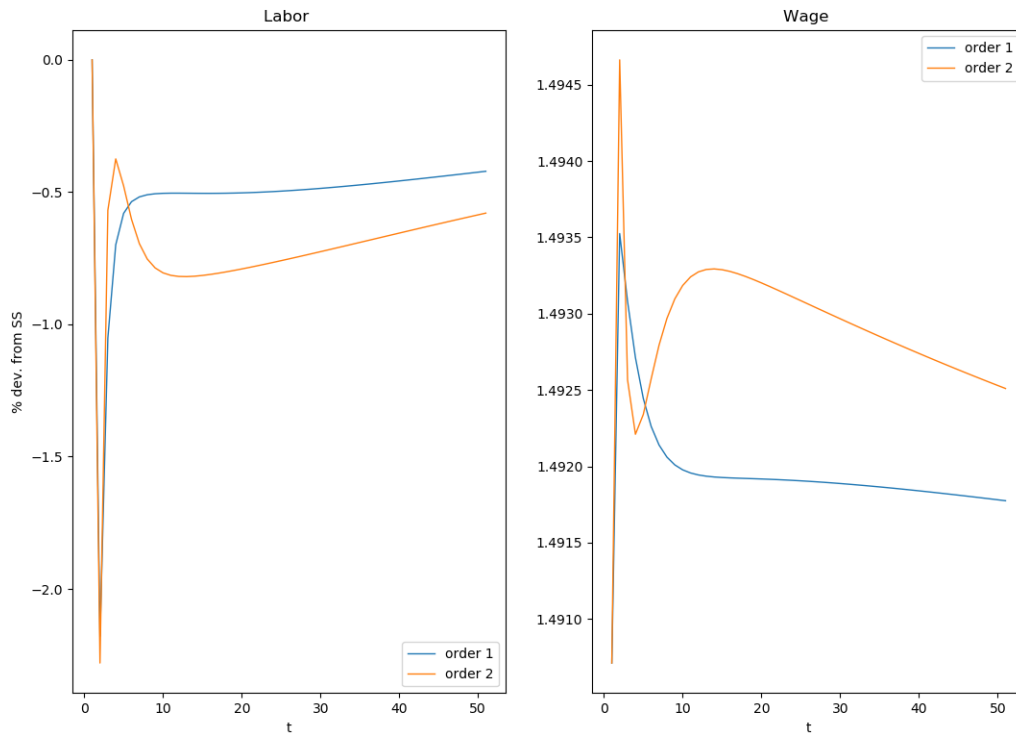
C.2 Impulse Response Functions

Figure A2: Household responses to aggregate productivity shock



Notes: Nonlinear impulse response functions to a one standard deviation shock to aggregate productivity e^z . Borrowing constraints (upper left panel) are shown as a fraction of annual steady state output for low to high income households (resp. lowest to highest line). The other five panels compare impulse response functions in the first-order (linear, in blue) and second-order perturbations (nonlinear, in orange) of the model. Initial period: steady state. One period is a quarter, variables are annualized.



Figure A3: Aggregate labor responses to borrowing constraint shocks



Notes: Nonlinear impulse response functions to a one standard deviation shock to borrowing constraints ϵ^ϕ , for aggregate employment and the equilibrium wage. The two panels compare impulse response functions in the first-order (linear, in blue) and second-order perturbations (nonlinear, in orange) of the model. Initial period: steady state. One period is a quarter, variables are annualized.

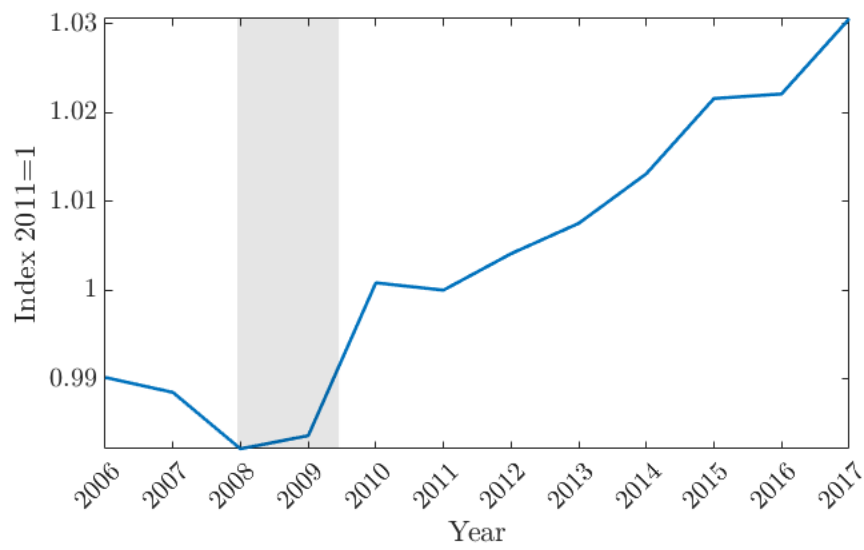
C.3 Comparison with Post-Great Recession Data

Figure A4: Lending standards for consumer credit in survey data



Source: Senior Loan Officer Opinion Survey on Bank Lending Practices, Federal Reserve Board, January 2024. One period is a quarter.

Figure A5: Aggregate productivity in the data



Source: Penn World Table 9.0. Aggregate productivity is measured as total factor productivity at constant national prices for the United States. The shaded area represents the NBER-dated recession.