Financial Frictions and the Wealth Distribution

Jesús Fernández-Villaverde

University of Pennsylvania, NBER, and CEPR

Samuel Hurtado — Banco de España

Galo Nuño — Banco de España*

June 17, 2019

Abstract

We postulate a continuous-time model of financial frictions *à la* Basak and Cuoco (1998) and Brunnermeier and Sannikov (2014), but with a non-trivial distribution of wealth among households. We show how such a model can be efficiently computed, despite its substantial non-linearities and time-varying wealth distributions, using tools from machine learning. We illustrate how the model can be structurally estimated with a likelihood function with aggregate and micro observations, using tools from inference with diffusions. Regarding results, we document, first, the strong non-linearities created by financial frictions. Second, we report the existence of multiple stochastic steady-states with properties that differ from the deterministic steady state along important dimensions. Third, we illustrate how the generalized impulse-response functions of the model are highly state-dependent. In particular, we find that the recovery after a negative aggregate shock is more sluggish when the financial expert is more leveraged. Fourth, we prove that wealth heterogeneity matters in this economy because of the asymmetric responses of household consumption decisions to aggregate shocks.

Keywords: Heterogeneous agents; aggregate shocks; continuous-time; machine learning; neural networks; structural estimation; likelihood functions.

JEL codes: C45, C63, E32, E44, G01, G11.

*We thank Emmanuel Farhi, Xavier Gabaix, Mark Gertler, Ben Moll, Gianluca Violante, Ivan Werning, and participants at numerous seminars for pointed comments. The views expressed in this manuscript are those of the authors and do not necessarily represent the views of the European Central Bank or the Bank of Spain.

1 Introduction

This paper argues that, in a standard heterogeneous agents model with financial frictions, idiosyncratic individual shocks interact with exogenous aggregate shocks to generate endogenously time-varying levels of aggregate risk. In particular, such an economy displays highly nonlinear behavior, with i) multiple stochastic steady states; ii) multimodal, skewed, and fat-tailed ergodic distributions of endogenous variables; and iii) strong state-dependence on the responses of endogenous variables to aggregate shocks.

To do so, we build, compute, and estimate using maximum likelihood and U.S. data a continuous-time neoclassical growth model with heterogeneous households subject to labor productivity shocks as in Aiyagari (1994) and Krusell and Smith (1998). We enrich the model with a financial expert (a stand-in for banks or financial intermediaries), limited financial markets participations as in Basak and Cuoco (1998) and Brunnermeier and Sannikov (2016), and an aggregate shock to capital. Households save in a noncontingent bond to self-insure against their labor productivity shocks and the variations in wages and interest rates induced by aggregate risk. The financial expert cannot issue state-contingent assets (i.e., outside equity), but it can issue a noncontingent bond to leverage its equity and accumulate physical capital that is rented to a representative firm. Only the expert can hold physical capital.

The interaction between the demand for bonds by the households and the supply of bonds by the financial expert begets, for parameter values that match important aspects of the U.S. data, multiple stochastic steady states (SSSs). This multiplicity occurs even despite the model having a unique deterministic steady state (DSS).[1] In particular, we will have a high-leverage SSS and a low-leverage SSS, each of them with their own basin of attraction and endogenous levels of aggregate risk (there is a third, unstable SSS that we do not need to discuss).

The intuition for the existence of multiple SSSs is as follows. When the financial expert is highly levered, negative aggregate shocks have dire consequences: the financial expert net wealth is considerably eroded and the economy suffers a deep and prolonged recession. This recession translates into persistently low wages. Since the households want to self-insure against these low wages, their demand for bonds is high: this SSS has a higher share of wealthy households and more wealth and income inequality than the DSS. This high demand for bonds translates into a low risk-free interest rate for the bond, a high expected excess return for productive capital, and, therefore, justifies the high leverage by the financial expert. High-leverage regions of the economy are regions with high aggregate volatility, even when the variance of the aggregate shock is the same than in the low-leverage region. In other words, the optimal responses of agents in the presence of financial frictions generate high endogenous aggregate risk.

[1] A SSS is a fixed-point of the equilibrium conditions of the model when the *realization* of the aggregate shock is zero. A DSS is a fixed-point of the equilibrium conditions of the model when the *volatility* of the aggregate shock (but not of idiosyncratic shocks) is zero.

To document the higher persistence of shocks when leverage is high, we compute the generalized impulse response functions and the distributional impulse response functions to a negative capital shock in different points of the state space. The responses are very similar on impact, but the ensuing recession is more persistent when leverage is high. This higher persistence is due to the dynamics of aggregate household consumption. In a high-leverage economy, the decline in consumption of wealthy households is less severe. This milder decline produces a slower path of capital accumulation and, hence, creates a slow recovery path. From the households' perspective, the attenuation in the decline of consumption is consistent with the expected path of interest rates, which is more persistent in the high-leverage economy (thus, inducing less intertemporal substitution of consumption) due to capital dynamics.

In comparison, when the financial expert is lowly levered, the recessions after a negative aggregate shock are mild. Thus, the demand for bonds by the household is low, the risk-free rate high, the expected excess return low, and these prices sustain the low leverage by the financial expert. With low leverage, the endogenous aggregate risk of the economy is smaller: the economy responds more mutedly to the same aggregate shock than when leverage is high.

For our baseline parameter values, the economy will spend more time, on average, around the high-leverage SSS than around the low-leverage SSS. We will show, however, how this finding varies as we change the volatility of idiosyncratic and aggregate shocks.

The multiplicity of SSSs does not imply a multiplicity of equilibria: in our model, we find a unique stochastic equilibrium. Whether the economy is a point of high or low leverage is a consequence of the sequence of aggregate shocks in the past. Sometimes, while the economy is traveling in the basin of attraction of the high-leverage SSS, a large aggregate shock will move it to the basin of attraction of the low-leverage SSS (and vice versa). Thus, the interaction of financial frictions with the consumption-saving decisions made by the agents generates endogenous time-variation in aggregate risk: after some sequences of shocks, the economy will be fragile and prone to severe recessions due to high leverage and, after other sequences, the economy will be more resilient and less volatile thanks to low leverage.

The intuition above also explains why heterogeneity at the household level is crucial for our argument and why the "quasi-aggregation result" of Krusell and Smith (1998) breaks down in our environment. As we reduce idiosyncratic labor productivity risk, our model collapses to a version of the model in Brunnermeier and Sannikov (2016). In the absence of idiosyncratic risk, the precautionary saving motive is smaller and the high-leverage SSS disappears: the households are not willing to demand all the bonds required to sustain the high leverage SSS. The presence of a unique SSS does not preclude the existence of long-lived transient dynamics after an aggregate shock (and their associated multimodal ergodic distributions of aggregate variables) as those reported by Brunnermeier and Sannikov (2014), but it fundamentally changes the structure of fluctuations with respect to the case where we have two SSSs.

This result suggests that the rise in wealth inequality among capital owners documented by Alvaredo et al. (2017) before the financial crisis of 2007 is directly linked with the increase in debt and leverage witnessed during the same period (Adrian and Shin, 2010, and Nuño and Thomas, 2017). Thus, changes in micro turbulence (e.g., more volatile labor markets) can translate into higher aggregate volatility even when the variance of aggregate shocks is constant. This mechanism is, as far as we know, new in the literature. It is also different from the "paradox of volatility" in Brunnermeier and Sannikov (2014). In their model, a low volatility of aggregate shocks leads to higher leverage by the financial expert and, thus, deep recessions when a large shock hits the economy. In our model, a high volatility of idiosyncratic shocks leads to higher leverage and, thus, larger recessions on average, even when the aggregate shocks are not particularly large.

Conversely, when we increase the level of idiosyncratic labor productivity risk in the model, the low-leverage SSS disappears and only the high-leverage SSS survives. In this case, households are so concerned about their own idiosyncratic risk that they demand enough bonds as to push the risk-free rate sufficiently low that only high levels of leverage can be a SSS (paradoxically, increasing the aggregate risk they face).

The documentation of the importance of individual heterogeneity in models with financial frictions is a novel contribution of our paper. The insight that, in the presence of financial frictions, the wealth distribution is a state of the economy is not new. The idea is already discussed by Bernanke et al. (1999) and Kiyotaki and Moore (1997). However, the previous literature has focused on the case where there is *between-agents* heterogeneity, but no *within-agents* heterogeneity. *Between-agents* heterogeneity means that capital owners are different from experts. No *within-agents* heterogeneity means there is just one capital owner (a representative household) and a representative expert (or, perhaps, different capital owners and experts but where the heterogeneity is collapsed into an economic-wide average of leverage). Our paragraphs above document how *within-agents* is crucial to understand the aggregate consequences of financial frictions.

Researchers have avoided studying economies with *within-agents* heterogeneity because characterizing models with financial frictions and non-trivial wealth distributions is challenging: since the latter is an infinite-dimensional object, standard dynamic programming techniques cannot be employed. To overcome this problem, our paper also provides new methodological tools for the global, nonlinear solution and estimation of heterogeneous agent models with aggregate shocks.

More concretely, we rely on the machine learning literature and employ a neural network to obtain a flexible approximation of the perceived law of motion (PLM) of the cross-sectional distributions of assets (the expert's equity and households' bonds) with a finite set of moments in the spirit of Krusell and Smith (1998). Naturally, other machine learnings schemes may also

be proposed (or, for the matter, other nonlinear universal approximators such as series expansions or splines). However, our approach is particularly convenient, both regarding theoretical properties and practical considerations.

First, the *universal approximation theorem* (Hornik et al., 1989; Cybenko, 1989) states that a neural network can approximate any unknown Borel measurable function. Second, the neural network breaks the curse of dimensionality for a large class of approximated functions, which allows our method to be extended to much richer environments with many more state variables (or more moments). Third, the neural network can be efficiently trained using a combination of the *gradient descent* and the *back-propagation* algorithms. Fourth, not only is our algorithm efficient and easy to code, but also particularly amenable to massive parallelization in graphic or tensor processing units. Finally, our approach reflects in a fairly transparent way the self-justified equilibria nature of the "bounded rationality" solution of most heterogeneous agents models (Kubler and Scheidegger, 2018). The PLM is computed based on the samples drawn in the simulation of paths within the aggregate ergodic distribution. The agents employ the neural network to extrapolate the dynamics outside of the equilibrium region.

In comparison, in Krusell and Smith (1998) and most of the subsequent literature, the PLM of the aggregate variables is approximately linear in the endogenous state variables (but nonlinear in the exogenous states, since the coefficients of the regression are allowed to vary across shocks). This traditional PLM is a poor choice in our model, in which the nonlinearities of the endogenous state variables play a central role and where, because of the use of continuous time, exogenous states are incorporated into the endogenous states instantaneously. In fact, we will document how a naive implementation of the Krusell and Smith (1998) algorithm in our economy, and even of refinements such as using Chebyshev polynomials, deliver a much worse numerical performance. The use of neural networks allows us to be much more flexible and avoiding having to specify, ex-ante, any concrete nonlinear structure of the PLM.

Continuous time help us to characterize much of the equilibrium dynamics analytically and to worry only about local derivatives (instead of the whole shape of equilibrium functions) even when solving the model globally. However, nothing essential depends on this choice and we could replicate our approach (with somewhat heavier computational costs) in discrete time (see Achdou et al., 2017, and Nuño and Thomas, 2017, for a more general presentation of the advantages of continuous-time methods).

We also illustrate how the model can be structurally estimated with a likelihood function with aggregate and micro observations, using tools from inference with diffusions. Such likelihood is computationally straightforward once we have solved the model with the approach outlined above: it just amounts to transposing a matrix. Then, we take our model to the data by matching some features of the U.S. economy, such as the leverage of the corporate sector, and estimating the volatility of the aggregate shocks using maximum likelihood.

Our work relates to several important threads in macroeconomics. First, we follow the recent macro-finance literature pioneered by Basak and Cuoco (1998), Adrian and Boyarchenko (2012), He and Krishnamurthy (2012, 2013), and Brunnermeier and Sannikov (2014), among others. As we mentioned before, most of these papers only consider *between-agents* heterogeneity, but no *within-agents* heterogeneity. Instead, we deal with a more complex –and empirically relevant– form of heterogeneity, namely a non-trivial heterogeneous wealth distribution across households.

Our paper also contributes to the literature on global solution methods for heterogeneous agent models with aggregate shocks such as Den Haan (1996), Den Haan (1997), Algan et al. (2008), Reiter (2009, 2010), Den Haan and Rendahl (2010), Maliar et al. (2010), Sager (2014), and Pröhl (2015) (a recent survey of the field can be found in Algan et al. 2014). To the best of our knowledge, ours is the first paper to generalize the celebrated algorithm of Krusell and Smith (1998) to accommodate a universal nonlinear law of motion in the endogenous state variables.[2]

Finally, our paper builds on the nascent literature on the application of machine learning techniques to compute dynamic general equilibrium models. The proposed methods have so far been concerned with the solution of high dimensional dynamic programming problems. Scheidegger and Bilionis (2017) combine Gaussian process regression with an active subspace method to solve discrete-time stochastic growth models of up to 500 dimensions. Duarte (2018) instead employs a reinforcement learning algorithm together with a neural network to solve a two-sector model with 11 state variables. In contrast, our machine learning algorithm is used to provide a nonlinear forecast of aggregate variables within the model itself. In this respect, our paper is loosely related to an early literature using neural networks to model bounded rationality and learning, such as Barucci and Landi (1995), Cho (1995), and Cho and Sargent (1996).

Our methodology may also be useful to analyze other heterogeneous agents models with aggregate shocks. An obvious candidate is the heterogeneous agent New Keynesian (HANK) model with a zero lower bound (ZLB) on the nominal interest rates, such as Auclert (2016), Gornemann et al. (2012), Kaplan et al. (2018), Luetticke (2015), and McKay et al. (2016). The ZLB introduces a nonlinearity in the state space of aggregate variables that cannot be addressed either with local methods or with global methods based on linear laws of motion. Other potential candidates include any model in which the "quasi-aggregation" result breaks down and, thus, we require higher-order moments to provide an accurate characterization of the equilibrium dynamics of the agents' distribution.

[2] Ahn et al. (2017) introduce a related method to compute the solution to heterogeneous agent models with aggregate shocks in continuous time. However, theirs is a local solution, based on first-order perturbation around the DSS and, thus, unable to analyze the class of nonlinearly-related questions posed in our paper.

2 Model

We postulate a continuous-time, infinite-horizon model. Three types of agents populate our economy: a representative firm, a representative expert, and a continuum of households. There are two assets: a risky asset, which we identify as capital, and a risk-free one, which we call bonds. Only the expert can hold the risky asset. In the interpretation implicit in our terminology, this is because the expert is the only agent with knowledge in accumulating capital. However, other interpretations, such as the expert standing in for banks or other financial intermediaries in the economy, are possible. In contrast, households can lend to the expert at the riskless rate, but cannot hold capital themselves, as they lack the required skill to handle it. The expert cannot issue outside equity, but she can partially finance her holdings of the risky asset by issuing bonds to households. Together with market clearing, our assumptions imply that the economy has a risky asset in positive net supply, capital, and a risk-free asset in zero net supply, bonds. As will become clear below, there is no need to separate between the representative firm and expert, and we could write the model consolidating both agents in a single type. Keeping both agents separate, though, clarifies the exposition at the cost of a little additional notation. Also, we only introduce heterogeneity on the side of the households (and not among the experts or the firms) because, in our environment, it is the type of heterogeneity that generates the most interesting aggregate consequences.

2.1 The firm

A representative firm rents aggregate capital, K_t, and aggregate labor, L_t, to produce output with a Cobb-Douglas technology:

$$Y_t = F(K_t, L_t) = K_t^{\alpha} L_t^{1-\alpha}.$$

Since input markets are competitive, wages, w_t, are equal to the marginal productivity of labor:

$$w_t = \frac{\partial F\left(K_t, L_t\right)}{\partial L_t} = (1-\alpha)\,\frac{Y_t}{L_t} \tag{1}$$

and the rental rate of capital, rc_t, is equal to the marginal productivity of capital:

$$rc_t = \frac{\partial F\left(K_t, L_t\right)}{\partial K_t} = \alpha \frac{Y_t}{K_t}. \tag{2}$$

During production, capital depreciates at a constant rate δ and receives a growth rate shock Z_t that follows a Brownian motion with volatility σ. Thus, aggregate capital evolves according

to:

$$\frac{dK_t}{K_t} = (\iota_t - \delta)\, dt + \sigma dZ_t, \tag{3}$$

where ι_t is the reinvestment rate per unit of capital that we will characterize below. The capital growth rate shock is the only aggregate shock to the economy. Following convention, the rental rate of capital rc_t is defined over the capital contracted, K_t, and not over the capital returned after depreciation and the growth rate shock. Thus, we define the instantaneous return rate on capital dr_t^k as:

$$dr_t^k = (rc_t - \delta)\, dt + \sigma dZ_t.$$

The coefficient of the time drift, $rc_t - \delta$, is the profit rate of capital, equal to the rental rate of capital less depreciation. The volatility σ is the capital gains rate.

2.2 The expert

The expert holds capital $\widehat{K}_t$ (we denote variables related to the expert with a caret). She rents this capital to the representative firm. And, to finance her holding of $\widehat{K}_t$, the expert issues risk-free debt $\widehat{B}_t$ at rate r_t to the households. The financial frictions in the model come from the fact that the expert cannot issue state-contingent claims (i.e., outside equity) against $\widehat{K}_t$. In particular, the expert must absorb all the risk from holding capital.

The net wealth (i.e., inside equity) of the expert, $\widehat{N}_t$, is the difference between her assets (capital) and her liabilities (debt):

$$\widehat{N}_t = \widehat{K}_t - \widehat{B}_t.$$

We allow $\widehat{N}_t$ to be negative, although this would not occur along the equilibrium path.

Let $\widehat{C}_t$ be the consumption of the expert. Then, the dynamics of $\widehat{N}_t$ are given by:

$$\begin{aligned} d\widehat{N}_t &= \widehat{K}_t dr_t^k - \widehat{B}_t r_t dt - \widehat{C}_t dt \\ &= \widehat{\omega}_t \widehat{N}_t dr_t^k + \left[(1 - \widehat{\omega}_t)\, \widehat{N}_t r_t - \widehat{C}_t \right] dt \\ &= \left[(r_t + \widehat{\omega}_t\, (rc_t - \delta - r_t))\, \widehat{N}_t - \widehat{C}_t \right] dt + \sigma \widehat{\omega}_t \widehat{N}_t dZ_t, \end{aligned} \tag{4}$$

where $\widehat{\omega}_t \equiv \frac{\widehat{K}_t}{\widehat{N}_t}$ is the leverage ratio of the expert. The term $r_t + \widehat{\omega}_t\, (rc_t - \delta - r_t)$ is the deterministic return on net wealth, equal to the return on bonds, r_t, plus $\widehat{\omega}_t$ times the excess return on leverage, $rc_t - \delta - r_t$. The term $\sigma \widehat{\omega}_t \widehat{N}_t$ reflects the risk of holding capital induced by the capital growth rate shock.

The previous expression allows us to derive the law of motion for $\widehat{K}_t$:

$$\begin{aligned} d\widehat{K}_t &= d\widehat{N}_t + d\widehat{B}_t \\ &= \left[\left(r_t + \widehat{\omega}_t \left(rc_t - \delta - r_t\right)\right)\widehat{N}_t - \widehat{C}_t\right] dt + \sigma\widehat{\omega}_t \widehat{N}_t dZ_t + d\widehat{B}_t. \end{aligned}$$

The expert's preferences over $\widehat{C}_t$ are representable by:

$$\widehat{U}_j = \mathbb{E}_j \left[\int_j^\infty e^{-\widehat{\rho}(t-j)} \log(\widehat{C}_t) dt\right], \tag{5}$$

where $\widehat{\rho}$ is her discount rate. Using a log utility function will make our derivations below easier, but could be easily generalized to the class of recursive preferences introduced by Duffie and Epstein (1992).

The expert decides her consumption levels and leverage ratio to solve the problem:

$$\max_{\{\widehat{C}_t, \widehat{\omega}_t\}_{t\geq 0}} \widehat{U}_0, \tag{6}$$

subject to evolution of her net wealth (4), an initial level of net wealth N_0, and the No-Ponzi-game condition:

$$\lim_{T\to\infty} e^{-\int_0^T r_\tau d\tau} B_T = 0. \tag{7}$$

2.3 Households

There is a continuum of infinitely-lived households with unit mass. Households are heterogeneous in their wealth a_m and labor supply z_m for $m \in [0, 1]$. The distribution of households at time t over these two individual states is $G_t(a, z)$. To save on notation, we will drop the subindex m when no ambiguity occurs.

Each household supplies z_t units of labor valued at wage w_t. Idiosyncratic labor productivity evolves stochastically following a two-state Markov chain: $z_t \in \{z_1, z_2\}$, with $0 < z_1 < z_2$. The process jumps from state 1 to state 2 with intensity λ_1 and vice versa with intensity λ_2. The ergodic mean of z is 1. As in Huggett (1993), we identify state 1 with unemployment (where z_1 is the value of leisure and home production) and state 2 with working. We will follow this assumption when the model faces the data, but nothing essential depends on it. Also, increasing the number of states of the chain is trivial, but notationally cumbersome.

Households can save an amount a_t in the riskless debt issued by the expert at interest rate r_t. Hence, a household's wealth follows:

$$da_t = (w_t z_t + r_t a_t - c_t) dt = s(a_t, z_t, K_t, G_t) dt, \tag{8}$$

where the short-hand notation $s(a_t, z_t, K_t, G_t)$ denotes the drift of the wealth process. The first two variables, a_t and z_t, are the household individual states, the next two, K_t and G_t, are the aggregate state variables that determine the returns on its income sources (labor and bonds). All four variables pin down the optimal choice, $c_t = c(a_t, z_t, K_t, G_t)$, of the control. The households also faces a borrowing limit that prevents them from shorting bonds:

$$a_t \geq 0. \tag{9}$$

Households have a CRRA instantaneous felicity function over consumption flows $c(\cdot)$:

$$u(c_t) = \frac{c_t^{1-\gamma} - 1}{1-\gamma}$$

discounted at rate $\rho > 0$. As before, we could substitute the CRRA felicity function with a more general class of recursive preferences.

Two points are worth discussing here. First, we have a CRRA felicity function to allow different risk aversions in the households and the expert. Second, we make the households less patient than the expert, $\rho > \widehat{\rho}$. We will show later how the risk-free rate in the DSS (recall, the deterministic steady state) is pinned down by the discount factor of the expert, i.e., $r = \widehat{\rho}$ (we drop the subindex when we denote a variable evaluated at the DSS). But, if $\rho \leq r = \widehat{\rho}$, the households would want to accumulate savings without bounds to self-insure against idiosyncratic labor risk (Aiyagari, 1994). Hence, we can only have a DSS –and an associated ergodic distribution of individual endogenous variables– if we increase the households' discount rate above the expert's.[3]

In summary, households maximize

$$\max_{\{c_t\}_{t\geq 0}} \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} \frac{c_t^{1-\gamma} - 1}{1-\gamma} dt \right], \tag{10}$$

subject to the budget constraint (8), initial wealth a_0, and the borrowing limit (9).

[3]This property of our economy stands in contrast with models à la Bernanke et al. (1999), where borrowers are more impatient than lenders to prevent the former from accumulating enough wealth as to render the financial friction inoperative. But in these models, borrowers are infinitesimal and subject to idiosyncratic risk, and the lenders' discount rate determines the DSS risk-free rate. In our model, the situation is reversed, with the lenders being infinitesimal and subject to idiosyncratic risk and the borrower's discount rate controlling the DSS risk-free rate. We have framed our discussion for the case without aggregate shocks since we want to ensure the existence of a DSS. The characterization of the admissible region for ρ in relation with $\widehat{\rho}$ when we only care about the properties of the economy with aggregate shocks is beyond the scope of our paper.

2.4 Market clearing

There are three market clearing conditions. First, the total amount of debt issued by the expert must equal the total amount of households' savings:

$$B_t \equiv \int a dG_t\left(da, dz\right) = \widehat{B}_t, \tag{11}$$

which also implies $dB_t = d\widehat{B}_t$.

Second, the total amount of labor rented by the firm is equal to labor supplied:

$$L_t = \int z dG_t.$$

Due to the assumption about the ergodic mean of z, we have that $L_t = 1$. Then, total payments to labor are given by w_t. If we define total consumption by households as

$$C_t \equiv \int c\left(a_t, z_t, K_t, G_t\right) dG_t\left(da, dz\right),$$

we get:

$$d\widehat{B}_t = dB_t = \left(w_t + r_t B_t - C_t\right) dt, \tag{12}$$

which tells us that the evolution of aggregate debt is the labor income of households (w_t) plus its debt income ($r_t B_t$) minus their aggregate consumption C_t.

Third, the total amount of capital in this economy is owned by the expert,

$$K_t = \widehat{K}_t$$

and, therefore, $dK_t = d\widehat{K}_t$ and $\widehat{\omega}_t = \frac{K_t}{N_t}$, where $N_t = \widehat{N}_t = K_t - B_t$. With these results, we derive

$$\begin{aligned} dK_t &= \left(\left(r_t + \widehat{\omega}_t\left(rc_t - \delta - r_t\right)\right)\widehat{N}_t - \widehat{C}_t\right) dt + \sigma\widehat{\omega}_t\widehat{N}_t dZ_t + d\widehat{B}_t \\ &= \left(\left(rc_t - \delta\right)K_t + w_t - C_t - \widehat{C}_t\right) dt + \sigma K_t dZ_t \\ &= \left(Y_t - \delta K_t - C_t - \widehat{C}_t\right) dt + \sigma K_t dZ_t, \end{aligned} \tag{13}$$

where the last line uses the fact that, from competitive input markets and constant-returns-to-scale, $Y_t = rc_t K_t + w_t$. Recall, from equation (3), that

$$dK_t = \left(\iota_t - \delta\right) K_t dt + \sigma K_t dZ_t.$$

Then, equating (13) and (3) and cancelling terms, we get

$$\iota_t = \frac{Y_t - C_t - \widehat{C}_t}{K_t},$$

i.e., the reinvestment rate is output less aggregate consumption divided by aggregate capital.

2.5 Density

The households distribution $G_t(a, z)$ has a density on assets a, $g_{it}(a)$, conditional on the labor productivity state $i \in \{1, 2\}$. The density satisfies the normalization

$$\sum_{i=1}^{2} \int_0^\infty g_{it}(a) da = 1.$$

The dynamics of this density conditional on the realization of aggregate variables are given by the Kolmogorov forward (KF) equation:

$$\frac{\partial g_{it}}{\partial t} = -\frac{\partial}{\partial a}\left(s\left(a_t, z_t, K_t, G_t\right) g_{it}(a)\right) - \lambda_i g_{it}(a) + \lambda_j g_{jt}(a), \ i \neq j = 1, 2. \tag{14}$$

Reading equation (14) is straightforward: the density evolves according to the optimal consumption-saving choices of each household plus two jumps corresponding to households that circulate out of the labor state i ($\lambda_i g_{it}(a)$) and the households that move into state j ($\lambda_j g_{jt}(a)$).

3 Equilibrium

An equilibrium in this economy is composed by a set of prices $\left\{w_t, rc_t, r_t, r_t^k\right\}_{t\geq 0}$, quantities $\left\{K_t, N_t, B_t, \widehat{C}_t, c_{mt}\right\}_{t\geq 0}$ and a density $\left\{g_{it}(\cdot)\right\}_{t\geq 0}$ for $i \in \{1, 2\}$ such that:

1. Given w_t, r_t, and g_t, the solution of household m's problem (31) is $c_{mt} = c\left(a_t, z_t, K_t, G_t\right)$.
2. Given r_t^k, r_t, and N_t, the solution of the expert's problem (6) is $\widehat{C}_t$, K_t, and B_t.
3. Given K_t, the firm maximizes their profits and input prices are given by w_t and rc_t and the rate of return on capital by r_t^k.
4. Given w_t, r_t, and c_t, g_{it} is the solution of the KF equation (14).
5. Given r_t, g_{it}, and B_t, the debt market (11) clears and $N_t = K_t - B_t$.

3.1 Equilibrium characterization

Several properties of the equilibrium are characterized with ease. We proceed first with the expert's problem. The use of log-utility implies that the expert consumes a constant share $\widehat{\rho}$ of her net wealth and chooses a leverage ratio proportional to the difference between the expected return on capital and the risk-free rate:

$$\widehat{C}_t = \widehat{\rho} N_t \tag{15}$$

$$\omega_t = \widehat{\omega}_t = \frac{1}{\sigma^2}\left(rc_t - \delta - r_t\right). \tag{16}$$

Second, rewriting the latter result, we get that the excess return on leverage,

$$rc_t - \delta - r_t = \sigma^2 \frac{K_t}{N_t},$$

depends positively on the variance of the aggregate shock, σ^2, and the leverage of the economy $\frac{K_t}{N_t}$. The higher the volatility or the leverage ratio in the economy, the higher the excess return that the expert requires to isolate households from dZ_t. A positive capital growth rate shock, by increasing N_t relatively to K_t, lowers the excess return. Analogously, a higher volatility of the aggregate shock increases the excess return.

Third, we can use the values of rc_t, L_t, and ω_t in equilibrium to get the wage $w_t = (1-\alpha) K_t^{\alpha}$, the rental rate of capital $rc_t = \alpha K_t^{\alpha-1}$, and the risk-free interest rate:

$$r_t = \alpha K_t^{\alpha-1} - \delta - \sigma^2 \frac{K_t}{N_t}. \tag{17}$$

Since $K_t = N_t + B_t$, these three equations depends only on the expert's net wealth N_t and debt B_t.

Fourth, we can describe the evolution of N_t:

$$\begin{aligned} dN_t &= \left[\left(r_t + \omega_t\left(rc_t - \delta - r_t\right)\right) N_t - \widehat{C}_t\right] dt + \sigma \omega_t N_t dZ_t \\ &= \left(\alpha K_t^{\alpha-1} - \delta - \widehat{\rho} - \sigma^2\left(1 - \frac{K_t}{N_t}\right)\frac{K_t}{N_t}\right) N_t dt + \sigma K_t dZ_t \end{aligned} \tag{18}$$

as a function only of N_t, B_t, and dZ_t. Equation (18) shows the nonlinear dependence of dN_t on the leverage level $\frac{K_t}{N_t}$. We will stress this point in the next pages repeatedly. For convenience, some times we will write

$$dN_t = \mu^N(B_t, N_t)dt + \sigma^N(B_t, N_t)dZ_t,$$

where $\mu^N(B_t, N_t) = \left(\alpha K_t^{\alpha-1} - \delta - \widehat{\rho} - \sigma^2\left(1 - \frac{K_t}{N_t}\right)\frac{K_t}{N_t}\right) N_t$ is the drift of N_t and $\sigma^N(B_t, N_t) = \sigma K_t$ its volatility.

Fifth, we have from equation (12):

$$\begin{aligned} dB_t &= (w_t + r_t B_t - C_t)\, dt \\ &= \left((1-\alpha)\, K_t^{\alpha} + \left(\alpha K_t^{\alpha-1} - \delta - \sigma^2 \frac{K_t}{N_t}\right) B_t - C_t\right) dt, \end{aligned} \tag{19}$$

which depends nonlinearly on N_t and B_t and linearly on aggregate household consumption, C_t, an endogenous variable we must determine.

We can stack all the equilibrium conditions (except the optimality condition for households, which we will discuss in the next section) in two blocks. The first block includes all the variables that depend directly on N_t, B_t, and dZ_t:

$$w_t = (1-\alpha)\, K_t^{\alpha} \tag{20}$$

$$rc_t = \alpha K_t^{\alpha-1} \tag{21}$$

$$r_t = \alpha K_t^{\alpha-1} - \delta - \sigma^2 \frac{K_t}{N_t} \tag{22}$$

$$dr_t^k = (rc_t - \delta)\, dt + \sigma dZ_t \tag{23}$$

$$dN_t = \left(\alpha K_t^{\alpha-1} - \delta - \widehat{\rho} - \sigma^2\left(1 - \frac{K_t}{N_t}\right)\frac{K_t}{N_t}\right) N_t dt + \sigma K_t dZ_t. \tag{24}$$

The second block includes the equations determining the aggregate consumption of the households, dB_t, dK_t, and $\frac{\partial g_{it}}{\partial t}$:

$$C_t \equiv \sum_{i=1}^{2} \int c\,(a_t, z_t, K_t, G_t)\, g_{it}\,(a)\, da \tag{25}$$

$$dB_t = \left((1-\alpha)\, K_t^{\alpha} + \left(\alpha K_t^{\alpha-1} - \delta - \sigma^2 \frac{K_t}{N_t}\right) B_t - C_t\right) dt \tag{26}$$

$$dK_t = dN_t + dB_t \tag{27}$$

$$\frac{\partial g_{it}}{\partial t} = -\frac{\partial}{\partial a}\left(s\,(a_t, z_t, K_t, G_t)\, g_{it}(a)\right) - \lambda_i g_{it}(a) + \lambda_j g_{jt}(a),\ i \neq j = 1, 2. \tag{28}$$

The second block shows i) how the density $\{g_{it}\,(\cdot)\}_{t\geq 0}$ for $i \in \{1, 2\}$ matters to determine C_t, ii) that C_t pins down dB_t, and iii) that once we have dB_t, we can calculate dK_t. Therefore, in practice, computing the equilibrium of this economy is equivalent to finding C_t. Once C_t is known, all other aggregate variables follow directly.

3.2 The DSS of the model

Before we explain our general solution in the next section, we use the previous equations to analyze three interesting cases. In this subsection, we describe the DSS of the model where there are no capital growth rate shocks, but we still have idiosyncratic household shocks. In the next two subsections, we will look at the SSS (recall, a stochastic steady state, also known as the risky steady state) of our economy and the representative household version of the model.

To study the DSS, we can go to the law of motion for the expert net wealth (18), set $\sigma = 0$, and get:

$$dN_t = \left(\alpha K_t^{\alpha-1} - \delta - \widehat{\rho}\right) N_t dt. \tag{29}$$

Since the drift of N_t, $\mu^N(B, N) = (\alpha K^{\alpha-1} - \delta - \widehat{\rho})\, N$, must be zero in a DSS (remember that we drop the t subindex to denote the DSS value of a variable), we get

$$K = \left(\frac{\widehat{\rho} + \delta}{\alpha}\right)^{\frac{1}{\alpha-1}}.$$

With this result, the DSS risk-free interest rate (17) equals the return on capital and the rental rate of capital less depreciation:

$$r = r_t^k = rc_t - \delta = \alpha K_t^{\alpha-1} - \delta = \widehat{\rho}. \tag{30}$$

As mentioned above, this condition forces us to have $\widehat{\rho} < \rho$. Otherwise, the households would accumulate too many bonds and the DSS would not be well-defined.

Finally, the dispersion of the idiosyncratic shocks determines the DSS expert's net wealth:

$$N = K - B = K - \int a dG\left(da, dz\right),$$

a quantity that, unfortunately, we cannot compute analytically.

3.3 The SSS of the model

A SSS in our model is formally defined as a density $g^{SSS}(\cdot)$ and equity N^{SSS} that remain invariant in the absence of aggregate shocks. Let $\Gamma_\sigma(g(\cdot), N, W)$ be the law of motion of the economy given an aggregate capital volatility σ and a realization of the Brownian motion W. More precisely, $\Gamma_\sigma(\cdot,\cdot,\cdot)$ is an operator that maps income-wealth densities $g(\cdot)$ and equity levels N into changes in these variables:

$$\lim_{\Delta t \to 0} \frac{1}{\Delta t} \begin{bmatrix} g_{t+\Delta t}(\cdot) - g_t(\cdot) \\ N_{t+\Delta t} - N_t \end{bmatrix} = \Gamma_\sigma(g_t(\cdot), N_t, W_t).$$

The SSS, therefore, solves:

$$\Gamma_\sigma(g^{SSS}(\cdot), N^{SSS}, 0) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

In general, we will have multiple SSSs that solve the previous functional equation. We will document below how several of them appear in our quantitative exercise.

The difference between the SSS and the DSS is that the former is the steady state of an economy where individual agents make their decisions taking into account aggregate risks ($\sigma > 0$) −using equation (33)− but no shock arrives along the equilibrium path, whereas in the latter agents live in an economy without aggregate risks ($\sigma = 0$) and arrange their consumption paths accordingly. The DSS is then formally defined, using our notation above, as

$$\Gamma_0(g^{DSS}(\cdot), N^{DSS}) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

3.4 The representative household version of the model

In this subsection, we eliminate the idiosyncratic labor risk of households. In that way, we come up with a representative household version of the model where we keep the *between-agents* heterogeneity (we still have an expert and a household), but we do not have any *within-agents* heterogeneity (there are one representative expert and one representative household).

Equations (20) to (24) of the equilibrium, as well as equations (26) and (27), remain unchanged. Instead of equation (25), the consumption rule of the household $C_t = c(B_t, N_t)$ depends now just on aggregate debt and capital and it solves the problem:

$$\max_{\{C_t\}_{t\geq 0}} \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} \frac{C_t^{1-\gamma} - 1}{1-\gamma} dt \right], \tag{31}$$

subject to the budget constraint (26) (where input prices have already been substituted in), initial wealth B_0, and the borrowing limit $B_t \geq 0$. Because of the absence of *within-agents* heterogeneity, there is no need for an analogous of equation (28) in this version of the model.

4 Solution

Our previous discussion highlighted the role of finding, in the general version of the model, the households aggregate consumption, C_t, to compute the equilibrium of the economy given some structural parameter values:

$$\Psi = \{\alpha, \delta, \sigma, \widehat{\rho}, \rho, \gamma, z_1, z_2, \lambda_1, \lambda_2\}.$$

To do so, we follow Krusell and Smith (1998) and assume that, when forming their expectations, households only use a finite set of n moments of the cross-sectional distribution of assets instead of the complete distribution. In contrast to Krusell and Smith (1998), in which the income-wealth distribution is the only endogenous state variable, here the expert's net wealth N_t is also a state variable. At the same time, we do not have any exogenous state variable, as $K_t = N_t + B_t$ instantaneously incorporates the capital growth rate shocks. For ease of exposition, we discuss the case with $n = 1$. All the techniques can be trivially extended to the case with $n > 1$ at the cost of heavier notation.

More concretely, households consider a *perceived law of motion* (PLM) of aggregate debt:

$$dB_t = h\left(B_t, N_t\right) dt, \tag{32}$$

where $h\left(B, N\right)$ is the conditional expectation of dB_t given available information (B_t, N_t):

$$h\left(B_t, N_t\right) = \frac{\mathbb{E}\left[dB_t | B_t, N_t\right]}{dt},$$

instead of the exact law of motion (19). We borrow the term PLM from the learning literature (Evans and Honkapohja, 2001). Our choice of words accentuates that we allow $h\left(\cdot, \cdot\right)$ be a general function, and not just a linear or polynomial function of its arguments (perhaps with state-dependent coefficients). In fact, our methodology will let the PLM approximate, arbitrarily well, equation (19).

In the original Krusell and Smith (1998) approach, the PLM is assumed to be log-linear in the endogenous state variables and nonlinear in the exogenous state by making the coefficients of the log-linear specification dependent on the realization of the aggregate shock. As explained in Subsection 4.3, we propose a more flexible methodology, in which the functional form $h\left(\cdot, \cdot\right)$ is not specified, but obtained from simulated data by employing machine learning techniques. This extra flexibility is key given the complex nonlinearities present in laws of motion of N_t, equation (18), and B_t, equation (19).

Given the PLM, the household's problem has an associated Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{aligned} \rho V_i\left(a, B, N\right) &= \max_c \frac{c^{1-\gamma} - 1}{1-\gamma} + s\frac{\partial V_i}{\partial a} + \lambda_i\left[V_j(a, B, N) - V_i(a, B, N)\right] \\ &\quad + h\left(B, N\right)\frac{\partial V_i}{\partial B} + \mu^N\left(B, N\right)\frac{\partial V_i}{\partial N} + \frac{\left[\sigma^N\left(B, N\right)\right]^2}{2}\frac{\partial^2 V_i}{\partial N^2}, \end{aligned} \tag{33}$$

$i \neq j = 1, 2$, and where we use the shorthand notation $s = s\left(a, z, N + B, G\right)$. Notice how the HJB incorporates $h\left(B, N\right)$. Equation (33) complements the equilibrium conditions (20)-(28) by making the problem of the household explicit.

4.1 An overview of the algorithm

The algorithm to find $h(B, N)$ in (32) proceeds according to the following iteration:

1) Start with h_0, an initial guess for h.

2) Using current guess for h, solve for the household consumption, c_m, in the HJB equation (33). This solution can be obtained by using an upwind finite differences scheme described in Appendix A (although other numerical algorithms such as complex differentiation can be applied).

3) Construct a time series for B_t by simulating the cross-sectional distribution over time. Given B_t, we can find N_t and K_t using equations (18) and (27).

4) Use a universal nonlinear approximator to obtain h_1, a new guess for h.

5) Iterate steps 2)-4) until h_n is sufficiently close to h_{n-1} given some pre-specified norm and tolerance level.

Steps 1)-5) show that our solution has two main differences with respect to the original Krusell-Smith algorithm: the use of continuous time and our employment of a universal nonlinear approximator to update the guess of the PLM. Both differences deserve some explanation.

4.2 Continuous time

Krusell and Smith (1998) wrote their model in discrete time. Our continuous-time formulation, while not changing any fundamental feature of the model, enjoys several numerical advantages with respect to discrete time (Achdou et al., 2017). First, continuous time naturally generates sparsity in the matrices characterizing the transition probabilities of the discretized stochastic processes. Intuitively, continuously moving state variables such as wealth only drift an infinitesimal amount in an infinitesimal unit of time. Therefore, in an approximation that discretizes the state space, households reach only states that directly neighbor the current state. Second, the optimality characterizing consumption has a simpler structure than in discrete time:

$$c_i^{-\gamma} = \frac{\partial V_i}{\partial a}. \tag{34}$$

Third, it is easier to capture occasionally binding constraints such as equation (9) in continuous time than in discrete time as the optimality condition (34) for consumption holds with equality everywhere in the interior of the state space. Fourth, the dynamics of the cross-sectional wealth distribution are characterized by the KF equation (14). The discretization of this equation yields an efficient way to simulate a time series of the cross-sectional distribution (although this

can also be performed in discrete time, as in Ríos-Rull 1997, Reiter 2009, and Young 2010, at some additional cost). Appendix A provides further details.

Regarding the generation of data, we simulate T periods of the economy with a constant time step Δt. We start from the initial income-wealth distribution at the DSS (although we could pick other values). A number of initial samples is discarded as a burn-in. If the time step is small enough, we have

$$B_{t_j+\Delta t} = B_{t_j} + \int_{t_j}^{t_j+\Delta_t} dB_s = B_{t_j} + \int_{t_j}^{t_j+\Delta t} h\left(B_s, N_s\right) ds \approx B_{t_j} + h\left(B_{t_j}, N_{t_j}\right) \Delta t.$$

Our simulation $\left(\mathbf{S}, \widehat{\mathbf{h}}\right)$ is composed by a vector of inputs $\mathbf{S} = \{\mathbf{s}_1, \mathbf{s}_2, ..., \mathbf{s}_J\}$, where $\mathbf{s}_j = \left\{s_j^1, s_j^2\right\} = \left\{B_{t_j}, N_{t_j}\right\}$ are samples of aggregate debt and expert's net wealth at J random times $t_j \in [0, T]$, and a vector of outputs $\widehat{\mathbf{h}} = \left\{\widehat{h}_1, \widehat{h}_2 ..., \widehat{h}_J\right\}$, where

$$\widehat{h}_j \equiv \frac{B_{t_j+\Delta t} - B_{t_j}}{\Delta t}$$

are samples of the growth rate of B_t. The evaluation times t_j should be random and uniformly distributed over $[0, T]$ as, ideally, samples should be independent.

4.3 A universal nonlinear approximator

In the original Krusell-Smith algorithm, the law of motion linking the mean of capital tomorrow and the mean of capital today is log-linear, with the coefficients in that function depending on the aggregate shock. This approximation is highly accurate due to the near log-linearity of their models in the vicinity of the DSS. Indeed, in such a model, the DSS and SSS almost coincide. But, as shown in equations (18) and (19), this linearity of the law of motion of the endogenous variables with respect to other endogenous variables does not carry out to our model.

This nonlinear structure causes two problems. First, we face the *approximation* problem: we need an algorithm that searches for an unknown nonlinear functional instead of a simple linear regression with aggregate-state-dependent coefficients. Second, we need to tackle the *extrapolation* problem. While the theoretical domain of B_t and N_t is unbounded, practical computation requires to limit it to a compact subset of $\mathbb{R}^2$ large enough as to prevent boundary conditions from altering the solution in the subregion where most of the ergodic distribution accumulates. But precisely because we deal with such a large area, the simulation in step 3) of the algorithm in Subsection 4.1 never visits an ample region of the state space. Thus, the approximation algorithm should not only provide an accurate nonlinear approximation in the visited region, but also a "reasonable" extrapolation to the rest of the state space. We will

return to what "reasonable" means in this context momentarily.

To address these two problems, we propose to employ a nonlinear approximation technique based on neural networks. Our approach has four crucial advantages. First, the *universal approximation theorem* (Hornik et al. 1989; Cybenko 1989) states that a neural network with at least one hidden layer can approximate any Borel measurable function mapping finite-dimensional spaces arbitrarily well. In particular, the theorem does not require that the approximated function be differentiable and can handle cases with kinks and occasionally binding constraints.[4]

Second, the neural network coefficients can be efficiently estimated using gradient descent methods and back-propagation. This allows for an easier coding and shorter implementation time than other approaches.

Third, neural networks are more economical, for middle and high dimensions, than other approximators. More concretely, Barron (1993) shows that a one-layer neural network achieves, for functions on the first moment of the magnitude distribution of the Fourier transform, integrated square errors of order $O(1/n)$, where n is the number of nodes. In comparison, for series approximations (polynomials, spline, and trigonometric expansions), the integrated square error is of order $O(1/(n^{2/d})$ where d is the dimensions of the function to be approximated. In other words: the "curse of dimensionality" does not apply to neural networks that approximate functions of the very wide class considered by Barron (1993).[5] This advantage is not present in our baseline model, with $d = 2$, but will appear in any extension with additional aggregate state variables. Even going to $d = 3$ or $d = 4$ saturates alternatives such as Chebyshev polynomials.[6]

Fourth, neural networks extrapolate outstandingly. This is, in practice, key. Neural networks have well-behaved shapes outside their training areas. In contrast, Chebyshev polynomials (or other series) more often than not display explosive behaviors outside the fitted area that prevent the algorithm from converging. Figures 21 and 22 in Appendix B show this disappointing behavior of an approximation to the PLM in our model with Chebyshev polynomials. The two figures document how, within the area of high density of the ergodic distribution, Chebyshev polynomials approximate the law of motion for aggregate debt fairly (compare them with panel c) in Figure 2, obtained with our neural network). But Chebyshev polynomials start oscillating as soon as we abandoned the well-traveled area of the simulation.

[4]Recall that Lusin's theorem tells us that every measurable function is a continuous function almost everywhere. Thus, we can approximate jumps in a finite number of points, but not functions with extremely intricate shapes. Those complex functions, however, are unlikely to be of much relevance in solving standard dynamic equilibrium models.

[5]In fact, we can rely on the more general theorems shown by Bach (2017) that cover non-decreasing positively homogeneous activation functions like the rectified linear unit and that show, beyond the break of the curse dimensionality, approximation and the estimation errors.

[6]Similarly, approaches, such as Smolyak interpolation, that alleviate the "curse of dimensionality" in standard problems are harder to apply here because we deal with shapes of the ergodic distribution are hard to characterize ex-ante. Neural networks are more resilient to sparse initial information regarding such shapes.

4.4 More on the neural network approximator

Here we briefly describe our neural network approximator of the PLM in more detail. For excellent introductory treatments of this material, see Bishop (2006) and Goodfellow et al. (2016).

A single hidden layer neural network $h(\mathbf{s};\theta)$ is a linear combination of Q fixed nonlinear basis (i.e., activation) functions $\phi(\cdot)$:

$$h(\mathbf{s};\theta) = \theta_0^2 + \sum_{q=1}^{Q} \theta_q^2 \phi\left(\theta_{0,q}^1 + \sum_{i=1}^{2} \theta_{i,q}^1 s^i\right), \tag{35}$$

where $\mathbf{s}$ is an two-dimensional input and θ a vector of coefficients (i.e., weights):

$$\theta = \left(\theta_0^2, \theta_1^2, ..., \theta_Q^2, \theta_{0,1}^1, \theta_{1,1}^1, \theta_{2,1}^1, ..., \theta_{0,Q}^1, \theta_{1,Q}^1, \theta_{2,Q}^1\right).$$

Note how we call θ "coefficients," as they represent a numerical entity, in comparison with the structural parameters, Ψ, that have a sharp economic interpretation.

The neural network provides a flexible parametric function h that determines the growth rate of aggregate debt:

$$\widehat{h}_j = h(\mathbf{s}_j;\theta), \ j = 1, .., J,$$

and that satisfies the properties of universal approximation, breaking of the curse of dimensionality, good extrapolative behavior, and easy implementation we discussed above.

Different alternatives are available for the activation function. For our model, we choose a *softplus* function, $\phi(x) = \log(1 + e^x)$ for a given input x. The softplus function has a simple sigmoid derivative, which avoids some of the problems caused by the presence of a kink in rectified linear units, while keeping an excellent efficient computation and gradient propagation.

The neural network (35) can be generalized to include many hidden layers, stacked one after the other. In that case, the network is called a *deep* neural network. However, for the particular problem of approximation a two-dimensional function a single layer is enough. The size of the hidden layer is determined by Q. This hypercoefficient can be set by regularization or simply by trial-and-error in relatively simple problems, such as the one presented here. In our case, we set $Q = 16$ because the cost of a larger hidden layer is small.

The vector of coefficients θ is selected to minimize the quadratic error function $\mathcal{E}\left(\theta;\mathbf{S},\widehat{\mathbf{h}}\right)$

given a simulation $\left(\mathbf{S}, \widehat{\mathbf{h}}\right)$:

$$\begin{aligned}\theta^* &= \arg\max_{\theta} \mathcal{E}\left(\theta; \mathbf{S}, \widehat{\mathbf{h}}\right) \\ &= \arg\max_{\theta} \sum_{j=1}^{J} \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right) \\ &= \arg\max_{\theta} \frac{1}{2} \sum_{j=1}^{J} \left\| h\left(\mathbf{s}_j; \theta\right) - \widehat{h}_j \right\|^2.\end{aligned}$$

A standard approach to perform this minimization in neural networks is the *stochastic gradient descent* algorithm. The stochastic gradient descent begins by drawing a random initialization of θ_0 from a known distribution Θ, typically a Gaussian or uniform:

$$\theta_0 \sim \Theta. \tag{36}$$

Then, θ is recursively updated according to $\theta_{m+1} = \theta_m - \epsilon_m \nabla \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)$, where:

$$\nabla \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right) \equiv \left[\frac{\partial \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)}{\partial \theta_0^2}, \frac{\partial \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)}{\partial \theta_1^2}, ..., \frac{\partial \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)}{\partial \theta_{2,Q}^1}\right]^\top$$

is the gradient of the error function with respect to θ evaluated at $\left(\mathbf{s_j}, \widehat{h}_j\right)$. To improve performance, it is typical to group several points of the simulation in each evaluation step –instead of just one point– in what is known as a *minibatch* algorithm. This takes advantage of the fast convergence of the gradient (think about it as a score of the error function) towards its value with an infinite sample, but at a much lower computational cost. The step size $\epsilon_m > 0$ is selected in each iteration according to a *line-search* algorithm in order to minimize the error function in the direction of the gradient. The algorithm is run until the distance between θ_{m+1} and θ_m is below a threshold ε:

$$\|\theta_{m+1} - \theta_m\| < \varepsilon.$$

An advantage of neural networks is that the error gradient can be efficiently evaluated using a *back-propagation* algorithm, originally developed by Rumelhart et al. (1986), which builds on the chain rule of differential calculus. See 8 for details.

One concern with neural networks is that the algorithm might converge to a local minimum. A way of coping with it is to implement a Monte Carlo multi-start. We select P initial vectors θ_0^p, with $p = 1, ..., P$ from (36). For each of these vectors, we run the stochastic gradient descent until convergence. Once we achieve convergence, we select the θ_m^p that yields the minimum

error across all the trials. Furthermore, since we are interested in approximating an unknown function, not clearing a market or satisfying an optimality condition, local minima that are close to a global minimum are acceptable solutions to the approximation problem.

Finally, notice that the algorithm is massively parallel, either in CPUs or GPUs (and, in the middle-run, in the new generation of Tensor Processing Units or TPUs), a most convenient feature for scaling and estimation. See Appendix xxx for details.

5 Estimation

Once we have solved the model given some structural parameter values Ψ, the next step is to take the model to the data, to let observations determine the values of Ψ and perform inference on them. We will proceed in two stages. First, we will discuss the simple case where the econometrician has access to output data and wants to build the likelihood associated with it. After we understand how to construct the likelihood function in this situation, we will move to describe how to add microeconomic observations from the cross-sectional distribution of assets and evaluate a more sophisticated likelihood. We will close this section by showing the results of our estimation with real data.

5.1 Building the likelihood function

Let us assume that the econometrician has access to $D+1$ observations of output, Y_t, at fixed time intervals $[0, \Delta, 2\Delta, .., D\Delta]$:

$$Y_0^D = \{Y_0, Y_\Delta, Y_{2\Delta}, ..., Y_D\}.$$

The derivations below would be very similar for observables other than output. More important, though, is that since we have one aggregate shock in the model (to capital), we can only use one observable in our likelihood. Otherwise, we would suffer from stochastic singularity. If we wanted to have more observables, such other aggregate variables or the states of model, we would need to either enrich the model with more shocks or to introduce measurement shocks in the observables. In those more complex situations, we might need to resort to a sequential Monte Carlo approximation to the filtering problem described by the associated Kushner-Stratonovich equation of our dynamic system (see, for a related approach in discrete time, Fernández-Villaverde and Rubio-Ramírez, 2007).

The likelihood function $\mathcal{L}_D\left(Y_0^D|\Psi\right)$ for our observations of output has the form:

$$\mathcal{L}_D\left(Y_0^D|\Psi\right) = \prod_{d=1}^{D} p_Y\left(Y_{d\Delta}|Y_{(d-1)\Delta}; \Psi\right),$$

where $p_Y\left(Y_{d\Delta}|Y_{(d-1)\Delta};\Psi\right)$, the conditional density function of $Y_{d\Delta}$ given $Y_{(d-1)\Delta}$, is equal to:

$$p_Y\left(Y_{d\Delta}|Y_{(d-1)\Delta};\Psi\right)=\int f_{d\Delta}(Y_{d\Delta},B)dB$$

given a density function for output and debt, $f_{d\Delta}(Y_{d\Delta},B)$, implied by the solution of the model.

Our task is, therefore, to compute the sequences of conditional densities $p_Y\left(Y_{d\Delta}|Y_{(d-1)\Delta};\Psi\right)$ at the fixed time intervals $[0,\Delta,2\Delta,..,D\Delta,]$. To do so, first, we obtain the diffusion of $Y_t=(B_t+N_t)^{\alpha}$.

Applying Itô's lemma, we get:

$$\begin{aligned}
dY_t &= \frac{\partial\left(B+N\right)^{\alpha}}{\partial B}dB_t+\frac{\partial\left(B+N\right)^{\alpha}}{\partial N}dN_t+\frac{1}{2}\frac{\partial^2\left(B+N\right)^{\alpha}}{\partial N^2}\sigma^2\left(B+N\right)^2dt\\
&= \alpha K_t^{\alpha-1}\left[\left(h(B_t,N_t)+\mu_t^N\right)dt+\sigma K_t dZ_t\right]+\frac{1}{2}\alpha\left(\alpha-1\right)K_t^{\alpha-2}\sigma^2K_t^2dt\\
&= \alpha Y_t^{\frac{\alpha-1}{\alpha}}\left[h(B_t,Y_t^{\frac{1}{\alpha}}-B_t)+\alpha Y_t-\delta Y_t^{\frac{1}{\alpha}}-\left(\alpha Y_t^{\frac{\alpha-1}{\alpha}}-\delta-\sigma^2\frac{Y_t^{\frac{1}{\alpha}}}{Y_t^{\frac{1}{\alpha}}-B_t}\right)B_t-\hat{\rho}\left(Y_t^{\frac{1}{\alpha}}-B_t\right)\right]dt\\
&\quad+\frac{\sigma^2}{2}\alpha\left(\alpha-1\right)Y_tdt+\sigma\alpha Y_tdZ_t\\
&= \mu^Y\left(B_t,Y_t\right)dt+\sigma_t^Y\left(Y_t\right)dZ_t, \qquad (37)
\end{aligned}$$

where

$$\begin{aligned}
\mu^Y\left(B_t,Y_t\right) &= \alpha Y_t^{\frac{\alpha-1}{\alpha}}*\left[h(B_t,Y_t^{\frac{1}{\alpha}}-B_t)+\alpha Y_t+\left[\frac{\left(\alpha-1\right)\sigma^2}{2}-\delta\right]Y_t^{\frac{1}{\alpha}}\right.\\
&\quad\left.-\left(\alpha Y_t^{\frac{\alpha-1}{\alpha}}-\delta-\sigma^2\frac{Y_t^{\frac{1}{\alpha}}}{Y_t^{\frac{1}{\alpha}}-B_t}\right)B_t-\hat{\rho}\left(Y_t^{\frac{1}{\alpha}}-B_t\right)\right],
\end{aligned}$$

and

$$\sigma^Y\left(Y_t\right)=\alpha\sigma Y_t.$$

With the diffusion for Y_t, (37), the density $f_t^d(Y,B)$ follows the Kolmogorov forward (KF) equation in the interval $[(d-1)\Delta,d\Delta]$:

$$\begin{aligned}
\frac{\partial f_t}{\partial t} &= -\frac{\partial}{\partial Y}\left[\mu^Y(Y,B)f_t(Y,B)\right]-\frac{\partial}{\partial B}\left[h(B,Y^{\frac{1}{\alpha}}-B)f_t^d(Y,B)\right]\\
&\quad+\frac{1}{2}\frac{\partial^2}{\partial Y^2}\left[\left(\sigma^Y(Y)\right)^2f_t(Y,B)\right]. \qquad (38)
\end{aligned}$$

The initial condition at the beginning of the interval is

$$f_{(d-1)\Delta}(Y,B)=\delta\left(Y-Y_{(d-1)\Delta}\right)f_{(d-2)\Delta}(B|Y_{(d-1)\Delta}),$$

where $f_{(d-1)\Delta}(B|Y_{(d-1)\Delta})$ is the probability of B conditional on $Y = Y_{(d-1)\Delta}$:

$$f_{(d-2)\Delta}(B|Y_{(d-1)\Delta}) = \frac{f_{(d-2)\Delta}(Y_{(d-1)\Delta}, B)}{f_{(d-2)\Delta}(Y_{(d-1)\Delta})} = \frac{f_{(d-2)\Delta}(Y_{(d-1)\Delta}, B)}{\int f_{(d-2)\Delta}(Y_{(d-1)\Delta}, B)dB)},$$

if $d \geq 2$, $f_{-1}(B) = f(B)$ is the ergodic distribution of B, and $\delta(\cdot)$ is the Dirac delta function.

Lo (1988) pioneered the estimation of the likelihood of a continuous-time stochastic process on discrete-time samples using the Kolmogorov forward equation to characterize the transition density functions. That paper provides a number of technical assumptions that has to be satisfied for the estimation. In our model they are satisfied provided that $h(B, N)$ is twice continuously differentiable in B and N and three times continuously differentiable in Ψ, which is guaranteed if it is computed using a neural network with activation functions that meet this requirements –as it happens in our case– and that Ψ lies in the interior of a finite- dimensional closed and compact parameter space.

A fundamental property of the operator in the KF equation (38) is that it is the adjoint of the infinitesimal generator employed in the HJB. The intuition is that one can think about the dynamic choices of the agents implied by the HJB as a probability distribution of their future choices, and, hence, a distribution on observables, such as output, induced by the stochastic shocks of the model. There is, in other words, an intimate link between optimal choices and likelihood functions.

This result is remarkable since it means that the solution of the KF equation amounts to transposing an inverting a sparse matrix that has already been computed when we solved the HJB. This provides a highly efficient way of evaluating the likelihood after the model is solved.[7]

5.2 Adding micro observations

A promising avenue to improve the estimation is to add micro observations, which bring much additional information and help integrating different levels of aggregation to assess the empirical validity of the model. More concretely, let $X_t \equiv [g_t(a, z); N_t]'$, be a vector of observations on the asset holdings of agents in this economy (households, $g_t(a, z)$, and the expert, N_t). Imagine, as before, that we have $D + 1$ observations of X_t at fixed time intervals $[0, \Delta, 2\Delta, .., D\Delta,]$:

$$X_0^D = \{X_0, X_\Delta, X_{2\Delta}, ..., X_D\}.$$

At this moment, we need to assume –as it is typically done in models with heterogeneous agents and aggregates shocks– that the *conditional no aggregate uncertainty* (CNAU) condition

[7] If the KF would become numerically cumbersome in more general models, we could construct Hermite polynomials expansions of the (exact but unknown) likelihood as in Aït-Sahalia (2002). We could also consider methods of moments in continuous time such as those pioneered by Andersen and Lund (1997) and Chacko and Viceira (2003).

holds. See, for instance, Miao (2006), following Bergin and Bernhardt (1992). This condition implies that if households are distributed on the interval $I = [0, 1]$ according to the Lebesgue measure Φ, then

$$G_t(A \times Z) = \Phi\left(i \in I : \left(a_t^i, z_t^i\right) \in A \times Z\right),$$

for any subsets $A \subset [0, \infty)$, $Z \subset \{z_1, z_2\}$. That is, the probability under the conditional distribution is the same as the probability according to the Lebesgue measure across the interval I.

The likekihood that an individual agent $i \in I$ at time $t = d\Delta$ is at state $(a_{d\Delta}^i, z_{d\Delta}^i, B_{d\Delta}, N_{d\Delta})$ is $f_{d\Delta}^d(a_{d\Delta}^i, z_{d\Delta}^i, B_{d\Delta}, N_{d\Delta})$. The log-likelihood is then $\log\left[f_{d\Delta}^d(a_{d\Delta}^i, z_{d\Delta}^i, B_{d\Delta}, N_{d\Delta})\right]$. Notice that this log-likelihood is a function of i.

The conditional aggregate log-likelihood across all agents is

$$\log p_X\left(X_{d\Delta}|X_{(d-1)\Delta}; \Psi\right) = \int \log\left[f_{d\Delta}^d(a_{d\Delta}^i, z_{d\Delta}^i, B_{d\Delta}, N_{d\Delta})\right] \Phi(di),$$

and, taking into account the CNAU condition, we get:

$$\begin{aligned}\int \log\left[f_{d\Delta}^d(a_{d\Delta}^i, z_{d\Delta}^i, B_{d\Delta}, N_{d\Delta})\right] \Phi(di) &= \int \log\left[f_{d\Delta}^d(a, z, B_{d\Delta}, N_{d\Delta})\right] G_{d\Delta}(da, dz) \\ &= \sum_{i=1}^{2} \int_0^\infty \log\left[f_{d\Delta}^d(a, z_i, B_{d\Delta}, N_{d\Delta})\right] g_{d\Delta}(a, z) da,\end{aligned}$$

where, in the second line, we have applied the definition of the Radon-Nikodym derivative to get the differential in a.

The density $f_t^d(a, z, B, N)$ follows the KF equation:

$$\begin{aligned}\frac{\partial f_t^d}{\partial t} = &-\frac{\partial}{\partial a}\left(s_t(a, z_i) f_t^d(a, z_i, B, N)\right) - \lambda_i f_t^d(a, z_i, B, N) + \lambda_j f_t^d(a, z_j, B, N) \\ &-\frac{\partial}{\partial B}\left[h(B, N) f_t^d(a, z_i, B, N))\right] - \frac{\partial}{\partial N}\left[\mu_t^N(B, N) f_t^d(a, z_i, B, N)\right] \\ &+\frac{1}{2}\frac{\partial^2}{\partial N^2}\left[\left(\sigma_t^N(B, N)\right)^2 f_t^d(B, N)\right], \quad (i \neq j = 1, 2)\end{aligned} \tag{39}$$

where

$$f_{(d-1)\Delta}^d = g_{(d-1)\Delta}(a, z)\delta\left(B - B_{(d-1)\Delta}\right)\delta\left(N - N_{(d-1)\Delta}\right),$$

which, following the same reasoning than the one in the previous subsection, is easy to evaluate.

More concretely, we use the notation $f_{i,j,l,m}^d \equiv f_i^d(a_j, B_l, N_m)$ and define a time step $\Delta t = \frac{\Delta}{S}$, where $1 << S \in \mathbb{N}$ is a constant. If we solve the KF equation (39) using a finite difference

scheme, we have, for $t = (d-1)\Delta$ and $s = 1, .., S-1$, where

$$\begin{aligned} \mathbf{f}^d_{t+s\Delta t} &= \left(\mathbf{I} - \Delta t \mathbf{A}^{\mathbf{T}}\right)^{-1} \mathbf{f}^d_{t+(s-1)\Delta t}, \\ \mathbf{f}^d_t &= \mathbf{g}_t \delta_{N_{(d-1)\Delta}} \delta_{B_{(d-1)\Delta}}, \end{aligned}$$

where δ is the Kronecker delta and $\mathbf{f}^d_t$ is defined as

$$\mathbf{f}^d_t = \begin{bmatrix} f_{1,1,1,t} \\ g_{1,1,1,2,t} \\ \vdots \\ g_{2,J,L,M,t} \end{bmatrix}.$$

The conditional density $p_X\left(X_{d\Delta}|X_{(d-1)\Delta};\gamma\right)$ can be approximated by

$$p_X\left(X_{d\Delta}|X_{(d-1)\Delta};\gamma\right) = \sum_{i=1}^{2}\sum_{j=1}^{J}\sum_{l=1}^{L} f^d_{i,j,l^d,m} g^d_{i,j} \Delta a \Delta B,$$

where $f^d_{i,j,l^d,m}$ is the density evaluated at the observed equity point $N_{d\Delta}$, $f^d_i(a_j, B_l, N = N_{d\Delta})$ and $g^d_{i,j}$ are the elements of the observed distribution $\mathbf{g}_{d\Delta}$.

5.3 Maximizing the likelihood

Once we have evaluated the likelihood, we can either maximize it or perform Bayesian inference relaying on a posterior sampler. In this paper, for clarity of exposition, we follow the former approach. Also, since we are dealing with a rather novel approach to solution and estimation of models with heterogeneous agents, we want to keep the estimation relatively simple, and we fix most of the structural parameters at conventional calibrated values for the U.S. economy and use only aggregate data. This choice will help when we conduct robustness exercises below.

As our data, we will rely on U.S. quarterly output observations for 1984Q1-2017Q4, with bandpass filter keeping frequencies between 20 and 60 quarters (between 5 and 15 years). We start in 1984, as often done in the literature, to focus in capturing the dynamics that have governed aggregate fluctuations in the U.S. after the arrival of the great moderation (Galí and Gambetti, 2009).[8] We bandpass the data as to eliminate long-run trends and to skip the business cycle frequencies caused by productivity shocks and monetary policy shocks our model is not designed to account for. Note, however, that our methodology does not depend on this

[8]See also the updated evidence in Liu et al. (2018), who document how the great moderation has clearly survived the financial crisis of 2008.

filtering and that a richer model could be estimated with raw data without theoretical problems.

We tweak and tune the standard neural network training scheme in several ways to meet the specific demands of the problem at hand.

First, it needs to give a good approximation in a very consistent way: a not-good-enough approximation in any of the dozens of iterations of the algorithm can make it break and not deliver a result. This is the reason we switched to line-search instead of using a constant or adaptative learning rate: it prevents bad steps in the minimization algorithm.

Second, it can't introduce big amounts of noise in the progressive approximation to the solution of the model, otherwise that noise can mask or prevent convergence (strict convergence criteria can only be met by chance, if at all). This is the reason we use batch gradient descent instead of stochastic or mini-batch gradient descent (i.e. all training points are used in every gradient calculation and line-search step): the random choice of traning points in each step, common in the machine learning literature, would make progress stochastic and introduces noise.

Reducing stochastic elements in the solution method is also the reason we train the model not with the noisy simulated points, but with a grid approximation to that that clears out noise: we define a 101x101 grid over the B,N support, and assign each simulated point to one of the knots in that grid; with all the points assigned to each knot we run a linear regression, and use that to estimate the height of the PLM at that knot. This grid could later be used to solve the model using interpolation (e.g. with splines, or with natural neighbor interpolation) and linear extrapolation (and we do that as a robustness check, finding similar results to those of the solution with the neural network), but on a 2D surface such as our PLM that extrapolation tends to generate ridges (because of the amplification of sample noise in far extrapolations), which could prevent convergence at the HJB step. Instead, we use those knots to train the neural network, which provides a good-enough fit in the visited area and a much smoother extraplation to the non-visited area of the B,N support.

Finally, the need to avoid stochastic elements that introduce noise in the algorithm, plus a desire for fast running times, made us change the usual initialization of the neural network parameters for all iterations except the first one: in the first iteration of the algorithm, we do ten random initializations of the parameters of the neural network and ten subsequent training sessions, later choosing the best-performing trained network across those ten training sessions. From the second iteration onwards, the neural network is initialized using weights that were found to be optimal in the previous iteration, and a single training session is carried out. This avoids re-introducing a stochastic element on each step of the algorithm that would greatly reduce its ability to meet a strict convergence criterion (or, more accurately, it would require much longer training sessions in order to clear out that noise and guarantee convergence). Using a small regulaization parameter in the PLM update step (it starts at 0.30 and exponentially

decays towards 0.05, reaching 0.20 after 5 iterations and 0.10 after 16) makes the convergence of the full algorithm slower but smoother, with slow updates to the optimal neural network that help this non-random initialization work well.

We report rates at an annual term. The capital share parameter, α, is taken to be 0.35 and the depreciation rate of capital, δ, is 0.1. The discount rate ρ, is set to 0.05. The risk aversion of the households γ is set to 2. All these are standard values in the business cycle literature to match aggregate observations such as investment-output ratios and the rates of return on capital.

The idiosyncratic income process parameters are calibrated following our interpretation of state 1 as unemployment and state 2 as employment. The transition rates between unemployment and employment (λ_1, λ_2) are chosen such that (i) the unemployment rate $\lambda_2/(\lambda_1+\lambda_2)$ is 5 percent and (ii) the job finding rate is 0.3 at monthly frequency or $\lambda_1 = 0.986$ at annual frequency.[9] These numbers describe the 'US' labor market calibration in Blanchard and Galí (2010). We normalize average income $\bar{y} = \frac{\lambda_2}{\lambda_1+\lambda_2} y_1 + \frac{\lambda_1}{\lambda_1+\lambda_2} y_2$ to 1. We also set z_1 equal to 71 percent of z_2, as in Hall and Milgrom (2008). Both targets allow us to solve for y_1 and y_2. We set the experts' discount rate $\widehat{\rho}$ to ensure that the leverage ratio K/N in the most-visited SSS is 2, which is roughly the average leverage from a Compustat sample of non-financial corporations. Table 1 summarizes our baseline calibration.

Table 1. Baseline parametrization

Parameter	Value	Description	Source/Target
α	0.35	capital share	standard
δ	0.1	capital depreciation	standard
γ	2	risk aversion	standard
ρ	0.05	households' discount rate	standard
λ_1	0.986	transition rate unemp.-to-employment	monthly job finding rate of 0.3
λ_2	0.052	transition rate employment-to-unemp.	unemployment rate 5 percent
z_1	0.72	income in unemployment state	Hall and Milgrom (2008)
z_2	1.015	income in employment state	$\mathbb{E}(y) = 1$
$\widehat{\rho}$	0.0497	experts' discount rate	$K/N = 2$

We solve the model according to the algorithm in Section 4. Then, we evaluate the likelihood of the observations on U.S. output for different values of σ, the volatility of the aggregate shock and, therefore, the most interesting parameters in terms of the properties of the model. While doing so, we keep all the other parameter values fixed at their calibrated quantities. We

[9] Analogously to Blanchard and Galí (2010, footnote 20), we compute the equivalent annual rate λ_1 as $\lambda_1 = \sum_{i=1}^{12} (1-\lambda_1^m)^{i-1} \lambda_1^m$, where λ_1^m is the monthly job finding rate.

maximize the likelihood function by searching on a grid between 0.008 and 0.018 with a step 0.00025.

We plot the resulting log-likelihood in Figure 1. The point estimate, 0.0142, is drawn as a vertical discontinuous red line, with a standard error TBC, computed by the local derivative of the function. The smoothness of the plot confirms, also, that our algorithm has successfully converged and that changes in one parameter value do not lead us into substantially different numerical solutions.

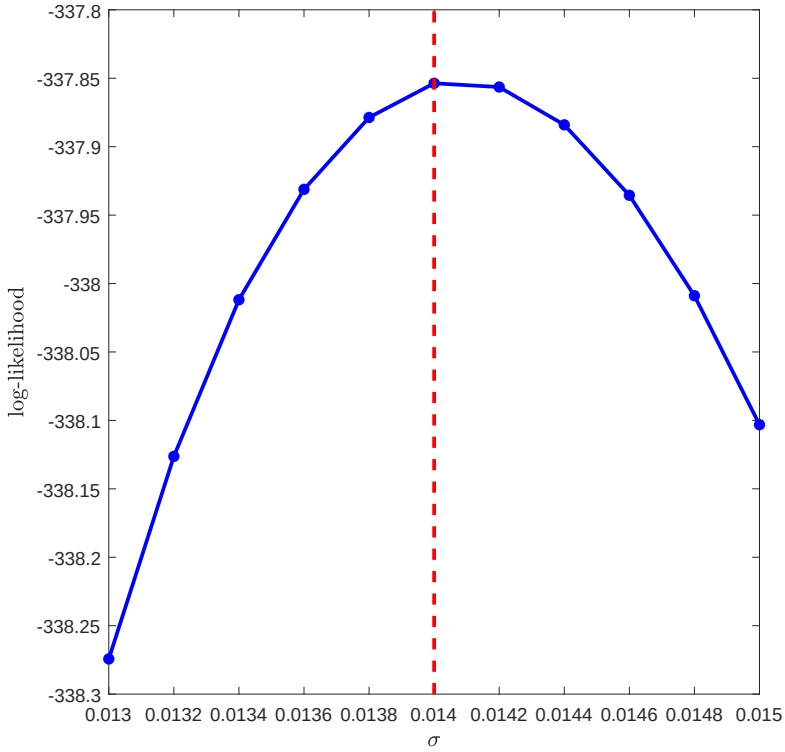


Figure 1: Log-likelihood for different values of σ and point estimate.

6 Quantitative results

In this section, we present the quantitative results generated by our solution algorithm with the parameter values from Section 5. We will report, first, the PLM for aggregate debt, $h(B, N)$. We will also assess our approximation for $dh(B, N)$ and evaluate the accuracy of our solution. Next, we will explore the phase diagram of the model to understand why we find several SSS(s) and analyze the role of the value of σ in determining those SSS(s). Then, we will move to document the aggregate ergodic distribution of debt and equity and the wealth distributions in the DSS and SSS(s). We will close by looking at the dynamic responses of the model to an aggregate shock in terms of aggregate variables at different points of the state space.

6.1 The PLM

The resulting PLM for aggregate debt, $h(B, N)$, is reported in Figure 2. Panel (a), at the top left, displays three transversal cuts of $h(B, N)$ along a range of values of equity (N). The first cut is fixes B at its baseline SSS value ($B = 1.9669, N = 1.7442$, with $K = 3.7111$ and $\frac{K}{N} = 2.1277$), the second cut fixes B at a selected high-leverage point ($B = 2.15, N = 1.5$, with $K = 3.65$ and $\frac{K}{N} = 2.4333$), and the third cut fixes B at a SSS with low-leverage point ($B = 1.0259, N = 2.6714$, with $K = 3.6973$ and $\frac{K}{N} = 1.3840$). The thicker part of the lines indicate the regions of the state space in which the ergodic distribution of aggregates variables, to be described below, is nonzero. The white point indicates, in the first cut, the baseline SSS; in the second cut, the high leverage point described above; in the third cut, the SSS with low leverage point. Panel (b), at the top right, follows the same pattern that panel (a), but switching the roles of equity (N) and debt (B). Finally, panel (c), at the bottom, shows the complete three-dimensional representation of the PLM. The shaded area in this bottom panels highlights the region of the PLM visited in the ergodic distribution with (non-trivial) positive probability. The thin red line is the "zero" level intersected by the PLM: to the right of the line, aggregate debt falls, and to the left, it grows.

Figure 2 demonstrates the non-linearity of $h(B, N)$ even within the area of the ergodic distribution that has positive mass. The agents in our economy expect very different growth rates of B_t in each region of the state space, with the function switching from concavity to convexity along the state space. While this argument is clear from the general shape of panel (c), it encodes rich dynamics that deserve further explanation. For example, panel (b) shows how, as leverages increases, $h(B, N)$ becomes stepper and, in the ergodic distribution, more concave. Given the same level of debt, a higher level of leverage induces larger changes in the level of aggregate debt as the financial expert is exposed to comparatively more capital risk. This result will resurface when we discuss the dynamics of the economy after an aggregate shock. The general shape of panel (c) also shows that, as intuition suggests, $h(B, N)$ is generally

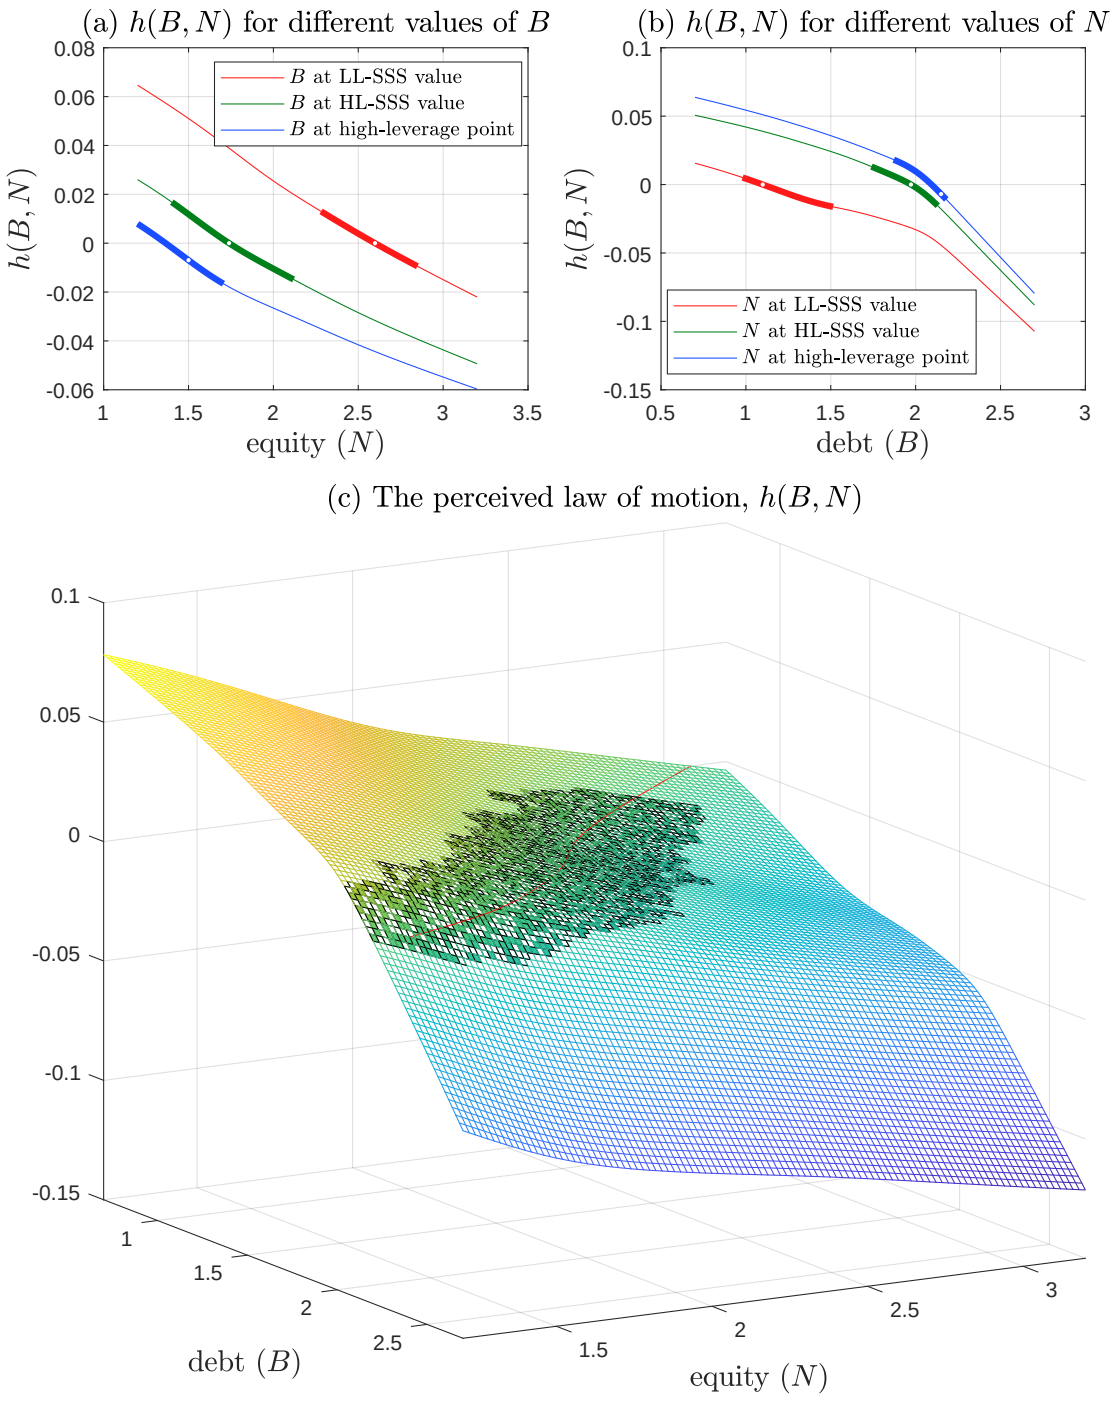


Figure 2: The PLM $h(B,N)$.

Note: White points in panels (a) and (b) indicate the most-visited SSS and selected points in the high- and low-leverage regions of the ergodic distribution. The thicker part of the lanes in panels (a) and (b) and the shaded area in panel (c) displays the region of the PLM visited in the ergodic distribution. The thin red line is the "zero" level intersected by the PLM.

decreasing in debt and equity.

Figure 3 reproduces the same three panels of Figure 2, except for $dh(B,N)$ once we have

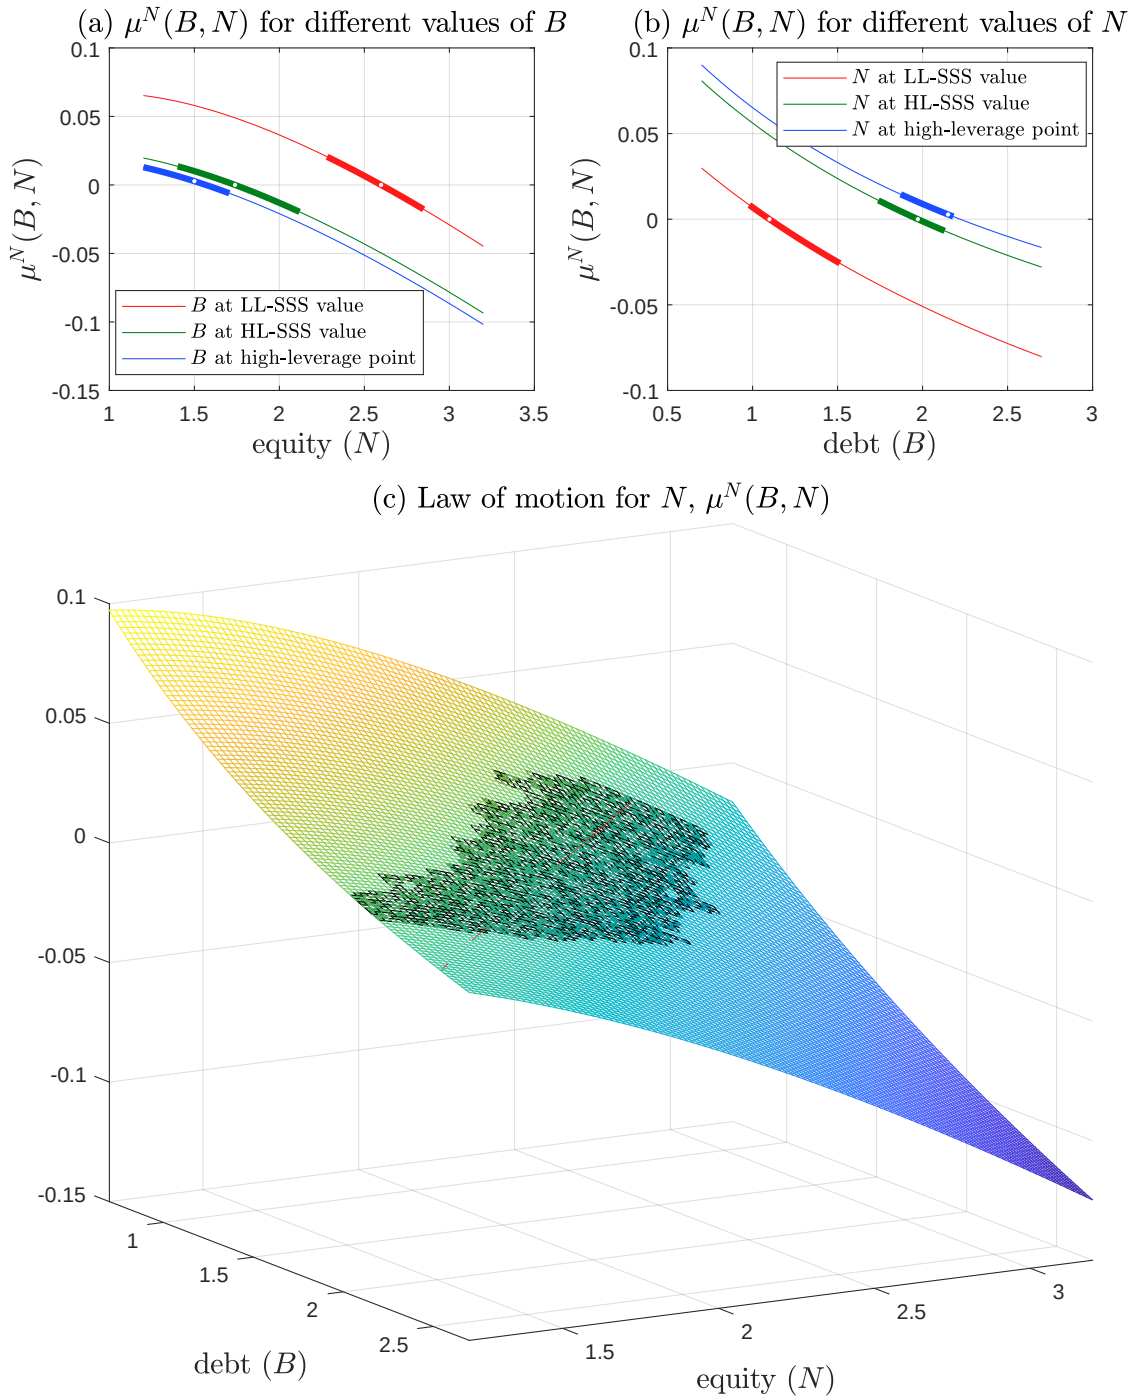


Figure 3: The law of motion $dh(B,N)$.

Note: White points in panels (a) and (b) indicate the most-visited SSS and selected points in the high- and low-leverage regions of the ergodic distribution. The thicker part of the lanes in panels (a) and (b) and the shaded area in panel (c) displays the region of $dh(B,N)$ actually in the ergodic distribution. The thin red line is the "zero" level intersected by $dh(B,N)$.

used equation (32). Similar comments regarding the non-linear structure of the solution apply here. For example, now, $dh(B,N)$ becomes higher as a function of equity as the level of leverage

falls.

The non-linearity of the PLM confirms our conjecture that more traditional solution methods that rely on linear structures (conditional on aggregate shocks) might not be appropriate for solving the class of models we are exploring. We can illustrate this argument more formally by looking at the forecasting capability of our PLM. The R^2 associated with the PLM we compute using our neural network is 0.9922, with an RMSE of 0.0004. The forecasting errors, furthermore, are nicely clustered around zero, with a mode roughly equal to zero. To compare it with the standard Krusell and Smith (1998) algorithm of finding an OLS over a linear regression on endogenous state variables, we have recomputed our model using the latter approach. The R^2, in that case, is 0.8275, considerably lower than typical values reported in the literature for more standard variations of the stochastic neoclassical growth model with heterogeneous agents, and with an RMSE of 0.0021.

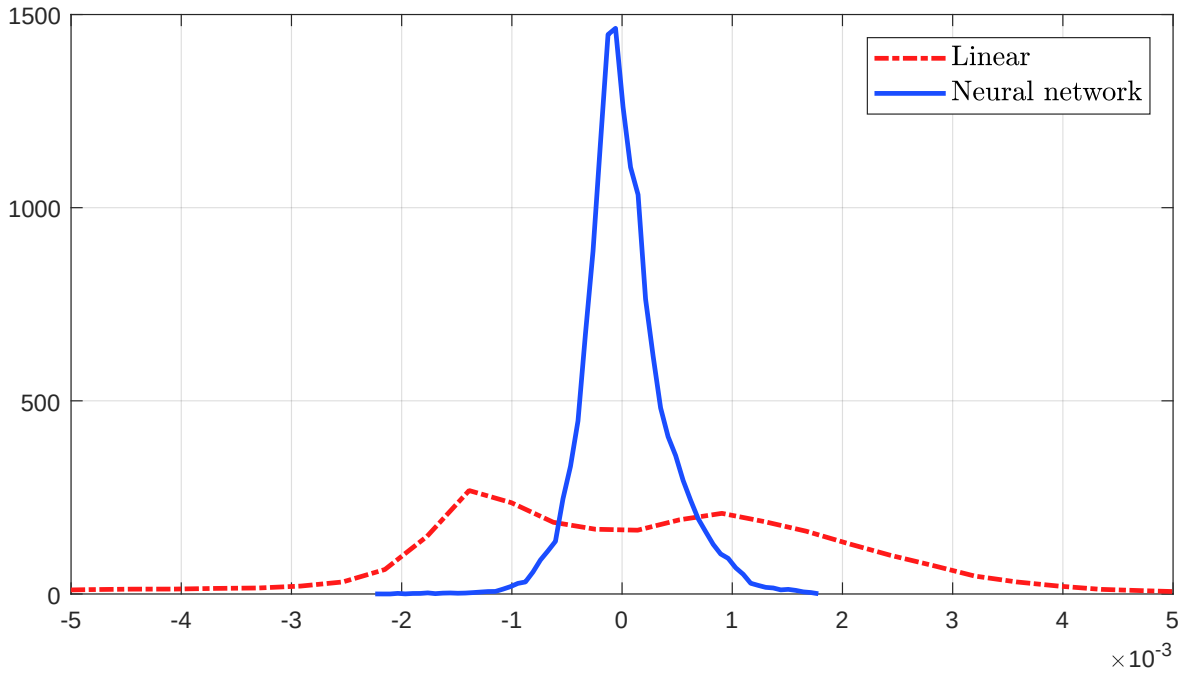


Figure 4: Forecasting error distribution at a 1-month horizon, linear PLM (left) and nonlinear (right).

Figure 4 plots the histogram of forecasting errors at a 1-month horizon (the time step selected in the simulation step), with a continuous blue line representing the errors from our algorithm and the discontinuous red line the errors from a linear-in-endogenous variables Krusell-Smith algorithm. The latter approach produces more volatile forecasting errors, which are also skewed to the right and without a mode at zero. These results back the importance of taking into account the nonlinearities of the model when computing the PLM to achieve the required

accuracy.[10]

6.2 The phase diagram

Figure 5 plots the phase diagram of our model along the aggregate debt (B) on the x-axis and equity (N) on the y-axis.

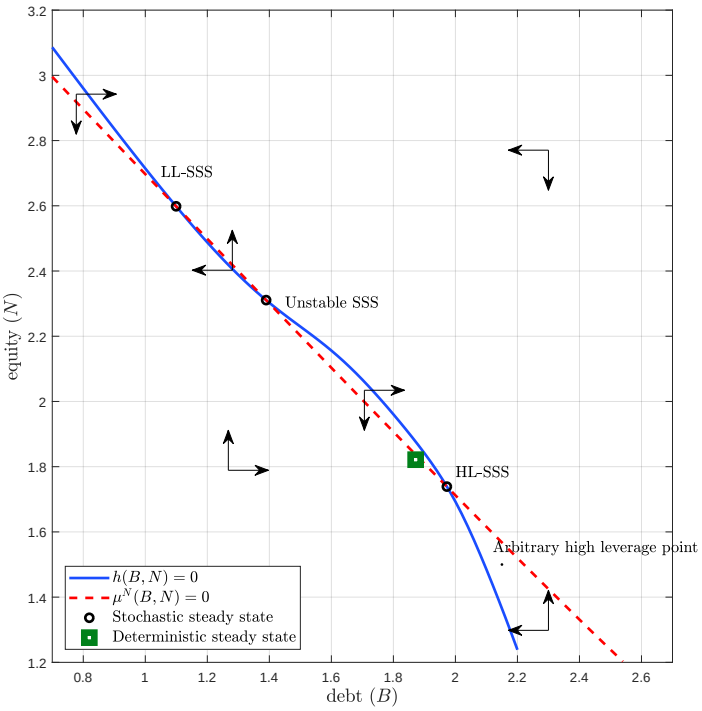


Figure 5: Phase diagram, DSS, and SSS(s).

The blue line in Figure 5 represents the loci of zero changes in debt, $h(B,N)=0$. The line inherits the non-linear dependence on B and N of the right-hand side of equation (19), the object $h(B,N)$ approximates. In particular, there is a convex segment for low levels of debt and a concave segment for high levels of debt. The discontinuous red line represents the loci of zero changes in the equity of the expert, $dN(B,N)=0$. These two lines are the intersections of zero with $h(B,N)$ and $dN(B,N)$ in Figures 2 and 3, which we represented –in those figures– with a thin red line. The arrows in Figure 5 indicate the movement of debt, B, and equity, N, when we are away from the blue and red lines. For completeness, we also plot the point of high leverage that we use in Figures 2 and 3 (and later, when we compute the GIRFs of the model) to show the behavior of the economy when leverage is high.

The two lines intersect three times, defining three SSS(s). From the bottom right, the first intersection is the baseline, stable SSS (recall, with $B=1.9669, N=1.7442$, with $K=3.7111$

[10]In the Appendix, we discuss other alternatives to the standard Krusell and Smith (1998) algorithm and argue that our method has advantages over them as well. Also, we checked that adding additional moments to the OLS regression do not help much concerning accuracy.

and $\frac{K}{N} = 2.1277$). We call this SSS "baseline" because of two reasons. First, this SSS (and its neighborhood) is the most visited one in the ergodic distribution (see Figure 12 below). Second, this SSS is the closest to the DSS (green square). In comparison with the DSS ($B = 1.8718, N = 1.8215$, with $K = 3.6933$ and $\frac{K}{N} = 2.0276$), the baseline SSS has 0.5% more capital, 5.1% more debt, and 4.2% less equity.

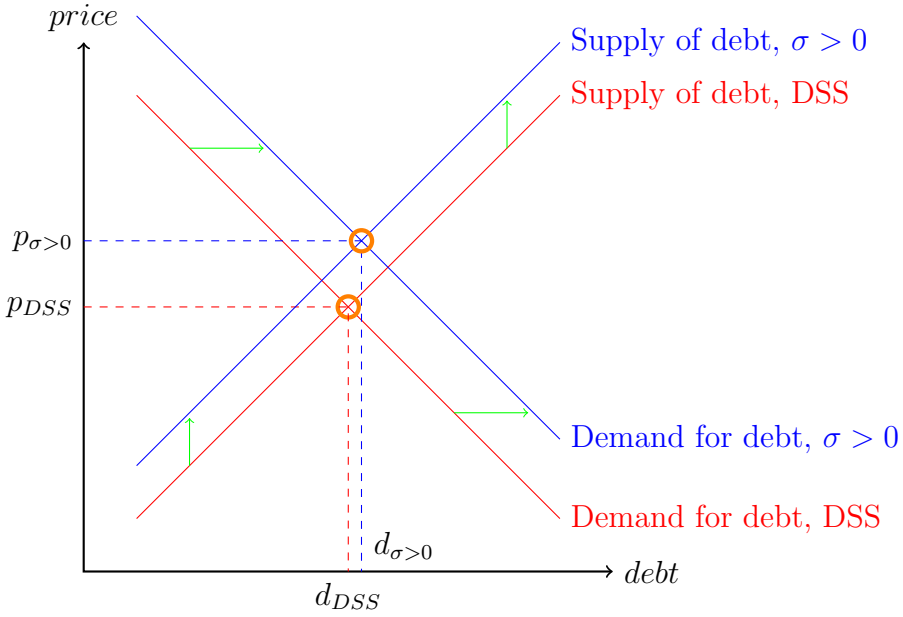


Figure 6: DSS vs. $\sigma > 0$.

Figure 6 illustrates the price mechanism that leads to more debt and less equity. In red lines, we plot the demand curve for debt by households and the supply curve of debt by the expert at the DSS as a function of the price of debt. These curves take $\sigma = 0$ and the wealth of each agent as given. Then, we look at the effect of making $\sigma > 0$. Everything else equal, this aggregate uncertainty moves to the left the supply curve of debt (increasing blue line) because it generates risk for the expert's capital holdings. The expert lowers her demand curve for capital and her willingness to issue debt to finance it through leverage. Aggregate uncertainty also creates additional consumption risk for the households due to the stochastic variation of wages and the risk-free rate (in the DSS, households are already exposed to idiosyncratic risk regarding their labor productivity). The precautionary motive induced by the aggregate risk shifts the household's demand for debt to the right (the decreasing blue line in Figure 6). Given a lower supply and higher demand curve for debt, the risk-free rate (i.e., the inverse of the price of debt) must fall to clear the debt market.

While the effect on the risk-free interest rate of $\sigma > 0$ around the DSS is unequivocal, the impact on quantities depends on whether the supply or the demand curve move more. When $\sigma = 0.014$, the movement to the right of the demand curve is larger than the movement to the left of the supply curve, and the new market-clearing level of debt is higher than in the DSS. The

lower risk-free rate persuades the expert to rely more on debt and less on equity and gather the larger excess returns from leverage until the economy reaches the baseline SSS (recall the phase arrows in Figure 5; Figure 6 only performs static comparative statics on σ around the DSS). The dynamic effect of this excess return is so strong that capital ends up being slightly higher in the baseline SSS than in the DSS. Hence, at this SSS, the expert absorbs much of the aggregate risk in the economy. We will see below how different values of σ change the net effects of the moves of the supply and demand curves of debt and, with them, the dynamic consequences of variations on σ on quantities and the distribution of aggregate risk in the economy.

Mechanically, the risk-free interest rate in the baseline SSS is given by equation (17):

$$r^{SSS} = \alpha \left(K^{SSS}\right)^{\alpha-1} - \delta - \sigma^2 \frac{K^{SSS}}{N^{SSS}},$$

whereas in the DSS is given by equation (30):

$$r^{DSS} = \alpha \left(K^{DSS}\right)^{\alpha-1} - \delta.$$

From these two equations and our parameter values, $K^{SSS} = 3.7111 > K^{DSS} = 3.6933$, we get $r^{DSS} = 0.0497 > r^{SSS} = 0.0488$. This lower risk-free rate is due to two terms. First, a higher capital lowers the marginal productivity of capital from 0.1497 to 0.1492. Second, the term $\sigma^2 \frac{K^{SSS}}{N^{SSS}}$ pushes down the risk-free rate an additional 4.17 basis points (the average excess return).

The second intersection in Figure 5 is at a middle SSS with less debt and more equity ($B = 1.3800, N = 2.3204$, with $K = 3.7004$). This SSS is, however, unstable, and the dynamics of the economy quickly move away from it. Thus, we will not discuss it further.

Finally, the third intersection in Figure 5, at the top left, is a stable SSS with low leverage: debt is much smaller ($B = 1.0259$) and equity much higher ($N = 2.6714$) than in the baseline stable SSS, yielding $K = 3.6973$. Figure 7 helps explaining why we have this low-leverage SSS by performing a comparative statics exercise. The red lines are the same supply and demand curves at the DSS than in Figure 6. Imagine now that, starting at the DSS, we implement a large redistribution of wealth from the households to the expert, while keeping everything else equal. The negative wealth effect for the household moves their demand for debt (decreasing blue line) to the left. Conversely, the higher equity for the expert moves its supply curve of debt (increasing blue line) to the left since she does not need much additional capital.[11] The market clearing quantity of debt is much lower than the debt at the DSS.

[11] We need the redistribution of wealth to be large to get the financial expert's equity sufficiently close to the total optimal amount of capital. When the redistribution around the DSS is small, the financial expert still wants takes advantage of the high excess returns of capital by issuing more debt. See, again, the phase arrows in Figure 5 to visualize these dynamics.

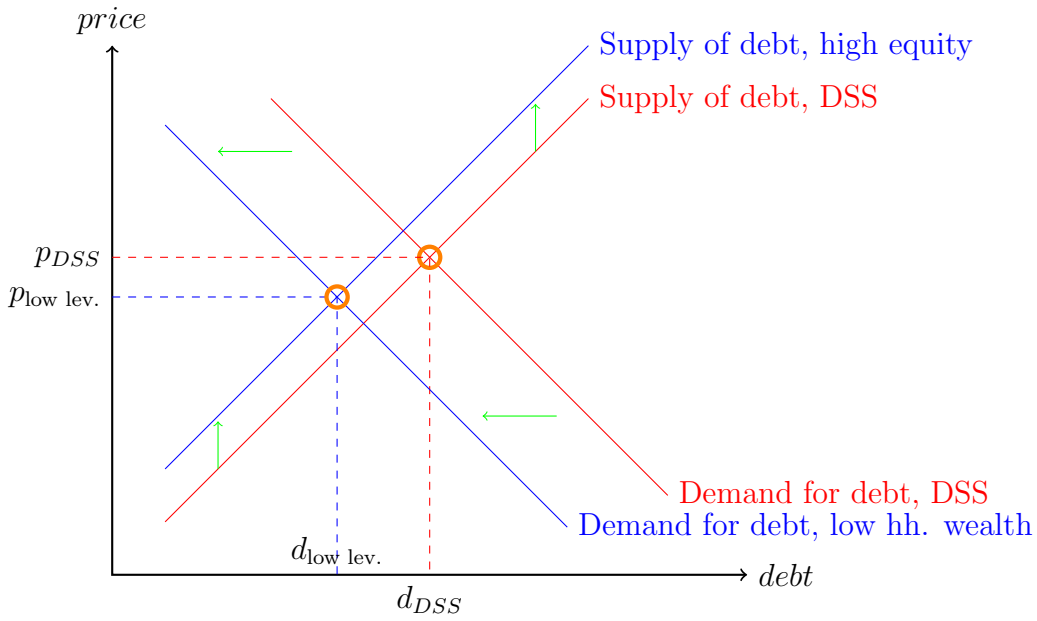


Figure 7: DSS vs. low-leverage.

If the role of σ around the DSS was clear for prices and ambiguous for quantities, now we face the reverse situation: we know that quantities fall, but the evolution of the risk-free rate will depend on the particular value of σ. For our point estimate of $\sigma = 0.014$, the low leverage translates into a high risk-free rate, 0.0493, and a low average excess return, 3.77 basis points. This low excess return ensures the local stability of this SSS, as the expert has a much lower incentive to gather excess returns by increasing her leverage (notice, from equation (16), that the expert's desire leverage is quite sensitive to the excess return since this return is divided by σ^2, a small number).

Due to the low debt of this SSS, households are less capable of smoothing consumption after an aggregate shock. Loosely speaking, at the low-leverage SSS, there is a "shortage" of debt from the perspective of the households that does not exist in the baseline SSS. Because of this latent desire of households for much higher savings as a response to relatively small changes in prices or wealth, the distance with the unstable SSS of this low-leverage SSS is lower than the distance between the unstable SSS and the baseline SSS. Shocks are more likely to push the economy from a point to the left of the unstable SSS toward the right of the unstable SSS than the other way around. Consequently, as we will document below, this low-leverage SSS occurs much less often than the baseline SSS.

We can document how the economy evolves around each SSS by plotting, in Figure 8, the basins of attraction of each SSS. In red, we plot the paths around the low-leverage SSS and, in blue, the paths around the baseline SSS. When the shocks push the economy from one basin to the other, the economy will tend to move, on average, toward the corresponding SSS.

How can we have two such different stable SSS(s)? The basic intuition is that, in our economy

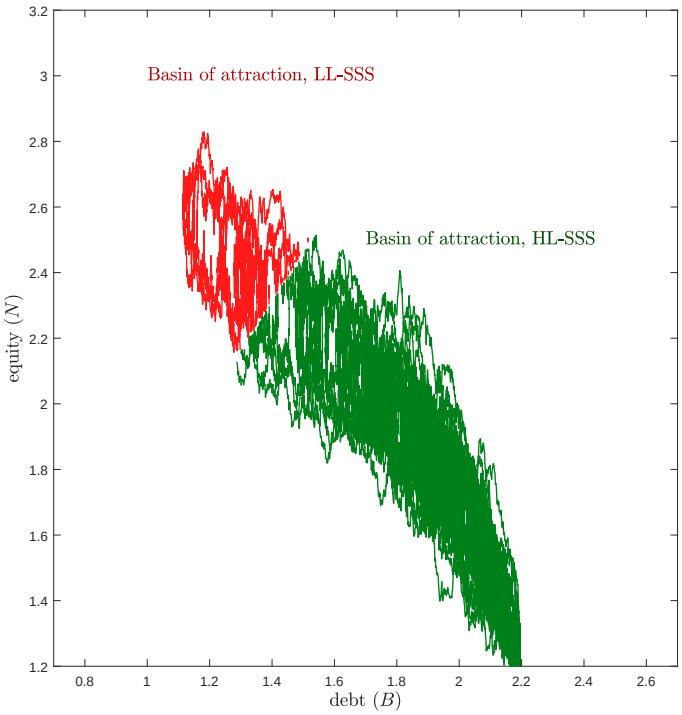


Figure 8: Simulation of equilibrium paths.

with financial frictions, we can divide the (approximately) same total wealth in multiple ways between the financial expert and the household. Since each division of wealth distributes the aggregate risks in different ways among the agents, there will be, in general, multiple leverage levels that would induce no further changes in equity and debt in the absence of shocks (i.e., the multiple SSS(s)) through the variations in excess returns. The impossibility of issuing state-contingent debt by the expert and the high sensitivity of leverage to excess returns eliminates the "concavity" forces that lead to a unique SSS in a neoclassical setup through a more complete allocation of the risk.

6.3 The role σ in determining the different SSS(s)

The previous discussion highlighted how the equilibrium risk-free rate distributes the consequences of aggregate risk among agents. We can push this argument further by exploring the role of σ in determining the different SSS(s) in Figure 9 through its effect on the risk-free rate and excess returns. Each panel plots, for a different value of σ, the phase diagram of the economy following the same convention than Figure 5. For low values of σ, the financial experts can absorb much of the aggregate risk and, therefore, the baseline SSS is to the right of the DSS, i.e., the baseline SSS has more debt and less equity than the DSS (we still have, nevertheless, the second stable SSS with much more equity and less debt). However, as σ increases, the expert reduces its desired level of leverage and the baseline SSS gets closer to the DSS, until it

crosses it for σ around 0.023 and we end up with a baseline SSS with less debt and more equity. Also, as σ rises, $h(B, N) = 0$ becomes curvier and the additional SSS(s) disappear.

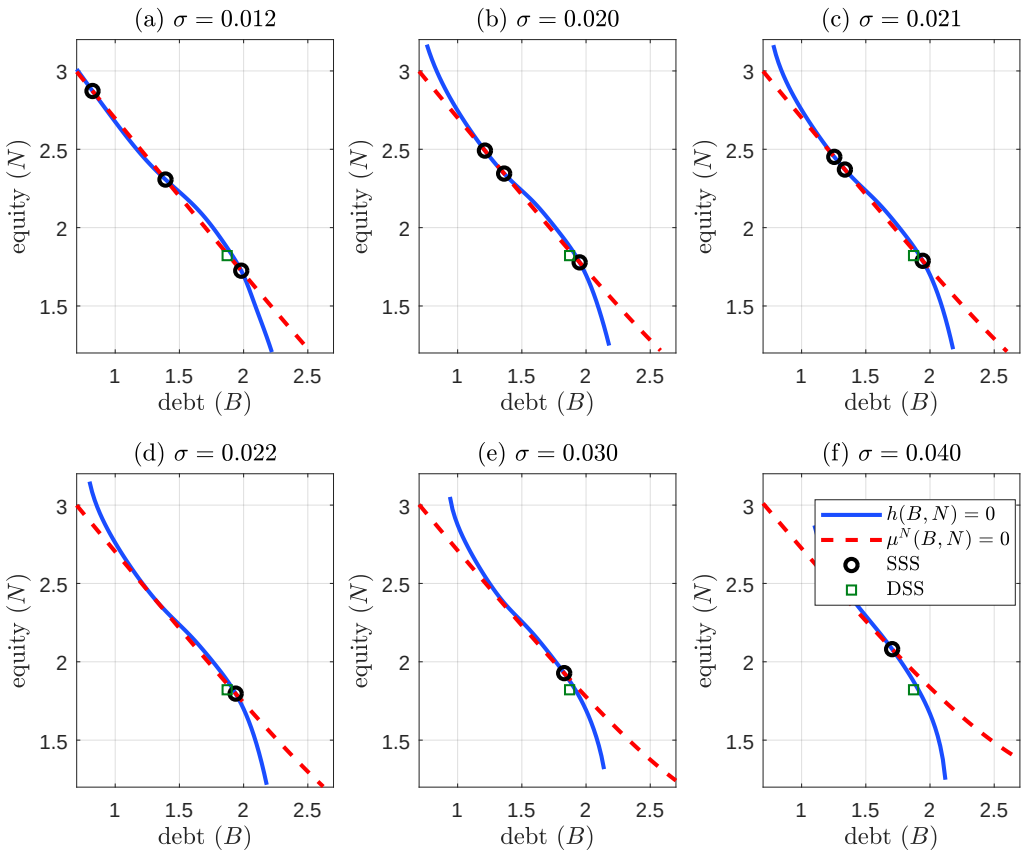


Figure 9: Phase diagram as a function of σ.

Figure 10 complements the previous results by plotting the values of the low-leverage SSS, the unstable SSS, and the baseline SSS (plus, for reference, the DSS) as a function of σ. We can see how the leverage in the baseline SSS is a negative function of σ, a roughly constant function in the unstable SSS, and an increasing function in the low-leverage SSS (until these additional SSS disappear). The mechanism for these three slopes is the same than the one we highlighted above. In the baseline SSS, as σ grows, the financial expert wants to unload some of the capital risk by reducing its leverage. In comparison, in the low-leverage SSS, the households demand more debt as σ increases (recall that at low-leverage, there is a "shortage" of debt that does not exist in the baseline SSS).

6.4 Convergence to the SSS(s)

Does the economy converge to any of the SSS(s)? The state space $(g(\cdot), N)$ is infinite dimensional and, hence, we cannot numerically analyze convergence for any possible initial state. Instead, we analyze convergence for densities visited in the aggregate ergodic distribution. In

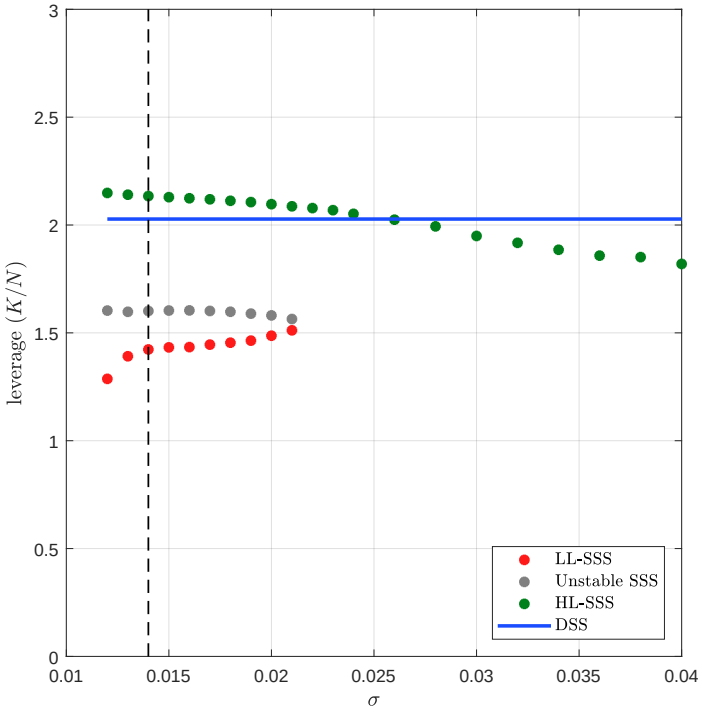


Figure 10: SSS as a function of σ.

Figure 11, we consider an array of different initial values, i.e., income-wealth densities and equity levels $(g_0(\cdot), N_0)$, selected from the simulations employed to compute the aggregate ergodic distribution above and analyze the transitional dynamics in the absence of aggregate shocks (even if agent continue forming their expectations assuming the continuous arrival of aggregate shocks with $\sigma > 0$). In all the cases considered the economy converges to the SSS. The convergence path is very slow: for initial values in the high- or low-leverage regions convergence may take several centuries.

Notwithstanding, we cannot rule out that for other densities than the ones in the aggregate ergodic distribution the model would not converge to the SSS. This limitation is related to the self-confirming nature of the solution. The PLM is computed based on the income-wealth distributions visited along the paths of the ergodic distribution. Our algorithm could potentially find a distribution with disparate first moments that would lead to alternative dynamics.

An interesting feature of the transitional dynamics is the curvature of the paths that lie in regions with relatively low levels of high capital. In these cases, the paths are bent because debt adjusts faster than equity, the latter mainly depending on the leverage ratio. If equity and aggregate capital are below (above) their long-run values households (de-)accumulate wealth increasing their leverage. Once they reach a point with still low (high) equity, but a capital level closer to that in the SSS, there is a progressive redistribution of wealth from households to the expert (or the other way around in the case of initial high equity and capital). In the next subsection, we discuss the model's ergodic distribution in more detail.

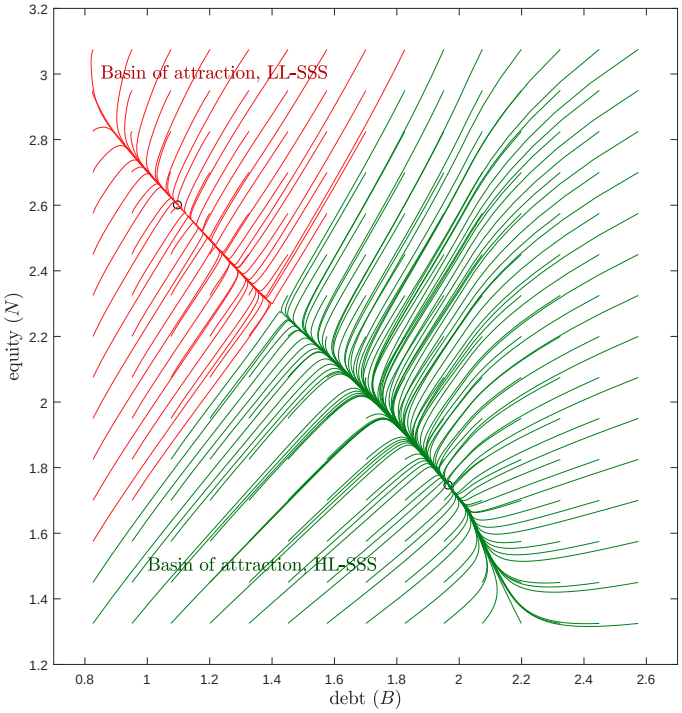


Figure 11: Convergence paths.

6.5 The aggregate ergodic distribution of debt and equity

Panel (a) of Figure 12 displays the aggregate ergodic joint distribution of debt and equity $F(B,N)$. This distribution is defined as

$$\mathbb{P}\left\{(B,N)\in\Omega\right\}=\int_{\Omega}dF,$$

for any subset Ω of the state space.

The ergodic distribution is not obtained directly from the PLM, but from the simulation of the paths of the income-wealth distribution and equity. Naturally, the PLM is employed in the HJB equation (33) to obtain the optimal consumption policies of individual households. We use a Monte Carlo simulation of 5,400 years at a monthly frequency. We initialize the model at the DSS, and we disregard the first 400 samples as a burn-in.[12]

Panels (b) and (c) of Figure 12 plot the marginal ergodic distributions for debt and eq-

[12] As we mentioned in Section 4, here we face the problem of extrapolation. We compute the PLM based on simulations on a relatively limited region of the state space. However, when forming expectations, households evaluate the PLM over the entire state space. Thus, the PLM is extrapolated over the regions of the state space not included in the support of the ergodic distribution. There is no guarantee that the dynamics of the model in the extrapolated region coincide with the ones expected in the PLM. Thus, exactly as it happens in Krusell and Smith (1998), our approximation can be understood as a particular instance of a self-confirming equilibrium (SCE). See Brumm and Scheidegger (2017) and Piazzesi and Schneider (2016) for related discussions. In a SCE, households' beliefs about the PLM coincide with the actual law of motion only in the equilibrium paths. Off-equilibrium they may diverge, but households never discover it as this region is never visited.

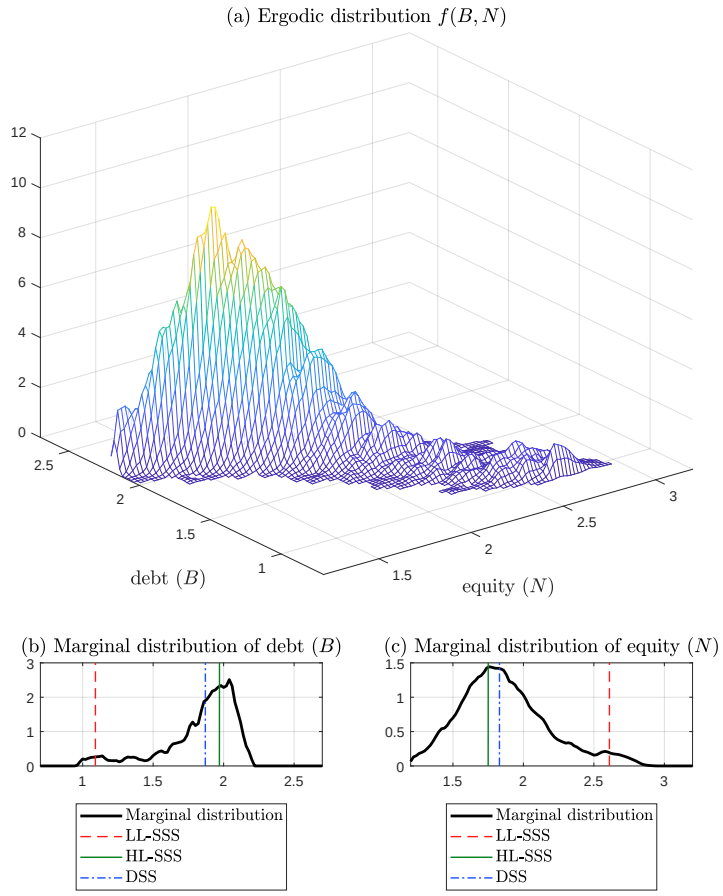


Figure 12: Ergodic distribution $F(B, N)$. Lighter colors indicate higher probability.

uity. These two panels show how the economy spends most of the time in a region with debt levels between 1.5 and 2.2 and equity between 1 and 2.2. Those values of debt and equity correspond to the neighborhood of the baseline SSS ($B = 1.9669, N = 1.7442$). In comparison, the neighborhood of the low-leverage SSS ($B = 1.3800, N = 2.3204$) appears much less often. However, we can see a small peak in the ergodic distribution slightly to the left of this SSS (i.e., less debt and more equity). This small peaks documents how, after the right sequence of shocks, the economy can linger in that region from some time, although less so than in models such as Brunnermeier and Sannikov (2014), which present a clearer bimodality of the ergodic distribution. As we explained above, the desire of households of accumulating more debt lowers the height of this second peak as there is a strong force for reversion toward higher levels of leverage.

Note, as well, how the marginal distribution of debt is much more concentrated than the marginal distribution of equity, reflecting how . The ergodic distribution has a substantial tail in areas of high equity and low debt, creating an interesting asymmetry. Finally, the total amount of capital ($K = B + N$) is on average larger in the low leverage region than in the high leverage one.

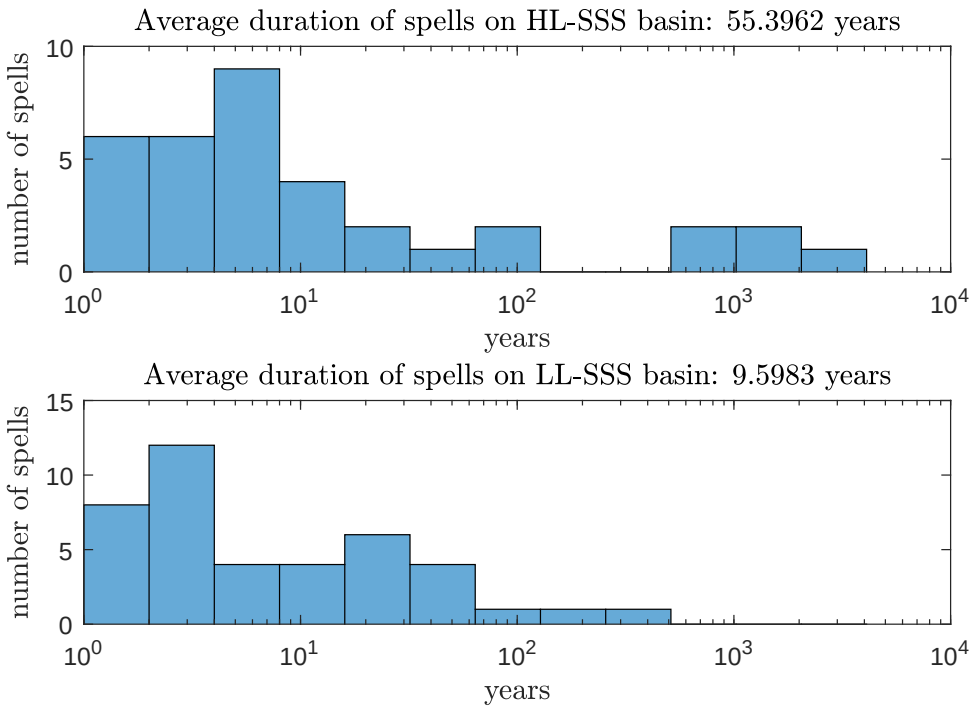


Figure 13: Spell durations at each SSS.

We complement the previous discussion with Figure 13, which plots the histogram of the duration of spells of the economy around the baseline SSS and the low-leverage SSS. The average duration of a spell of the economy around the baseline SSS is 48.69 years, a long period of relatively high leverage. On the other hand, the average duration of a spell of the economy around the low-leverage SSS is 4.75 years, a much shorter length, reflecting the strong push effect of high risk-free rates on household accumulation decisions (and, therefore, on the re-building of leverage). Nevertheless, the accumulated effects of repeated shocks also mean that the spells around this low-leverage SSS can be much longer than 4.75 years, with some lasting as long as a century. Figure 13 shows how the interaction of financial frictions and heterogeneity can lead to rather complex middle- and long-run dynamics that are not present in more conventional business cycle models.

6.6 Wealth distributions in the DSS and SSS(s)

Figure 14 compares the wealth distribution in the DSS (discontinuous blue line), the baseline SSS (continuous green line), and the low-leverage SSS (discontinuous red line). The wealth distribution slightly shifts to the left as we move from the DSS to the baseline SSS, and moves rather dramatically, also to the left, as we move from the baseline SSS to the low-leverage SSS.

As we saw before, everything else equal, the presence of $\sigma > 0$ lowers the risk-free interest rate. Thus, in the baseline SSS, the assets of the household are pushed to the left of where they

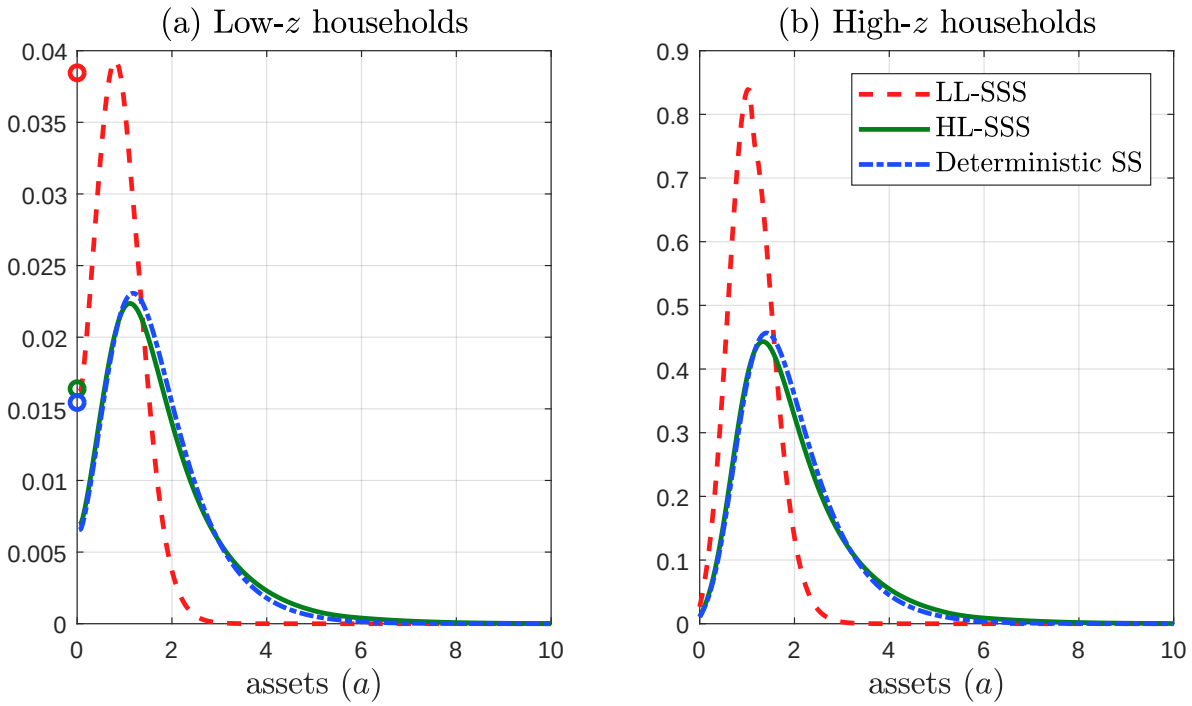


Figure 14: Wealth distribution in the DSS and SSS.

would be in the DSS. However, this move is small as the baseline SSS displays higher wages and lower risk-free rates due to the higher capital level and the higher leverage ratio. The net effect is that savings almost coincide for low-wealth households in the DSS and baseline SSS, but are higher for high-wealth households in the DSS. There is, however, a thicker right tail in the baseline SSS. Therefore, the presence of aggregate risk in this economy produces an allocation with a higher level of total savings and more wealth inequality, together with higher output and more leverage, compared to the allocation without aggregate risk. In the low-leverage SSS, in contrast, there is little debt (recall our comment above about the "debt shortage"), and most households accumulate closer to zero assets than in the baseline SSS.

As in Krusell and Smith (1998), our solution only considers the first set of moments of the wealth distribution to approximate the PLM. This "approximate aggregation" holds well in our model due to the linearity of the consumption policy. Only agents close to the borrowing limit face a nonlinear consumption policy but, being close to zero assets, they contribute relatively little to the aggregate dynamics of capital (and, as shown by Figure 14, there are not that many of them in any case). However, and this point is vital to understand our results, the quasi-linearity of the consumption function is only with respect to the household state variables, not with respect to the aggregate state variables. This is why we can still have highly non-linear behavior for the economy as a whole while we can approximate well the PLM as a non-linear function of only one moment.

6.7 Dynamics after an aggregate shock

The analysis above has shown how the introduction of aggregate risk modifies the distribution of wealth across households and between households and the expert. In this subsection, we analyze whether those different wealth distributions change the transmission of aggregate shocks.

Figure 15 displays the generalized impulse response functions (GIRFs) to a one-standard-deviation negative capital shock. The GIRF is defined as the difference between the transition path if an initial shock hits the economy and that if no shock arrives. We consider initial states $(g_0(\cdot), N_0)$ in the high- and low-leverage regions and the SSS.[13]

We study first the GIRFs to one-standard-deviation negative capital shock that hits the economy at the baseline SSS (continuous green lines). Time units are years. The shock destroys capital (left panel in the central row) and, as the expert absorb all the risk in a world where debt is not state-contingent, lowers equity by a considerable degree (right panel in the central row). The reduction in capital produces a contraction in output (left panel in the top row) and a decline in wages (left panel in bottom row). The lower level of capital increases the return on capital and the risk-free interest rate (center and right panels in bottom row). Over time, the expert rebuilds its equity by consuming less (note the significant drop in the expert's consumption at impact, right panel in top row), borrowing more (central panel in the central row), and accumulating excess returns.

A central property of the GIRFs is the persistence they show: even after 40 years, the economy is still around half a point percentage below its pre-shock level. The dynamics of equity accumulation (see, again, right panel in the central row) propagate aggregate shocks in ways that are simply not present in models without financial frictions.

The GIRFs when the shock hits at the low-leverage SSS (discontinuous red lines) and the high-leverage point we used before ($B = 2.15, N = 1.5$; discontinuous blue lines) are similar, but with a few noticeable differences. For example, when leverage is low, the expert reduces her consumption much less and issues debt much faster, which allows for quicker recovery in capital and output. In the opposite case of the high-leverage point, the output drop at impact is acuter and recovery of the economy even slower than in the baseline SSS (also, recall that the high-leverage point is not a fixed point of the equilibrium dynamics and that the economy will not return to it in the absence of aggregate shocks). This slow recovery is because, in the high-leverage case, the decline in consumption in the first decade after the shock is less pronounced. We will see, later, when we report the comparison of the time-varying consumption decision

[13]The word "generalized" comes because, in comparison with linear models, impulse responses in non-linear models are state dependent, that is, they are not independent of the point where the economy is when the shock arrives, their size depends non-linearly in the size of the shock, and they are not symmetric (i.e., the response to a negative shock is not the mirror image of the response to a positive shock. Thus, we need to specify the size and sign of the shock and when this shock occurs.

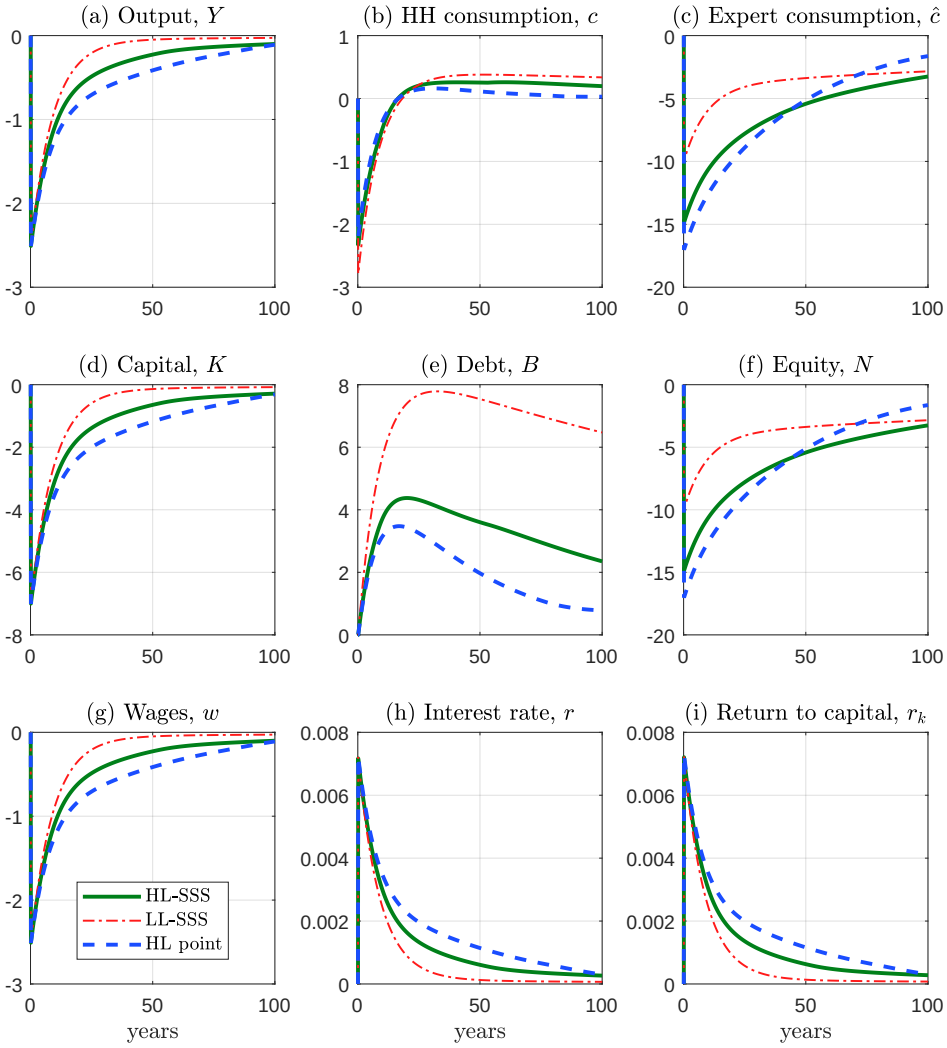


Figure 15: GIRFs for different initial states.

rules in Figure 19, how wealthy households mainly drive this phenomenon.

The differences among the GIRFs depending on the state at which the aggregate shock hits demonstrate, once more, the critical non-linearities of our model. Note, also, that the one-standard-deviation shock is not large enough to send the economy away from the basins of attraction of each of the two SSS. An even larger shock, by inducing a switch between the basins, will have even more dramatic non-linear consequences.

7 The role of heterogeneity

After presenting the quantitative results in the previous section, there is still one fundamental question we must address: Does heterogeneity matter? In other words, does the introduction

of heterogeneous households change the predictions of the model with respect to the representative household version outlined in Subsection 3.4?

We will answer this question in four steps. First, we will vary the level of idiosyncratic risk among the households. Second, we will report that the SSS of the economy with a representative household is quite different. Third, we will show how the GIRFs of the model are indeed different when we have a representative household or heterogenous households. Indeed, the response of the economy to the same aggregate shock may be strongly modified depending on the initial wealth allocation among households: either disperse with heterogeneity or collapsed without it. Second, we will explore how the differences in consumption decisions and the move in distributions over time account for the differences in the GIRFs.

7.1 Varying idiosyncratic risk

In Figure 16, we plot the phase diagram of the model for six different values of z_1 (still keeping the ergodic mean of z equal to 1). The first panel in the top row is the benchmark case of $z_1 = 0.72$ and, therefore, reproduces Figure 5. As we increase z_1, the DSS moves toward the left top corner rather quickly (low B and high N): with less idiosyncratic risk, households have less incentive to accumulate assets and the level of debt in the economy plummets. Because of the same reason, the high leverage SSS starts moving to the left until it disappears by the time we reach $z_1 = 0.9$.

In Figure 17, we plot the ergodic distribution of equity and debt. The ergodic distribution starts spreading toward the upper left corner first slowly, but gathering considerable steam by the time we reach $z1 = 0.8$. At this level, there is a bifurcation and region around the low leverage SSS becomes the predominant one, with the higher leverage region eventually disappearing. This change in the ergodic distribution is crucial for aggregate fluctuations since, as we saw before, the response of the economy to a capital shock is very different with low leverage than with leverage.

7.2 Different SSS(s)

When we have a representative household, the SSS of the model ($B = 3.4099, N = 0.5947$, with $K = 4.0046$ and $\frac{K}{N} = 6.7339$) is far away from any of the SSS(s) of the model with heterogeneous households. In particular, the level of leverage, $\frac{K}{N} = 6.7339$, is much higher

7.3 Different GIRFs: heterogenous vs. representative households

Figure 18 reports the GIRFs of the heterogenous households version of the model (HA; continuous green line) to a one-standard-deviation negative capital shock when the economy is

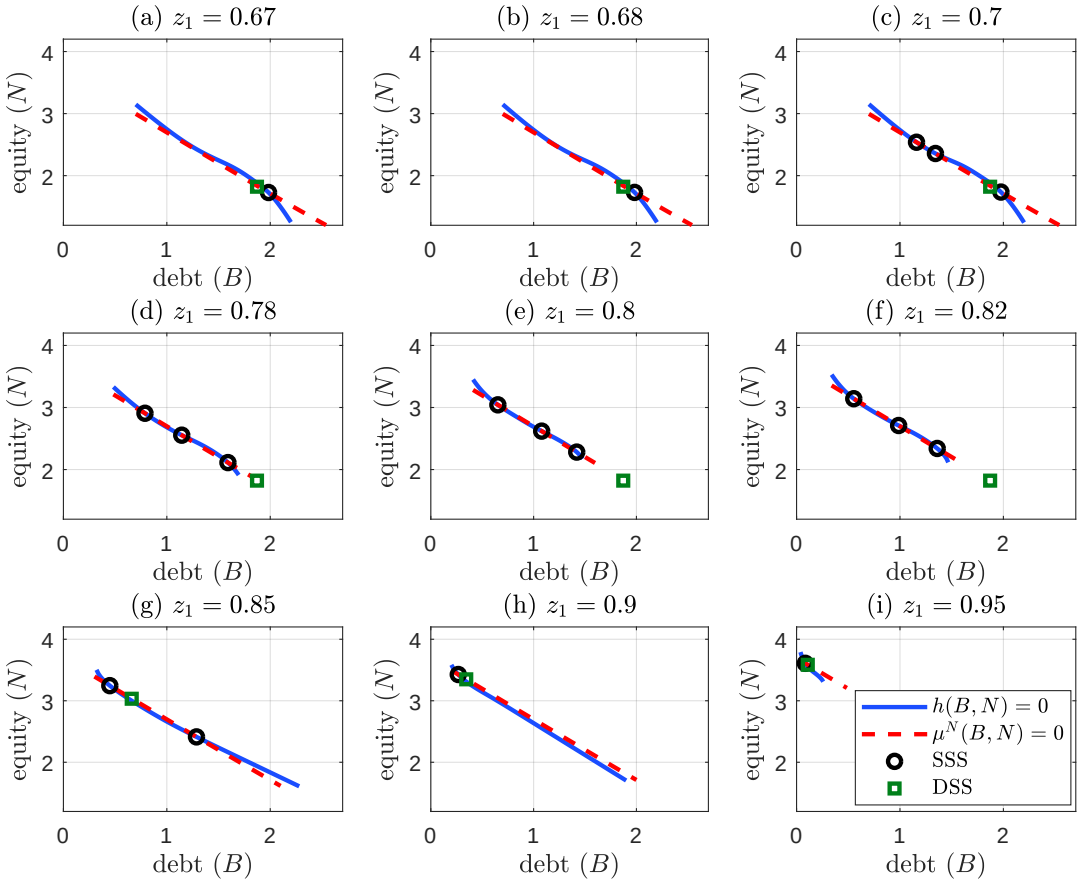


Figure 16: Phase diagram as a function of z_1.

at the baseline SSS. These GIRFs are, by construction, identical to the GIRFs in continuous green lines in Figure 15. Figure 18 also plots the GIRFs of the representative household version of the model (RA model, HA calibration; discontinuous blue line).[14]

Both sets of GIRFs display considerable differences. The fall in output (left panel in the top row) is considerably more persistent when we have a representative household. We will show below how this is related with the dynamics of consumption and wealth accumulation by rich households. The reduction of total consumption of the households (central panel in the top row) is lower at impact in the case of the representative household but its recovery is also slower. In comparison, the consumption of the expert drops much more when we have a representative household. This is because the expert starts with a higher leverage (see in right panel of the central row the large decrease in equity) and it cannot issue much additional debt for a long time (central panel of the central row).

[14]We build this representative household version using the same parameter values than in , except for the labor market parameters, which disappear with the absence of household heterogeneity. We use, as well, the same value of σ that we estimated for the heterogeneous household model. We will return to this point below in the main text.

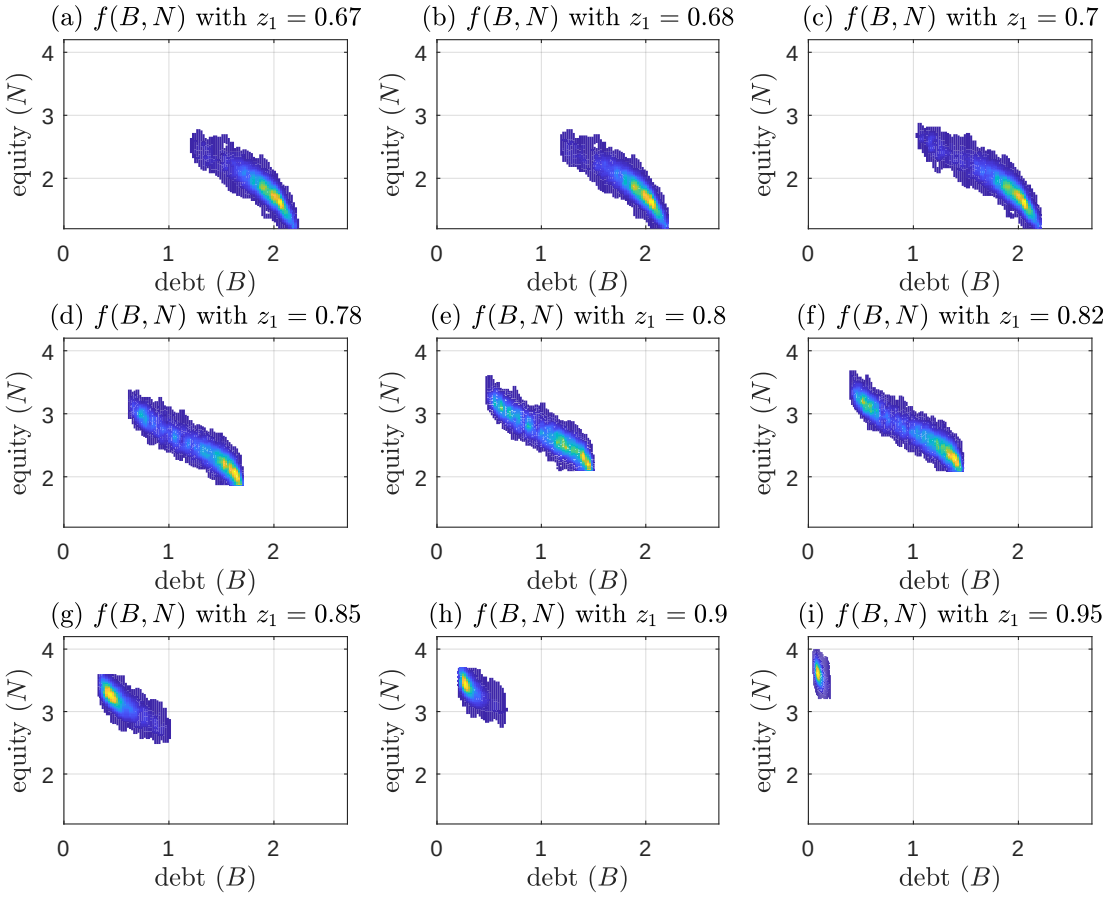


Figure 17: Ergodic distribution $F(B,N)$ as a function of z_1. Lighter colors indicate higher probability.

Figure 18 also plot the GIRFs of the model when $z_1 = 0.98$. While this is not exactly the representative household version of the model (which would required a different computational approach), it is extremely close: the household only faces a 2% reduction on wages with a low probability. Thus, we call this version of the model a quasi-representative household.

7.4 Differences across households

How can we understand the differences in GIRFs from the previous subsection? The key is to explore how the reduction in households' total income and consumption (see the central panel in the top row of Figure 18) is distributed among households.

Recall that, after the negative aggregate shock, wages decrease and the risk-free rate increases. Poorer households, which are mainly dependent on their labor earnings, loose much. In comparison, richer households, which depend much more on income from their assets, loose less or, if they are sufficiently wealthy, can even gain. Even more, these rich households can

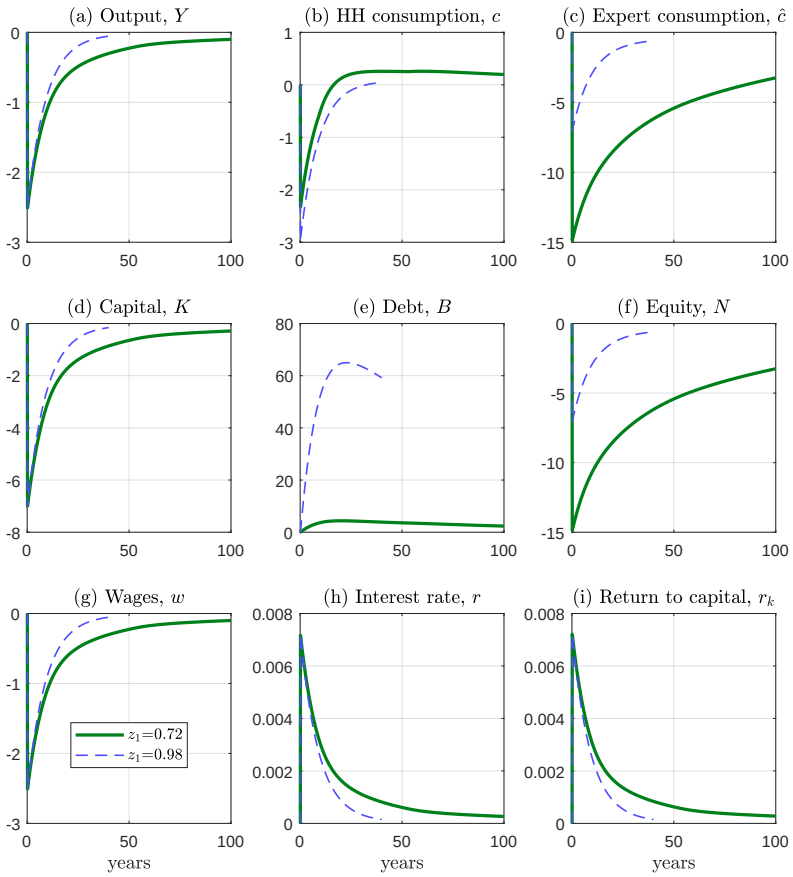


Figure 18: GIRFS, heterogeneous households vs. (quasi-)representative households.

take advantage of a higher risk-free rate to accumulate wealth fast. They reduce their consumption not as much as a consequence of lower income (wealth effect), but as a consequence of better rewards for savings (substitution effect). Moreover, the sluggish aggregate dynamics of the model means that wealthy households expect the higher risk-free rate to last for a long time. The integral of these consumption responses across households is different, because of the marginal decreasing utility of consumption, than the response of a household with average wealth. Hence, the responses of total consumption and the evolution of total debt are also different.

The heterogeneity of consumption effects is documented by Figure 19, which displays the consumption decision rules for high-productivity households ($z = z_2$) along the asset axis (results for low-income households are qualitatively similar) one-standard-deviation negative capital shock that hits the economy at the baseline SSS (continuous green line). To facilitate interpretation, we plot the *difference* in consumption decision rules with respect to the case without the shock, instead of the level. Furthermore, we plot the difference in consumption decision rules at different points in time after the shock: at impact (left panel at the top row);

after 5 years (right panel at the top row); after 10 years (left panel at the bottom row); and after 20 years (right panel at the bottom row)

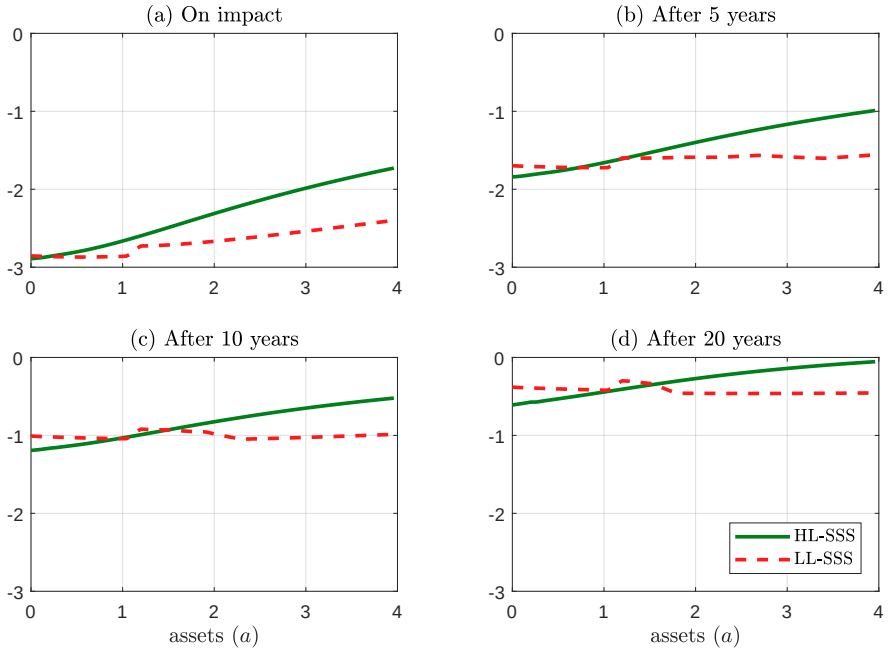


Figure 19: Difference in consumption decision rules at different points in time after the shock.

At impact, all households reduce their consumption, but poorer households do so by a larger amount, reflecting their lower income. However, richer households (level of assets of 4) reduce their consumption around one-third less than poor households (level of assets of 0). The same difference in consumption reduction survives over time.

The asymmetry in the consumption response is much smaller when we have a one-standard-deviation negative capital shock at the low-leverage SSS. Since in this case even the relatively rich households want to accumulate more debt for consumption smoothing in the future, they also reduce their consumption. Consequently, the expert can issue additional debt and the economy recovers fast.

The asymmetric consumption responses have a direct impact on how the wealth distribution evolves over time. To illustrate this point, we plot, in Figure (20), what we call the *distributional impulse response functions* or DIRFs. A DIRF is the natural analog of a GIRF except that, instead of plotting the evolution of an aggregate variable such as output or wages, we plot the evolution of the wealth density $g_t(\cdot)$. More concretely, we the difference between the density after the shock and that in the transition path. Time, in years, is plotted in the z-axis, assets in the x-axis, and the DIFRs in the y-axis. A positive value of the DIRF at a given asset level and point of time should be read as the density being higher at that asset level and point of

time than it would have been in the absence of a shock. A negative value has the opposite interpretation.

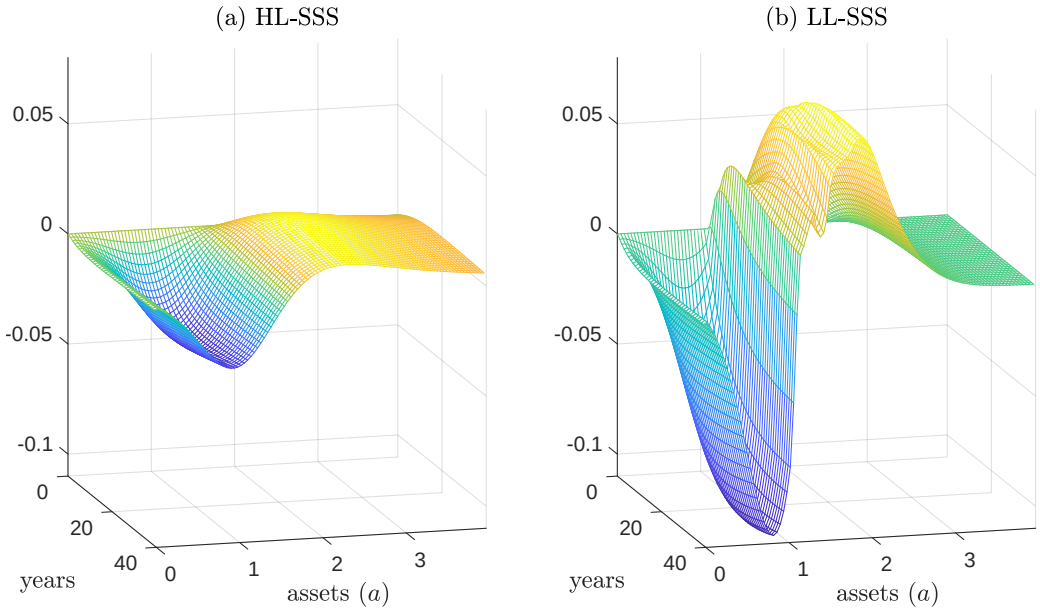


Figure 20: Distributional impulse response at the SSS.

In the left panel of Figure (20), we plot the DIRF to a one-standard-deviation negative capital shock when the economy is at the baseline SSS. In the right panel, we plot the DIRF to a one-standard-deviation negative capital shock when the economy is at the low-leverage SSS.

In the left panel, we can see how households with low assets must draw from their wealth to smooth consumption (even if consumption still drops) to compensate for lower income. This mechanism makes the DIRF negative in that region. In comparison, households with higher assets are reducing their consumption to respond to a temporarily higher risk-free rate and accumulate wealth. Thus, the DIRF is positive in the region of high assets. These effects are even more pronounced in the right panel. In the low-level SSS, poor households have too little debt to smooth consumption, and wealthy households have a sizeable latent valuation of additional debt when the risk-free changes.

In summary: heterogeneity matters in this economy. It makes the responses of aggregate variables entirely different. The mechanism is through the heterogeneity in the consumption response of poor vs. wealthy households and the disparate impact of the wealth and substitution effect when wages fall and the risk-free rate increases.

8 Conclusion

In this paper, we have postulated, computed, and estimated a continuous-time model of financial frictions *à la* Basak and Cuoco (1998) and Brunnermeier and Sannikov (2014), but with a non-trivial distribution of wealth among households. Such exercise has allowed us to uncover critical non-linear features of the model, such as the presence of multiple SSS(s) and the strong state-dependence of the GIRFs and to document the importance of household heterogeneity in the presence of financial frictions.

While we stand by the importance *per se* of our quantitative results, this paper can also be understood as a "proof of concept" of how to efficiently compute and estimate a continuous-time model with heterogeneous agents. For the computation, we have exploited tools borrowed from machine learning. For the estimation, we have built on contributions from inference with diffusions. The importance of heterogeneity for many questions in macroeconomics suggests that there are multiple potential applications for the methodological tools that we presented and we hope to show, in the close future, how to extend our current results in richer environments.

References

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2017). Income and wealth distribution in macroeconomics: A continuous-time approach. Working Paper 23732, National Bureau of Economic Research.

Adrian, T. and Boyarchenko, N. (2012). Intermediary Leverage Cycles and Financial Stability. Staff Reports 567, Federal Reserve Bank of New York.

Adrian, T. and Shin, H. S. (2010). Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437.

Ahn, S., Kaplan, G., Moll, B., Winberry, T., and Wolf, C. (2017). When inequality matters for macro and macro matters for inequality. In *NBER Macroeconomics Annual 2017, Vol. 32*. University of Chicago Press.

Aït-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica*, 70(1):223–262.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684.

Algan, Y., Allais, O., and Den Haan, W. J. (2008). Solving Heterogeneous-Agent Models with Parameterized Cross-Sectional Distributions. *Journal of Economic Dynamics and Control*, 32(3):875–908.

Algan, Y., Allais, O., Den Haan, W. J., and Rendahl, P. (2014). Solving and simulating models with heterogeneous agents and aggregate uncertainty. In Schmedders, K. and Judd, K. L., editors, *Handbook of Computational Economics*, volume 3, pages 277–324. Elsevier.

Alvaredo, F., Chancel, L., Piketty, T., Saez, E., and Zucman, G. (2017). Global inequality dynamics: New findings from wid.world. Working Paper 23119, National Bureau of Economic Research.

Andersen, T. G. and Lund, J. (1997). Estimating continuous-time stochastic volatility models of the short-term interest rate. *Journal of Econometrics*, 77(2):343–377.

Auclert, A. (2016). Monetary policy and the redistribution channel. Stanford University.

Bach, F. (2017). Breaking the curse of dimensionality with convex neural networks. *Journal Machine Learning Research*, 18(1):629–681.

Barron, A. R. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945.

Barucci, E. and Landi, L. (1995). Non-parametric versus linear learning devices: a procedural perspective. Technical report, University of Florence.

Basak, S. and Cuoco, D. (1998). An equilibrium model with restricted stock market participation. *Review of Financial Studies*, 11(2):309–341.

Bergin, J. and Bernhardt, D. (1992). Anonymous sequential games with aggregate uncertainty. *Journal of Mathematical Economics*, 21(6):543–562.

Bernanke, B. S., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, pages 1341–1393. Elsevier.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer-Verlag.

Blanchard, O. and Galí, J. (2010). Labor markets and monetary policy: A new keynesian model with unemployment. *American Economic Journal: Macroeconomics*, 2(2):1–30.

Brumm, J. and Scheidegger, S. (2017). Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica*, 85(5).

Brunnermeier, M. and Sannikov, Y. (2016). Macro, money, and finance: A continuous-time approach. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 1497–1545. Elsevier.

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

Candler, G. V. (1999). Finite difference methods for continuous time dynamic programming. In Marimón, R. and Scott, A., editors, *Computational Methods for the Study of Dynamic Economies*, pages 172–194. Oxford University Press.

Chacko, G. and Viceira, L. M. (2003). Spectral gmm estimation of continuous-time processes. *Journal of Econometrics*, 116(1):259–292.

Cho, I.-K. (1995). Perceptrons Play the Repeated Prisoner's Dilemma. *Journal of Economic Theory*, 67(1):266–284.

Cho, I.-K. and Sargent, T. J. (1996). Neural networks for encoding and adapting in dynamic economies. In Rust, H. A. D. K. J., editor, *Handbook of Computational Economics*, volume 1, pages 441–470. Elsevier.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314.

Den Haan, W. J. (1996). Heterogeneity, Aggregate Uncertainty, and the Short-Term Interest Rate. *Journal of Business & Economic Statistics*, 14(4):399–411.

Den Haan, W. J. (1997). Solving Dynamic Models with Aggregate Shocks and Heterogeneous Agents. *Macroeconomic Dynamics*, 1(02):355–386.

Den Haan, W. J. and Rendahl, P. (2010). Solving the Incomplete Markets Model with Aggregate Uncertainty Using Explicit Aggregation. *Journal of Economic Dynamics and Control*, 34(1):69–78.

Duarte, V. (2018). Sectoral reallocation and endogenous risk-aversion: Solving macro-finance models with machine learning. MIT Sloan School of Management.

Duffie, D. and Epstein, L. (1992). Stochastic differential utility. *Econometrica*, 60(2):353–94.

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.

Fernández-Villaverde, J. and Rubio-Ramírez, J. F. (2007). Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies*, 74(4):1059–1087.

Galí, J. and Gambetti, L. (2009). On the sources of the great moderation. *American Economic Journal: Macroeconomics*, 1(1):26–57.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. http://www.deeplearningbook.org.

Gornemann, N., Kuester, K., and Nakajima, M. (2012). Monetary policy with heterogeneous agents. Working Papers 12-21, Federal Reserve Bank of Philadelphia.

Hall, R. E. and Milgrom, P. R. (2008). The limited influence of unemployment on the wage bargain. *American Economic Review*, 98(4):1653–74.

He, Z. and Krishnamurthy, A. (2012). A model of capital and crises. *Review of Economic Studies*, 79(2):735–777.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–70.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kiyotaki, N. and Moore, J. (1997). Credit cycles. *Journal of Political Economy*, 105(2):211–248.

Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Kubler, F. and Scheidegger, S. (2018). Self-justified equilibria: Existence and computation. Working paper, DBF, University of Zurich.

Liu, P., Theodoridis, K., Mumtaz, H., and Zanetti, F. (2018). Changing macroeconomic dynamics at the zero lower bound. *Journal of Business & Economic Statistics*, 0(0):1–14.

Lo, A. (1988). Maximum likelihood estimation of generalized itô processes with discretely sampled data. *Econometric Theory*, 4(2):231–247.

Luetticke, R. (2015). Transmission of monetary policy and heterogeneity in household portfolios. UCL.

Maliar, L., Maliar, S., and Valli, F. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell-smith algorithm. *Journal of Economic Dynamics and Control*, 34(1):42–49.

McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–58.

Miao, J. (2006). Competitive equilibria of economies with a continuum of consumers and aggregate shocks. *Journal of Economic Theory*, 128(1):274–298.

Nuño, G. and Thomas, C. (2017). Bank leverage cycles. *American Economic Journal: Macroeconomics*, 9(2):32–72.

Piazzesi, M. and Schneider, M. (2016). Housing and macroeconomics. Working Paper 22354, National Bureau of Economic Research.

Pröhl, E. (2015). Approximating equilibria with ex-post heterogeneity and aggregate risk. Research Paper 17-63, Swiss Finance Institute.

Reiter, M. (2009). Solving Heterogeneous-Agent Models by Projection and Perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665.

Reiter, M. (2010). Solving the Incomplete Markets Model with Aggregate Uncertainty by Backward Induction. *Journal of Economic Dynamics and Control*, 34(1):28–35.

Ríos-Rull, J.-V. (1997). Computation of equilibria in heterogeneous agent models. Staff Report 231, Federal Reserve Bank of Minneapolis.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323:533–536.

Sager, E. (2014). Solving the incomplete markets model with aggregate uncertainty: The method of mixtures. Bureau of Labor Statistics, Price Research Division.

Salmon, M. (1995). Bounded rationality and learning: Procedural learning. In Kirman, A. P. and Salmon, M., editors, *Learning and Rationality in Economics*, pages 236–275. Blackwell Publishing.

Scheidegger, S. and Bilionis, I. (2017). Machine learning for high-dimensional dynamic stochastic economies. UCL.

Young, E. R. (2010). Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-Stochastic Simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41.

Appendix

A. Numerical algorithm

We describe the numerical algorithm used to jointly solve for the equilibrium value function, $v(a, z, B, N)$, the density $g(a, z, B, N)$ and the aggregate debt B and equity N. The algorithm proceeds in 3 steps. We describe each step in turn.

Step 1: Solution to the Hamilton-Jacobi-Bellman equation The HJB equation is solved using an *upwind finite difference* scheme similar to Candler (1999) and Achdou et al. (2017). It approximates the value function $v_i(a, B, N)$, $i = 1, 2$ on a finite grid with steps Δa, ΔB, $\Delta N : a \in \{a_1, ..., a_J\}, B \in \{B_1, ..., B_L\}, N \in \{N_1, ..., N_M\}$, where

$$\begin{aligned} a_j &= a_{j-1} + \Delta a = a_1 + (j-1)\Delta a, \ 2 \leq j \leq J, \\ B_l &= B_{l-1} + \Delta B = B_1 + (l-1)\Delta L, \ 2 \leq l \leq L, \\ N_m &= N_{m-1} + \Delta N = N_1 + (m-1)\Delta N, \ 2 \leq m \leq M. \end{aligned}$$

The lower bound in the wealth space is $a_1 = 0$, such that $\Delta a = a_J/(J-1)$. We use the notation $v_{i,j,l,m} \equiv v_i(a_j, B_l, N_m)$, and similarly for the policy function $c_{i,j,l,m}$. The derivatives are evaluated according to

$$\begin{aligned} \frac{\partial_i v(a_j, B_l, N_m)}{\partial a} &\approx \partial_f v_{i,j,l,m} \equiv \frac{v_{i,j+1,l,m} - v_{i,j,l,m}}{\Delta a}, \\ \frac{\partial_i v(a_j, B_l, N_m)}{\partial a} &\approx \partial_b v_{i,j,l,m} \equiv \frac{v_{i,j,l,m} - v_{i,j-1,l,m}}{\Delta a}. \\ \frac{\partial_i v(a_j, B_l, N_m)}{\partial B} &\approx \partial_B v_{i,j,l,m} \equiv \frac{v_{i,j,l+1,m} - v_{i,j,l,m}}{\Delta B}, \\ \frac{\partial_i v(a_j, B_l, N_m)}{\partial Z} &\approx \partial_N v_{i,j,l,m} \equiv \frac{v_{i,j,l,m+1} - v_{i,j,l,m}}{\Delta N}, \\ \frac{\partial_i^2 v(a_j, B_l, N_m)}{\partial N^2} &\approx \partial_{NN}^2 v_{i,j,l,m} \equiv \frac{v_{i,j,l,m+1} + v_{i,j,l,m-1} - 2v_{i,j,l,m}}{(\Delta N)^2}. \end{aligned}$$

Note that at each point of the grid, the first derivative with respect to a can be approximated with a forward (f) or a backward (b) approximation. In an upwind scheme, the choice of forward or backward derivative depends on the sign of the *drift function* for the state variable, given by

$$s_{i,j,l,m} \equiv w_{l,m} z_i + r_{l,m} a_j - c_{i,j,l,m}, \tag{40}$$

where

$$
\begin{aligned}
c_{i,j,l,m} &= \left[\frac{\partial v_{i,j,l,m}}{\partial a}\right]^{-1/\gamma}, && (41)\\
w_{l,m} &= (1-\alpha)\, Z\left(B_l + N_m\right)^{\alpha}, && (42)\\
r_{l,m} &= \alpha Z\left(B_l + N_m\right)^{\alpha-1} - \delta - \sigma^2 \frac{(B_l + N_m)}{N_m}. && (43)
\end{aligned}
$$

Let superscript n denote the iteration counter. The HJB equation is approximated by the following upwind scheme,

$$
\begin{aligned}
\frac{v^{n+1}_{i,j,l,m} - v^{n}_{i,j,l,m}}{\Delta} + \rho v^{n+1}_{i,j,l,m} &= \frac{(c^n_{i,j,l,m})^{1-\gamma}}{1-\gamma} + \partial_f v^{n+1}_{i,j,l,m} s^n_{i,j,l,m,f} \mathbf{1}_{s^n_{i,j,n,m,f}>0} + \partial_B v^{n+1}_{i,j,l,m} s^n_{i,j,l,m,b} \mathbf{1}_{s^n_{i,j,l,m,b}<0}\\
&\quad + \lambda_i \left(v^{n+1}_{-i,j,l,m} - v^{n+1}_{i,j,l,m}\right) + h_{l,m}\partial_B v_{i,j,l,m} + \mu^N_{l,m}\partial_N v_{i,j,l,m}\\
&\quad + \frac{\left[\sigma^N_{l,m}\right]^2}{2}\partial^2_{NN} v_{i,j,l,m}
\end{aligned}
$$

for $i = 1, 2$, $j = 1, ..., J$, $l = 1, .., L$, $m = 1, ..., M$, where where $\mathbf{1}(\cdot)$ is the indicator function and

$$
\begin{aligned}
h_{l,m} &\equiv h\left(B_l, N_m\right),\\
\mu^N_{l,m} &\equiv \mu^N\left(B_l, N_m\right) = \alpha Z\left(B_l + N_m\right)^{\alpha} - \delta\left(B_l + N_m\right) - r_{l,m} B_l - \widehat{\rho} N_m,\\
\sigma^N_{l,m} &\equiv \sigma^N\left(B_l, N_m\right) = \sigma\left(B_l + N_m\right),\\
s^n_{i,j,l,m,f} &= w_{l,m} z_i + r_{l,m} a_j - \left[\frac{1}{\partial^n_f v_{i,j,l,m}}\right]^{1/\gamma},\\
s^n_{i,j,l,m,b} &= w_{l,m} z_i + r_{l,m} a_j - \left[\frac{1}{\partial^n_b v_{i,j,l,m}}\right]^{1/\gamma}
\end{aligned}
$$

Therefore, when the drift is positive ($s^n_{i,j,l,m,f} > 0$), we employ a forward approximation of the derivative, $\partial^n_f v_{i,j,l,m}$; when it is negative ($s^n_{i,j,l,m,b} < 0$), we employ a backward approximation, $\partial^n_b v_{i,j,l,m}$. The term $\frac{v^{n+1}_{i,j,l,m} - v^n_{i,j,l,m}}{\Delta} \to 0$ as $v^{n+1}_{i,j,l,m} \to v^n_{i,j,l,m}$.

Moving all terms involving v^{n+1} to the left hand side and the rest to the right hand side, we obtain:

$$
\begin{aligned}
\frac{v^{n+1}_{i,j,l,m} - v^{n}_{i,j,l,m}}{\Delta} + \rho v^{n+1}_{i,j,l,m} &= \frac{(c^n_{i,j,n,m})^{1-\gamma} - 1}{1-\gamma} + v^{n+1}_{i,j-1,l,m}\alpha^n_{i,j,l,m} + v^{n+1}_{i,j,l,m}\beta^n_{i,j,l,m} + v^{n+1}_{i,j+1,l,m}\xi^n_{i,j,l,m}\\
&\quad + \lambda_i v^{n+1}_{-i,j,l,m} + v^{n+1}_{i,j,l+1,m}\frac{h_{l,m}}{\Delta B} + v^{n+1}_{i,j,l,m+1}\varkappa_{l,m} + v^{n+1}_{i,j,l,m-1}\varrho_{l,m} \qquad (44)
\end{aligned}
$$

where

$$\begin{aligned}
\alpha_{i,j}^{n} &\equiv -\frac{s_{i,j,B}^{n}\mathbf{1}_{s_{i,j,B}^{n}<0}}{\Delta a},\\
\beta_{i,j,l,m}^{n} &\equiv -\frac{s_{i,j,l,m,f}^{n}\mathbf{1}_{s_{i,j,n,mF}^{n}>0}}{\Delta a}+\frac{s_{i,j,l,m,b}^{n}\mathbf{1}_{s_{i,j,l,m,b}^{n}<0}}{\Delta a}-\lambda_{i}-\frac{h_{l,m}}{\Delta B}-\frac{\mu_{l,m}^{N}}{\Delta N}-\frac{\left(\sigma_{l,m}^{N}\right)^{2}}{\left(\Delta N\right)^{2}},\\
\xi_{i,j}^{n} &\equiv \frac{s_{i,j,F}^{n}\mathbf{1}_{s_{i,j,F}^{n}>0}}{\Delta a},\\
\varkappa_{l,m} &\equiv \frac{\mu_{l,m}^{N}}{\Delta N}+\frac{\left(\sigma_{l,m}^{N}\right)^{2}}{2\left(\Delta N\right)^{2}}=\frac{\left[\alpha Z\left(B_{l}+N_{m}\right)^{\alpha}-\delta\left(B_{l}+N_{m}\right)-r_{l,m}B_{l}-\widehat{\rho}N_{m}\right]}{\Delta N}+\frac{\sigma^{2}\left(B_{l}+N_{m}\right)^{2}}{2\left(\Delta N\right)^{2}},\\
\varrho_{l,m} &\equiv \frac{\left(\sigma_{l,m}^{N}\right)^{2}}{2\left(\Delta N\right)^{2}}=\frac{\sigma^{2}\left(B_{l}+N_{m}\right)^{2}}{2\left(\Delta N\right)^{2}}.
\end{aligned}$$

for $i=1,2$, $j=1,...,J$, $l=1,..,L$, $m=1,...,M$. We consider boundary state constraints in a ($s_{i,1,B}^{n}=s_{i,J,F}^{n}=0$). The boundary conditions in B and N are reflections.

In equation (44), the optimal consumption is set to

$$c_{i,j,n,m}^{n}=\left(\partial v_{i,j,l,m}^{n}\right)^{-1/\gamma}. \tag{45}$$

where

$$\partial v_{i,j,l,m}^{n}=\partial_{f}v_{i,j,l,m}^{n}\mathbf{1}_{s_{i,j,n,mF}^{n}>0}+\partial_{b}v_{i,j,l,m}^{n}\mathbf{1}_{s_{i,j,l,m,b}^{n}<0}+\partial\bar{v}_{i,j,l,m}^{n}\mathbf{1}_{s_{i,j,n,mF}^{n}\leq0}\mathbf{1}_{s_{i,j,l,m,b}^{n}\geq0}.$$

In the above expression, $\partial\bar{v}_{i,j,l,m}^{n}=(\bar{c}_{i,j,n,m}^{n})^{-\gamma}$ where $\bar{c}_{i,j,n,m}^{n}$ is the consumption level such that the drift is zero :

$$\bar{c}_{i,j}^{n}=w_{l,m}z_{i}+r_{l,m}a_{j}.$$

We define

$$\mathbf{A}^n_{l.m} = \begin{bmatrix} \beta^n_{1,1,l,m} & \xi^n_{1,1,l,m} & 0 & 0 & \cdots & 0 & \lambda_1 & 0 & \cdots & 0 \\ \alpha^n_{1,2,l,m} & \beta^n_{1,2,l,m} & \xi^n_{1,2,l,m} & 0 & \cdots & 0 & 0 & \lambda_1 & \ddots & 0 \\ 0 & \alpha^n_{1,3,l,m} & \beta^n_{1,3,l,m} & \xi^n_{1,3,l,m} & \cdots & 0 & 0 & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha^n_{1,J-1,l,m} & \beta^n_{1,J-1,l,m} & \xi^n_{1,J-1,l,m} & 0 & \cdots & \lambda_1 & 0 \\ 0 & 0 & \cdots & 0 & \alpha^n_{1,J,l,m} & \beta^n_{1,J,l,m} & 0 & 0 & \cdots & \lambda_1 \\ \lambda_2 & 0 & \cdots & 0 & 0 & 0 & \beta^n_{2,1,l,m} & \xi^n_{2,1,l,m} & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \lambda_2 & 0 & \cdots & \alpha^n_{2,J,l,m} & \beta^n_{2,J,l,m} \end{bmatrix},$$

$$\mathbf{v}^{n+1}_{l,m} = \begin{bmatrix} \mathbf{v}^{n+1}_{1,1,l,m} \\ \mathbf{v}^{n+1}_{1,2,l,m} \\ \vdots \\ \mathbf{v}^{n+1}_{1,J,l,m} \\ \mathbf{v}^{n+1}_{2,1,l,m} \\ \vdots \\ \mathbf{v}^{n+1}_{2,J,l,m} \end{bmatrix}$$

and

$$\mathbf{A}^n_m = \begin{bmatrix} \mathbf{A}^n_{1,m} & \frac{h_{1,m}}{\Delta K}\mathbf{I}_{2J} & \mathbf{0}_{2J} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} \\ \mathbf{0}_{2J} & \mathbf{A}^n_{2,m} & \frac{h_{2,m}}{\Delta K}\mathbf{I}_{2J} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} \\ \mathbf{0}_{2J} & \mathbf{0}_{2J} & \mathbf{A}^n_{3,m} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ & & & \mathbf{0}_{2J} & \mathbf{A}^n_{L-1,m} & \frac{h_{L-1,m}}{\Delta K}\mathbf{I}_{2J} \\ \mathbf{0}_{2J} & \mathbf{0}_{2J} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} & \left(\mathbf{A}^n_{L,m} + \frac{h_{L,m}}{\Delta K}\mathbf{I}_{2J}\right) \end{bmatrix}, \quad \mathbf{v}^{n+1}_m = \begin{bmatrix} \mathbf{v}^{n+1}_{1,m} \\ \mathbf{v}^{n+1}_{2,m} \\ \vdots \\ \mathbf{v}^{n+1}_{L,m} \end{bmatrix},$$

where $\mathbf{I}_n$ and $\mathbf{0}_n$ are the identity matrix and the zero matrix of dimension $n \times n$, respectively.

We can also define

$$\mathbf{A}^n = \begin{bmatrix} (\mathbf{A}_1^n + \mathbf{P}_1) & \mathbf{X}_1 & \mathbf{0}_{2J\times L} & \cdots & \mathbf{0}_{2J\times L} & \mathbf{0}_{2J\times L} \\ \mathbf{P}_2 & \mathbf{A}_2^n & \mathbf{X}_2 & \cdots & \mathbf{0}_{2J\times L} & \mathbf{0}_{2J\times L} \\ \mathbf{0}_{2J\times L} & \mathbf{P}_3 & \mathbf{A}_3^n & \cdots & \mathbf{0}_{2J\times L} & \mathbf{0}_{2J\times L} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ & & & \mathbf{P}_{M-1} & \mathbf{A}_{M-1}^n & \mathbf{X}_{M-1} \\ \mathbf{0}_{2J\times L} & \mathbf{0}_{2J\times L} & \cdots & \mathbf{0}_{2J\times L} & \mathbf{P}_M & (\mathbf{A}_M^n + \mathbf{X}_M) \end{bmatrix}, \ \mathbf{v}^{n+1} = \begin{bmatrix} \mathbf{v}_1^{n+1} \\ \mathbf{v}_2^{n+1} \\ \vdots \\ \mathbf{v}_M^{n+1} \end{bmatrix},$$

$$\mathbf{X}_m = \begin{bmatrix} \varkappa_{1,m}\mathbf{I}_{2J} & \mathbf{0}_{2J} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} \\ \mathbf{0}_{2J} & \varkappa_{2,m}\mathbf{I}_{2J} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ & & \mathbf{0}_{2J} & \varkappa_{L-1,m}\mathbf{I}_{2J} & \mathbf{0}_{2J} \\ \mathbf{0}_{2J} & \mathbf{0}_{2J} & \mathbf{0}_{2J} & \mathbf{0}_{2J} & \varkappa_{L,m}\mathbf{I}_{2J} \end{bmatrix},$$

$$\mathbf{P}_m = \begin{bmatrix} \varrho_{1,m}\mathbf{I}_{2J} & \mathbf{0}_{2J} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} \\ \mathbf{0}_{2J} & \varrho_{2,m}\mathbf{I}_{2J} & \cdots & \mathbf{0}_{2J} & \mathbf{0}_{2J} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ & & \mathbf{0}_{2J} & \varrho_{L-1,m}\mathbf{I}_{2J} & \mathbf{0}_{2J} \\ \mathbf{0}_{2J} & \mathbf{0}_{2J} & \mathbf{0}_{2J} & \mathbf{0}_{2J} & \varrho_{L,m}\mathbf{I}_{2J} \end{bmatrix}, \mathbf{u}^n = \begin{bmatrix} \frac{(c_{1,1,1,1}^n)^{1-\gamma}-1}{1-\gamma} \\ \frac{(c_{1,2,1,1}^n)^{1-\gamma}-1}{1-\gamma} \\ \vdots \\ \vdots \\ \frac{(c_{2,J,L,M}^n)^{1-\gamma}-1}{1-\gamma} \end{bmatrix}.$$

Then, equation (44) is a system of $2 \times J \times L \times M$ linear equations which can be written in matrix notation as:

$$\frac{1}{\Delta}\left(\mathbf{v}^{n+1} - \mathbf{v}^n\right) + \rho\mathbf{v}^{n+1} = \mathbf{u}^n + \mathbf{A}^n\mathbf{v}^{n+1}.$$

The system in turn can be written as

$$\mathbf{B}^n\mathbf{v}^{n+1} = \mathbf{d}^n \tag{46}$$

where , $\mathbf{B}^n = \left(\frac{1}{\Delta} + \rho\right)\mathbf{I} - \mathbf{A}^n$ and $\mathbf{d}^n = \mathbf{u}^n + \frac{1}{\Delta}\mathbf{v}^n$.

The algorithm to solve the HJB equation runs as follows. Begin with an initial guess $v_{i,j,l,m}^0$. Set $n = 0$. Then:

1. Compute $c_{i,j,l,m}^n$, $i = 1, 2$ using (45).
2. Find $v_{i,j,l,m}^{n+1}$ solving the linear system of equations (46).
3. If $v_{i,j,l,m}^{n+1}$ is close enough to $v_{i,j,l,m}^n$, stop. If not, set $n := n + 1$ and proceed to step 1.

Most programming languages, such as `Julia` or `Matlab`, include efficient routines to handle sparse matrices such as $\mathbf{A}^n$.

Step 2: Solution to the Kolmogorov Forward equation The income-wealth distribution conditional on the current realization of aggregate debt $B = B_l$ and equity $N = N_m$ can be characterized by the KF equation:

$$\begin{aligned} \frac{\partial g}{\partial t} &= -\frac{\partial}{\partial a}\left[s_i\left(a, B, N\right) g_{i,t}(a)\right] - \lambda_i g_{i,t}(a) + \lambda_{-i} g_{-i,t}(a),\ i = 1, 2. \quad (47) \\ 1 &= \int_0^\infty g(a) da. \quad (48) \end{aligned}$$

If we define a time step Δt we also solve this equation using an finite difference scheme. We use the notation $g_{i,j} \equiv g_i(a_j)$. The system can be now expressed as

$$\begin{aligned} \frac{g_{i,j,t+1} - g_{i,j}}{\Delta t} &= -\frac{g_{i,j,t} s_{i,j,l,m,b} \mathbf{1}_{s_{i,j,l,m,b}>0} - g_{i,j-1,t} s_{i,j-1,l,m,f} \mathbf{1}_{s_{i,j-1,l,m,f}>0}}{\Delta a} \\ &\quad - \frac{g_{i,j+1,t} s_{i,j+1,l,m,b} \mathbf{1}_{s_{i,j+1,l,m,b}<0} - g_{i,j,t} s_{i,j,l,m,b} \mathbf{1}_{s_{i,j,l,m,b}<0}}{\Delta a} - \lambda_i g_{i,j,t} + \lambda_{-i} g_{-i,j,t}, \end{aligned}$$

In this case, let us define

$$\mathbf{g}_t = \begin{bmatrix} g_{1,1,t} \\ g_{1,2,t} \\ \vdots \\ g_{1,J,t} \\ g_{2,1,t} \\ \vdots \\ g_{2,J,t} \end{bmatrix},$$

as the density conditional on the current state of B_l and N_m.We assume that g_0 is the density in the deterministic steady state (which coincides with the standard Aiyagari (1994) economy), the update in the next time period is given by the KF equation:

$$\mathbf{g}_{t+1} = \left(\mathbf{I} - \Delta t \mathbf{A}_{l,m}^{\mathbf{T}}\right)^{-1} \mathbf{g}_t,$$

where $\mathbf{A}_{l,m}^{\mathbf{T}}$ is the transpose matrix of $\mathbf{A}_{l,m} = \lim_{n\to\infty} \mathbf{A}_{l,m}^n$, defined above.

Complete algorithm

We can now summarize the complete algorithm. We begin a guess of the PLM $h^0\left(B, N\right)$. Set $s := 1$:

Step 1: Household problem. Given $h^{s-1}\left(B, N\right)$, solve the HJB equation to obtain an estimate of the value function $\mathbf{v}$ and of the matrix $\mathbf{A}$.

Step 2: Distribution. Given **A**, simulate T periods of the economy using the KF equation and obtain the aggregated ebt $\{B_t\}_{t=0}^{T}$ and equity $\{N_t\}_{t=0}^{T}$. The law of motion of equity is

$$N_t = N_{t-1} + \left[\alpha Z \left(B_t + N_t\right)^{\alpha} - \delta \left(B_t + N_t\right) - r_t B_t - \widehat{\rho} N_t\right] \Delta t + \sigma \left(B_t + N_t\right) \sqrt{\Delta t} \varepsilon_t,$$

where $\varepsilon_t \overset{iid}{\sim} N(0,1)$.

Step 3: PLM. Update the PLM using a neural netwok: h^s. If $\|h^s - h^{s-1}\| < \varepsilon$, where ε is a small positive constant, then stop. if not return to Step 1.

B. Backpropagation formulae

In our case, this results in:

$$\begin{aligned}
\frac{\partial \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)}{\partial \theta_0^2} &= h\left(\mathbf{s}_j; \theta\right) - \widehat{h}_j \\
\frac{\partial \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)}{\partial \theta_q^2} &= \left(h\left(\mathbf{s}_j; \theta\right) - \widehat{h}_j\right) \phi\left(\theta_{0,q}^1 + \sum_{i=1}^{2} \theta_{i,q}^1 s_j^i\right), \text{ for } q = 1, ..., Q \\
\frac{\partial \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)}{\partial \theta_{0,q}^1} &= \theta_q^2 \left(h\left(\mathbf{s}_j; \theta\right) - \widehat{h}_j\right) \phi'\left(\theta_{0,q}^1 + \sum_{i=1}^{2} \theta_{i,q}^1 s_j^i\right), \text{ for } q = 1, ..., Q \\
\frac{\partial \mathcal{E}\left(\theta; \mathbf{s_j}, \widehat{h}_j\right)}{\partial \theta_{i,q}^1} &= s_j^i \theta_q^2 \left(h\left(\mathbf{s}_j; \theta\right) - \widehat{h}_j\right) \phi'\left(\theta_{0,q}^1 + \sum_{i=1}^{2} \theta_{i,q}^1 s_j^i\right), \text{ for } i = 1, 2 \text{ and } q = 1, ..., Q,
\end{aligned}$$

where $\phi'(x) = \frac{1}{(1+e^{-x})}$.

B. Chebyshev polynomials

Here we show that the approximation computed with Chebyshev polynomials is not satisfactory.

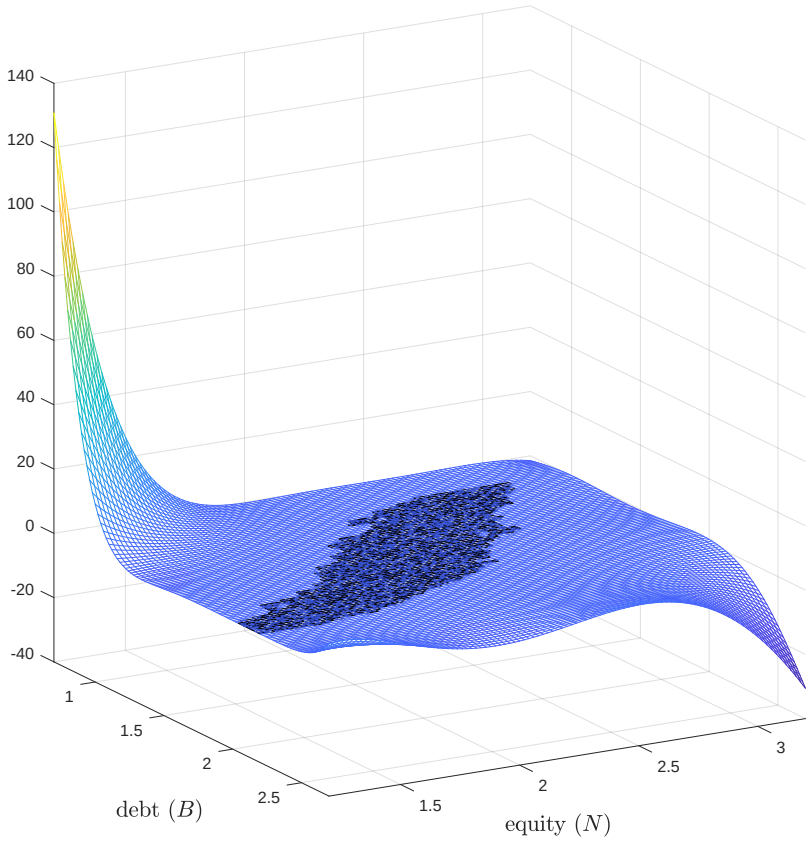


Figure 21: PLM with Chebyshev polynomials.

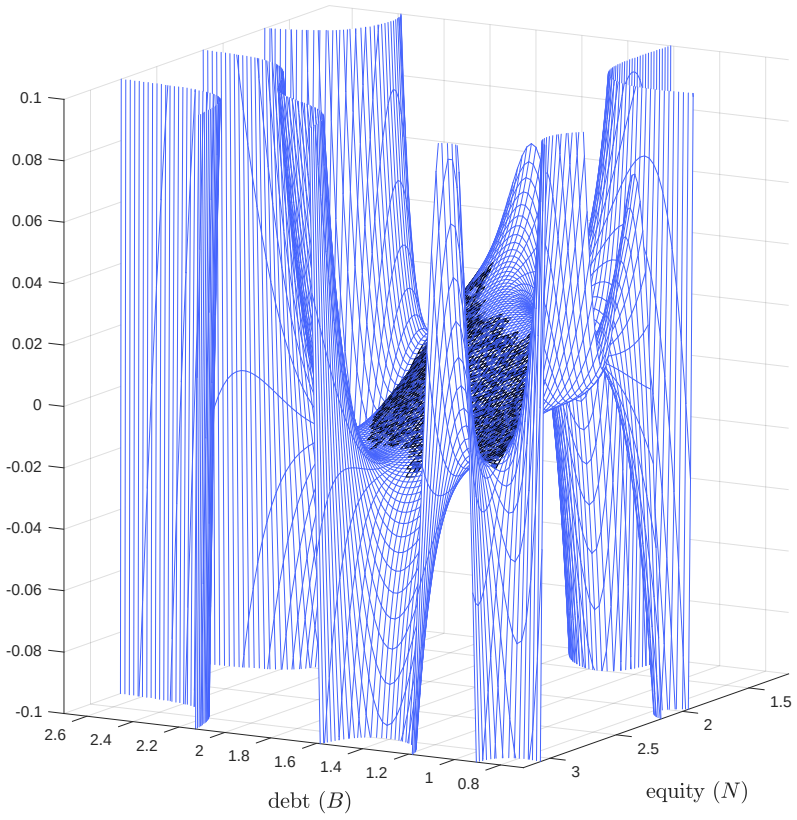


Figure 22: PLM with Chebyshev polynomials (zoom).