ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

https://www.econometricsociety.org/

Econometrica, Vol. 91, No. 3 (May, 2023), 869–901

FINANCIAL FRICTIONS AND THE WEALTH DISTRIBUTION

JESÚS FERNÁNDEZ-VILLAVERDE
Department of Economics, University of Pennsylvania, NBER, and CEPR

SAMUEL HURTADO
Banco de España

GALO NUÑO
Banco de España

Econometrica, Vol. 91, No. 3 (May, 2023), 869–901

FINANCIAL FRICTIONS AND THE WEALTH DISTRIBUTION

JESÚS FERNÁNDEZ-VILLAVERDE
Department of Economics, University of Pennsylvania, NBER, and CEPR

SAMUEL HURTADO
Banco de España

GALO NUÑO
Banco de España

We postulate a continuous-time heterogeneous agent model with a financial sector and households to study the nonlinear linkages between aggregate and financial variables. In our model, the interaction between the supply of bonds by the financial sector and the precautionary demand for bonds by households produces significant *endogenous aggregate risk*. This risk makes the economy transition between a high-leverage region and a low-leverage region, which, in turn, creates state dependence in impulse responses: the same shock starting from the high-leverage region gets propagated and amplified more than when the shock arrives when leverage is low. State dependence in impulse responses generates a time-varying aggregate precautionary savings motive that, by moving the risk-free rate, justifies the leverage level of the financial sector in each region. Finally, we illustrate the usefulness of neutral networks to solve for the nonlinear perceived law of motion of the model, and the importance of household heterogeneity in driving its quantitative properties.

KEYWORDS: Heterogeneous agents, wealth distribution, financial frictions, continuous time, neural networks, likelihood function.

1. INTRODUCTION

A GROWING BODY OF EMPIRICAL EVIDENCE has highlighted the nonlinear links between financial and aggregate variables. For example, Jordà, Schularick, and Taylor (2016) have gathered data from 17 advanced economies over 150 years to document how output growth, volatility, skewness, and tail events correlate with the levels of leverage in an economy. Similarly, Adrian, Boyarchenko, and Giannone (2019) have found how, in the United States, sharply negative output growth follows worsening financial conditions associated with leverage.

This paper shows how a continuous-time heterogeneous agent model can generate many of these nonlinear linkages. Our economy comprises a financial sector, modeled as a representative financial expert, and a continuum of households, subject to idiosyncratic labor productivity shocks. The economy is hit by aggregate shocks to the stock of capital and operates under two financial frictions. First, only the expert can hold the capital rented to a representative firm, and it cannot issue state-contingent assets (i.e., outside

Jesús Fernández-Villaverde: jesusfv@econ.upenn.edu
Samuel Hurtado: samuel.hurtado@bde.es
Galo Nuño: galo.nuno@bde.es

A `GitHub` repository with the codes for the paper and further examples can be found at https://github.com/jesusfv/financial-frictions. We thank three referees, Manuel Arellano, Alper Cenesiz, David Childers, Emmanuel Farhi, Xavier Gabaix, Lars Peter Hansen, Mark Gertler, Aubhik Khan, Davide Melcangi, Ben Moll, Chris Sims, Gianluca Violante, Ivan Werning, and participants at numerous seminars and conferences for comments. The views expressed in this manuscript are those of the authors and do not necessarily represent the views of the Eurosystem or the Bank of Spain.

 https://doi.org/10.3982/ECTA18180

equity) to finance it. Instead, the expert can issue a risk-free bond to leverage his equity and accumulate more capital; hence, he must absorb all the capital-return risk. Second, households can only save through the use of a risk-free bond to self-insure against idiosyncratic shocks and the wage and risk-free rate variations induced by aggregate fluctuations. The bond, in addition, cannot be shorted.

We include a financial sector in our model to capture the evolution of debt and leverage. We introduce household heterogeneity because, under the financial frictions we consider, such heterogeneity begets substantial endogenous aggregate risk. This aggregate risk derives from the interaction between the supply of bonds by the expert, determined by the excess return of capital over risk-free bonds, and the precautionary demand for bonds by households, driven by their income risk.

The endogenous aggregate risk induces an endogenous regime-switching process for output, the risk-free rate, excess returns, debt, and leverage. The economy transitions periodically between two different regions (technically, the basins of attraction of two different stochastic steady states, which we will define below). One region is characterized by high leverage and the other region is characterized by low leverage.

When the economy is in the high-leverage region (with a capital-over-equity ratio of around 2.1), the standard deviation of output is 14.2% higher than when the economy is in the low-leverage region (with a capital-over-equity ratio of around 1.4). This higher volatility is due to the higher persistence of the responses to the shock to capital. These changes in volatility account for 46% of the observed changes in volatility in U.S. output between the early 1990s and the mid-2000s reported in Fernández-Villaverde and Guerrón-Quintana (2020). Output is also less skewed when leverage is high than when it is low (0.0014 vs. 0.1186). Finally, the moves between high and low leverage resemble the debt supercycles (i.e., sustained periods of borrowing and deleveraging lasting for several decades) documented in Reinhart and Rogoff (2009) and Schularick and Taylor (2012).

We demonstrate the presence of endogenous aggregate risk by solving the model by obtaining a flexible, nonlinear approximation of the perceived law of motion (PLM) of aggregate debt as a function of the expert's equity and the cross-sectional distribution of households' bonds.[1]

First, we explain how we can approximate the PLM with high accuracy using just the first moment of the households' bond distribution—for example, aggregate debt—à la Krusell and Smith (1998) because the households' consumption decision rule is close to linear with respect to the household state variables (except for impoverished households). At the same time, we let the PLM be a nonlinear function of the aggregate states (given our assumption regarding the PLM: aggregate debt and equity) because households' consumption decision rule is sharply nonlinear with respect to them. Households understand the strongly state-dependent responses of wages and the risk-free rate with respect to the aggregate states and act accordingly.

In comparison, in Krusell and Smith (1998), the PLM of the aggregate variables is loglinear in the endogenous aggregate state variables.[2] We prove that a loglinear PLM is a poor choice in our model by computing a naive implementation of the Krusell and Smith (1998) algorithm. Fernández-Villaverde, Marbet, Nuño, and Rachedi (2022) found a similar result in a quite different heterogeneous agent environment (discrete time, nominal

[1] A `GitHub` repository with the codes for the paper and further examples of our methodology is available at https://github.com/jesusfv/financial-frictions.

[2] In Krusell and Smith (1998), the PLM is nonlinear in the exogenous states because the coefficients of the regression are allowed to vary across shocks. However, in our model, exogenous states are incorporated into the endogenous states instantaneously and this refinement is moot.

rigidities, etc.). Thus, we conjecture that this point applies to many heterogeneous agent economies: the PLM of the model can be well approximated by one or a few moments of the distribution and, yet, be fully nonlinear.

Second, we show how to employ neural networks to find the PLM as a nonlinear function of the aggregate state variables. Neural networks can capture nonlinearities without specifying, ex ante, any concrete structure for them. All linear models are alike, but every nonlinear model is nonlinear in its particular and unpredictable way. While other nonlinear approximation methods are possible, they often run into problems. For example, we show how a Chebyshev polynomial approximation delivers much worse accuracy than our approach.

Our approach generates a "bounded rationality" solution with the same pros and cons as the solution in Krusell and Smith (1998). On the positive side, since the forecast errors are small, this "bounded rationality" solution is a self-justified equilibrium (Kubler and Scheidegger (2018)). On the negative side, this equilibrium may or may not be close to the rational-expectations equilibrium, which we cannot compute. In particular, if we had more households clustered around their borrowing constraints, we would need to allow for more features of the wealth distribution beyond its first moment to enter into the PLM.

Also, given our PLM, we illustrate how a nonlinear model like ours can be taken to the data with a likelihood function employing aggregate and micro observations using inference with diffusions (Lo (1988)). The likelihood is computationally straightforward once we have solved the model with the approach outlined above: it simply amounts to transposing a matrix. Likelihood-based estimation is attractive in this context because methods of moments (either informal, such as calibration, or formal, such as the SMM) face difficulties when dealing with nonlinear DSGE models (Andreasen, Fernández-Villaverde, and Rubio-Ramírez (2018)).

Why does the nonlinear PLM create endogenous aggregate risk? For parameter values that match U.S. data and maximize the likelihood function, the curve within the PLM that implies zero changes in debt crosses the curve that implies zero changes in equity multiple times. Each of these crossings constitutes a stochastic steady state, or SSS(s): debt and equity accumulation are zero and, hence, so is capital accumulation. Interestingly, this occurs despite the model having a unique deterministic steady state (DSS).[3] In particular, we will have a high-leverage SSS (HL-SSS) and a low-leverage SSS (LL-SSS), each with its basin of attraction (there is a third, unstable SSS that we do not need to discuss).[4] Notice that the multiplicity of SSS(s) in our model is different from the multiplicity of equilibria: in our model, we find a unique equilibrium.

The intuition for the existence of two stable SSS(s) is as follows. In the basin of attraction of the HL-SSS, endogenous aggregate risk is high. After a negative aggregate shock, the economy suffers protracted recessions with persistently low wages (we explain below why). Due to precautionary behavior, households accumulate more bonds when risk is high. Higher savings have two consequences.

[3]An SSS, also known as a risky steady state (Coeurdacier, Rey, and Winant (2011)), 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. We care about SSSs because they are the points around which the ergodic distribution of the endogenous variables of the model is concentrated.

[4]The basin of attraction of an SSS is the set of starting points of debt and equity that converge to such an SSS in the absence of aggregate shocks.

First, higher savings increase wealth inequality among households: households with good realizations of their idiosyncratic shock accumulate much wealth. Second, higher savings lower the risk-free rate that clears the bond market and cause a high expected excess return for capital, pushing the expert to leverage aggressively, the defining feature of this basin of attraction. In comparison, aggregate risk is low in the basin of attraction of the LL-SSS. Thus, precautionary savings are lower, wealth inequality is smaller, the risk-free rate is high, and the expected excess return is reduced. These prices sustain the low leverage that characterizes this basin.

Why are recessions more severe in the basin of attraction of the HL-SSS than around the LL-SSS? When leverage is high, a negative aggregate shock to capital greatly erodes the expert's net wealth since his relatively small equity must absorb all the capital losses. Furthermore, in the aftermath of this negative shock, the expert is reluctant to take additional debt on too quickly: his leverage position is already quite risky. But since the expert does not issue too much debt, he also gets lower excess returns and accumulates equity sluggishly. Slow debt and equity accumulation lead to slow capital recovery. During this time, wages are persistently low and the risk-free rate is high.[5] By contrast, the recessions after a negative aggregate shock are mild when leverage is low. While the responses are similar on impact, the recovery is much faster because the expert experiences less reduction in his net wealth.

Whether the economy has high or low leverage is a consequence of past aggregate shocks. Sometimes, while the economy is traveling in the basin of attraction of the HL-SSS, a sequence of aggregate shocks will move it to the basin of attraction of the LL-SSS (and vice versa). For our baseline parameter values, the economy will spend more time, on average, in the basin of attraction of the HL-SSS than in the basin of attraction of the LL-SSS. This finding, however, varies as we change the volatility of idiosyncratic and aggregate shocks.

The transition between the two basins of attraction amounts to an endogenous regime-switching process for output, interest rates, debt, and leverage and generates a strong state-dependence in the impulse response functions of the model. This endogenous switching generates the time variation in the volatility and skewness of these variables and the long-lasting changes in the leverage and wealth inequality in the economy. Consider, for instance, the experience of the advanced economies between 1945 and today, as leverage was slowly rebuilt after World War II, pushing the economy toward higher levels of debt, wealth inequality, and financial fragility. Our model offers an intriguing interpretation of such a phenomenon.

Our findings are in stark contrast to the properties of the model when we move it toward its representative household version. As we reduce idiosyncratic risk (while keeping the volatility of the aggregate shock constant), the precautionary saving motive becomes smaller and the excess return is low. Eventually, the model encounters a bifurcation and the HL-SSS evaporates. Conversely, when we increase idiosyncratic risk, the LL-SSS vanishes. In this case, households are so concerned about their idiosyncratic risk that their demand for bonds propels the risk-free rate sufficiently low to allow only the existence of the HL-SSS (paradoxically, increasing the endogenous aggregate risk they face). Our findings also diverge from the version of the model where the expert can issue outside equity without constraints. In such a case, the expert and the households share the capital risk, and the large impact of leverage movements (and the associated nonlinearities) disappears.

[5] In our model, wages are equal to the marginal productivity of labor (lower when capital is low) and the risk-free interest rate depends on the marginal productivity of capital (higher when capital is low).

Thus, household heterogeneity allows us not only to generate nonlinear dynamics between aggregate variables and leverage, but also to think about ways in which (i) the rise in wealth inequality witnessed before the 2007–2009 financial crisis (Alvaredo, Chancel, Piketty, Saez, and Zucman (2017)); (ii) the increase in debt and leverage experienced during the same period (Adrian and Shin (2010), and Nuño and Thomas (2017)); and (iii) the low risk-free interest rates of the last two decades (Holston, Laubach, and Williams (2017)) can be linked. Furthermore, since household heterogeneity is at the core of our argument, changes in the forces behind precautionary savings affect the overall behavior of the economy. For example, an increase in idiosyncratic risk, such as the one documented for the United States since the late 1970s by Moffitt and Zhang (2018), translates in our model into higher macro volatility—even when the variance of aggregate shocks remains constant, a novel result in the literature.[6]

The papers closest to ours are Brunnermeier and Sannikov (2014) and He and Krishnamurthy (2012, 2013). Brunnermeier and Sannikov (2014) developed a model of financial frictions with *between-agents* heterogeneity: there is a representative expert and a representative household. Their model generates highly nonlinear amplification effects: sufficiently large shocks force the expert to sell his capital to the household, which operates it less efficiently. In comparison, the expert absorbs small shocks by adjusting payouts. Thus, as in our model, their economy transitions between regions of high and low leverage. He and Krishnamurthy (2012, 2013), also using a model of *between-agents* heterogeneity, focused on studying the dynamics of risk premia during crises in asset markets. There is an outside equity constraint. When the constraint is slack, aggregate risk aversion determines risk premia, but these increase sharply when the constraint is operative. Instead, the mechanism in our paper relies on *within-agents* heterogeneity: households never operate capital but have precautionary behavior due to aggregate and idiosyncratic shocks. With high leverage, households demand many bonds (justifying the high leverage in equilibrium), and, conversely, with low leverage, households demand few bonds. We see our paper as complementary to these early works.

Furthermore, Brunnermeier and Sannikov (2014) shared with our model a link between volatility and leverage, which they called the "paradox of volatility." In their model, 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, high volatility of idiosyncratic shocks leads to higher leverage and deeper recessions despite the volatility of aggregate shocks being constant.

Our paper also contributes to the literature on global solution methods for heterogeneous agent models with aggregate shocks. To the best of our knowledge, we are the first to generalize the celebrated algorithm of Krusell and Smith (1998) to accommodate a universal nonlinear law of motion in the endogenous state variables. We are also close to Ahn, Kaplan, Moll, Winberry, and Wolf (2017), who introduced a related method to compute the solution to heterogeneous agents 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 nonlinear dynamics posed by our paper. Similarly, Childers (2022) used linear projection methods (basis functions like histograms, splines, or polynomials) to solve discrete-time models locally. He described a setting in which the

[6]In the data, the economy is buffeted by many aggregate shocks. Some of them, such as monetary policy shocks, have become less volatile since the 1980s, thus lowering the overall volatility of the economy despite higher micro turbulence. Our claim is a comparative statics statement *with respect to* capital shocks.

model equilibrium conditions can be represented as a neural network (and so can be analyzed and implemented using neural network theory), but the paper does not solve for these equilibrium conditions.

Finally, our paper is related to the literature on the application of machine learning to compute dynamic models. The proposed methods have so far been concerned with the solution of high-dimensional dynamic programming problems. Examples include Scheidegger and Bilionis (2019), Duarte (2018), Maliar, Maliar, and Winant (2019), Azinović, Gaegauf, and Scheidegger (2022), and Ebrahimi Kahou, Fernández-Villaverde, Perla, and Sood (2021). Instead, our algorithm provides a nonlinear forecast of aggregate variables within the model.

While many nonlinear solution schemes are possible, our neural network approach is convenient, regarding both theoretical properties and practical considerations. First, the *universal approximation theorem* (Hornik, Stinchcombe, and White (1989), Bach (2017)) states that a neural network can approximate any unknown Borel measurable function. Second, the neural network can be efficiently trained using a combination of the gradient descent and the backpropagation algorithms. Third, our algorithm is easy to code and readily amenable to massive parallelization, and the computed solution generalizes well to simulation paths outside the aggregate ergodic distribution.

The Supplemental Material (Fernández-Villaverde, Hurtado, and Nuño (2023)) includes several appendices providing further details regarding the computation and estimation of the paper. Throughout the text, we will refer to this Supplemental Material as needed.

2. MODEL

We postulate a continuous-time DSGE model with a representative firm, a representative financial expert, and a continuum of households. There is one risky asset, capital, and a risk-free one, noncontingent bonds. Only the expert can hold capital because he is the only agent with the knowledge of how to accumulate it (we can also think about the expert as depicting the financial sector). In contrast, households can lend to the expert at the risk-free rate but cannot hold capital themselves, as they lack the required skill to handle it. The expert cannot issue outside equity, but can partially finance his capital holdings by issuing bonds to households. While in our model there is no need to distinguish between the firm and the expert, separating these two agents clarifies the exposition. We introduce heterogeneity on the side of the households—but not among the financial experts or the firms—because this heterogeneity triggers significant *endogenous aggregate risk*, the key feature of our investigation.

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 $Y_t = K_t^{\alpha} L_t^{1-\alpha}$. Since input markets are competitive, wages, w_t, are equal to the marginal productivity of labor:

$$w_t = (1-\alpha)\frac{Y_t}{L_t}, \tag{1}$$

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

$$rc_t = \alpha \frac{Y_t}{K_t}. \tag{2}$$

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

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

where ι_t is the investment rate per unit of capital. The capital shock Z_t is a common assumption in the financial frictions literature because it impacts the equity of the expert directly, making the effects of the shock transparent (Brunnermeier and Sannikov (2014)). The shock is a stand-in for any exogenous change in the efficiency units of capital. For example, it can represent the consequences of a new environmental regulation that lowers the profitability of existing capital or a technological innovation that renders already-installed capital economically obsolete, as in Greenwood, Hercowitz, and Krusell (1997).

We define the rental rate of capital rc_t over the capital contracted, K_t, and not over the capital returned after depreciation and the growth rate shock. Thus, the instantaneous return rate on capital is $dr_t^k = (rc_t - \delta)\,dt + \sigma\,dZ_t$.

2.2. *The Expert*

The representative expert holds capital $\widehat{K}_t$ and rents it to the firm (we denote variables related to the expert with a caret). To finance $\widehat{K}_t$, the expert issues risk-free debt $\widehat{B}_t$ at rate r_t to households. The financial friction in the model is 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 his assets (capital) and liabilities (debt), $\widehat{N}_t = \widehat{K}_t - \widehat{B}_t$. We allow $\widehat{N}_t$ to be negative, although this will not occur along the equilibrium path.

If $\widehat{C}_t$ is the consumption of the expert, $\widehat{N}_t$ evolves as

$$d\widehat{N}_t = \widehat{K}_t\,dr_t^k - \widehat{B}_t r_t\,dt - \widehat{C}_t\,dt = \big[\big(r_t + \widehat{\omega}_t(rc_t - \delta - r_t)\big)\widehat{N}_t - \widehat{C}_t\big]\,dt + \sigma\widehat{\omega}_t\widehat{N}_t\,dZ_t, \tag{4}$$

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 expected instantaneous 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 $\sigma\widehat{\omega}_t\widehat{N}_t$ is the risk of holding capital induced by the capital growth rate shock. Equation (4) allows us to derive the law of motion for $\widehat{K}_t$:

$$d\widehat{K}_t = d\widehat{N}_t + d\widehat{B}_t = \big[\big(r_t + \widehat{\omega}_t(rc_t - \delta - r_t)\big)\widehat{N}_t - \widehat{C}_t\big]\,dt + \sigma\widehat{\omega}_t\widehat{N}_t\,dZ_t + d\widehat{B}_t.$$

The expert's preferences over $\widehat{C}_t$ can be represented 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 his discount rate. The log felicity function will make our derivations below easier, but it could be generalized to recursive preferences à la Duffie and Epstein (1992).

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

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

subject to the evolution of N_t (4), an initial 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 $G_t(a, z)$. We will drop the subindex m when no ambiguity occurs to save on notation.

Each household supplies z_t units of labor valued at wage w_t. As in Huggett (1993), we assume that idiosyncratic labor productivity follows 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. When we face the model with the data, we will identify state 1 with unemployment (where z_1 is the value of leisure and home production) and state 2 with working.

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

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

where $s_i(a_t, G_t, N_t)$ is the drift of wealth for a household with labor supply z_i. The aggregate state variables G_t and N_t enter into this drift because they determine wages and the risk-free rate (as we will see in equation (11), G_t gives us B_t, and B_t, with N_t, gives us K_t). These variables pin down the optimal choice of consumption, $c_{it} = c_i(a_t, G_t, N_t)$. Households cannot short bonds:

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

Households have a CRRA instantaneous felicity function $u(c_t) = \frac{c_t^{1-\gamma}-1}{1-\gamma}$ discounted at rate $\rho > 0$. We pick this functional form to allow households and the expert to have different risk aversions, but we could substitute more general recursive preferences for it.

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).

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:

$$B_t \equiv \int a dG_t(a, z) = \widehat{B}_t, \tag{11}$$

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

Second, the total amount of labor rented by the firm is equal to labor supplied, $L_t = \int z dG_t$. Since the ergodic mean of z is 1, we have that $L_t = 1$ and total payments to labor are w_t.

Define households' total consumption as $C_t = C(G_t, N_t) \equiv \sum_{i=1}^{2} \int c_i(a, G_t, N_t)\, dG_t(a, i)$. Then, we can write

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

which tells us that the evolution of aggregate debt is the total income of households $w_t + r_t B_t$ minus their total 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

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

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 $dK_t = (\iota_t - \delta)K_t\, dt + \sigma K_t\, dZ_t$. Equating this result to (13) and canceling terms, we get $\iota_t = \frac{Y_t - C_t - \widehat{C}_t}{K_t}$, that is, the investment rate is output less aggregate consumption divided by aggregate capital. Also, we get the direct result:

$$dK_t = dN_t + dB_t. \tag{14}$$

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

$$\frac{\partial g_{it}(a)}{\partial t} = -\frac{\partial}{\partial a}\big(s_i(a, G_t, N_t) g_{it}(a)\big) - \lambda_i g_{it}(a) + \lambda_j g_{jt}(a), \quad i \neq j = 1, 2. \tag{15}$$

This density evolves according to the households' optimal consumption-saving choices plus the jumps corresponding to households that circulate out of $(\lambda_i g_{it}(a))$ and into $(\lambda_j g_{jt}(a))$ the labor state i. Also, the density satisfies the normalization $\sum_{i=1}^{2} \int g_{it}(a)\, da = 1$.

3. EQUILIBRIUM

Appendix A in the Supplemental Material defines an equilibrium for this economy and lists the equilibrium conditions. To characterize this equilibrium, we proceed first with the expert's problem. The use of log-utility implies that the expert consumes a constant share $\widehat{\rho}$ of his net wealth and chooses a leverage ratio proportional to the difference between the expected return on capital and the risk-free rate:

$$\widehat{C}_t = \widehat{\rho} N_t, \tag{16}$$

$$\omega_t = \widehat{\omega}_t = \frac{1}{\sigma^2}(rc_t - \delta - r_t). \tag{17}$$

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

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

depends positively on the variance of the aggregate shock, σ^2, and the leverage of the economy $\frac{N_t+B_t}{N_t}$. The higher the volatility or the leverage ratio, the higher the excess return the expert requires to isolate households from dZ_t.

Third, we can use the equilibrium values of rc_t, L_t, and ω_t 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:

$$r_t = \alpha K_t^{\alpha-1} - \delta - \sigma^2 \frac{K_t}{N_t} = \alpha(N_t + B_t)^{\alpha-1} - \delta - \sigma^2 \frac{N_t + B_t}{N_t}. \tag{18}$$

This equation will play a key role in explaining our quantitative results. Notice that equations (16)–(18) depend only on N_t and B_t.

Fourth, we use equations (4), (16), and (17) to describe the evolution of N_t as

$$\begin{aligned} dN_t &= \left[rc_t - \delta - \widehat{\rho} + (\omega_t - 1)(rc_t - \delta - r_t)\right] N_t\, dt + \sigma(N_t + B_t)\, dZ_t \\ &= \left[rc_t - \delta - \widehat{\rho} + \sigma^2(\omega_t - 1)\omega_t\right] N_t\, dt + \sigma(N_t + B_t)\, dZ_t \\ &= \left[\alpha(N_t + B_t)^{\alpha-1} - \delta - \widehat{\rho} + \sigma^2\left(\frac{N_t + B_t}{N_t} - 1\right)\frac{N_t + B_t}{N_t}\right] N_t\, dt + \sigma(N_t + B_t)\, dZ_t \\ &= \mu^N(B_t, N_t)\, dt + \sigma^N(B_t, N_t)\, dZ_t. \end{aligned} \tag{19}$$

The first line of equation (19) shows that the drift is the capital return plus the levered-up excess return. The second line uses the portfolio choice embedded in equation (17) to relate excess return to leverage. The third line demonstrates the nonlinear dependence of dN_t on N_t, B_t, and dZ_t. The fourth line writes the equation compactly in terms of a drift $\mu^N(B_t, N_t)$ and a volatility $\sigma^N(B_t, N_t)$.

Fifth, we have from equation (12):

$$\begin{aligned} dB_t &= (w_t + r_t B_t - C_t)\, dt \\ &= \left((1-\alpha)(N_t + B_t)^{\alpha} + \left(\alpha(N_t + B_t)^{\alpha-1} - \delta - \sigma^2 \frac{N_t + B_t}{N_t}\right) B_t - C_t\right) dt \\ &= \mu^B(B_t, G_t, N_t)\, dt, \end{aligned} \tag{20}$$

where the last line comes from recalling that $C_t = C(G_t, N_t)$. Equation (20) has a drift but no diffusion term because the shocks to capital are absorbed by the expert's equity. Thus, if we know $\mu^B(B_t, G_t, N_t)$, we get dB_t. With dB_t, we can calculate dK_t and all the other endogenous variables of the model follow. Section 4 will explain how we approximate $\mu^B(B_t, N_t, G_t)$ numerically given that we do not have a closed-form expression for C_t, the only unknown variable in equation (20).

Next, we define the DSS of the model. At the DSS, there are no capital growth rate shocks (i.e., $\sigma = 0$), but we still have idiosyncratic household shocks. From equation (19), we get

$$dN_t = \mu^N(B_t, N_t) = \left(\alpha(N_t + B_t)^{\alpha-1} - \delta - \widehat{\rho}\right) N_t\, dt. \tag{21}$$

Since, in the DSS, $dN_t = 0$, we find $K_{\text{DSS}} = N_{\text{DSS}} + B_{\text{DSS}} = \left(\frac{\widehat{\rho}+\delta}{\alpha}\right)^{\frac{1}{\alpha-1}}$. Thus, from equation (18), we get

$$r_{\text{DSS}} = r^k_{\text{DSS}} = rc_{\text{DSS}} - \delta = \alpha K_{\text{DSS}}^{\alpha-1} - \delta = \widehat{\rho}. \tag{22}$$

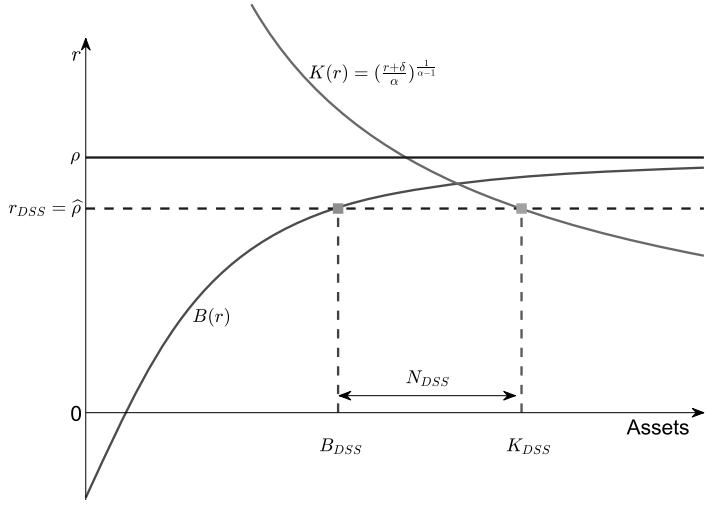

FIGURE 1.—DSS of the model.

This condition implies that we must have that $\widehat{\rho} < \rho$, that is, households must be less patient than the expert. Otherwise, households would want to accumulate savings without bounds to self-insure against idiosyncratic labor risk, and the DSS would not be well-defined.[7]

At r_{DSS}, there is demand for bonds by households, $B_{\text{DSS}} = \int a dG_{\text{DSS}}(a, z)$, a quantity we cannot compute analytically. We will impose parameter values in Section 5 that ensure that $B_{\text{DSS}} < K_{\text{DSS}}$ since the expert always wants to hold a positive net wealth ($N_{\text{DSS}} = K_{\text{DSS}} - B_{\text{DSS}}$) in equilibrium.

Figure 1 illustrates the DSS. The line $K(r) = (\frac{r+\delta}{\alpha})^{\frac{1}{\alpha-1}}$ is the capital demanded by the representative firm as a function of r. The crossing of the capital demand with $\widehat{\rho}$ determines r_{DSS} and K_{DSS}. The line $B(r)$ is the quantity of bonds demanded by households as a function of r. This function is monotonically increasing and asymptotes below ρ for the same reasons as in Aiyagari (1994). The crossing of the capital demand with $\widehat{\rho}$ pins down B_{DSS}. The difference between K_{DSS} and B_{DSS} gives us N_{DSS}.

In comparison, an SSS (stochastic steady state) is defined as a density $g_{\text{SSS}}(\cdot)$ and an equity N_{SSS} that remain invariant when the realization of the aggregate shock is zero. The difference between the SSS and the DSS is that, in the former, agents make their decisions taking into account aggregate risks ($\sigma > 0$), but no shock arrives along the equilibrium path, whereas, in the latter, $\sigma = 0$ and the agents are aware of it. While the DSS is unique in our model, we will find multiple SSSs in our quantitative exercise.

As mentioned in the Introduction, we care about these SSSs because they are the points around which the ergodic distribution of the variables of the model is concentrated (Schenk-Hoppé and Schmalfuss (2001)). For example, they are the points toward

[7]This property stands in contrast to models à la Bernanke, Gertler, and Gilchrist (1999), where borrowers are more impatient than lenders to prevent the former from accumulating enough wealth to render the financial friction inoperative. In these models, borrowers are infinitesimal and subject to idiosyncratic risk, and the lenders' discount rate determines the DSS risk-free rate. The situation is reversed in our model, where the households (lenders) must be more impatient than the expert (borrower).

which an economy returns in a generalized impulse response function (GIRF).[8] Hence, the SSSs are good summaries of many of the dynamic properties of the model.

4. SOLUTION

Our discussion of equation (20) highlighted the role of finding $\mu^B(B_t, N_t, G_t)$ to compute the equilibrium of the economy. In order to do so, we need to find the solution to the households' optimization problem (10) to get, by aggregating individual consumption choices, $C_t = \sum_{i=1}^{2} \int c_i(a, G_t, N_t)\, dG_t(a, i)$.

Using recursive notation, the households' $i \neq j = 1, 2$ problem can be expressed as a Hamilton–Jacobi–Bellman (HJB) equation:

$$\begin{aligned}\rho V_i(a, G, N) = \max_c \frac{c^{1-\gamma} - 1}{1-\gamma} + s_i(a, G, N)\frac{\partial V_i}{\partial a} + \lambda_i\big[V_j(a, G, N) - V_i(a, G, N)\big] \\ + \sum_{j=1}^{2} \int \frac{\delta V_i}{\delta G_j}(a, G, N; \widehat{a})\frac{\partial G_j}{\partial t}\, d\widehat{a} + \mu^N(B, N)\frac{\partial V_i}{\partial N} \\ + \frac{\big[\sigma^N(B, N)\big]^2}{2}\frac{\partial^2 V_i}{\partial N^2},\end{aligned} \tag{23}$$

where $\frac{\delta V_i}{\delta G_j}(\cdot; \widehat{a})$ is the Fréchet derivative of the value function V_i with respect to the distribution G. This is not a normal finite-dimensional derivative, as it reflects how the value function responds to changes in G, which is an infinite-dimensional object. Thus, equation (23) cannot be solved using standard tools for finite-dimensional partial differential equations.[9]

We tackle this challenge by making two assumptions. First, we follow Krusell and Smith (1998) and assume that, when forming their expectations, households only keep track of total debt, B_t, the first moment of G_t, instead of the entire distribution. As shown in Section 7.3, the first moment of G_t will numerically suffice because of the linearity of the household consumption policy with respect to individual wealth for most wealth levels (except those very close to the borrowing limit). Thus, instead of approximating $\mu^B(B_t, G_t, N_t)$, we approximate the much simpler conditional expectation of dB_t given (B_t, N_t), that is, $h(B_t, N_t) = \frac{\mathbb{E}[dB_t | B_t, N_t]}{dt}$. In other words, we assume that households consider a *perceived law of motion* (PLM) of B_t:

$$dB_t = h(B_t, N_t)\, dt, \tag{24}$$

instead of equation (20). We borrow the term PLM from the learning literature (Evans and Honkapohja (2001)) to accentuate that we allow $h(\cdot, \cdot)$ to be a general function.

[8]The word "generalized" in GIRF is used because, in comparison with linear models, impulse responses in nonlinear models are (i) state-dependent; (ii) a nonlinear function of the size of the shock; (iii) and nonsymmetric. Thus, we need to specify the size and sign of the shock and when this shock occurs.

[9]This difficulty does not emerge in the absence of aggregate risk, as the complete dynamics of G can be summarized by calendar time, and, hence, the problematic fourth term in equation (23) collapses to $\frac{\partial V}{\partial t}$. See Bilal (2021) for a derivation of the HJB in the presence of aggregate risk.

Given the PLM, the household's HJB (23) can be rewritten as

$$\rho V_i(a,B,N) = \max_c \frac{c^{1-\gamma}-1}{1-\gamma} + s_i(a,B,N)\frac{\partial V_i}{\partial a} + \lambda_i\big[V_j(a,B,N) - V_i(a,B,N)\big] + h(B,N)\frac{\partial V_i}{\partial B} + \mu^N(B,N)\frac{\partial V_i}{\partial N} + \frac{\big[\sigma^N(B,N)\big]^2}{2}\frac{\partial^2 V_i}{\partial N^2}, \tag{25}$$

where now both V_i and s_i depend on B instead of G. If we know $h(B_t, N_t)$, equation (25) can be solved using standard numerical techniques, such as the upwind finite differences scheme described in Appendix B in the Supplemental Material, or any other suitable algorithm (e.g., a deep neural network when the number of state variables is high).

Second, we approximate $h(\cdot,\cdot)$ by training a neural network on simulated data. Our algorithm will approximate the PLM arbitrarily well. This extra flexibility is key given the complex nonlinearities present in equations (19)–(20). By doing so, we will capture equilibrium dynamics missed if we approximated the PLM by a simple polynomial function (see Appendix G).

The ideas in the previous paragraphs are quite general. At the core of every heterogeneous agent model, we have laws of motion for aggregate variables dependent on the agents' distribution. These laws of motion can be approximated efficiently by neural networks. Appendix C presents a general algorithm to solve heterogeneous agent models using this approach. Our application in this paper is an instance of such a general algorithm.

An overview of the algorithm: Our algorithm to find $h(B,N)$ in (24) proceeds as follows:

(1) Start with h_0, an initial guess for h. Set $n := 0$.
(2) Using h_n, solve for household consumption, c, in the HJB equation (25).
(3) Construct a time series for B_t by simulating the cross-sectional distribution over time using the KF equation (15). Given B_t, we can find N_t and K_t using equations (19) and (14).
(4) Use a neural network to obtain $\tilde{h}$, a new guess for h.
(5) If $\|\tilde{h} - h_n\| < \epsilon$ given a norm $\|\cdot\|$ and tolerance level $\epsilon > 0$, stop. The solution is $\tilde{h}$. Otherwise, set $n := n+1$, make $h_n = \tilde{h}$, and return to step 2.

Steps 1–5 show that our solution has two differences with respect to the original Krusell–Smith algorithm: the use of continuous time and our employment of neural networks to update the guess of the PLM. Both differences deserve some explanation.

Simulating in continuous time: We use continuous time for convenience: it helps us to characterize much of the equilibrium dynamics analytically and to avoid dealing with expectations in the HJB equations (23) and (25) thanks to Itô's lemma. However, nothing essential in our analysis depends on it and Appendix C shows that the algorithm is easily adapted to discrete time without changing any important idea. See, for details, Achdou, Han, Lasry, Lions, and Moll (2021) and Nuño and Thomas (2016).

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 are discarded as a burn-in. Then, we have

$$B_{t_j} = \sum_{i=1}^{2} \int g_{it_j}(a; B_{t_j}, N_{t_j})\, da,$$

where the dynamics of the density are given by the KF equation (15).

Our simulation $(\mathbf{X}, \widehat{\mathbf{h}})$ is composed of a vector of inputs $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_J\}$, where $\mathbf{x}_j = \{x_j^1, x_j^2\} = \{B_{t_j}, N_{t_j}\}$ are samples of aggregate debt and the expert's net wealth at J random times $t_j \in [0, T]$, and a vector of outputs $\widehat{\mathbf{h}} = \{\widehat{h}_1, \widehat{h}_2 ..., \widehat{h}_J\}$, where $\widehat{h}_j \equiv \frac{B_{tj+\Delta t} - B_{t_j}}{\Delta t}$ 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.

Neural networks as universal nonlinear approximators: In Krusell and Smith (1998), the law of motion of the mean of capital is log-linear in capital, the endogenous variable, with state-dependent coefficients. This approximation is highly accurate due to the near log-linearity of their model in the vicinity of the DSS. Indeed, in their model, the DSS and SSS almost coincide. But, as shown in equations (19) and (20), this linearity does not extend to our model.

This nonlinearity causes two problems. First, we face the *approximation* problem: we need to search for an unknown nonlinear functional. Second, we need to fix the *generalization* problem. While the domain of B_t and N_t is unbounded, practical computation requires limiting it to a compact subset of $\mathbb{R}^2$ large enough to prevent boundary conditions from altering the solution in the subregion where the ergodic distribution accumulates. But this is such a large area that the simulation in step 3 of the algorithm above never visits much of it (even for extremely long simulations). Thus, the approximation algorithm should generalize in a "reasonable" fashion to the rest of the state space.[10]

To address these two problems, we employ neural networks to approximate the PLM. Our approach displays four strengths. First, the *universal approximation theorem* (Hornik, Stinchcombe, and White (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. The theorem applies to functions with nondifferentiabilities, kinks, and occasionally binding constraints.

Second, neural networks generalize outstandingly and are resilient to sparse initial information about the problem's solution. Neural networks have well-behaved shapes outside their training areas (the "reasonableness" we mentioned two paragraphs ago). In contrast, Chebyshev polynomials (or other series) often display explosive behaviors outside the fitted area. Indeed, Figure S3 in Appendix G shows the poor behavior of an approximation to the PLM in our model with Chebyshev polynomials.

Third, neural networks can be coded and trained using standard machine learning libraries. Furthermore, training a neural network is particularly well suited to massive parallelization in CPUs, GPUs, FPGAs, or in the dedicated AI accelerators such as TPUs specially designed for this class of problems (see, for FPGAs in economics, Cheela, DeHon, Fernández-Villaverde, and Peri (2022)).

Fourth (although this feature is not relevant for our model), neural networks can also tackle highly dimensional h_t's. Under technical conditions, neural networks break the "curse of dimensionality" (Barron (1993), and Bach (2017)). Researchers have found that this property holds very well in practice. For instance, Ebrahimi Kahou et al. (2021) have solved a model of industry dynamics with up to 10,000 state variables.

[10]The dynamics of the model in the generalized region might not coincide with the ones expected in the PLM. As in Krusell and Smith (1998), our approximation is a self-justified equilibrium in which households' beliefs about the PLM coincide with the actual law of motion (ALM) only in the equilibrium paths (Kubler and Scheidegger (2018)). Off-equilibrium, the PLM and the ALM may diverge, but households never discover it, as this region is never visited.

More concretely, a neural network with one hidden layer $h(\mathbf{x};\boldsymbol{\theta})$ is a linear combination of Q activation functions $\phi(\cdot)$:

$$h(\mathbf{x};\boldsymbol{\theta}) = \theta_0^2 + \sum_{q=1}^{Q} \theta_q^2 \phi\left(\theta_{0,q}^1 + \sum_{i=1}^{2} \theta_{i,q}^1 x^i\right), \tag{26}$$

where $\mathbf{x}$ is a two-dimensional input and $\boldsymbol{\theta} = (\theta_0^2, \theta_1^2, \ldots, \theta_{2,Q}^1)$ is a vector of weights. Thus, the neural network provides a flexible parametric function h that determines the growth rate of aggregate debt $\widehat{h}_j = h(\mathbf{x}_j;\boldsymbol{\theta})$, $j = 1, .., J$. The neural network (26) can be extended to include additional hidden layers. In that case, the network is called a *deep* neural network. However, for our problem, we checked that one layer was enough.

In terms of the network architecture, we pick a *softplus* activation function, $\phi(x) = \log(1 + e^x)$. This function has a simple sigmoid derivative, while keeping an efficient computation and gradient propagation. Through trial-and-error, we select a generous $Q = 16$ because the computational cost of a large hidden layer is small.[11]

The network weights θ are trained by minimizing a quadratic error function:

$$\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}} \mathcal{E}(\boldsymbol{\theta}; \boldsymbol{X}, \widehat{\mathbf{h}}) = \arg\max_{\boldsymbol{\theta}} \frac{1}{2} \sum_{j=1}^{J} \left\| h(\mathbf{x}_j;\boldsymbol{\theta}) - \widehat{h}_j \right\|^2$$

given a simulation $(\mathbf{X}, \widehat{\mathbf{h}})$. We solve this problem using the batch gradient descent algorithm. Appendix B describes the training of the network and how we handle local minima.

Why do neural networks approximate functions efficiently? Neural networks are a chain of geometric transformations. We take the vector of state variables, apply a linear transformation to it (in our case, $\theta_{0,q}^1 + \sum_{i=1}^{2} \theta_{i,q}^1 x^i$), and twist this transformation through an activation function ($\phi(\cdot)$). The process is repeated successively in each layer of the network. These geometric transformations move the functional approximation problem from lying in a difficult-to-characterize space to one living in an equivalent but much more convenient space.

A well-known example of what we mean by geometric transformations is as follows (Chollet (2021, p. 47)). Imagine that you have crumpled a sheet of paper into a ball. The paper ball is the function you want to approximate in the computer but doing so, with all its irregularities, is hard. That is why standard methods like splines or Chebyshev polynomials fail as soon as we deal with an irregular ball in four or five dimensions. The geometric transformations of the neural network expand the paper ball (the linear translation) and unfold it (the activation function) until it is a plain sheet of paper again. Approximating a sheet of paper (a linear plane) in a computer is straightforward.

5. TAKING THE MODEL TO THE DATA

We take our model to the data by mixing calibration and likelihood estimation. Table I summarizes the calibrated parameters. The capital share, α, is taken to be 0.35 and the depreciation rate of capital, δ, is 0.1 (all rates are annual). The discount rate, ρ, is set to 0.05. The risk aversion of the households γ is 2. These are standard values in the business cycle literature to match the investment-output ratio and the rate of return on capital.

[11] Overparameterizing the network with a Q too high is usually not a problem because of the "double descent phenomenon": even extremely overparameterized neural networks converge to a minimum-norm solution within the space of approximating functions. See Ebrahimi Kahou et al. (2021) for details.

TABLE I

BASELINE PARAMETERIZATION.

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%
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	$K/N=2$

The idiosyncratic income process is 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% and (ii) the job finding rate is 0.3 at a monthly frequency or $\lambda_1 = 0.986$ at an annual frequency. These numbers describe the 'U.S.' labor market calibration in Blanchard and Galí (2010).[12] We normalize average income $\bar{z} = \frac{\lambda_2}{\lambda_1+\lambda_2} z_1 + \frac{\lambda_1}{\lambda_1+\lambda_2} z_2$ to 1. We also set z_1 equal to 71% of z_2, as in Hall and Milgrom (2008). Both targets allow us to solve for z_1 and z_2. We set the expert's discount rate $\widehat{\rho}$ to ensure that the leverage ratio K/N in the DSS is nearly 2, which is roughly the average leverage from a Compustat sample of nonfinancial corporations.

Next, and keeping all the other parameters fixed at their calibrated quantities, we estimate σ—the key parameter in terms of the properties of the model—using maximum likelihood and U.S. quarterly observations for 1984.Q1–2017.Q4. We only estimate one parameter to make the exercise as transparent as possible. We start in 1984, as often done in the literature, to focus on the dynamics that have governed aggregate fluctuations in the U.S. during the last decades (Galí and Gambetti (2009)). We transform output using a bandpass filter output keeping frequencies between 20 and 60 quarters to eliminate long-run trends and to skip the business cycle frequencies caused by productivity and monetary policy shocks, which our model is not designed to account for. However, our methodology does not depend on this filtering, which is just done for convenience.

Our estimation illustrates an advantage of our solution method: it allows for a fast evaluation of the likelihood function implied by the model using ideas from inference with diffusions. Appendix D explains how to do so using aggregate data and Appendix E explores how to also add microeconomic observations from the cross-sectional distribution of assets.

We solve the model with the algorithm in Section 4. We use four Monte Carlo simulations of 2.5 million years, each at a monthly frequency. We initialize the model at the DSS, and we disregard the first 500 years as a burn-in. Then, we evaluate the likelihood for different values of σ. Our point estimate is 0.0142, with a standard error of 0.00011342. Appendix F provides details and a plot of the likelihood function.

[12] 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.

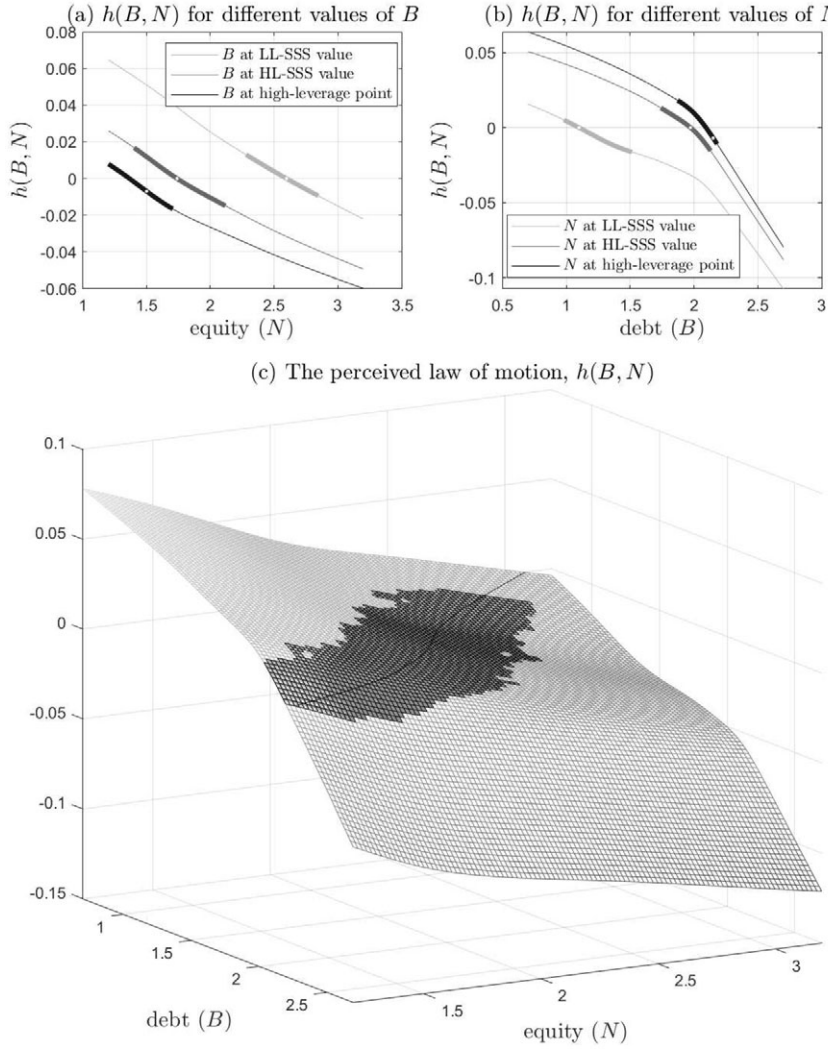

FIGURE 2.—The PLM $h(B, N)$ and transversal cuts.

6. QUANTITATIVE RESULTS

Figure 2 reports the PLM, $h(B, N)$, that we compute with our algorithm in Section 4 and the parameter values in Section 5. Panel (a), at the top left, displays three transversal cuts of $h(B, N)$ along with a range of values of equity (N). The first cut fixes B at the high-leverage SSS (HL-SSS) value ($B^{\mathrm{HL}} = 1.9641$, $N^{\mathrm{HL}} = 1.7470$, with $K^{\mathrm{HL}} = 3.7111$ and $\frac{K^{\mathrm{HL}}}{N^{\mathrm{HL}}} = 2.1243$), the second cut fixes B at an arbitrary 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 the low-leverage SSS (LL-SSS; $B^{\mathrm{LL}} = 1.0967$, $N^{\mathrm{LL}} = 2.6010$, with $K^{\mathrm{LL}} = 3.6977$ and $\frac{K^{\mathrm{LL}}}{N^{\mathrm{LL}}} = 1.4216$).

The thicker part of the lines indicates the regions of the state space in which the ergodic distribution of aggregate variables is nonzero. The white point indicates, in the first top cut, the LL-SSS; in the second cut, the HL-SSS; and in the third cut, the high-leverage point described above. Panel (b), at the top right, follows the same pattern as panel (a),

but switches the roles of equity (N) and debt (B). Finally, panel (c), at the bottom, shows the three-dimensional representation of the PLM. The shaded area highlights the region of the PLM visited in the ergodic distribution with nontrivial positive probability. The thin middle 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 nonlinearity of $h(B, N)$. The agents in our economy expect different growth rates of B_t in each region of the state space, with the function switching from concave to convex along the way. While this argument is clear from the shape of panel (c), it encodes rich dynamics. For instance, panel (b) shows how, as leverage increases, $h(B, N)$ becomes steeper 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 more capital risk. This result will resurface several times in future paragraphs. As intuition suggests, $h(B, N)$ generally decreases in debt and equity.

Figure 3 replicates Figure 2, except now for $\mu^N(B, N)$, the drift for equity. Similar comments regarding the nonlinear structure of the solution apply here. For example, now, $\mu^N(B, N)$ becomes higher as a function of equity as the level of leverage falls.

The nonlinearity of the PLM confirms that more traditional solution methods that rely on linear structures (conditional on aggregate shocks) might not be appropriate for this model. We check this argument by looking at the forecasting capability of our PLM and comparing it with the forecasting capability of an alternative PLM computed using the Krusell and Smith (1998) algorithm.[13]

Figure 4 plots the histogram of forecasting errors at a one-month horizon (the time step in the simulation), with a continuous line representing the errors from our algorithm and the discontinuous line the errors from a Krusell–Smith algorithm. The R^2 associated with the neural-network PLM 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. The Krusell–Smith algorithm produces an R^2 of 0.8275, considerably lower than typical values reported in the literature for standard heterogeneous agents models, and with an RMSE of 0.0021. Also, the forecasting errors are more volatile, skewed to the right, and without a mode at zero. We checked that adding more moments to the OLS regression does not help much in terms of accuracy.

But considering nonlinearities is not just about accuracy. More importantly, the model dynamics change dramatically if we do not consider them. Appendix G shows that, when we solve for the PLM using a Krusell–Smith algorithm, the multiplicity of SSSs at the core of our paper disappears. At the same time, when we solve the model using other nonlinear methods, such as splines, the multiplicity and location of SSS(s) remain unchanged.[14] Thus, our results arise from considering nonlinearities in the PLM, not using neural networks per se. Appendix G also discusses other alternative solution methods, such as Chebyshev polynomials, and argues that our method has advantages over them as well.

[13]More concretely, we approximate $dB_t = h(B_t, N_t)\,dt$ with a linear regression of $\frac{dB_t}{dt}$ on B_t and N_t. We checked that the results were nearly the same with a log-linear regression.

[14]Neural networks are not strictly needed to solve our model nonlinearly. We just argue that they are particularly convenient to do so.

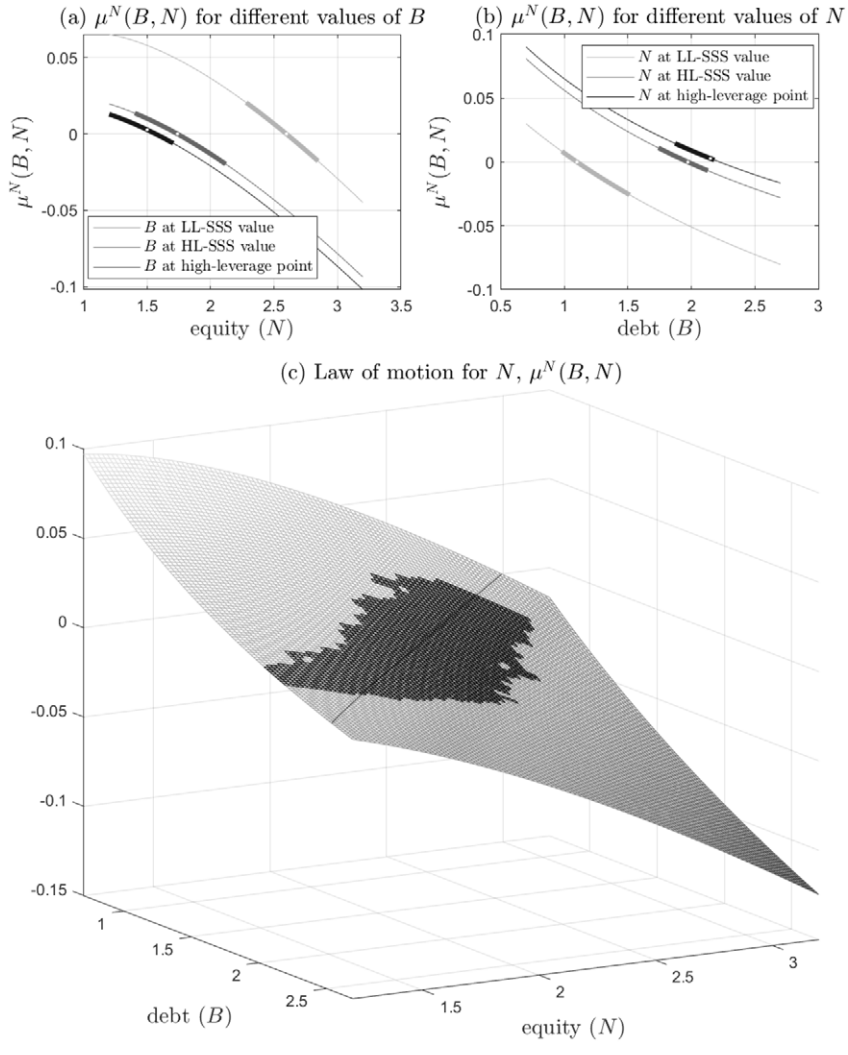

FIGURE 3.—The law of motion $\mu^N(B,N)$ and transversal cuts.

6.1. *The Phase Diagram*

Figure 5 plots the phase diagram of our model. The x-axis is aggregate debt, B. The y-axis is equity, N. The continuous line plots the pairs of B and N that make $h(B,N)=0$ in equation (24), that is, the combinations of debt and equity that imply a zero drift in the expert's debt. This line comes from the intersections of zero with $h(B,N)$ in Figure 2. The discontinuous line plots the combinations of debt and equity that make $\mu^N(B,N)=0$ in equation (19), that is, the combinations of debt and equity that imply a zero drift in the expert's net wealth. This second line comes from the intersections of zero with $\mu^N(B,N)$ in Figure 3. Both lines have a negative slope because the higher the debt, the lower the equity required to sum up the total capital that clears the demand of capital by the representative firm in equilibrium.

The intersections of the two lines—that is, the B and N such that $h(B,N)=\mu^N(B,N)=0$—define the SSS(s) of the model, that is, the combinations of debt and

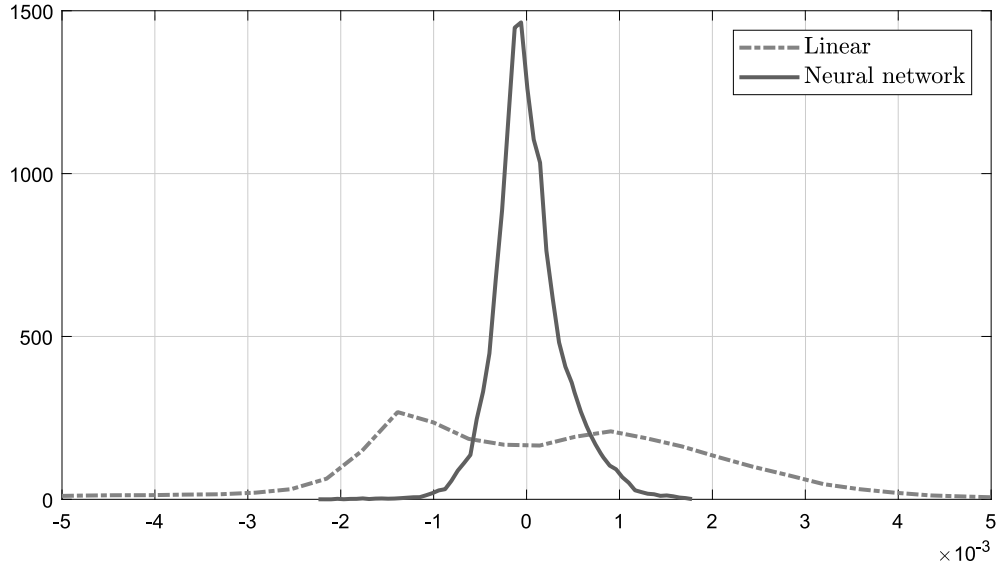

FIGURE 4.—Forecasting error distribution at a one-month horizon, linear PLM (left) and neural network (right).

equity that remain constant over time when the realization of the aggregate shock is zero. The arrows in Figure 5 indicate the movement of debt, B, and equity, N, when we are away from the two lines.

Given our parameter values, the two lines intersect three times, defining three SSS(s). From the bottom right, the first intersection is the stable HL-SSS ($B^{\text{HL}} = 1.9641$, $N^{\text{HL}} = 1.7470$, and $K^{\text{HL}} = 3.7111$). This SSS is the closest to the DSS (the square). In comparison with the DSS ($B = 1.8718$, $N = 1.8215$, and $K = 3.6933$), the HL-SSS has 0.5% more capital, 4.9% more debt, and 4.1% less equity. The second intersection is at a middle-leverage SSS with less debt and more equity ($B^{ML} = 1.3897$, $N^{ML} = 2.3108$, and $K^{ML} = 3.7005$). This SSS is, however, unstable, and the dynamics of the economy move away from it. Thus, we will not discuss it further. The final third intersection, at the top left, is the stable LL-SSS. Here, debt is much smaller ($B^{\text{LL}} = 1.0967$) and equity considerably higher ($N^{\text{LL}} = 2.6010$) than in the HL-SSS, yielding, however, a similar capital, $K^{\text{LL}} = 3.6977$. We also plot the point of high leverage that we use in Figures 2 and 3. Also, Appendix H shows the convergence to the SSS(s) for points visited in the ergodic distribution.

6.2. *Differences in Persistence at Each Stable SSS(s)*

To explain why we have multiple SSS(s), we must first understand the differences in the persistence of the dynamic responses to a capital shock in each basin of attraction. To do so, Figure 6 displays the GIRFs to a two-standard-deviations negative capital shock.

We plot three GIRFs: when the shock hits the economy at the HL-SSS (continuous lines); when the shock hits the economy at the LL-SSS (discontinuous dotted lines); and when the shock hits the economy at the arbitrary high-leverage point defined in Figure 5 (discontinuous lines). Time units are years.

In the three cases, the shock destroys the same amount of capital (panel d) and output falls, at impact, the same (panel a). However, the higher the leverage, the larger the reduction of equity at impact (panel f). To see why, rewrite the law of motion for the expert's

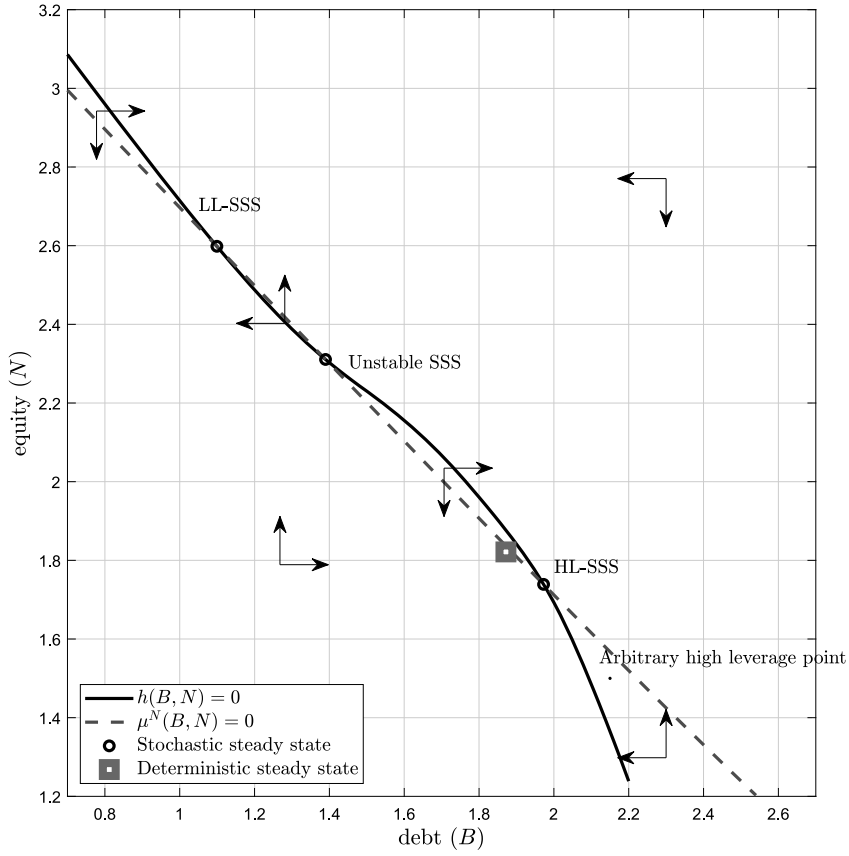

FIGURE 5.—Phase diagram, DSS, and SSS(s).

net wealth (19) as

$$\frac{dN_t}{N_t} = \left(\alpha K_t^{\alpha-1} - \delta - \widehat{\rho}\right) dt + \sigma^2 B_t \frac{K_t}{N_t^2}\, dt + \sigma \frac{K_t}{N_t}\, dZ_t. \tag{27}$$

The volatility term in equation (27) depends positively on leverage: the expert must absorb the capital losses from a negative shock using a small net wealth. At the same time, the recovery rate of the expert's net wealth is roughly independent of leverage: the first drift term, $(\alpha K_t^{\alpha-1} - \delta - \widehat{\rho})$, does not depend on leverage and the second drift term, $\sigma \frac{K_t}{N_t}\, dZ_t$, is negligible in size because it is multiplied by σ^2.

The expert also accumulates less debt after the negative shock when leverage is high (panel e). A negative shock to capital increases leverage above the level of the HL-SSS. While taking on more debt is attractive due to the excess return it delivers, it is also quite risky when leverage is so high. A possible future negative shock would be absorbed by a small buffer of net wealth, forcing the expert's consumption ($\widehat{C}_t = \widehat{\rho} N_t$) down too much. Thus, the expert only takes on more debt as he accumulates more equity. But since we just argued that accumulating equity is slow, the evolution of debt is also very persistent.

Interestingly, the GIRF for debt is hump-shaped. After many periods without an aggregate shock (recall, we are looking at a GIRF, where the realization of aggregate shocks

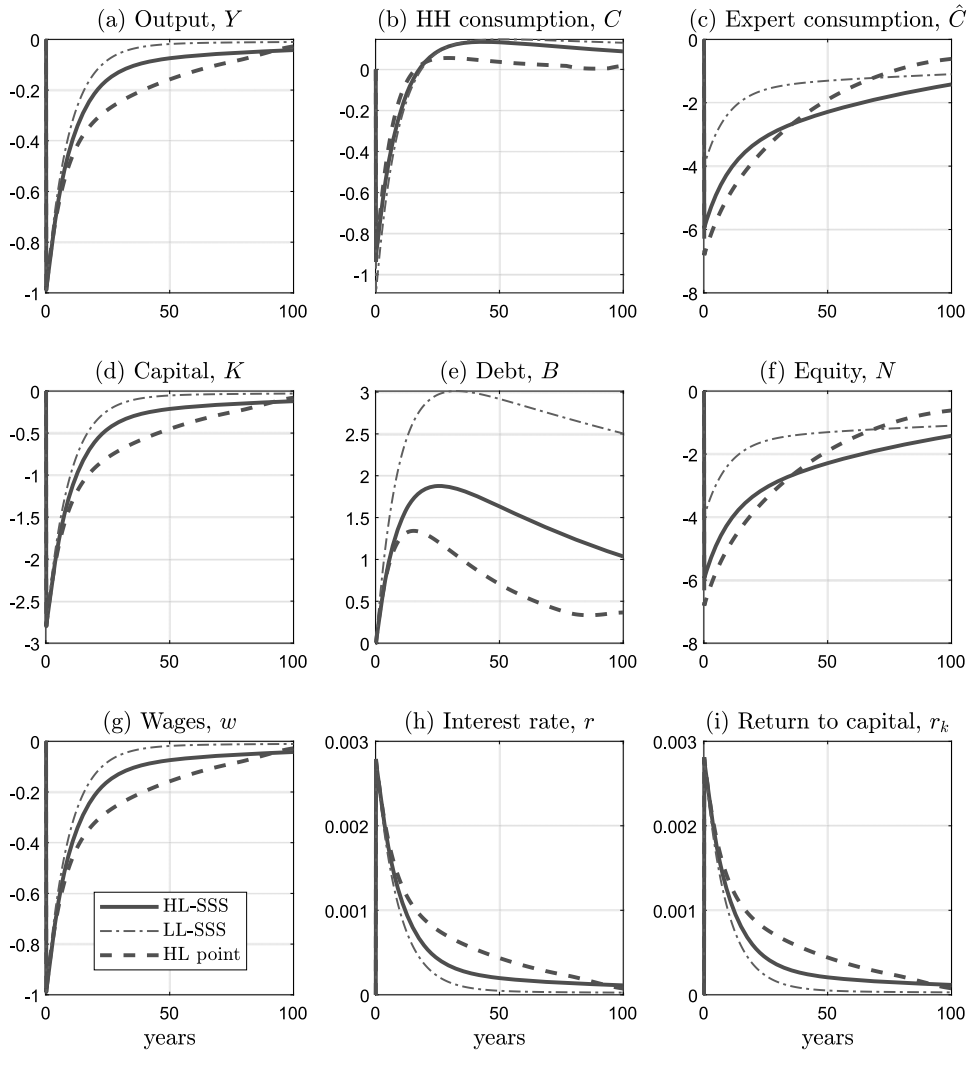

FIGURE 6.—GIRFs for different initial states.

is zero), the expert has accumulated enough capital so that we are very close to the SSS (panel d). But since the excess returns continue accruing, the expert can pay back his debt.

Combining the slow accumulation of debt when leverage is high with the finding that the growth rate of equity is roughly independent of leverage, we get that capital (panel d) and output (panel a) recover more slowly when leverage is high. This gives us persistently lower wages (panel g) and a higher risk-free rate (panel h). Section 7.2 explains how these changes affect the consumption of different households along the wealth distribution. Suffice it to say here that households experience a negative income effect (lower labor income), a positive income effect (higher bond income), and an intertemporal substitution effect (higher risk-free rate in their Euler equations). The net size of these opposite effects remains roughly unchanged as we vary the levels of leverage (panel b). For example, when leverage is high, the negative labor income effect is larger, but the positive risk-free rate income effect is bigger as well.

Figure 6 also explains why our economy diverges from a version of the model where the expert can issue outside equity without constraints. In such a case, the expert and households share the capital risk. Thus, after a negative capital shock, we would not observe the large leverage movements in panel (f) that drive the nonlinear responses of other aggregate variables.[15]

Three final points about the GIRFs deserve discussion. First, the GIRFs are highly persistent. When leverage is high, even after 40 years, the economy is still around half a percentage point below its pre-shock level. The dynamics of equity and debt accumulation propagate aggregate shocks in ways that are not present in models without financial frictions. Second, the state-dependence of the GIRFs induces time variation in the volatility of aggregate variables. Third, the two-standard-deviations shock is not large enough to send the economy away from the basins of attraction of each SSS. An even larger shock or a shock closer to the frontier between the two basins will induce a switch of basin, the endogenous regime-switching behavior we discussed in the Introduction. These larger shocks induce a sharp nonlinearity, as the GIRFs will converge to the SSS of the new basin, not their origin. We will return to this point in Section 6.4.

6.3. *Why Do We Have Two Stable SSS(s)?*

We are now ready to discuss why we have two stable SSS(s). The key is the evolution of wages (panel g) in Figure 6. Because of the arguments in the previous subsection, when leverage is high, wages fall for a much longer period than when leverage is low. In the HL-SSS, it takes 19.08 years for 75% of the effect on wages of a negative aggregate shock to dissipate (and 30 years at the high-leverage point). In comparison, it takes 13.33 years in the LL-SSS. Similarly, the standard deviation of wages when the economy fluctuates around the basin of attraction of the HL-SSS is 0.0125, but 0.0108 around the basin of attraction of the LL-SSS.

The forecast of a persistent fall in wages embedded in the PLM induces a stronger precautionary behavior by households. At the same time, the supply of debt given by equation (17) is

$$B_t = \left(1 + \frac{1}{\sigma^2}(rc_t - \delta - r_t)\right)N_t.$$

Thus, the debt market will clear with a reduction in r_t. At the HL-SSS, this reduction in r_t will also need to compensate for the small N_t (the HL-SSS is characterized by having low equity) and slightly lower rc_t (as total capital is a bit higher than at the LL-SSS).

Hence, we have shown why we have a fixed point at the HL-SSS. High leverage makes wages persistently lower after a negative capital shock. Persistently lower wages create a precautionary motive that lowers the risk-free rate and induces high leverage by the expert. We only need to reverse the argument to show why we also have a fixed point at the LL-SSS. Low leverage makes wages recover quickly after a negative capital shock. The lower associated precautionary behavior means the risk-free rate is relatively high and that sustains the low leverage of the expert.

The high sensitivity of leverage to excess returns induced by $\frac{1}{\sigma^2}$ in equation (17) accounts for the large differences between the two SSS(s) in terms of how a roughly equal

[15]The importance of the financial friction in our model would still go through qualitatively if the expert could issue some outside equity, but subject to some cap due to "skin-in-the-game" constraints. Quantitatively, though, the effect of capital shocks could be smaller.

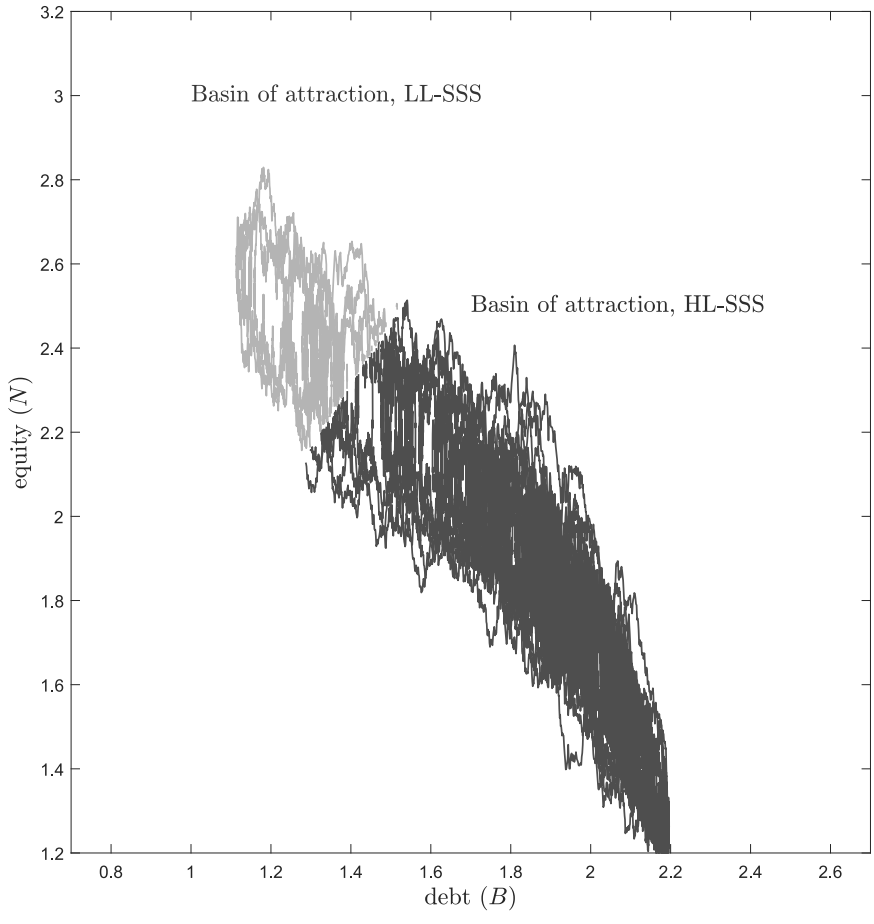

FIGURE 7.—Simulation of equilibrium paths.

total wealth ($K^{\text{HL}} = 3.7111$ vs. $K^{\text{LL}} = 3.6977$) is allocated between experts ($N^{\text{HL}} = 1.7470$ vs. $N^{\text{LL}} = 2.6010$) and households ($B^{\text{HL}} = 1.9641$ vs. $B^{\text{LL}} = 1.0967$). The excess return for the HL-SSS is 4.1636 basis points and, for the LL-SSS, 2.7864. This high sensitivity is the dual of the negligibility of the term multiplied by σ^2 in equation (27).

6.4. *Endogenous Aggregate Risk*

Figure 7 plots random equilibrium sequences for our economy corresponding to different sequences of aggregate shocks. In light gray, we plot the paths when they are in the basin of attraction of the LL-SSS. In darker gray, we plot the paths when they are in the basin of attraction of the HL-SSS. Along the equilibrium sequence, our economy has recurrent endogenous regime switches. A path in the basin of attraction of the LL-SSS can be pushed toward the basin of attraction of the HL-SSS by a series of shocks that reduce equity and increase debt. Conversely, a series of shocks that increase equity and lower

TABLE II

MOMENTS CONDITIONAL ON BASIN OF ATTRACTION.

	Mean	Standard deviation	Skewness	Kurtosis
$Y^{\text{basin HL}}$	1.5802	0.0193	0.0014	2.869
$Y^{\text{basin LL}}$	1.5829	0.0169	0.1186	3.0302
$r^{\text{basin HL}}$	4.92	0.3364	0.0890	2.866
$r^{\text{basin LL}}$	4.89	0.2947	−0.0282	3.0056
$w^{\text{basin HL}}$	1.0271	0.0125	0.0014	2.8691
$w^{\text{basin LL}}$	1.0289	0.0111	0.1186	3.0302

debt can push from the basin of attraction of the HL-SSS toward the basin of attraction of the LL-SSS.[16]

Table II reports the moments of the economy conditional on the basin of attraction. There are three takeaways from Table II. First, the mean of output is higher at the LL-SSS basin than at the HL-SSS, despite output being higher at the HL-SSS itself than at the LL-SSS ($Y^{HL} = 1.5824 > Y^{\text{LL}} = 1.5804$). The reversal is due to the different skewnesses of the distribution of output at each basin. The levels of low output at the HL-SSS basin occur in the aftermath of large negative shocks to capital.[17] Second, when leverage is high, the economy fluctuates more. The mechanism, outlined a few paragraphs back, is the higher persistence of capital in the basin of the HL-SSS after a shock. Third, there is a mild excess kurtosis of output at the HL-SSS. Similar points can be made about the risk-free rate and wages.

6.5. *Debt Supercycles*

Figure 8 plots the histogram of the duration of spells of the economy around each SSS on a log scale.[18] The interaction of financial frictions and precautionary behavior leads to extremely persistent middle- and long-run dynamics of borrowing and deleveraging that are not present in conventional DSGE models. The median duration of a spell around the HL-SSS is 4.167 years and around the LL-SSS, 3.917 years. However, the distribution of spells at the HL-SSS has a much thicker right tail, with many spells lasting centuries. Indeed, the economy spends 86.738% of its time (86.706% if we condition on spells lasting more than one year) in the basin of attraction of the HL-SSS.

The spells in Figure 8 resemble the debt supercycles documented by Reinhart and Rogoff (2009) and Schularick and Taylor (2012). Imagine, for instance, an economy that is pushed to the LL-SSS basin because of a significant deleveraging (this can be due to widespread defaults or high inflation). After a series of negative capital shocks that send the economy back to the HL-SSS basin, the economy will experience a protracted process

[16]Figure S4 in Appendix H plots convergence paths from a large number of initial conditions to show the deterministic dynamics of the model in the absence of shocks. Interestingly, these deterministic dynamics are often non-monotone and plodding, with extremely high persistence.

[17]See Fernández-Villaverde and Guerrón-Quintana (2020) and Dew-Becker, Tahbaz-Salehi, and Vedolin (2019) for evidence that aggregate activity is negatively skewed. Since our model suggests that over the last several decades we have been at the HL-SSS basin, this is what we would expect to find in the data.

[18]To make this section more informative, we eliminate from our statistics and Figure 8 the spells with a duration of less than one year. When the value of the state variables is close to the frontier between the two basins of attraction, there are many extremely short spells as the economy jumps from one basin to the other after a shock before the convergence dynamics of the model push the state variables closer to the SSS.

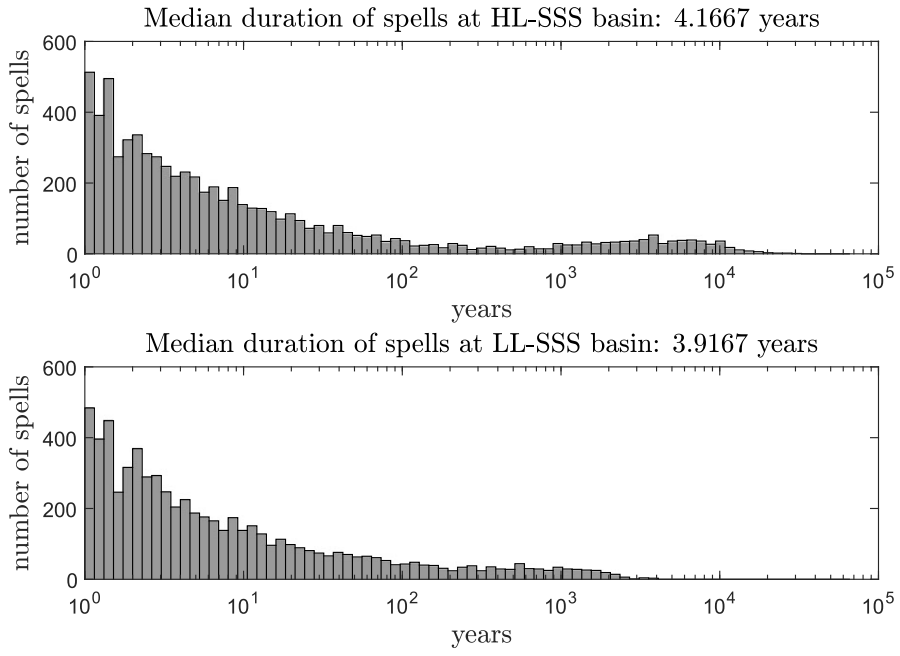

FIGURE 8.—Spell durations at each SSS.

of progressively higher wealth inequality, increasing leverage, falling risk-free rates, and larger responses to capital shocks.[19]

6.6. *The Aggregate Ergodic Distribution of Debt and Equity*

Panel (a) of Figure 9 displays the aggregate ergodic joint distribution of debt and equity $F(B, N)$. This distribution is defined, for any subset Ω of the state space, as $\mathbb{P}\{(B, N) \in \Omega\} = \int_\Omega dF$ and it can be computed by simulation.

Panels (b) and (c) of Figure 9 plot the marginal ergodic distributions for debt and equity. 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 correspond to the neighborhood of the HL-SSS. In comparison, the neighborhood of the LL-SSS appears much less often. As seen in Figure 8, the desire of households to accumulate more debt lowers the height of this second peak as there is a strong force for reversion toward higher levels of leverage. The marginal distribution of debt is much more concentrated than the marginal distribution of equity. The ergodic distribution has a substantial tail in areas of high equity and low debt, creating asymmetry in aggregate dynamics.

[19]This process qualitatively reminds us of the experience of many rich economies after 1945. For example, the United States came out of World War II with a low level of household and corporate debt, the fruit of pervasive defaults and the lack of new borrowing during the Great Depression and the war (McKinsey Global Institute (2010, pp. 71–73)). Our model focuses on private debt. But even if we consider government debt, the U.S. had low levels of aggregate debt at the end of 1945. This situation lasted until leverage started rising in the 1970s. Once the economy was in the new basin, higher leverage was accompanied by higher wealth inequality, falling risk-free rates, and a heightened fragility of the economy to adverse shocks.

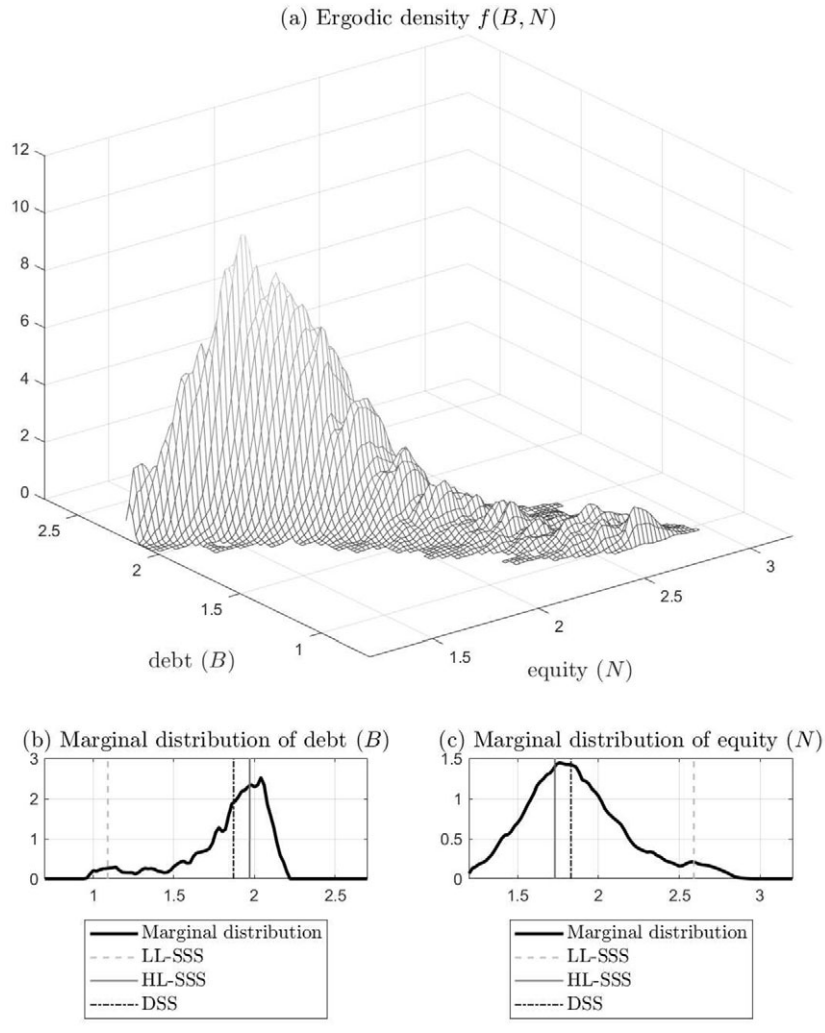

FIGURE 9.—Ergodic distributions. Lighter colors indicate higher probability.

7. THE ROLE OF HOUSEHOLD HETEROGENEITY

This section explores how heterogeneity among households drives our results. We start by comparing, in the top row of Figure 10, the wealth distribution in the DSS (discontinuous dotted line), the HL-SSS (continuous line), and the LL-SSS (discontinuous line). In the left panel, we plot the distributions for low-z households (with circles denoting the mass at zero assets). In the right panel, we plot the distributions for high-z households.

The wealth distribution slightly shifts to the left as we move from the DSS to the HL-SSS, but shifts rather dramatically, also to the left, as we travel from the HL-SSS to the LL-SSS. The reason behind these moves is the changes in endogenous aggregate risk we have discussed before. The higher volatility and persistence of wages in the HL-SSS generate a thicker right tail than at the DSS.

These movements in the distribution lead to substantial differences in the wealth Gini coefficient: 0.28977 in the DSS, 0.24318 in the LL-SSS, and 0.31935 in the HL-SSS. In

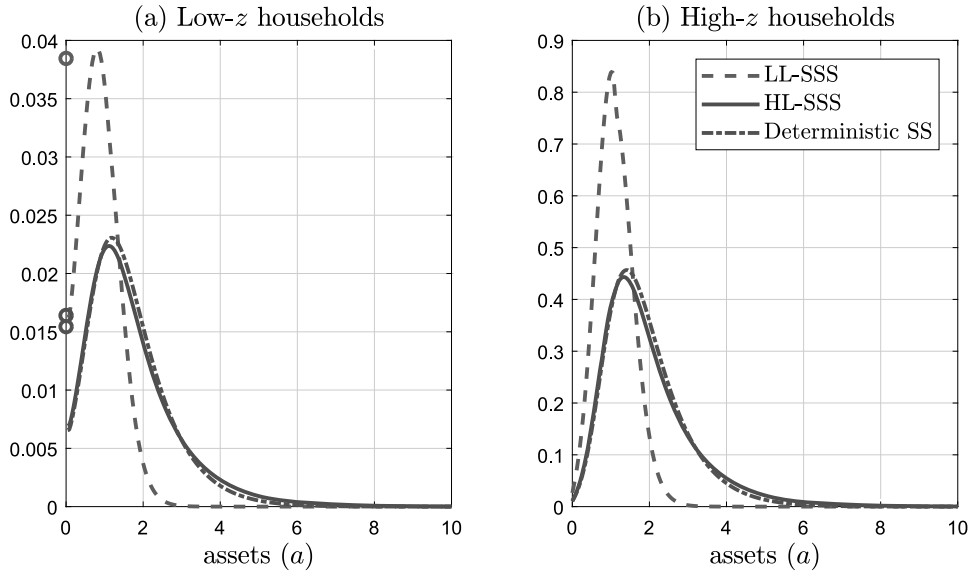

FIGURE 10.—Wealth distribution in the DSS and SSS.

comparison, the canonical Krusell–Smith model reported in Krueger, Mitman, and Perri (2016, p. 873, Table 7) has a wealth Gini of 0.350. If we read the evidence as suggesting that the U.S. economy has been at the basin of attraction of the HL-SSS over the last few decades, our model does quite well in comparison with the canonical Krusell–Smith environment despite the fact that we do not have correlated aggregate and idiosyncratic shocks. This result is not surprising since there is a fair amount of idiosyncratic risk in our model: adding aggregate risk only to the model increases the standard deviation of income from 0.09064 to 0.093115 when we are at the HL-SSS and from 0.072379 to 0.093115 when we are at the LL-SSS.

Having more wealth also means that the average welfare of households at the HL-SSS is higher. Households are willing to give up, on average, 0.034% of their consumption to (instantaneously) move from the DSS to the HL-SSS. However, we need to compensate them, also on average, 0.340% of their consumption to go to the LL-SSS. These numbers must be read carefully, though. The value functions of households at the LL-SSS are above the value functions of those at the HL-SSS for all asset values because the economy is less volatile and the utility function is concave (see the plot of the value functions in Figure S10 in Appendix J). However, we have more households with larger assets at the HL-SSS.

Next, we show how the results of Section 6 change as we modify the forces driving the wealth distributions in Figure 10. We will answer this question in three steps. First, we will vary the level of aggregate and idiosyncratic risk among households. Second, we will explore how the differences in consumption decisions and the variation in distributions over time account for our results. Third, we will conclude with a brief comment on our approximation of the PLM.

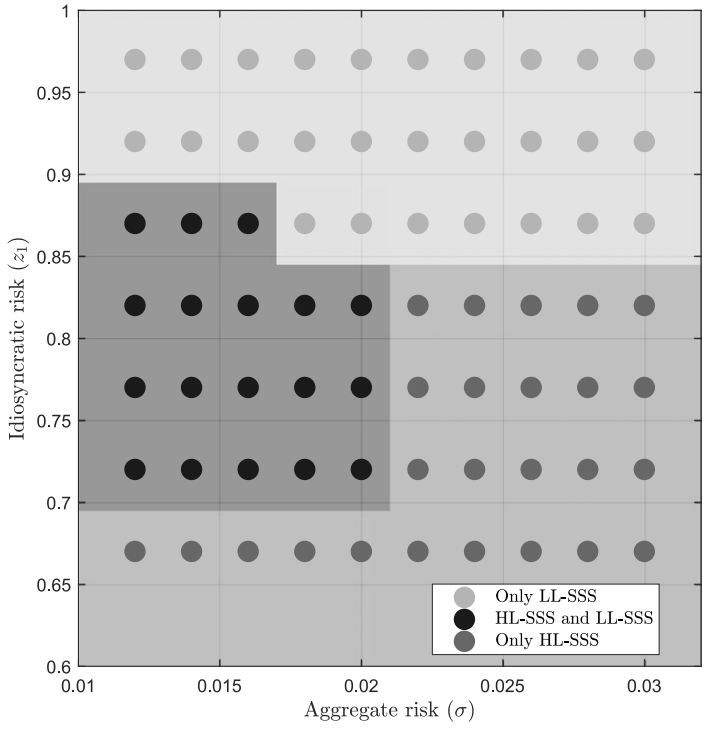

FIGURE 11.—Multiplicity of SSS(s) as a function of aggregate and idiosyncratic volatility.

7.1. *The Interaction of Aggregate and Idiosyncratic Risk*

Figure 11 shows the consequences of moving aggregate and idiosyncratic risk simultaneously and, as such, it is a good summary of most of our results.[20] Each point in the figure represents a different combination of values of z_1 and σ, with the associated colors of intermediate values computed with a nearest-neighbors algorithm.

There are three regions in the figure. For high levels of idiosyncratic risk (i.e., low z_1), and due to the subsequent high precautionary behavior of households, there is only one HL-SSS. The region becomes larger (i.e., for higher z_1) as we increase σ. For intermediate values of z_1 (between 0.7 and 0.9) and moderate levels of aggregate risk (σ below 0.02; recall the maximum likelihood point estimate of 0.0142 using U.S. data), we have an HL-SSS and an LL-SSS, exactly as discussed in previous pages. The region shrinks as aggregate risk increases: households exhibit more precautionary behavior and the HL-SSS disappears. Finally, for high values of z_1, the level of idiosyncratic risk is sufficiently low that only the LL-SSS exists. One can think about the top row of points, when $z_1 = 0.97$, as a version of the model with very low household heterogeneity. While this is not strictly the representative household version of the model, idiosyncratic risk is so small (low productivity is just 3% below the ergodic mean, and only 5% of households are in such a situation) that we approximate very well the limit case of a representative household.[21]

[20] Appendix I documents how the properties of the model change when we vary either *only* aggregate risk or *only* idiosyncratic risk.

[21] We do not compute the exact representative household version of the model because it would require a different algorithm than the one we use for the heterogeneous case. Consequently, there would be numerical

Figure 11 demonstrates the importance of household heterogeneity in terms of the existence of different SSS(s). If we look at a version of the model with low or no household heterogeneity (the top region), we only discover a unique SSS and the model has quite disparate properties (the next subsection will illustrate some of them). Such a model would need different parameter values to match the U.S. data, biasing the usefulness of the model for counterfactual analysis and welfare evaluation.

Appendix K complements this discussion by comparing the GIRFs of the baseline high household heterogeneity version of the model with the GIRFs of the low household heterogeneity version. These sets of GIRFs differ in important ways, further illustrating the importance of household heterogeneity.

7.2. *Inspecting the Heterogeneity Mechanism*

The key to understanding Figure 11 is to analyze how the reduction in households' total consumption after a negative aggregate shock is distributed. After this shock, wages decrease and the risk-free rate increases. Poorer households, mainly dependent on labor earnings, lose much. In comparison, richer households, with more income from their assets, lose less or, if they are sufficiently wealthy, can even gain (see, in the third row of Figure S10 in Appendix J, how the value function improves for households with high assets, in particular in the HL-SSS).

Appendix L plots the consumption decision rules of households along their asset axis and shows how the wealth distribution evolves after a shock. Rich households reduce their consumption less at the HL-SSS than in the LL-SSS. At the HL-SSS, rich households reduce their consumption not so much as a consequence of lower labor income (wealth effect) but as a consequence of better rewards for savings (substitution effect). Moreover, the sluggish aggregate dynamics of the model mean that wealthy households expect a higher risk-free rate to last for a long time. Because of the marginal decreasing utility of consumption, the integral of these consumption responses across households is different than the response of a household with average wealth. Hence, the responses of total consumption and the evolution of total debt are also different.

7.3. *On the Approximation of the PLM*

But if household heterogeneity matters in this economy, why can we approximate our PLM well with just one moment of the distribution? Because the consumption decision rule of households is close to linear with respect to the household state variables. Only agents close to the borrowing limit face a nonlinear consumption decision rule but, being close to zero assets, they contribute little to the aggregate dynamics of capital.

Interestingly, the consumption decision rule of households is close to linear with respect to the household state variables even though it is strongly nonlinear with respect to the aggregate state variables (B, N). Households understand the nonlinear responses of wages and the risk-free rate they receive from labor and savings with respect to (B, N) and behave accordingly. See the different decision rules in Figure 12.

This point is of general applicability: we can have models with agents' heterogeneity that imply nearly linear decision rules of the agents with respect to individual state variables and very nonlinear decision rules with respect to aggregate state variables. In that

chatter between both solutions that might complicate the comparison. Hence, we prefer to set $z_1 = 0.97$ (as high as our algorithm can go before breaking down).

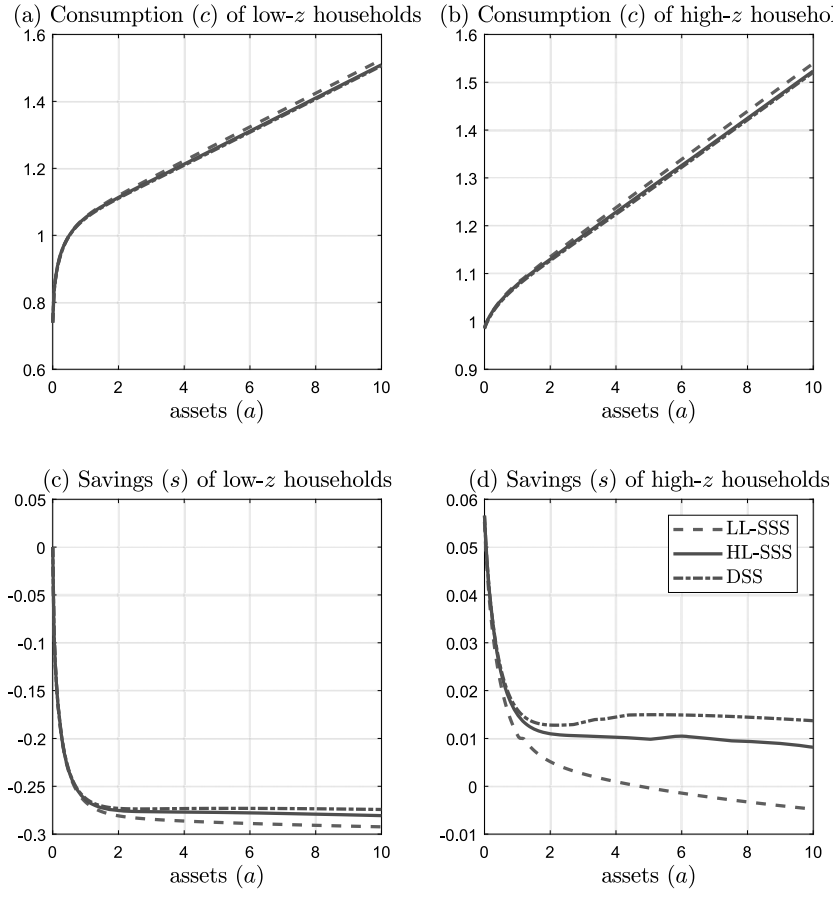

FIGURE 12.—Consumption and saving functions.

case, we will need solution methods that can capture such a nonlinearity and neural networks provide the required flexibility.

8. CONCLUSION

In this paper, we have analyzed a continuous-time DSGE model with a financial sector and heterogeneous households. This exercise has allowed us to argue that the wealth distribution among households is central to studying the effects of financial frictions. More concretely, we have shown how the total level of endogenous aggregate risk creates an endogenous regime-switching process for output, interest rates, debt, and leverage. The regime switching generates strong state dependence on the GIRFs of the model.

Methodologically, we also show how to efficiently compute and estimate a general class of models 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. Our methodology can help analyze heterogeneous agent models with aggregate shocks that display significant nonlinear features. An obvious candidate is the heterogeneous agent New Keynesian (HANK) model with a zero lower bound (ZLB) on the nominal interest rates explored by Fernández-Villaverde et al. (2022). The

ZLB introduces a nonlinearity in the state space of aggregate variables that local or global methods based on linear laws of motion cannot address. We hope this HANK model will be the first of many other related papers.

REFERENCES

ACHDOU, YVES, JIEQUN HAN, JEAN-MICHEL LASRY, PIERRE-LOUIS LIONS, AND BENJAMIN MOLL (2021): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *Review of Economic Studies*, 89, 45–86. [881]

ADRIAN, TOBIAS, AND HYUN SONG SHIN (2010): "Liquidity and Leverage," *Journal of Financial Intermediation*, 19, 418–437. [873]

ADRIAN, TOBIAS, NINA BOYARCHENKO, AND DOMENICO GIANNONE (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–1289. [869]

AHN, SEHYOUN, GREG KAPLAN, BENJAMIN MOLL, THOMAS WINBERRY, AND CHRISTIAN WOLF (2017): "When Inequality Matters for Macro and Macro Matters for Inequality," in *NBER Macroeconomics Annual 2017*, Vol. 32. University of Chicago Press. [873]

AIYAGARI, S. RAO (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 659–684. [879]

ALVAREDO, FACUNDO, LUCAS CHANCEL, THOMAS PIKETTY, EMMANUEL SAEZ, AND GABRIEL ZUCMAN (2017): "Global Inequality Dynamics: New Findings from WID.world," *American Economic Review*, 107, 404–409. [873]

ANDREASEN, MARTIN, JESÚS FERNÁNDEZ-VILLAVERDE, AND JUAN RUBIO-RAMÍREZ (2018): "The Pruned State-Space System for Non-Linear DSGE Models: Theory and Empirical Applications," *Review of Economic Studies*, 85, 1–49. [871]

AZINOVIĆ, MARLON, LUCA GAEGAUF, AND SIMON SCHEIDEGGER (2022): "Deep Equilibrium Nets," *International Economic Review*, 63, 1471–1525. [874]

BACH, FRANCIS (2017): "Breaking the Curse of Dimensionality With Convex Neural Networks," *Journal of Machine Learning Research*, 18, 629–681. [874,882]

BARRON, ANDREW R. (1993): "Universal Approximation Bounds for Superpositions of a Sigmoidal Function," *IEEE Transactions on Information Theory*, 39, 930–945. [882]

BERNANKE, BEN S., MARK GERTLER, AND SIMON GILCHRIST (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," in *Handbook of Macroeconomics*, Vol. 1, ed. by John B. Taylor and Michael Woodford. Elsevier, 1341–1393. [879]

BILAL, ADRIEN (2021): "Solving Heterogeneous Agent Models With the Master Equation," Tech. rep, University of Chicago. [880]

BLANCHARD, OLIVIER, AND JORDI GALÍ (2010): "Labor Markets and Monetary Policy: A New Keynesian Model With Unemployment," *American Economic Journal: Macroeconomics*, 2, 1–30. [884]

BRUNNERMEIER, MARKUS K., AND YULIY SANNIKOV (2014): "A Macroeconomic Model With a Financial Sector," *American Economic Review*, 104, 379–421. [873,875]

CHEELA, BHAGATH, ANDRÉ DEHON, JESÚS FERN ÁNDEZ-VILLAVERDE, AND ALESSANDRO PERI (2022): "Programming FPGAs for Economics: An Introduction to Electrical Engineering Economics," Working Paper 29936, National Bureau of Economic Research. [882]

CHILDERS, DAVID (2022): "Automated Solution of Heterogeneous Agent Models," Carnegie Mellon. [873]

CHOLLET, FRANÇOIS (2021): "Deep Learning With Python, Manning Publication," (Second Ed.). Manning Publication. [883]

COEURDACIER, NICOLAS, HÉLÈNE REY, AND PABLO WINANT (2011): "The Risky Steady State," *American Economic Review*, 101, 398–401. [871]

CYBENKO, GEORGE (1989): "Approximation by Superpositions of a Sigmoidal Function," *Mathematics of Control, Signals and Systems*, 2, 303–314. [882]

DEW-BECKER, IAN, ALIREZA TAHBAZ-SALEHI, AND ANDREA VEDOLIN (2019): "Macro Skewness and Conditional Second Moments: Evidence and Theories," Northwestern University. [893]

DUARTE, VICTOR (2018): "Sectoral Reallocation and Endogenous Risk-Aversion: Solving Macro-Finance Models With Machine Learning," MIT Sloan School of Management. [874]

DUFFIE, DARRELL, AND LARRY EPSTEIN (1992): "Stochastic Differential Utility," *Econometrica*, 60, 353–394. [875]

EBRAHIMI KAHOU, MAHDI, JESÚS FERNÁNDEZ-VILLAVERDE, JESSE PERLA, AND ARNAV SOOD (2021): "Exploiting Symmetry in High-Dimensional Dynamic Programming," Working Paper 28981, National Bureau of Economic Research. [874,882,883]

EVANS, GEORGE W., AND SEPPO HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press. [880]

FERNÁNDEZ-VILLAVERDE, JESÚS, AND PABLO GUERRÓN-QUINTANA (2020): "Uncertainty Shocks and Business Cycle Research," *Review of Economic Dynamics*, 37, S118–S146. [870,893]

FERNÁNDEZ-VILLAVERDE, JESÚS, SAMUEL HURTADO, AND GALO NUÑO (2023): "Supplement to 'Financial Frictions and the Wealth Distribution'," *Econometrica Supplemental Material*, 91, https://doi.org/10.3982/ECTA18180. [874]

FERNÁNDEZ-VILLAVERDE, JESÚS, JÖEL MARBET, GALO NUÑO, AND OMAR RACHEDI (2022): "Inequality and the Zero Lower Bound," University of Pennsylvania, available at https://www.sas.upenn.edu/∼jesusfv/inequality_ZLB.pdf. [870,899]

GALÍ, JORDI, AND LUCA GAMBETTI (2009): "On the Sources of the Great Moderation," *American Economic Journal: Macroeconomics*, 1, 26–57. [884]

GREENWOOD, JEREMY, ZVI HERCOWITZ, AND PER KRUSELL (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87, 342–362. [875]

HALL, ROBERT E., AND PAUL R. MILGROM (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98, 1653–1674. [884]

HE, ZHIGUO, AND ARVIND KRISHNAMURTHY (2012): "A Model of Capital and Crises," *Review of Economic Studies*, 79, 735–777. [873]

——— (2013): "Intermediary Asset Pricing," *American Economic Review*, 103, 732–770. [873]

HOLSTON, KATHRYN, THOMAS LAUBACH, AND JOHN C. WILLIAMS (2017): "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, 108, S59–S75. [873]

HORNIK, KURT, MAXWELL STINCHCOMBE, AND HALBERT WHITE (1989): "Multilayer Feedforward Networks Are Universal Approximators," *Neural Networks*, 2, 359–366. [874,882]

HUGGETT, MARK (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969. [876]

JORDÀ, ÒSCAR, MORITZ SCHULARICK, AND ALAN M. TAYLOR (2016): "Macrofinancial History and the New Business Cycle Facts," in *NBER Macroeconomics Annual 2016*, Vol. 31, University of Chicago Press, 213–263. [869]

KRUEGER, DIRK, KURT MITMAN, AND FABRIZIO PERRI (2016): "Macroeconomics and Household Heterogeneity," in *Handbook of Macroeconomics*, Vol. 2, ed. by John B. Taylor and Harald Uhlig. Elsevier, 843–921. [896]

KRUSELL, PER, AND ANTHONY A. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896. [870,871,873,880,882,886]

KUBLER, FELIX, AND SIMON SCHEIDEGGER (2018): "Self-Justified Equilibria: Existence and Computation," Working Paper, DBF, University of Zurich. [871,882]

LO, ANDREW (1988): "Maximum Likelihood Estimation of Generalized itô Processes With Discretely Sampled Data," *Econometric Theory*, 4, 231–247. [871]

MALIAR, LILIA, SERGUEI MALIAR, AND PABLO WINANT (2019): "Will Artificial Intelligence Replace Computational Economists Any Time Soon?" CEPR Discussion Papers DP14024, CEPR. [874]

MCKINSEY GLOBAL INSTITUTE (2010): "Debt and Deleveraging: The Global Credit Bubble and Its Economic Consequences," Tech. rep, McKinsey Global Institute. [894]

MOFFITT, ROBERT, AND SISI ZHANG (2018): "Income Volatility and the PSID: Past Research and New Results," *AEA Papers and Proceedings*, 108, 277–280. [873]

NUÑO, GALO, AND CARLOS THOMAS (2016): "Optimal Monetary Policy With Heterogeneous Agents," Working Paper 1624, Banco de España. [881]

——— (2017): "Bank Leverage Cycles," *American Economic Journal: Macroeconomics*, 9, 32–72. [873]

REINHART, CARMEN M., AND KENNETH S. ROGOFF (2009): *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press. [870,893]

SCHEIDEGGER, SIMON, AND ILIAS BILIONIS (2019): "Machine Learning for High-Dimensional Dynamic Stochastic Economies," *Journal of Computational Science*, 33, 68–82. [874]

SCHENK-HOPPÉ, KLAUS REINER, AND BJÖRN SCHMALFUSS (2001): "Random Fixed Points in a Stochastic Solow Growth Model," *Journal of Mathematical Economics*, 36, 19–30. [879]

SCHULARICK, MORITZ, AND ALAN M. TAYLOR (2012): "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008," *American Economic Review*, 102, 1029–1061. [870,893]

Co-editor Dave Donaldson handled this manuscript.

Manuscript received 3 March, 2020; final version accepted 23 January, 2023; available online 23 February, 2023.