Asset Pricing in a Low Rate Environment*

Preliminary and Incomplete

Marlon Azinovic†
University of Pennsylvania

Harold Cole
University of Pennsylvania

Felix Kubler
University of Zurich

February 28, 2023

Abstract

We examine asset prices in environments where the risk-free rate lies considerably below the growth rate. To do so, we introduce a tractable model of a production economy featuring heterogeneous trading technologies, as well as idiosyncratic and aggregate risk. We show that allowing for the possibility of firms exiting is crucial for matching key macroeconomic moments and, simultaneously, the risk-free rate, the market price of risk, and price-earnings ratios. In particular, our model allows us to consider calibrations that match the high observed market price of risk and average interest rates that can be as low as 2-3 percent below the average growth rate. High values for risk aversion or non-standard preferences are not necessary for this. We use the model to examine under which conditions realistic calibrations allow for an infinite rollover of government debt. For our benchmark calibration, rollover is impossible even if the average risk-free rate lies three percent below the average growth rate.

*We are grateful for the helpful comments we received from Nir Jaimovich.

†Azinovic is grateful to the Economics Department at the University of Pennsylvania for their kind hospitality and acknowledges support from the Swiss National Science Foundation under grant P500PS_210799, "Macro-Finance with Heterogeneous Households".

1 Introduction

We develop a tractable model of a production economy that matches key statistics of both asset prices and business cycle dynamics: it produces a standard market price of risk, a low and smooth risk-free rate, realistic price-earnings ratios as well as volatilities of aggregate output and consumption in line with business cycle statistics. We do this within a heterogeneous agents economy with a realistic degree of risk aversion by incorporating the life-and-death cycle of firms. Our model imposes restrictions on the asset pricing kernel. These restrictions allow us to make predictions about the impact of lower interest rates on the prices of long-run assets.

Since 1980, interest rates on U.S. government bonds have steadily decreased. They are now lower than the nominal growth rate. Figure 1 depicts different measures of the US real risk-free rate over the last 30 years. Depending on the choice of the time horizon, the average annual real rate lies somewhere between 0.5 and 0.9 percent (see Mehra and Prescott (1985) and Beeler and Campbell (2012)). The growth rate of per capita output has been roughly 2 percent. At the same time, the price-earnings ratio of S&P500 stocks lies around 25, and Shiller's home price index, deflated by the CPI, which captures real land rents, only rose by 60 percent. This is while the real 10-year rate in the figure fell from around 4 percent in 1990 to averaging less than 0.5 percent between 2011-2021.

The positive average gap between the U.S. growth rate and the real interest rate on U.S. Treasuries has sparked a large literature on the question of whether an infinite roll-over of government debt is possible and whether deficit finance has no fiscal cost (see, e.g., Blanchard (2019), Mian et al. (2021), Aguiar et al. (2021)). One important message of the literature on this debate is that the average interest rate alone contains little information on whether infinite debt rollover is possible (see, e.g. (Kocherlakota (2022), Bloise and Reichlin (2022)). Intuitively, the reason for this

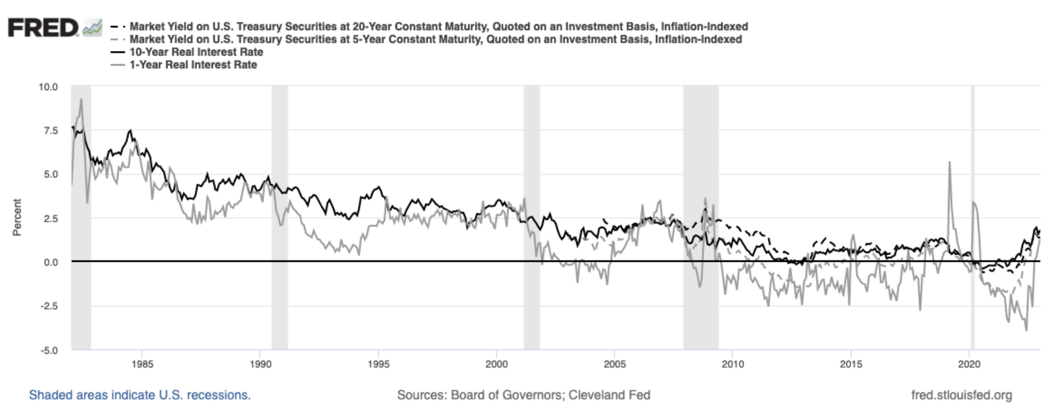

Figure 1: Real interest rates

ambiguity stems from the fact that when $r - g$ goes from 0.01 to 0.0001, the share of the present value of a constant stream realized within the first 100 years falls from 60 percent to 1 percent. This extreme reliance on the tail outcomes implies that if there is one state where the interest rate is positive and that occurs with probability 1 in finite time, and if in this state all state prices for transitions into different states are zero, then this will effectively truncate the future and deliver a finite present value. Of course, while the existence of one state with a positive interest rate is very likely, zero state prices are impossible in a model without satiation. But this highlights the importance of using a structural model of the macro-economy to determine endogenously determine returns and the variation in the pricing kernel.

While simple consumption-based asset pricing models cannot reproduce observed returns on equity and the risk-free rate (see, e.g., Weil (1989)) it is now well understood how the joint assumption of agents' heterogeneity and financial frictions can enrich the asset pricing implications of the model (see Constantinides and Duffie (1996)). Our model focuses on two dimensions of heterogeneity: first, we model two types of households with differential access to trading technologies, as in Chien et al. (2011). The first type of households, which we call advanced traders, can trade a full set of aggregate-state contingent securities. The second type of traders, which we call non-participants, only trade a one-period risk-free bond. The second dimension of heterogeneity we focus on is wealth heterogeneity, which arises endogenously from uninsurable idiosyncratic risk. Uninsurable idiosyncratic risk gives rise to a distribution of wealth across and within types and realizations of the idiosyncratic shocks. Due to the heterogeneity in trading technologies, the wealth distribution moves strongly in response to aggregate shocks. The large movements in the wealth distribution drive the high market price of risk. At the same time, we restrict ourselves to household preferences that exhibit a degree of risk aversion of five and do not include a habit factor that can lead to much higher local degrees of aversion to risk. We take this to be an important component of generating a reasonable pricing kernel off of the marginal investor.

As in Aiyagari (1994) and Huggett (1993), idiosyncratic income risk, together with borrowing constraints, are the key drivers of a low risk-free rate in our model. The presence of adjustment costs on the firms' investment implies that the infinitely lived firms pay an infinite stream of non-zero dividends. When interest rates are very low, the market price of the firms can potentially become very large. To match observed price-dividend ratios, we introduce firm-exit into our model: Every period, a fixed fraction of firms exit the market and are replaced by new firms founded by households. We argue that models without firm exit are unlikely to jointly match the risk-free rate, the market price of risk, and price-earnings ratios when we take the real interest rate to be very small. It is, of course, well-documented that every year a substantial fraction of firms and establishments in the US "die"(see, e.g. Crane et al. (2022)) and we calibrate the exit rate in our model to observed average death rates.

Even for high interest rates, computing equilibria for our model turns out to be substantially more difficult than for the model considered in Krusell and Smith (1998). First, in order to forecast the next period's endogenous variables, we need to introduce another aggregate variable, the wealth share of the advanced traders, in addition to aggregate capital. One aggregate moment does not suffice to capture the strong nonlinearities and movements in the wealth distribution. Second, in order to solve for equilibrium prices of the Arrow securities we need to employ approximation methods for the consumption functions which have good extrapolation properties, in the sense that the gradients do not change too much when moving outside the domain on which the approximating function was previously fitted. Advanced methods used in the literature such as Gaussian processes (Scheidegger and Bilionis (2019)) or neural nets (Azinovic et al. (2022)) performed worse in our experiments.

We employ a simulation-based method to solve for the equilibrium of our model, in which we alternate between simulating and updating. Hence, our equilibrium functions must be able to extrapolate well. To forecast the equilibrium outcomes of our model, we take as the aggregate state: the level of capital, the wealth share of the advanced traders, and the exogenous aggregate productivity shock. This state is sufficient to accurately forecast its own update. As is standard in Euler equation methods, we use this policy function (along with a forecast of next period's advance trader wealth share) to determine the policy of the firm next period and then solve jointly for next period's capital stock and Arrow prices so as to equate asset demand to asset supply. This critically involves determining the overall asset demand of the households. The natural state for households includes, in addition to the aggregate state, their type, idiosyncratic shock, and their asset level. Given this household state, we again use an Euler equation method to solve for current asset demand given tomorrow's consumption policy function and current Arrow prices. Because the response of individual consumption with respect to changes in the aggregate capital stock and wealth share of the advanced traders differs markedly in their own wealth level, fitting this directly would require a very flexible functional form, which typically would not extrapolate well. To deal with this we posit separate simple consumption policy functions for our households on a discretized grid of the aggregate shock, the idiosyncratic shock, and the households' asset holdings.

An important aspect of our approach is that we can compute equilibria for our model even when prices are not summable. The value of the firms remains finite even if the present value of a strictly positive stream of future payments explodes. Borrowing constraints for households ensure that even though the present value of future labor endowments might be infinite, the household's optimization problem still has a solution. Of course, computing an equilibrium of our model when prices are nearly summable or nearly non-summable is computationally challenging both because of the rich state space and the impact of low interest rates, which effectively lengthens pricing horizons. This reinforces the need for computational methods that are tailored precisely to our economic problem.

As an application, we investigate whether a consol with a dividend growing at the long-run growth rate has a finite value. If this is the case, it is impossible to roll over debt infinitely. We find that as long as the parameterization of our model matches a realistically high market price of risk, average interest rates can become very small (with a growth rate of 2 % we consider average interest rates of up to -1 %), while the value of the consol remains finite.

This paper builds on the literature studying limited stock market participation as a central channel to explain asset prices. The relevance of this channel to explain asset returns is established by Vissing-Jørgensen (2002) and Vissing-Jørgensen and Attanasio (2003). In line with these findings, Parker and Vissing-Jorgensen (2009) show that the consumption growth of high-consumption households is substantially more exposed to aggregate shocks than for average households. Guo (2004) and Chien et al. (2011) model an endowment economy with limited stock market participation and show that they can obtain a low risk-free rate and a large and counter-cyclical equity premium in line with the data. In difference to these papers, we model a production economy. Guvenen (2009) models two traders, which are heterogeneous in their trading technology as well as their intertemporal elasticity of substitution (IES) in a production economy. He obtains a large and counter-cyclical risk premium as well as a low and smooth risk-free rate. While the main mechanism in his paper, the heterogeneity in trading technologies, is similar to ours, there are several important differences in our paper. First, there is only one representative agent per type in his model, *i.e.* there is no uninsurable idiosyncratic risk. Consequently, his model is not suitable for studying issues related to the wealth distribution. Moreover, one cannot study environments with very low average risk-free rates. Second, the more advanced type of traders in his model can only trade a stock and a bond, while the advanced traders in our model can trade a complete set of aggregate-state contingent securities as in Chien et al. (2011). This simplification is not innocuous in the sense that it has significant effects on the resulting portfolio choice. Third, the between-type heterogeneity in our model is driven exclusively by their access to trading technologies. At the same time, Guvenen (2009) models a difference in preferences (the IES) between the two agents. While this does not seem to be crucial for his key results it does add an additional layer of complexity to his model.

Production-based asset pricing models with uninsurable idiosyncratic risks, maybe because of the associated computational difficulties, remain understudied. An exception in the context of limited stock market participation, which is closely related to our paper, is Favilukis (2013). Favilukis (2013) models a production economy with an OLG structure featuring idiosyncratic risk with counter-cyclical variations on the household side. Households have to pay two types of fixed costs to participate in the stock market. Similarly to Guvenen (2009), there are only two assets, a stock, and a bond. Like our model, his model generates asset pricing and business cycle statistics, which are roughly in line with the data. To the best of our knowledge, there is no work on production economies with limited stock-market participation, which incorporates firms' exit and entry and allows for the possibility of very low risk-free rates.

The paper is organized as follows. In Section 2, we introduce the model; Section 3 describes the calibration and key aspects of the quantitative implications of our model. In Section 4, we discuss the possibility of debt rollover, and Section 5 concludes. In the Appendix, we explain our computational approach in detail.

2 Model

We model a production economy with segmented financial markets and infinitely lived heterogeneous households that face aggregate and idiosyncratic risk. The existing firms have an exogenous probability of exit, with new firms entering to replace them. Time is discrete and indexed by $t = 0,\ 1,\ 2, \ldots$. We denote the aggregate shock in period t by $z_t \in \mathcal{Z} = \{1, \ldots, Z\}$ and assume that aggregate shocks follow a Markov chain with transition π. We use $\pi(z_{t+1}|z_t)$ for the probability to transit into aggregate shock z_{t+1} when currently in aggregate shock z_t. We denote a history of aggregate shock by z^t and we use $\pi(z^t)$ to denote the unconditional probabilities of shock sequences z^t. In addition, households face idiosyncratic risks which are assumed to cancel out in the aggregate and which we discuss below.

2.1 Firms

We model a continuum of firms that take capital and labor as input and produce the single consumption good with identical Cobb-Douglas production technology. The production function of firm i is

$$y_t^i = \xi(z_t)(k_t^i)^\alpha (A_t l_t^i)^{1-\alpha}, \tag{1}$$

where $\xi(z_t)$ is the stochastic productivity level, A_t grows deterministically at a fixed rate, k_t^i and l_t^i denote the amount of capital and labor firm i chooses in period t.

In every period, a firm exits with an exogenous probability $1-\Gamma$, where Γ denotes the exogenous survival probability of firms. If firms exit, they do not pay dividends, and their capital evaporates.[1] In each period, a measure of $(1-\Gamma)/\Gamma$ new firms enter the economy, and upon entry, new firms invest an exogenous amount $\bar{i}_t > 0$. Conditional on surviving, new firms entering in period t start producing in period $t+1$ with capital $k_{t+1}^i = \bar{i}_t$. The difference between the investment of new and incumbent firms is that the investment of new firms is exogenously chosen and not subject to adjustment costs.

Conditional on surviving, incumbent firm i's capital accumulates according to

$$k_{t+1}^{i*} = k_t^i(1-\delta) + i_t^i. \tag{2}$$

δ denotes the depreciation rate of capital and k_{t+1}^{i*} denotes firm i's capital in the next period conditional on surviving. Throughout the paper, x_{t+1}^{i*} denotes quantity x for firm i at time $t+1$ conditional on surviving from period t to $t+1$. Firms producing at time t pay dividends

$$d_t^i = y_t^i - \omega_t l_t^i - i_t^i - \psi(k_{t+1}^{i*}, k_t^i), \tag{3}$$

where ω_t denotes the wage at time t. Adjustment costs are given by

$$\psi(k', k) := \xi^{\text{adj}} k \left(\frac{k'}{k} - (1 - \delta + x^{\text{target}}) \right)^2. \tag{4}$$

In this specification of adjustment costs, the parameter ξ^{adj} denotes the level of the adjustment costs, and x^{target} denotes the level of the investment to capital ratio for which no adjustment costs have to be paid.

Firms have access to financial markets and trade in a complete set of aggregate-state contingent securities, *i.e.* Arrow securities for the aggregate states. The price of an Arrow security that pays at node z^{t+1}, when traded at node z^t, is denoted by $p(z^{t+1}|z^t)$. From this, we can price all aggregate date-events — we denote the price in period t of consumption at some future date-event z^τ by $p_t(z^\tau)$.

Since Modigliani and Miller (1958) result holds here, the firm's capital structure does not affect its value, so we introduce the firm's problem without the consideration of financial policy. A firm's Bellman equation is given by

$$v_t^i = \max_{k_{t+1}^{i*}} d_t^i + \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^{i*}, \tag{5}$$

[1] In a more detailed model, the capital of exiting firms would have some resale value, but this is typically thought of as being very low and cyclical, see, e.g., Shleifer and Vishny (2011).

where k_{t+1}^{i*} and v_{t+1}^{i*} denote the firm's capital and value in the next period conditional on surviving.

The birth and death cycle of firms combined with adjustment costs means that young firms can differ in size relative to older firms. However, since the probability of death is independent of age or size, these firms effectively aggregate, with each firm having the same capital-to-labor rate, and the same investment-to-capital rate. Thus, effectively there is a representative surviving firm and a representative new firm. This is the sense in which the economy aggregates. To establish this formally, note that each firm i takes the wage as given and chooses l_t^i such that

$$\omega_t = \xi_t A_t (1-\alpha) \left(\frac{k_t^i}{A_t l_t^i} \right)^{\alpha} \tag{6}$$

$$\Leftrightarrow \frac{k_t^i}{A_t l_t^i} = \left(\frac{\omega_t}{A_t \xi_t (1-\alpha)} \right)^{\frac{1}{\alpha}} =: \mathcal{K}_t, \tag{7}$$

where optimal labor choice implies that $\mathcal{K}_t$ is constant across all operating firms. Therefore we obtain that the return to capital across firms must be identical and that

$$w_t = A_t (1-\alpha) \xi_t \mathcal{K}_t^{\alpha} \tag{8}$$

$$r_t = r_t^i = \alpha \xi_t \mathcal{K}_t^{\alpha - 1}. \tag{9}$$

Since firms don't face idiosyncratic risk, all incumbent firms face the same first-order condition for optimality which reads as follows.

$$1 + \underbrace{\frac{\partial \psi(k_{t+1}^{i*}, k_t^i)}{\partial k_{t+1}^{i*}}}_{=2\xi^{\text{adj}}\left(\frac{k_{t+1}^{i*}}{k_t^i} - (1-\delta+x^{\text{target}})\right)} = \sum_{z_{t+1}} p_{z_{t+1}}^t \left(r_{t+1} + 1 - \delta + \underbrace{\frac{\partial \psi(k_{t+2}^{i**}, k_{t+1}^{i*})}{\partial k_{t+2}^{i**}}}_{-\xi^{\text{adj}}\left(\left(\frac{k_{t+2}^{i**}}{k_{t+1}^{i*}}\right)^2 - (1-\delta+x^{\text{target}})^2\right)} \right) \tag{10}$$

Therefore each incumbent firm chooses the same growth rate, which, for a given period, t, we denote by $(g^k)_t^i = g_t^k$.

We define K_t to be the aggregate capital in period t. There is a measure 1 of firms, but only a measure Γ owns capital since startups have to invest to have capital and produce in the next period. Hence the average capital per producing firm is K_t/Γ. For the aggregate capital in period $t+1$ we have

$$K_{t+1} = \underbrace{\Gamma \int_{i \in \mathcal{P}_t} k_{t+1}^{i*} di}_{\text{capital of survivors}} + \underbrace{\Gamma \frac{1-\Gamma}{\Gamma} \bar{i}_t}_{\text{capital of surviving time } t \text{ entrants}} \tag{11}$$

$$\Leftrightarrow \frac{K_{t+1}}{\Gamma} - \frac{1-\Gamma}{\Gamma} \bar{i}_t = \int_{i \in \mathcal{P}_t} k_{t+1}^{i*} di, \tag{12}$$

where $\mathcal{P}_t$ is the set of indexes for firms, which are producing in period t, and where $\bar{i}_t$ denotes the capital level of time t startups in period $t+1$. We define V_t as the value of all firms producing in period t and obtain

$$
\begin{aligned}
V_t &= \int_{j\in\mathcal{P}_t} v_t^j dj \\
&= \int_{j\in\mathcal{P}_t} \left[\begin{array}{c} y_t^j - \omega_t l_t^j - i_t^j - k_t^j \ \xi^{\text{adj}} \left(g_t^k - (1-\delta + x^{\text{target}})\right)^2 \\ +\Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^{j*} \end{array} \right] dj \\
&= Y_t - \omega_t L_t - \int_{j\in\mathcal{P}_t} i_t^j dj - K_t \xi^{\text{adj}} \left(g_t^k - (1-\delta + x^{\text{target}})\right)^2 \\
&+ \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} \int_{j\in\mathcal{P}_t} v_{t+1}^{j*} dj.
\end{aligned}
\tag{13}
$$

For aggregate investment, which includes the investment by startups as well as incumbents, we have that

$$
I_t = \int_{j\in\mathcal{P}_t} i_t^j dj + \frac{1-\Gamma}{\Gamma}\bar{i}_t \quad \Leftrightarrow \int_{j\in\mathcal{P}_t} i_t^j dj = I_t - \frac{1-\Gamma}{\Gamma}\bar{i}_t \tag{14}
$$

The value of a new firm entering in time t is given by

$$
\bar{v}_t := \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^{t\text{-entrant}*}, \tag{15}
$$

where $v_{t+1}^{t\text{-entrant}*}$ denotes the value of an idiosyncratic firm in period $t+1$ with capital $k_{t+1}^j = \bar{i}_t$. Since the value of an idiosyncratic firm is proportional to its capital, we have

$$
v_{t+1}^{t\text{-entrant}*} = V_{t+1}\frac{\bar{i}_t}{K_{t+1}} \tag{16}
$$

While the net value of creating a new firm is simply $\bar{v}_t - \bar{i}_t$.

The value of firms producing next period is given by the value of surviving firms plus firms created today,

$$
V_{t+1} = (1-\Gamma)V_{t+1}\frac{\bar{i}_t}{K_{t+1}} + \int_{j\in\mathcal{P}_t} \Gamma v_{t+1}^{j*} dj
$$

thus the value of the surviving producing firms is given by

$$
\int_{j\in\mathcal{P}_t} \Gamma v_{t+1}^{j*} dj = V_{t+1}\left(1 - (1-\Gamma)\frac{\bar{i}_t}{K_{t+1}}\right) \tag{17}
$$

The recursion for the firm value, expressed in aggregates, is given by

$$
\begin{aligned}
V_t = Y_t - \omega_t L_t &- \left(\left(\frac{K_{t+1}}{\Gamma} - (1-\delta)K_t\right) - \frac{1-\Gamma}{\Gamma}\bar{i}_t\right) \\
&- K_t \xi^{\text{adj}} \left((g^k)_t - (1-\delta + x^{\text{target}})\right)^2 \\
&+ V_{t+1}\left(1 - (1-\Gamma)\frac{\bar{i}_t}{K_{t+1}}\right).
\end{aligned}
\tag{18}
$$

Aggregate dividends are given by

$$D_t = Y_t - \omega_t L_t - (K_{t+1} - (1-\delta)K_t) - K_t \xi^{\text{adj}} \left(g_t^k - (1 - \delta + x^{\text{target}}) \right)^2 . \tag{19}$$

2.1.1 Life-Cycle of Firms

As pointed out above, all surviving firms grow at the same rate, independent of size. To the extent that new entrants are smaller than the average firm size, a version of Zipf's law holds. The new entrants must grow faster than the growth rate of aggregate capital. To better understand this, we abstract from shocks and note that the "detrended" level of aggregate capital in steady state must satisfy

$$\hat{K}_{t+1} = \underbrace{\Gamma \frac{1+g^k}{1+g} \hat{K}_t}_{\text{surviving period t incumbents}} + \underbrace{\Gamma \frac{1-\Gamma}{\Gamma} s \hat{K}_{t+1}}_{\text{surviving period t startups}} . \tag{20}$$

where $\hat{K}_t = \frac{K_t}{(1+g)^t}$ denotes the detrended level of aggregate capital, g^k denotes the growth rate of capital for surviving incumbent firms, and $\frac{1-\Gamma}{\Gamma} s \hat{K}_{t+1}(1+g)$ denotes the period t investment into startups. Since, in steady state, $\hat{K}_t = \hat{K}_{t+1} = \hat{K}^{ss}$, we must have

$$\hat{K}^{ss} = \underbrace{\Gamma \frac{1+g^k}{1+g} \hat{K}^{ss}}_{\text{surviving period t incumbents}} + \underbrace{(1-\Gamma) s \hat{K}^{ss}}_{\text{surviving period t startups}} \tag{21}$$

$$\Rightarrow 1 = \left(\Gamma \frac{1+g^k}{1+g} + (1-\Gamma)s \right) . \tag{22}$$

As one can see from inspection, $g^k > g$ to the extent that $s < 1$. The relative share of employment in new vs. existing firms is given by $(1-\Gamma)s$, and that in exiting firms by $\Gamma \frac{1+g^k}{1+g}$. Thus we can calibrate the entering size with the shares of (employment) capital at exiting and entering firms, or to the growth rate of employment at existing firms.

We have assumed that new entrants are immediately included in the overall value of firms once they start producing. In the data, the stock market includes only publicly traded firms, which means that only some of the firms are not counted and enter with potentially a much longer delay. However, since the value of a firm is proportional to installed capital, and since this is equal for all firms that start producing at a given date, which firms are "included" in the market will only affect the value of these firms relative to overall output and not any of their standard statistics like the price-earnings ratio.

2.2 Households

We model a unit measure of households. The households are of two types ex-ante and there is ex-post heterogeneity within type due to idiosyncratic productivity risk. The types of ex-ante heterogeneous households are based on Chien et al. (2011): a measure $0 < \mu \leq 1$ of households are *advanced traders* and a measure $1-\mu$ of households are *non-participants*. Advanced traders can trade the full set of aggregate-state contingent securities, while the non-participants only trade the risk-free one-period bond. Both types of agents cannot insure against idiosyncratic risk and face borrowing constraints in the form of a zero lower bound on asset holdings.

Each agent experiences one of two idiosyncratic shocks in every period – we denote a realization of the idiosyncratic shock with $\eta_t \in \mathcal{N}$ and a history of idiosyncratic shocks with η^t. We assume that both, the aggregate as well as the idiosyncratic shock, are first-order Markov and evolve independently. We use $\pi_\eta(\eta_{t+1}|\eta_t)$ for the probability to transition into idiosyncratic shock η_{t+1} when currently in idiosyncratic shock η_t. Every period each agent receives a stochastic labor endowment $l(\eta^t) = l(\eta_t)$, which depends on its current idiosyncratic shock. We use A to denote the advanced traders and B to denote the non-participants.[2] We assume that agents supply their labor inelastically at the equilibrium wage.

2.2.1 Preferences

Agents have identical, times-separable, von Neumann-Morgenstern utility functions.

$$U((c_t)_{t=0}^{\infty}) = E_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \tag{23}$$

where instantaneous utility is given by

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma} \tag{24}$$

where we refer to γ as the coefficient of relative risk aversion and to β as the time-preference parameter.

2.2.2 Advanced traders

We denote the asset holding of an agent at node $x^t = (z^t, \eta^t)$ of aggregate-state contingent securities with a payout of one only at node $z^\tau \succ z^t$ by $a^{z^\tau}(x^t)$. In addition to trading in financial markets, advanced traders start up new firms. For each startup, advanced traders have to invest the amount $\bar{i}_t$ and then own a start-up of value $\bar{v}_t$ We assume that the firm will only start producing in the next period and the start-up owners can trade the firm in financial markets. Therefore a start-up opportunity simply means an exogenous cash transfer of value $\bar{v}_t - \bar{i}_t$. This value is mostly positive, but may also become negative, in particular in economic downturns. We assume the $(1-\Gamma)/\Gamma$ startup opportunities to be equally distributed across the advanced traders. Each trader hence starts $(1-\Gamma)/(\Gamma\mu)$ startups, where μ denotes the measure of advanced traders. We denote the cash flow received from a startup opportunity as $\daleth(z^t) = \bar{v}_t - \bar{i}_t$. The Bellman equation for advanced

[2] We use the letter B, because it goes well with A and because type B only trades a **b**ond.

traders is given by

$$V(z^t,\eta_t,a_{t-1}^{z_t}) = \max_{\{a_t^{z_{t+1}}\geq \underline{a}\}_{z_{t+1}\in\mathcal{Z}}} u(c_t)$$
$$+\beta\sum_{\tilde{z}_{t+1}}\sum_{\tilde{\eta}_{t+1}}\pi_z(\tilde{z}_{t+1}|z_t)\pi_\eta(\tilde{\eta}_{t+1}|\eta_t)V((z^t,\tilde{z}_{t+1}),\tilde{\eta}_{t+1},a_t^{\tilde{z}_{t+1}}), \tag{25}$$

where

$$c_t = l^A(\eta_t)\omega(z^t)+\frac{1-\Gamma}{\Gamma\mu}\daleth(z^t)+a_{t-1}^{z_t}-\sum_{z_{t+1}\in\mathcal{Z}}a_t^{z_{t+1}}p^{z_{t+1}}(z^t). \tag{26}$$

and where the price of an aggregate-state contingent bond is given by

$$p^{z_{t+1}}(z^t)=p(z^{t+1}|z^t),\text{ for } z^{t+1}=(z^t,z_{t+1})\succ z^t. \tag{27}$$

2.2.3 Non-participants

For non-participants a sufficient statistic for the agents' problem at node $x^t=(z^t,\eta^t)$ is given by

$$s_t^B := (z^t,\eta_t,b_{t-1}). \tag{28}$$

The Bellman equation of non-participants is given by

$$V(z^t,\eta_t,b_{t-1})=\max_{b_t\geq\underline{b}}u(c_t)+\beta\sum_{\tilde{z}_{t+1}\in\mathcal{Z}}\sum_{\tilde{\eta}_{t+1}\in\mathcal{N}}\pi_z(\tilde{z}_{t+1}|z_t)\pi_\eta(\tilde{\eta}_{t+1}|\eta_t)V((z^t,\tilde{z}_{t+1}),\eta_{t+1},b_t), \tag{29}$$

where

$$c_t = l^B(\eta_t)\omega(z^t)+b_{t-1}-b_tp^b(z^t). \tag{30}$$

2.3 Market clearing

There are $Z+1$ relevant market-clearing conditions in this economy. One for labor and one for each Arrow security. Consumption markets clear by Walras' law, and throughout, we normalize spot prices of consumption to one. Since households supply their labor inelastically and since we do not model population growth, the total amount of labor supplied by the households is given by

$$L_t^{\text{households}} := L_t^A+L_t^B,\text{ where} \tag{31}$$
$$L_t^A := \mu\sum_{\eta_t\in\mathcal{N}}\left(\int_{\underline{a}}^\infty l^A(\eta_t)\rho_t^A(\eta_t,x)\mathrm{d}x\right) \tag{32}$$
$$L_t^B := (1-\mu)\sum_{\eta_t\in\mathcal{N}}\left(\int_{\underline{b}}^\infty l^B(\eta_t)\rho_t^B(\eta_t,x)\mathrm{d}x\right) \tag{33}$$

In our calibration, we choose initial conditions to ensure that $L_t^{\text{households}} = 1$ for all t. Taking an equilibrium wage ω_t as given, the labor demand by the firm is given by

$$L_t^{\text{firm}} := \left(\frac{\omega_t}{(1-\alpha) A_t^{1-\alpha} \xi_t K_t^{\alpha}} \right)^{-\frac{1}{\alpha}}. \tag{34}$$

Labor market-clearing implies $L_t^{\text{firm}} = L_t^{\text{households}}$, hence the market-clearing wage on labor is given by

$$\omega_t = (1-\alpha)\xi_t K_t^{\alpha} L_t^{-\alpha}. \tag{35}$$

The firms are owned by an intermediary that issues Arrow securities that, in each aggregate shock next period, are collateralized by the firm's value. We do not introduce this intermediary formally and simply note that asset markets clear in period t if at each z_{t+1} the total financial wealth in the economy equals the value of the firms. The total financial wealth owned by households is given by

$$w_t^{\text{households}} := w_t^A + w_t^B, \text{ where} \tag{36}$$

$$w_t^A := \mu \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{a}}^{\infty} x \rho_t^A(\eta_t, x) \mathrm{d}x \right) \tag{37}$$

$$w_t^B := (1-\mu) \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{b}}^{\infty} x \rho_t^B(\eta_t, x) \mathrm{d}x \right). \tag{38}$$

Financial markets clear when

$$V_t^{\text{firm}} = w_t^{\text{households}}. \tag{39}$$

Note that V_t^{firm} as well as $w_t^{\text{households}}$, conditional on the aggregate shock $z_t \in \mathcal{Z}$, are predetermined at time $t-1$. Hence, at $t-1$, there are $|\mathcal{Z}|$ market-clearing conditions for the financial markets.

3 Calibration

We distinguish between exogenous parameters that we take from the existing literature or estimate from the appropriate data and endogenous parameters chosen to match the market price of risk and the average real rate in our model to values that are typically considered realistic. We also compare the consumption volatility, the volatility of the real rate, and the average price-earnings ratio in our model to values from the data.

3.1 Exogenous parameters

We take the capital share in the production function, Equation (1), to be $\alpha = 0.33$ and assume that (yearly) depreciation is $\delta = 0.1$. The adjustment cost parameter, ξ^{adj}, is taken to be an endogenous parameter and discussed in the next subsection. In our benchmark model, we set the exit-entry rate of firms to 2.5 percent. As Crane et al. (2022) report, the Census Bureau publishes both firm and establishment exit data through the annual Business Dynamics Statistics (BDS) product. Over recent decades, the employment-weighted firm death rate has been about 2.5 percent (the establishment death rate is much higher at roughly 4.5 percent). We assume that investment in new firms is proportional to K_{t+1} and take as our benchmark the special case $\bar{i}_t = K_{t+1}$, which simplifies the equations for the firm side. The employment share of firms who die is roughly 2.5 percent in the data. At the same time, the employment share of new entrants is reported to be 1 percent in the first year,[3] while Luttmer (2011) reports that firms' employment growth is roughly 1 percent on average. Both of these values suggest that $\bar{i}_t/\hat{K}_{t+1}$ is roughly 50 percent, given the firm death rate. With this, results do not change significantly. In our benchmark calibration, we assume that startups enter at average size, $\bar{i}_t/\hat{K}_{t+1} = 1$. On the balanced growth path, the capital of surviving incumbent firms hence grows at the deterministic growth rate $g^k = g$ and the employment share of new firms is $1 - \Gamma$.

We choose $x^{\text{target}} = \delta + g - 1$ such that no adjustment costs have to be paid if $\frac{K_{t+1}}{K_t} = (1 - \delta + x^{\text{target}}) = g$. We assume a trend growth $g = 2\%$ for labor augmenting productivity.

Following Chien et al. (2011), we take the share of advanced traders to be $\mu = 0.1$. We take the coefficient of relative risk aversion to be 5. This is roughly in line with the estimates in Calvet et al. (2021) and well within the range considered realistic by Mehra and Prescott (1985)[4]. The time preference parameter, β, is an endogenous parameter.

We model shocks to total factor productivity as a 3-state discretized AR(1) process for deviations from the deterministic trend. For the AR(1) process we have

$$\log(\xi_t) = \rho^{\text{tfp}} \log(\xi_{t-1}) + \sigma^{\text{tfp}} \epsilon_t^{\xi}, \tag{40}$$

where $\epsilon_t^{\xi} \sim \mathcal{N}(0, 1)$. We choose the auto-correlation to be $\rho^{\text{tfp}} = 0.8145$, as in Guvenen (2009), and the standard deviation of innovations of $\sigma^{\text{tfp}} = 0.0247$ to match the standard deviation of output growth we measure in the data (2.6%).[5] We discretize the AR(1) process using the method from Rouwenhorst (1995).

Because of our 3-state process, note that a simple bond and stock portfolio will not span the space of aggregate outcomes. As a result, there is a distinct advantage to being able to trade a richer set of assets; as our advanced traders do. The richer set of assets then also insures that the value of the firm is well defined.

In addition to the aggregate shocks, we assume that individuals face two idiosyncratic shocks. For simplicity, we assume that both types face the same idiosyncratic risk. An agent's labor productivity given shock η is given by $l(\eta) = X_\eta$ To match the persistence and standard deviation of the idiosyncratic income process we approximate the process from Storesletten et al. (2004), abstracting from to co-movement of idiosyncratic risk with the aggregate state of the economy.

[3] See BLS online publication: Business Employment Dynamics by Age and Size of Firms: Spotlight on Statistics: U.S. Bureau of Labor Statistics.

[4] They consider values below 10, but as they point out, Arrow (1974) argues that the coefficient should not be much larger than one.

[5] We take the series "A939RX0Q048SBEA" (Real gross domestic product per capita, Chained 2012 Dollars, Quarterly, Seasonally Adjusted Annual Rate) from FRED between 1947 and 2008, compute the quarterly growth rates and aggregate them to a yearly frequency.

Details can be found in Appendix B. We obtain

$$X = \begin{pmatrix} 0.46 \\ 1.54 \end{pmatrix}, \quad \pi^X = \begin{pmatrix} 0.89 & 0.11 \\ 0.11 & 0.89 \end{pmatrix} \tag{41}$$

The resulting cross-sectional standard deviation of log earnings is 0.60 and matches the standard deviation of the process simulated in the Appendix. It is in the ballpark of values typically used in the heterogeneous agents literature (see, *e.g.* Auclert et al., 2019).

3.2 Matching Moments

We have two remaining parameters, which we calibrate inside the model: the adjustment cost parameter ξ^{adj} and the time-preference parameter, β. We choose these parameters to match two key asset pricing facts, namely the low real average interest rate and the high *market price of risk.*

As a benchmark, we match an average yearly interest rate of 0.8%. Depending on estimating the mean with quarterly (1947.2-2008.4) or with yearly (1930-2008) data, Beeler and Campbell (2012) obtain values of 0.89 and 0.56, respectively. Mehra and Prescott (1985) estimate the average real return on a "relatively riskless security"from 1889 to 1978 to be 0.8 %. If we consider the monthly reported 1-year real interest rates[6] for the periods 1991-2016, we also obtain an average of 0.8. However, if we expand the interval to 1991-2022 it falls to 0.5. So clearly, real rates are low and seemingly lower of late.

Our second target is the market price of risk (MPR). Since there is a complete set of aggregate Arrow securities in our model, we can define it within the model as the standard deviation of the stochastic discount factor (that can be constructed from the observed Arrow security prices) normalized by its mean. Formally we have that at a given node z^t, the market price of risk is given by

$$\text{MPR}(z^t) = \frac{\sqrt{\sum_{z_{t+1} \in \mathcal{Z}} \pi(z_{t+1}|z_t)(\frac{p^{z_{t+1}}}{\pi(z_{t+1}|z_t)} - p^b(z^t))^2}}{p^b(z^t)}$$

The market price of risk (also referred to as Hansen-Jagannathan bound, Hansen and Jagannathan (1991)) exceeds the Sharpe ratio attained by any portfolio. Specifically, given the observed Sharpe ratio, the bound tells us that the SDF must be at least just as volatile. We prefer to target the market price of risk rather than targeting the Sharpe ratio directly because the latter depends on the firms' debt policies on which we do not want to take a stand. In their segmented market exchange economy Chien et al. (2011) target moments of equity returns and report an MPR of 45 percent. This is the value we target in our benchmark calibration. However, several studies report a Sharpe ratio in annual US equity returns that lies significantly above 45 percent (see, e.g., Lo (2002) or Lustig and Verdelhan (2012)). In Section 4 below, we explain why, for the exercise there, we want to employ a calibration with a relatively low (but still realistic) MPR. As we show below, we can obtain an MPR of 50 percent (the value cited by Cochrane (2009)) in our economy if we choose different values for the time preference parameter and adjustment costs.

[6] We take the "1-Year Real Interest Rate, Percent, Monthly, Not Seasonally Adjusted" time-series from FRED ("REAINTRATREARAT1YE"). We form simple averages from January 1991 to December 2016 and December 2022 respectively.

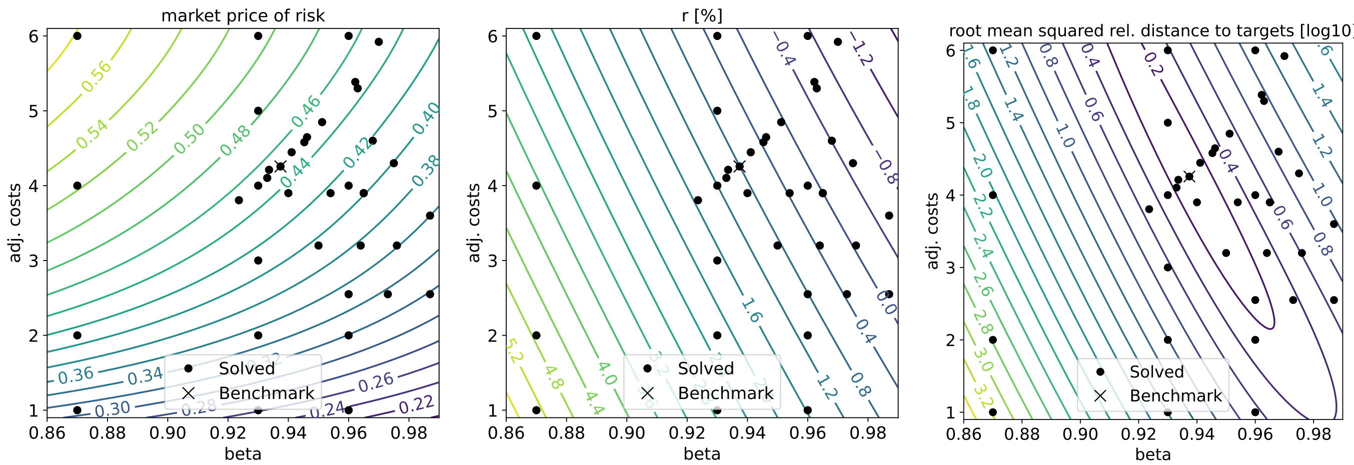

Figure 2: Dependence of the market price of risk (left panel), the interest rate (middle panel), and the root mean squared relative distance from the calibration targets (right panel) on the time-preference parameter β and the adjustment cost parameter ξ^{adj}.

3.2.1 Matching the two targets

For our calibration exercise, we need a mapping from the two free parameters, the time-preference parameter and the adjustment costs, to the moment of interest. It turns out that in our model many mappings from exogenous parameters to endogenous quantities of interest (most importantly the mapping to the mean market price of risk and the mean interest rate, which we use in our calibration exercise) are smooth. Hence we can approximate these mappings very well with a *surrogate model*, which is orders of magnitude cheaper to evaluate (see, *e.g.* Scheidegger and Bilionis, 2019; Catherine et al., 2022). Following Scheidegger and Bilionis (2019) we fit a Gaussian Process to obtain such a surrogate model.

Let r^{target} and $\text{mpr}^{\text{target}}$ denote our targets for the average interest rate and the average market price of risk. Similarly, let $\overline{r}(\beta, \xi^{\text{adj}})$ and $\overline{\text{mpr}}(\beta, \xi^{\text{adj}})$ denote the mean interest rate and the mean market price of risk in our economy with given parameters β and ξ^{adj}. To find parameters that match these two targets, we numerically minimize the average root mean squared error of the model implied moments relative to the targets[7]

$$\ell(\beta, \xi^{\text{adj}}) := \left(\frac{1}{2} \left(\frac{\overline{\text{mpr}}(\beta, \xi^{\text{adj}}) - \text{mpr}^{\text{target}}}{\text{mpr}^{\text{target}}} \right)^2 + \frac{1}{2} \left(\frac{\overline{r}(\beta, \xi^{\text{adj}}) - r^{\text{target}}}{\max\{0.01, r^{\text{target}}\}} \right)^2 \right)^{\frac{1}{2}}. \tag{42}$$

In appendix A, we discuss both the basic computational method to solve for equilibrium as well as the computational method used to match moments.

The third panel in Figure 2 shows the value of our objective function, (42), for different combinations of β and ξ^{adj}. For our Benchmark model, we chose the values which allow us to match both targets precisely. This is the case for $\beta = 0.9347$ and $\xi^{adj} = 4.2565$. As can be seen in the figure, the model matches the targets for the interest rate and the market price of risk (almost) exactly, and our two parameters are well-identified. By construction, our model matches observed

[7]Since our model matches both targets (almost) exactly, the obtained parameters are not sensitive to the specific functional form.

\	Model	Data	Rep. agent
E(r)	0.8 %	0.8 %	17 %
MPR	45 %	45 %	11 %
Vol. output	2.6 %	2.6 %	2.6 %
std r	2.0 %	1.8-2.9 %	2.7 %
Vol. cons.	2.0 %	1.4-2.0 %	2.2 %
log V/E	2.99	2.8-2.99	1.9

Table 1: Key moments in model and data.

output volatility, the observed average risk-free rate, and an MPR of 45 percent.

The first two panels of 2 show how the two targets vary with different values of the parameters. Note that an MPR of 50 percent or higher is still attainable if one chooses higher values for the adjustment cost parameter.

3.2.2 Other moments

Without targeting these moments, we also compare the results in our model to the data for the volatility of the real rate, the volatility of aggregate consumption, and the price-earnings ratio.

Beeler and Campbell (2012) estimate the (annual) volatility of the risk-free rate to be 2.89 in yearly data (1930-2008) and 1.82 in quarterly data (1947.2-2008.4). We find the volatility of nondurable consumption to be 2.0% in the data.[8] The average of the log price-earnings ratio in the data lies somewhere between 2.8 and 2.99 depending on whether we consider the "Cyclically Adjusted Price Earnings Ratio P/E10" or if the "Cyclically Adjusted Total Return Price Earnings Ratio P/E10".[9] Table 1 summarizes key statistics from our benchmark calibration of the model and compares them to the data and a representative agent specification.[10]

The first two targets are met almost exactly by construction. The volatility of the growth rates of real consumption per capita is 2 percent which appears very reasonable, the volatility of the real interest rate and the average (log) value earnings ratios are all well in the range of what can be considered realistic. Our model performs well in all these dimensions. If we were to target a significantly higher MPR (say 60 %) or higher, a higher risk aversion would be needed.

In comparison, the representative agent model (with the same preference parameters and production function) fails to reproduce average interest rates and the MPR. This is not a surprise (it has been pointed out many times in the literature, e.g. Weil (1989)) but the numbers are provided for comparison.

It is worth noting that the high market price of risk in our model is generated by the same

[8]We take the series "A796RX0Q048SBEA " (Real personal consumption expenditures per capita: Services: Nondurable goods, Chained 2012 Dollars, Quarterly, Seasonally Adjusted Annual Rate) and the sum of "A796RX0Q048SBEA " and "A797RX0Q048SBEA " (Real personal consumption expenditures per capita: Goods: Nondurable goods, Chained 2012 Dollars, Quarterly, Seasonally Adjusted Annual Rate) from FRED between 1947 and 2008, compute the quarterly growth rates and aggregate them to a yearly frequency.

[9]We divide the real price from Shiller's monthly stock market data (the "U.S. Stock Markets 1871-Present and CAPE Ratio" available at http://www.econ.yale.edu/∼shiller/data/ie_data.xls, last accessed: January 2023) by the real earnings, take logs, and compute the average between 1947 and 2008.

[10]The representative agent model is, for comparison, solved with identical parameters but without firm-exit.

mechanism already described in Chien et al. (2011). By construction, the entire aggregate risk in asset returns is held by the advanced traders. Since these only constitute 10 % of all agents in the economy their individual consumption will be very volatile. In our benchmark calibration, the consumption of the group of advanced traders is roughly four times more volatile than aggregate consumption. The assets are priced off the individual consumption of unconstrained advanced traders, which generates the high market price of risk in our model. The advanced traders, on average, hold more wealth and attain a higher consumption level.

3.3 Mapping parameters to moments

As mentioned above, the first two panels in figure 2 show how the market price of risk, the interest rate, and the root mean squared relative error vary with the parameters β and ξ^{adj}. High adjustment costs are a crucial part of our model's success. If adjustment costs are only 2 percent, there is no way to get a sizable market price of risk.

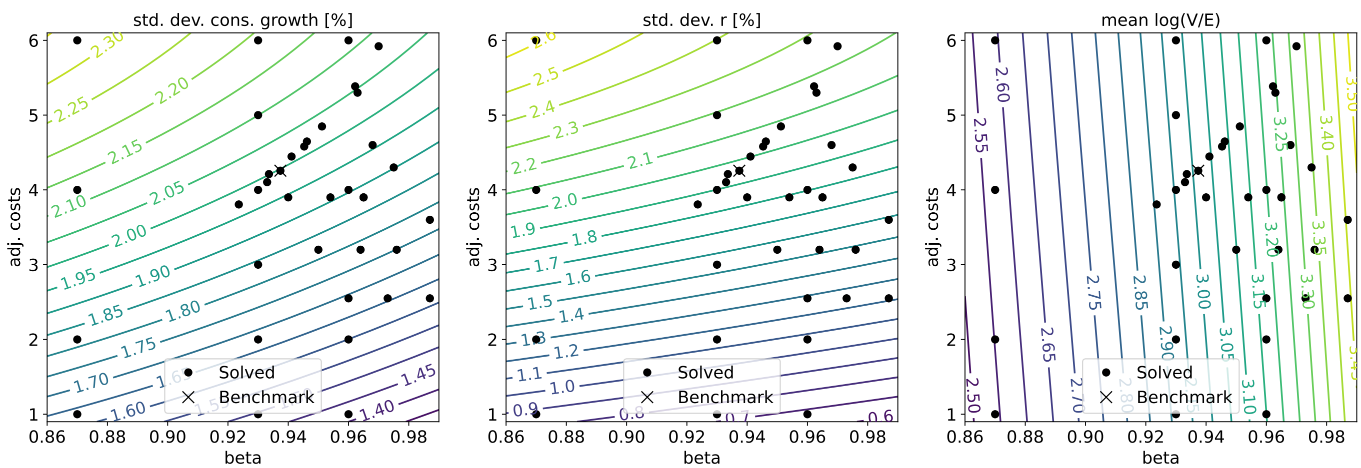

Figure 3: Standard deviation of consumption growth (left), standard deviation of the interest rate (middle), and the price-earnings ratio (right) for different values of the time-preference parameter (vertical axis), and the adjustment costs parameter (horizontal axis).

Figure 3 shows how the volatility of consumption growth and the volatility of interest rates, as well as the price-earnings ratios, depend on the time preference parameter and the adjustment costs. As one would expect, the time preference parameter has a large impact on the real rate and, therefore, on price-earnings ratios. Adjustment costs are the main determinant of the volatilities of both consumption growth and the risk-free rate. A higher MPR would imply higher consumption volatility.

As we show in Section 4.1 below, our model setup allows us to match a wide range of possible interest rates while keeping the MPR constant and 45 percent. For this, three modeling ingredients play an important role, idiosyncratic income risk, borrowing constraints, and firms' death. Ever since Aiyagari (1994) the first two are standard in dynamic models with heterogeneous agents. To the best of our knowledge, our model is the first to allow for firms' death in this setting.

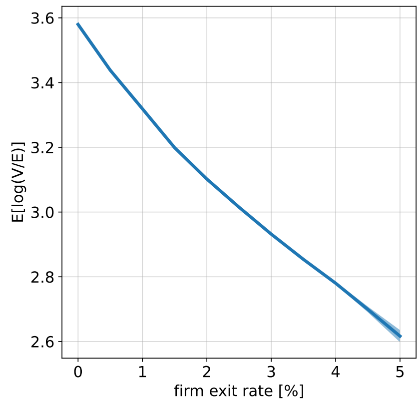

Figure 4: Mean log price-earnings ratio for different values of the firm exit rate. For a given firm exit rate, the adjustment costs and the patience are recalibrated to maintain an average market price of risk of 45% and an average interest rate of 0.8%. The shaded area shows the confidence interval of the prediction by the surrogate model.

3.4 The role of firms' exit

As explained in the introduction, one of the main innovations in our model is the feature that firms die and are replaced by new firms that are initially financed by households. The rate of firms exiting plays an important role in a realistic price-earnings (as well as price-dividend) ratio. In a low risk-free rate environment the current value of a discounted infinite stream of future payments can obviously be very large (we discuss this point in detail in the next section).

We alternatively solve the model for different firm exit rates and a variety of combinations for the adjustment costs and the patience to obtain a mapping from the exit rate, the adjustment costs and the patience to the moments of interest (see appendix A.2 for details). Figure 4 shows, based on the surrogate model, how the mean price-earnings ratio in our model depends on the firm exit rate. For each firm exit rate, we recalibrate the adjustment costs and the patience to maintain a market price of risk of 45% and a mean interest rate of 0.8%.

As explained above, the mean (log) V/E ratio for historical US data lies between 2.8 and 2.99. Since we are considering a log-scale a price-earnings ratio of 3.3 or higher is clearly at odds with the data. As figure 4 shows, the price-earnings ratio in our model without firm death, or with a death-rate below 1 % would be substantially too high. In order to examine low average interest rates firm exit must be a crucial ingredient of the model.

Note that when the interest rate is very low, dividends paid out in the far future greatly impact prices today. In the next section, we consider average interest rates as low as -1 %. For these calibrations, it takes more than 700 periods for the firm's value to converge, implying that the value of long-lived assets is crucially dependent on their payoffs 200 to 600 years in the future. This seems clearly counterfactual.

4 Low real rates and debt roll-over

As we noted earlier, the positive average gap between the U.S. growth rate and the real interest rate on U.S. Treasuries has sparked a large literature on the question of whether an infinite roll-over of government debt is possible and whether deficit finance has no fiscal cost (see Blanchard (2019), Mian et al. (2021) or Kocherlakota (2022), Bloise and Reichlin (2022)). Our model is ideally suited to answer this question because it produces realistic asset prices in a production economy and allows parametrizations of the model where infinite rollover is possible.[11]

In a deterministic model, debt rollover is possible if and only if the real risk-free rate, r, is smaller than (or equal to) the growth rate, g. It follows from Kocherlakota (2022) and Bloise and Reichlin (2022) that under uncertainty, debt rollover might be possible if $r > g$, or it might be impossible even if $r < g$. This is easy to understand if one considers a simple non-parametric setting (very similar to the setting in Kocherlakota (2022)). Suppose asset prices follow an S-state Markov chain with transition π, which is assumed to have a unique stationary distribution π^*. We define an $S \times S$ matrix Q of Arrow-prices, $q_{i,j}$ denoting the Arrow-price for state j in state i. The interest rate in state s is then $\frac{1}{R_s} = \sum_{i=1}^{S} q_{s,i}$, the average interest rate is $\bar{R} = \sum_{s=1}^{S} \pi_s^* R_s$. We abstract from growth. It follows from Aiyagari and Peled (1991) (see also Kocherlakota (2022), Proposition 1) that debt rollover is impossible if and only if the largest eigenvalue of the matrix Q is less than 1. It is easy to see that prices of infinite payments that are bounded away from zero are finite under the exact same condition. As $n \to \infty$ the series

$$(I + Q + \ldots + Q^n) \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$$

converges if and only if the largest eigenvalue is less than 1. We refer to this as *summable prices*, and, in what follows, will refer to no rollover and summable prices exchangeably. For debt rollover to be possible, there cannot be assets traded that pay off a strictly positive amount for the infinite future. In our economic model, households' future labor income cannot be traded because households face borrowing constraints. However, firms are claims to an infinite stream of future dividends. When debt rollover is possible, the value of firms can be finite only if dividend payments are negative in some states or if firms die at a fast enough rate. As explained above, in our economic model, we chose to introduce firm death.

Requirements on the largest eigenvalue of Q impose no restrictions on the interest rates R_s, $s = 1, \ldots, S$, except that (i) for debt rollover to be possible, there must be one state s for which $R_s < 1$ and ii) for debt rollover to be impossible there must be a state s' for which $R_{s'} > 1$. Obviously, these assumptions impose no restrictions on the average risk-free rate, $\bar{R}$. To see why this is the case, define a matrix

$$Q_\epsilon = \begin{pmatrix} \frac{1}{R_1} & \epsilon & \ldots & \epsilon \\ \frac{1}{R_2} & \epsilon & \ldots & \epsilon \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{R_S} & \epsilon & \ldots & \epsilon \end{pmatrix}$$

It is easy to see that for $\epsilon = 0$, the largest eigenvalue of this matrix is $\frac{1}{R_1} < 1$ and hence has nothing to do with $\bar{R}$. Moreover, for sufficiently small $\epsilon > 0$, the implicit function theorem can be applied, and it follows that the largest eigenvalue varies smoothly with ϵ.

[11] We do not address the perhaps more important question of whether government debt can be Pareto-improving in such a setting (see, e.g. Amol and Luttmer (2022) or Brumm et al. (2021))

β	ξ^{adj}	E(r)	MPR	Vol. output	std r	Vol. cons.	E(log V/E)	E(p^{console})
0.9700	5.921	-0.94 %	44.9 %	2.6 %	2.3 %	2.1 %	3.32	637
0.9622	5.385	-0.69 %	44.5 %	2.6 %	2.2 %	2.1 %	3.26	266
0.9512	4.848	0.03 %	44.9 %	2.6 %	2.1 %	2.0 %	3.14	138
0.9453	4.580	0.49 %	45.2 %	2.6 %	2.0 %	2.0 %	3.07	107
0.9374	4.257	0.80 %	45.0 %	2.6 %	2.0 %	2.0 %	3.01	84
0.9330	4.103	0.95 %	44.9 %	2.6 %	1.9 %	2.0 %	2.98	75
0.9236	3.805	1.57 %	45.4 %	2.6 %	1.8 %	2.0 %	2.90	59
0.978	2.663	0.0 %	33.0 %	2.6 %	1.3 %	1.6 %	3.36	∞

Table 2: Parameters and resulting moments in the model with different mean interest rates.

However, once one imposes an equilibrium model with agents that maximize an expected utility function, the above construction imposes strong assumptions on the movements of individuals' consumption. If debt rollover is impossible and in state 1 the interest rate is high, $1/R_1 < 1$, it must be the case that individual consumption in the next period drops sharply if the economy does not stay in the same state. In all other states, individual consumption must increases when moving to the high-interest state in the next period. This narrative fits with our model if the market price of risk is sufficiently high. Since shocks are temporary, the interest rate is high in the state where aggregate consumption is high, and with enough variability, consumption falls steeply when moving to worse states.

It is then a quantitative question of how low interest rates can become with prices still being summable. Next, we use our equilibrium model to investigate this question. We show that a gap between the growth rate and the average interest rate of 3 percent might still not allow for debt rollover if one keeps the MPR at 45 percent.

4.1 Summable Prices in Our Model

We check numerically whether our computed equilibrium prices are summable. Although there are only three exogenous shocks, state prices obviously take on infinitely many values since they depend on the (endogenously changing) wealth distribution. While the largest eigenvalue of the matrix of average state prices typically gives a good indication of whether prices are summable, we actually compute the price of a (fictitious) consol (an asset that pays a $(1+g)^t$ risk-free at any time t in the future). Note that there are always cases where it is not decidable whether prices are summable - one cannot insure the convergence or divergence of an infinite sum by checking finitely many terms. However, when we report below that prices are summable (or not) prices converge (or diverge) sufficiently fast to give us high confidence in our results. Table 2 shows the resulting moments in our model for different interest rates, keeping the market price of risk at 45%. In the last column, E(p^{console}) denotes the expected price for a console with payout g^t in every period. To attain a higher mean interest rate while keeping the targeted market price of risk, as well as all other parameters, constant, our model requires lower patience and adjustment costs parameters. A higher mean interest rate leads to a lower volatility of the interest rate as well as a lower price-earnings ratio.

We first note that for our benchmark calibration, prices are summable (and therefore, debt rollover is impossible). As explained above, an average interest rate of 0.8 % matches US averages well. However, it is instructive to consider a much lower rate as well. We, therefore, recalibrate the model to match an average interest rate of 5 %, 0 %, and -1 %. The latter certainly seems to be a lower bound given the available data (although there have been historical episodes where the return on government bonds lied below -1 percent, this is certainly not a realistic average return.). As Table 2 shows, as long as the MPR is held fixed at 45 %, prices remain summable (i.e. debt rollover remains impossible) even when average interest rates are 2.7 percent below the average growth rate. This is consistent with our non-parametric toy model in the previous section, the volatility of interest rates is sufficiently high to ensure that they are above the growth rate for a substantial fraction of time.

A realistically[12] high MPR turns out to be crucial for this result. For example, with a coefficient of relative risk aversion of 5, $\beta = 0.978$, and a relatively low adjustment cost parameter of 2.663, we obtain a market price of risk of 34.4 %. The average real rate is zero and its standard deviation is 1.3 %. In this economy, prices are not summable, and therefore debt rollover is possible.

To understand better the interacting effects of the average interest rate and the market price on risk we construct a surrogate model that maps average interest rates and the average market price of risk into the price of the consol. We fix all parameters except for the patience and the adjustment costs at fixed values; hence other moments, such as the standard deviation of the risk-free rate, also vary as the market price of risk varies. Other parameters might influence the consol price that we do not consider here.

Figure 5 shows the price of the console for different combinations of the interest rate and the MPR. For lower interest rates computing the consol's price becomes numerically difficult before a clear divergence of the price can be observed. We hence mark the area in which we cannot observe either clear divergence or convergence of the price with orange. Further, it should be noted that the value of the mean price of the console, which we obtain, comes with an error due to Monte Carlo simulations.

Infinite debt rollover is possible if and only if the value of the console is infinite. We can see that in regions where prices are summable, the average price of the consol is increasing faster than exponentially as the mean interest rate decreases. To see this, observe that the distance between the contour lines is decreasing, despite the figure showing the log of the average price of the console. A higher market price of risk lowers the price of the console and hence allows for summable prices with lower interest rates.

In summary, as it has been pointed out (e.g. Kocherlakota (2022) or Bloise and Reichlin (2022)) average low interest rates are neither necessary nor sufficient for the possibility of debt rollover. In our calibrated model, we find that debt rollover is impossible even if average interest rates are more than 2.5 percent below the average growth rate. A realistically high MPR is crucial for this result. In our calibrations, we chose to match an MPR that is rather on the low side, given the estimates in the literature.

[12]As mentioned above, we chose our calibration target for the MPR to be on the low side relative to estimates in the literate.

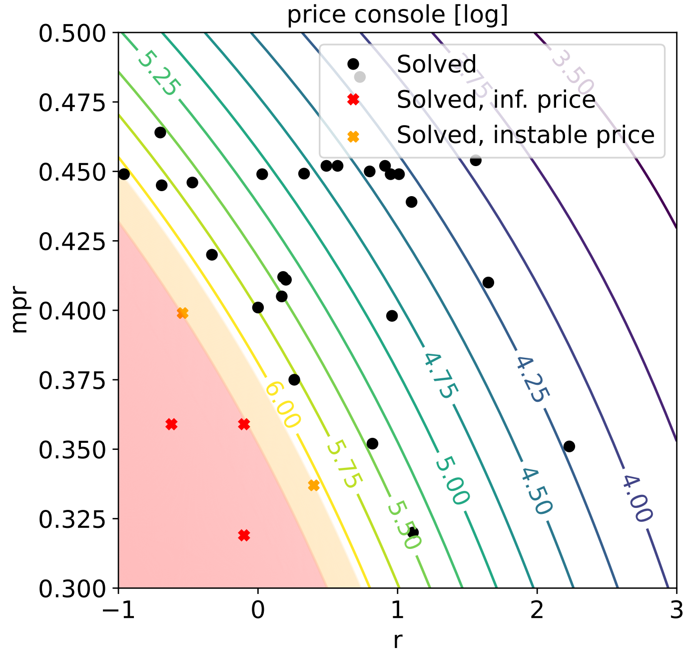

Figure 5: Mean price of a console, with a payout equal to $(1+g)^t$ in logs, plotted against the mean interest rate and the mean market price of risk.

5 Conclusion

We make three distinct contributions to the literature. Firstly, we introduce a tractable model with heterogeneous trading technologies, idiosyncratic and aggregate risk in a production economy, which produces asset pricing and business cycle statistics that are in line with the data. Secondly, we move the literature on limited stock market participation forward by relaxing the commonly used restriction that advanced traders can only choose from a pure bond and aggregate stock portfolio.[13] We show that relaxing this restriction has significant implications regarding the exposure of households' wealth to specific aggregate shocks. Lastly, we contribute to the development of computational methods by laying out how this model can be solved efficiently. Combining the ideas in Carroll (2006); Young (2010); Kubler and Scheidegger (2018), we efficiently solve the model exclusively on simulated aggregate states, which greatly ameliorates the problems resulting from extrapolation in higher dimensions. We show that, in addition to the discrete aggregate shocks, aggregate capital together with the wealth-share of advanced traders provide a good and tractable approximation of the aggregate state along the lines of Krusell and Smith (1998).

[13] *e.g.* Guo (2004), Guvenen (2009) and Favilukis (2013).

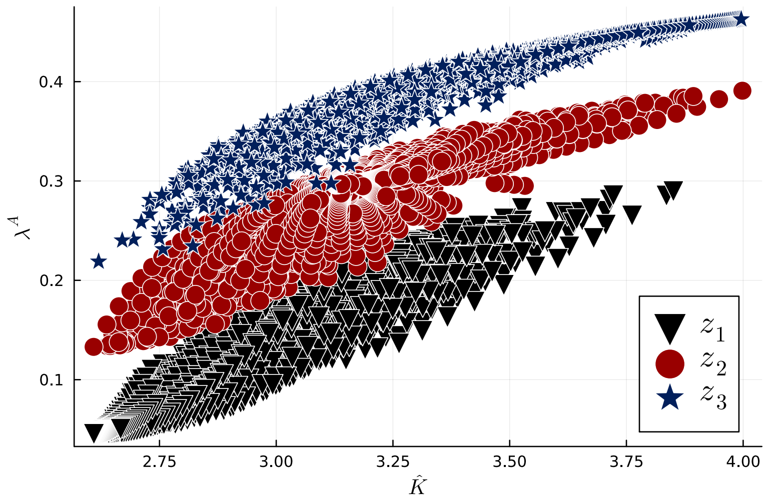

Figure 6: Realizations of the two endogenous aggregate variables, capital (horizontal axis) and the wealth share (vertical axis) of the advanced traders, for each of the exogenous aggregate shocks on a simulated path of 10000 periods for the benchmark model.

Appendix

A Computational Method

A.1 Solution Method

Our computational method is broadly based on the method developed by Krusell and Smith (1998) and more specifically on the algorithm developed by Kubler and Scheidegger (2018). We extend the classical method by Krusell and Smith (1998) in two important dimensions that allow us to solve the model efficiently and exclusively on simulated path of the economy. While we describe the method in the context of our model, it is generic and can be viewed as an extension of the general method in Krusell and Smith (1998).

More specifically, we solve all three problems, the household problem, the firm problem, and market-clearing simultaneously and only on simulated paths of the aggregate state-space. Since endogenous aggregate variables, in our case, aggregate capital and the wealth-share of the advanced traders, are usually very correlated a solution method requiring the household problem to be solved on a hypercubic domain, such as for the classical Krusell and Smith (1998) method or any approximation method based on a hypercubic domain,[14] would be problematic for two main reasons, which we illustrate in figure 6. Figure 6 shows realizations of the two endogenous aggregate variables, capital (horizontal axis) and the wealth share (vertical axis) of the advanced traders, for each of the exogenous aggregate shocks. The first problem is that the computational effort spent to solve the model for capital wealth-share combinations in the upper left and lower right corners would be completely wasted, since the model never reaches those states. The second problem is that, since the correct shock-contingent transitions for the endogenous aggregate variables are only observed on the simulated path, there is no obvious way how to extend a fitted forecasting

[14] e.g. adaptive sparse grids as in Brumm and Scheidegger (2017).

rule to an hypercubic domain. Fernández-Villaverde et al. (2019) illustrate the severity of this challenge and mitigate it by using neural networks to fit the forecasting rules and extrapolate to an hypercubic domain. Since our method allows us to solve the model exclusively on simulated paths, we circumvent the problem completely.

The second main feature of our method is that we extend the endogenous grid method by Carroll (2006) to be applicable in our setting. For a given aggregate state and a given exogenous idiosyncratic shock, we approximate the household consumption policies, which are all we need from the household side, as a piecewise linear function of the households asset holdings, rendering it very flexible and allowing us to take advantage of efficient interpolation schemes (see *e.g.* Druedahl (2021)). For a given value of asset holdings, which lies on a fixed asset grid, and given exogenous shocks (aggregate and idiosyncratic), we approximate the households' consumption functions as polynomials in capital and the wealth share of the advanced traders. Crucially. we fit a separate polynomial for each grid-value of the asset holdings and each combination of exogenous shocks, rendering our functional form very flexible and able to provide an excellent fit, despite the strong nonlinearities in our model.

A.1.1 Solution Algorithm

The exact aggregate state of the economy is given by its exogenous aggregate shock z_t and the joint distribution of households across types, idiosyncratic shocks, and asset holdings. In the spirit of Krusell and Smith (1998), we find a lower dimensional sufficient statistic to approximate the dependence of aggregate quantities on the wealth distribution. The two quantities are aggregate capital K_t, as in Krusell and Smith (1998), as well as the wealth share held by advanced traders, which we denote with λ_t^A. Aggregate capital is obviously important because, together with the exogenous aggregate shock, it characterizes the wages in the economy. The wealth distribution in our model moves a lot. The majority of households can only hold the bond while the small share of advanced traders hence must hold all the financial risk. As a result, the share of financial wealth held by the group of advanced traders will increase substantially following a good aggregate shock and decrease following a bad aggregate shock. Therefore aggregate capital alone would not be enough to forecast prices. We find that adding the wealth share held by Arrow traders to aggregate capital, provides for a very good sufficient statistic to predict endogenous aggregate quantities as well as their evolution, contingent on the exogenous shocks.

Our method is simulation-based. In every simulation step, we jointly solve for the households' policies, the firm's policy, and market clearing prices. Then, we draw a new aggregate shock and simulate the economy one period forward. To approximate the wealth distribution across trader types and idiosyncratic shocks, we use the non-stochastic simulation method developed in Young (2010). After collecting a sequence of simulated states, we use the computed prices and policies to update our approximating functions.

Set of functions we need to approximate We need to approximate five types of functions. First, we approximate the (aggregated) firm's policy function which determines the next period's capital as a function of the approximate aggregate state. Second, we approximate the (aggregated) firm's value function which gives the firm's value as a function of the approximate aggregate state. Third, we approximate a forecasting function to approximate the wealth share held by Arrow traders in the next period conditional on the next period's exogenous aggregate shock. The remaining two functions approximate the consumption functions of the advanced traders and non-participants as a function of the approximate aggregate state, the idiosyncratic shock, and the idiosyncratic asset holding. We denote our approximations with $g_x(\cdot)$, where the index x denotes what is approximated. The approximations we need are

$$g_K\left(z_t, \hat{K}_t, \lambda_t^A\right) \approx \hat{K}_{t+1} \tag{43}$$

$$g_{V^{\text{firm}}}\left(z_t, \hat{K}_t, \lambda_t^A\right) \approx \hat{V}_t^{\text{firm}} \tag{44}$$

$$g_{\lambda^A}\left(z_t, \hat{K}_t, \lambda_t^A, z_{t+1}\right) \approx \lambda_{t+1|z_{t+1}}^A \tag{45}$$

$$g_{c^A}\left(z_t, \hat{K}_t, \lambda_t^A, \eta_t, \frac{\hat{a}_{t-1}^{z_t}}{g}\right) \approx \hat{c}_t^A \tag{46}$$

$$g_{c^B}\left(z_t, \hat{K}_t, \lambda_t^A, \eta_t, \frac{\hat{b}_{t-1}}{g}\right) \approx \hat{c}_t^B. \tag{47}$$

Simulating the economy while solving for prices, policies, and values Given the set of approximating functions, we simulate the economy for T periods. We track the wealth distribution following the non-stochastic simulation method by Young (2010). At time step t we solve a system of four nonlinear equations for the prices of the three Arrow securities and the next period's capital. The four equations are the market-clearing conditions for each of the Arrow securities as well as the firm's first-order condition. Simultaneously, we obtain the households' policies implied by the prices and the transition of the aggregate summary statistics. Given a guess for prices and the next period's capital, we evaluate the equations the following way.

The firm's Euler equation is satisfied if $\hat{K}_{t+1}$ fulfills the aggregated version of equation (10). Where z_t, $\hat{K}_t$, and λ_t^A are given and $\hat{K}_{t+2}$ is obtained from the approximating function $\hat{K}_{t+2} = g_K\left(z_{t+1}, K_{t+1}, \lambda_{t+1}^{A\text{ forecast}}\right)$. The wealth share of advanced traders in period $t+1$ is obtained by using the forecasting function $\lambda_{t+1}^{A\text{ forecast}} = g_{\lambda^A}\left(z_t, K_t, \lambda_t^A, z_{t+1}\right)$.

To evaluate whether the market clearing conditions are satisfied, we need to know the financial wealth in the economy, *i.e.* the value of producing firms, which has to equal the asset demand by households. To obtain the value of operating firms in period $t+1$, we use the approximating function $\hat{V}_{t+1}^{\text{firm forecast}} = g_{V^{\text{firm}}}\left(z_{t+1}, \hat{K}_{t+1}, \lambda_{t+1}^{A\text{ forecast}}\right)$. To obtain the asset demand of households, we use the endogenous grid method by Carroll (2006). Despite its speed and its ability to exploit very efficient interpolation methods, a further advantage of using the endogenous grid method is that we only need to evaluate the approximating functions g_{c^A} and g_{c^B} on a pre-specified asset grid, corresponding to pre-specified asset positions in period $t+1$. Hence we can achieve a highly flexible functional form for function approximation by fitting several separate functions $g_{c^A}^{z_t,\eta_t,\hat{a}_{t-1}^{z_t}/g}\left(\hat{K}_t, \lambda_t^A\right)$ for each combination of the discretized exogenous shocks z_t and η_t, and each asset holding on the pre-specified grid $\hat{a}_{t-1}^{z_t}/g \in \mathcal{A}^{\text{grid}}$, where $\mathcal{A}^{\text{grid}}$ denote the set of points on the pre-specified asset grid for advanced traders. Analogously $\mathcal{B}^{\text{grid}}$ denotes the set of points on the pre-specified asset grid for bond traders. Each of the $g_{c^A}^{z_t,\eta_t,\hat{a}_{t-1}^{z_t}/g}\left(\hat{K}_t, \lambda_t^A\right)$, $g_{c^B}^{z_t,\eta_t,\hat{b}_{t-1}^{z_t}/g}\left(\hat{K}_t, \lambda_t^A\right)$, only have to capture the variation of consumption with aggregate capital and the wealth-share owned by advanced traders for fixed exogenous shocks and a fixed amount of financial wealth.

We choose the same grids for our histogram to track the wealth distribution by trader type and idiosyncratic shock. We denote the vector of masses of bond traders with idiosyncratic shock η and asset holdings $\hat{b}^{z_t}_{t-1}/g$ on the asset grid $\mathcal{B}^{\text{grid}}$ with $\mu^{B,\eta}_t$ and we denote masses of advanced traders with idiosyncratic shock η and asset holdings $\hat{a}^{z_t}_{t-1}/g$ on the asset grid $\mathcal{A}^{\text{grid}}$ with $\mu^{A,\eta}_t$. We separately compute the asset demand by bond and by advanced traders. In order to obtain the bond demand by bond traders, we follow the following steps:[15]

1. For all possible z_{t+1} and η_{t+1} we obtain the forecasted consumption of bond traders for a pre-specified grid of asset holdings $\hat{b}^{z_t}_{t+1} \in \mathcal{B}^{\text{grid}}$

$$\hat{c}^{B,z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}_{t+1} = g^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}_{c^B}\left(\hat{K}_{t+1}, \lambda^{A,\text{ forecast}}_{t+1}\right) \tag{48}$$

2. We obtain the associated marginal utility of purchasing the bond in period t for given exogenous shocks in $t+1$ and asset holdings

$$(V^{B\prime})^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}_{t+1} = u'(\hat{c}^{B,z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}})u'(g) \tag{49}$$

3. For all idiosyncratic shocks in period t, we compute the expected marginal utility from purchasing the bond when the asset holdings are on the grid $\mathcal{B}^{\text{grid}}$.

$$(EV^{B\prime})^{\eta_t,\mathcal{B}^{\text{grid}}}_t = \sum_{z_{t+1},\eta_{t+1}} \pi^z(z_t, z_{t+1})\pi^\eta(\eta_t, \eta_{t+1})(V^{B\prime})^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}_{t+1} \tag{50}$$

4. Compute the consumption consistent with the households Euler equation and the households savings choice lying on the grid $\mathcal{B}^{\text{grid}}$.

$$\mathcal{C}^{\eta_t,\mathcal{B}^{\text{grid}}}_t = \left(u'\right)^{-1}\left(\frac{\beta}{p^{\text{bond}}_t}(EV^{B\prime})^{\eta_t,\mathcal{B}^{\text{grid}}}_t\right) \tag{51}$$

5. The obtained tuple, $\left(\mathcal{C}^{\eta_t,\mathcal{B}^{\text{grid}}}_t, \mathcal{B}^{\text{grid}}\right)$ provides us with a mapping from consumption to the optimal (unconstrained) policy. In order to simulate the economy forward, however, we would like to know the households' policies when the period t asset holdings, not the period t asset choice, lies on the pre-specified grid, we choose the same as the grid for the asset histogram. Using the households budget constraint, we back out the asset holdings at the beginning of period t, which would be consistent with the consumption and savings policies $(\hat{c}^{B,\eta_t}_t, \hat{b}^{\eta_t}_{t+1}) \in \left(\mathcal{C}^{\eta_t,\mathcal{B}^{\text{grid}}}_t, \mathcal{B}^{\text{grid}}\right)$

$$\hat{b}^{\eta_t}_t = \hat{c}^{B,\eta_t}_t + p^{\text{bond}}_t \hat{b}^{\eta_t}_{t+1} g - \hat{w}_t \eta_t \tag{52}$$

   The resulting tuple $\left(\hat{b}^{\eta_t}_t, \hat{b}^{\eta_t}_{t+1} \in \mathcal{B}^{\text{grid}}\right)$ maps period t asset holdings to the savings choice, which lies exactly on the pre-specified asset grid.

6. We obtain the unconstrained savings choice for period t asset holdings lying on the pre-

[15]The steps exactly the canonical application of the endogenous grid method, repeated here for convenience.

specified grid with piecewise linear interpolation. We then ensure the borrowing constrained are satisfied and obtain our new mapping $\left(\hat{b}_t^{\eta_t} \in \mathcal{B}^{\text{grid}}, \hat{b}_{t+1}^{\text{new},\eta_t}\right)$.

7. Using the computed savings policy we compute an updated consumption policy $\left(\hat{b}_t^{\eta_t} \in \mathcal{B}^{\text{grid}}, \hat{c}_t^{B,\text{new},\eta_t}\right)$.
8. The asset demand by bond traders is then given by

$$d_t^B = \sum_{\eta_t} \sum_{\hat{b}_t^{\eta_t} \in \mathcal{B}^{\text{grid}}} \mu^{B,\eta_t}(\hat{b}_t^{\eta_t}) \cdot \hat{b}_{t+1}^{\text{new},\eta_t}(\hat{b}_t^{\eta_t}) g \tag{53}$$

The procedure to obtain the advanced traders' asset demand is similar, with some modification because they can purchase the aggregate-state contingent securities, each associated with its own Euler equation, as we outline below.

1. For all possible z_{t+1} and η_{t+1} we obtain the forecasted consumption of bond traders for a pre-specified grid of asset holdings $\hat{a}_{t+1}^{z_{t+1}} \in \mathcal{A}^{\text{grid}}$

$$\hat{c}_{t+1}^{A,z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} = g_{c^A}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} \left(\hat{K}_{t+1}, \lambda_{t+1}^{A,\text{ forecast}}\right) \tag{54}$$

2. We obtain the associated marginal utility of purchasing each of the aggregate-state contingent Arrow securities in period t for given exogenous shocks in $t+1$ and asset holdings

$$(V^{A'})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} = u'(\hat{c}^{A,z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}})u'(g) \tag{55}$$

3. For all idiosyncratic shocks in period t, and all possible shocks z_{t+1}, we compute the expected marginal utility from purchasing the corresponding Arrow security[16] when the asset holdings are on the grid $\mathcal{A}^{\text{grid}}$. In difference to the problem for the bond traders, the sum goes only over idiosyncratic shocks in time $t+1$, and we are computing separate terms for each of the aggregate shocks z_{t+1}.

$$(EV^{A'})_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}} = \sum_{\eta_{t+1}} \pi^{\eta}(\eta_t, \eta_{t+1})(V^{A'})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} \tag{56}$$

4. Compute the consumption consistent with the households Euler equation and the households savings choice for each of the Arrow securities lying on the grid $\mathcal{A}^{\text{grid}}$. In difference to the bond traders, the corresponding consumption now depends on the aggregate shock z_{t+1}.

$$\mathcal{C}_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}} = (u')^{-1}\left(\frac{\beta}{p_t^{z_{t+1}}}(EV^{A'})_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}}\right) \tag{57}$$

5. We now obtained a tuple, $\left(\mathcal{C}_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}}, \mathcal{A}^{\text{grid}}\right)$ provides us with a mapping from consumption to the optimal (unconstrained) policy for each of the aggregate state-contingent securities. Since the asset holdings at $t+1$ lie on the pre-defined grid for each of the Arrow securities, they are not consistent with each other and we hence can't straightforwardly use the budget-

[16] The marginal utility from purchasing the other Arrow securities, available for trade at time t, is zero once the aggregate shock z_{t+1} has realized.

constrained to back out the time t asset holdings, which would be consistent with those choices.

Instead, we choose a larger and denser consumption grid, which we denote by $\hat{c}_t^{\text{common}}$. We then interpolate the mappings $\left(\mathcal{C}_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}}, \mathcal{A}^{\text{grid}}\right)$ and apply the borrowing constraint to obtain mappings $(\hat{c}_t^{\text{common}}, \hat{a}_{t+1}^{z_{t+1},\text{common},\eta_t})$ for each asset (*i.e.* each z_{t+1} and each idiosyncratic shock η_t). We then obtain a mapping from the common consumption grid to total savings, $(\hat{c}_t^{\text{common}}, \hat{s}_t^{\text{common},\eta_t})$, for each idiosyncratic shock η_t, where

$$\hat{s}_t^{\text{common},\eta_t} = \sum_{z_{t+1}} p_t^{z_{t+1}} \hat{a}_{t+1}^{z_{t+1},\text{common},\eta_t} g. \tag{58}$$

With that savings function, we can now use the budget constraint to compute the asset holdings consistent with the consumption choice $\hat{c}_t^{\text{common}}$.

$$\hat{a}_t^{\text{common},\eta_t} = \hat{c}_t^{\text{common}} + \hat{s}_t^{\text{common},\eta_t} - \hat{w}_t\eta_t - \frac{1-\Gamma}{\Gamma\mu}\hat{\beth}_t, \tag{59}$$

where $\hat{\beth}_t$ denotes the payoff of creating a startup and the $\frac{1-\Gamma}{\Gamma\mu}$ denotes the startups per advanced trader.

As a result, we have two maps: the tuple $(\hat{a}_t^{\text{common},\eta_t}, \hat{c}_t^{\text{common}})$ maps asset holdings at the beginning of period t to period t consumption for each idiosyncratic shock in period t. The previously obtained tuples $\left(\mathcal{C}_t^{\eta_t,z_{t+1},\hat{a}_{t+1}^{z_{t+1}}\in\mathcal{A}^{\text{grid}}}, \mathcal{A}^{\text{grid}}\right)$ map period t consumption to the (unconstrained) asset choice for each idiosyncratic shock and each of the three assets.

Next, we use those mapping to obtain period t consumption and portfolio choice for the beginning of period asset holdings, $\hat{a}_t^{z_t,\eta_t} \in \mathcal{A}^{\text{grid}}$, lying on the predefined grid. Interpolating $(\hat{a}_t^{\text{common},\eta_t}, \hat{c}_t^{\text{common}})$ for values $\hat{a}_t^{z_t,\eta_t} \in \mathcal{A}^{\text{grid}}$, we obtain the new consumption policies $\left(\hat{a}_t^{z_t,\eta_t} \in \mathcal{A}^{\text{grid}}, \hat{c}_t^{A,\text{new},\eta_t}\right)$. Next we interpolate $\left(\mathcal{C}_t^{\eta_t,z_{t+1},\hat{a}_{t+1}^{z_{t+1}}\in\mathcal{A}^{\text{grid}}}, \mathcal{A}^{\text{grid}}\right)$ for values $\hat{c}_t^{A,\text{new},\eta_t}$ and apply the borrowing constraints to obtain new policy functions $\left(\hat{a}_t^{z_t,\eta_t} \in \mathcal{A}^{\text{grid}}, \hat{a}_{t+1}^{\text{new},z_{t+1},\eta_t}\right)$.

6. The asset demand by advanced traders is then given by

$$d_t^{A,z_{t+1}} = \sum_{\eta_t} \sum_{\hat{a}_t^{z_t,\eta_t}\in\mathcal{A}^{\text{grid}}} \mu^{A,\eta_t}(\hat{a}_t^{z_t,\eta_t}) \cdot \hat{a}_{t+1}^{\text{new},z_{t+1}\eta_t}(\hat{a}_t^{z_t,\eta_t}) g \tag{60}$$

For markets to clear, we need that for all z_{t+1}, we have

$$d_t^{A,z_{t+1}} + d_t^B = g_{V^{\text{firm}}}\left(z_{t+1}, \hat{K}_{t+1}, \lambda_{t+1}^{A \text{ forecast}}\right). \tag{61}$$

Together with the optimality condition for firm investment (*i.e.* equation (10)), this gives us four nonlinear equations to solve for the three prices and the next period's capital. Once the equation is solved, we recorded the new consumption function and the updated firm value, which we will use to update the consumption and firm value approximations, and simulate the economy one period forward using the non-stochastic simulation method by Young (2010). We also record the resulting wealth share held by advanced traders for each possible shock in the next period, which we will use to update the forecasting functions for the transition of their wealth share.

Repeating this step for T periods, we obtain the following newly computed sequences

1. aggregate shocks $\{z_t\}_{t=0}^T$
2. aggregate capital $\{\hat{K}_t\}_{t=0}^T$
3. wealth distribution of advanced traders $\{\rho_t^A\}_{t=0}^T$
4. wealth distribution of bond traders $\{\rho_t^B\}_{t=0}^T$
5. wealth share owned by advanced traders $\{\lambda_t^A\}_{t=0}^T$
6. prices for the aggregate-state contingent securities $\{p_t^{z_{t+1}}\}_{t=0}^T$
7. wealth share of advanced traders in the next period for each possible aggregate shock, as implied by the households investment policies $\{\lambda_{t+1}^{A,z_{t+1}}\}_{t=0}^T$
8. aggregate dividends paid by the operating firms $\{\hat{D}_t\}_{t=0}^T$
9. aggregate value of operating firms $\{\hat{V}_t\}_{t=0}^T$
10. consumption by advanced traders for each idiosyncratic shock and each grid point on the asset grid $\{\hat{c}_t^A\}_{t=0}^T$
11. consumption by bond traders for each idiosyncratic shock and each grid point on the asset grid $\{\hat{c}_t^B\}_{t=0}^T$

Updating the approximating functions

Updating the firm's policy We use the sequence of exogenous aggregate shocks, capital, and the wealth share of the advanced traders to update the firms policy $g_K^{z_t}\left(\hat{K}_t, \lambda_t^A\right)$. For each exogenous shock z_t, we predict $\hat{K}_{t+1}^{\text{pred}} = g_K^{z_t}\left(\hat{K}_t, \lambda_t^A\right)$. We then adjust the parameters of the function approximator to fit a weighted average between the predicted capital sequence and the newly obtained sequence $\hat{K}_{t+1}$, such that

$$g_K^{z_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1 - \alpha^{\text{update}})\hat{K}_{t+1}^{\text{pred}} + \alpha^{\text{update}}\hat{K}_{t+1}. \tag{62}$$

An $\alpha^{\text{update}} > 0$ dampens the updating of the policy function. Slow updating is useful in simulation-based methods to ensure that the domain of the functions (here aggregate capital and the wealth share of advanced traders) is not changing too quickly between subsequent iterations. We choose $\alpha^{\text{update}} = 0.05$.

Updating the forecasting function for the wealth share of advanced traders We use the sequence of exogenous aggregate shocks, capital, the wealth share of the advanced traders, and the shock contingent wealth share of advanced traders in the next period to update the forecasting functions $g_{\lambda^A}^{z_t,z_{t+1}}\left(\hat{K}_t, \lambda_t^A\right)$. For each pair of exogenous shock (z_t, z_{t+1}), we predict $\lambda_{t+1}^{A,z_{t+1},\text{pred}} = g_{\lambda^A}^{z_t,z_{t+1}}\left(\hat{K}_t, \lambda_t^A\right)$. Similar to capital, we fit the new parameters to a convex combination of the old prediction and the new values, such that

$$g_{\lambda^A}^{z_t,z_{t+1},\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1 - \alpha^{\text{update}})\lambda_{t+1}^{A,z_{t+1},\text{pred}} + \alpha^{\text{update}}\lambda_{t+1}^{A,z_{t+1}}. \tag{63}$$

We choose $\alpha^{\text{update}} = 0.05$.

Updating the households' consumption functions For the households' consumption functions, we follow the same procedure. The main difference is that we fit separate functions for each combination of the aggregate shock, the idiosyncratic shock, and each grid point on the individual asset grid. Hence we allow the dependence of the households' consumption on capital and wealth to be different depending on the households' wealth level. This is, in particular, suitable for our method, where we take the extremely correlated endogenous aggregate summary variables, *i.e.* capital and the wealth share of the advanced traders, from a simulated path, while having a grid for the idiosyncratic asset holding. Again we dampen the update by fitting the approximating function to a weighted average between the newly computed consumption and the consumption predicted by the previous function approximator.

$$g_{c^A}^{\eta_t, z_t, \hat{a}_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\hat{c}_t^{A,\eta_t,z_t,\hat{a}_t,\text{pred}} + \alpha^{\text{update}}\hat{c}_t^{A,\eta_t,z_t,\hat{a}_t}, \tag{64}$$

$$g_{c^B}^{\eta_t, z_t, \hat{b}_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\hat{c}_t^{B,\eta_t,z_t,\hat{b}_t,\text{pred}} + \alpha^{\text{update}}\hat{c}_t^{B,\eta_t,z_t,\hat{b}_t} \tag{65}$$

We choose $\alpha^{\text{update}} = 0.2$.

Updating the firm value To update the firm value we use the computed sequences of exogenous aggregate shocks, the wealth share of the advanced traders, the aggregate-state contingent prices, and the paid dividends. First, we compute the firm value resulting from iterating on the firm's Bellman equation, which is given in equation (18), while repeatedly fitting a function to the updated values. We denote the resulting firm value by $\hat{V}_t^{\text{div}}$. As for aggregate capital, we also predict the firm value based on the old policy approximation and dampen the update, so that the new function approximator fits

$$g_V^{z_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\hat{V}_t^{\text{pred}} + \alpha^{\text{update}}\hat{V}_t^{\text{div}}. \tag{66}$$

We choose $\alpha^{\text{update}} = 0.025$.

Functional form for the function approximators We approximate each of the approximating functions as

$$g(\hat{K}, \lambda^A) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{67}$$

and obtain the parameters $\{\theta_1, \theta_2, \theta_3, \theta_4, \theta_5\}$ by least-squares fitting the data.[17] Since we fit a different such function for each (pair of) discrete shocks and each asset grid point, this functional form provides us with sufficient flexibility to obtain a good fit to the data.

[17] We use the `curve_fit` command from the Julia library `LsqFit.jl`.

	$\hat{c}^A$ [%]	$\hat{c}^B$ [%]	$\hat{c}^A$ [%] η_1	$\hat{c}^B$ [%] η_1	$\hat{c}^A$ [%] η_2	$\hat{c}^B$ [%] η_2	$\hat{V}^{\text{firm}}$ [%]
Mean	0.01	0.00	0.01	0.00	0.01	0.00	0.01
90th percentile	0.02	0.00	0.02	0.00	0.02	0.00	0.01
99th percentile	0.06	0.01	0.08	0.01	0.04	0.01	0.02

Table 3: Statistics for the absolute difference between the values predicted by the approximating functions for the households' consumption and the firm value and the corresponding updated values obtained from a simulated path of 10000 periods for the benchmark model. The numbers denote the errors relative to the updated values and are expressed in %.

	$\hat{K}^A$ [%]	λ^A [%]	λ^A [%] $z_1 \to z_1$ $z_1 \to z_2$ $z_1 \to z_3$	λ^A [%] $z_2 \to z_1$ $z_2 \to z_2$ $z_2 \to z_3$	λ^A [%] $z_3 \to z_1$ $z_3 \to z_2$ $z_3 \to z_3$
Mean	0.00	0.01	0.01 0.01 0.01	0.01 0.01 0.01	0.01 0.01 0.01
90th percentile	0.00	0.02	0.02 0.02 0.02	0.02 0.02 0.02	0.01 0.01 0.01
99th percentile	0.00	0.03	0.04 0.03 0.03	0.04 0.04 0.03	0.03 0.03 0.03

Table 4: Statistics for the forecasting errors for aggregate capital and the wealth-share of the advanced traders on a simulated path of 10000 periods for the benchmark model. The errors for the capital forecast are relative errors in %, and the errors for the wealth share forecasts are absolute errors in %.

A.1.2 Accuracy of the solution method

To investigate the accuracy of our solution, we asses the *out of sample error* of our function approximators, *i.e.* the difference between new values collected on the simulated path of 10000 periods and the corresponding value predicted by our function approximators, which we obtained in the previous iteration. Table 3 shows that the remaining errors are low, with the 99th percentile well below 0.1%. Similarly, Table 4 shows the accuracy of the forecasting functions for the evolution of the two endogenous aggregate quantities that summarize the distribution in the spirit of Krusell and Smith (1998). As we can see both forecasting rules are accurate.

A.1.3 Necessity of adding the wealth share held by advanced traders

In the standard implementation of Krusell and Smith (1998) capital alone would be used to forecast prices. In our model capital remains important, since it pins down wages, but it is not enough. To illustrate this, figure 7 shows a scatter plot of capital against the state price for shock z_3, when the economy is in shock z_1 for 300 realizations along the simulated path in the benchmark model, normalized by the transition probability $\pi_{z_1 \to z_3}$. The circles show a prediction of the price when fitting a polynomial of the form $\theta_1 + \theta_2 \hat{K} + \theta_3 \hat{K}^2 + \theta_4 \lambda^A + \theta_5 (\lambda^A)^2$, where $\hat{K}$ denotes aggregate capital and λ^A denotes the wealth share of advanced traders. First, we can see that aggregate capital alone is not enough to predict the price accurately. Second, we can see that the prediction

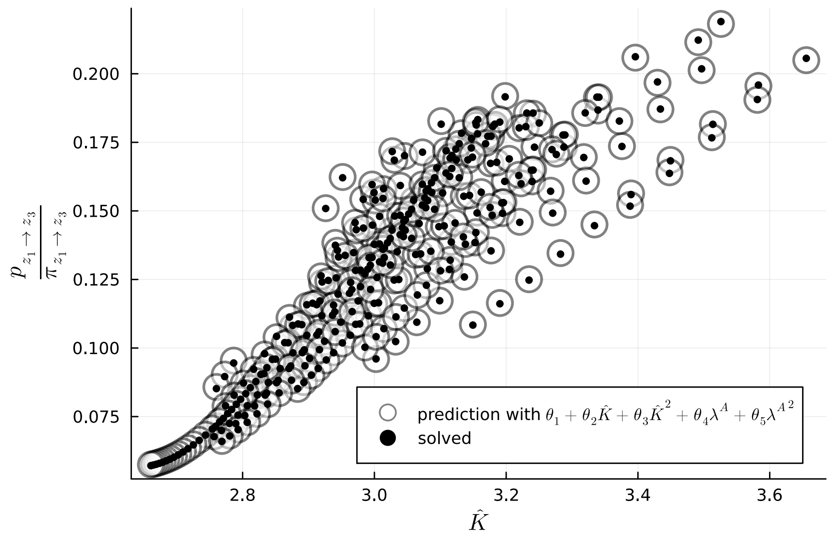

Figure 7: Realizations for the price of the aggregate-state contingent security for shock z_3, when the economy is in shock z_1, against the corresponding values of aggregate capital. The grey circles show the predictions of a fitted polynomial of the form $\theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2$.

based on capital and the simple functional form we are using, provides a very good fit to the data.[18]

Dependence of households' consumption on aggregates for different wealth levels

As described above, we approximate the households' consumption functions with a different function for each aggregate shock, idiosyncratic shock, and the households' asset level on a pre-specified grid. This allows us to fit a simple functional form for each of the functions.

$$g_{c^A}^{\eta_t, z_t, \hat{a}_t, \text{ new}}\left(\hat{K}_t, \lambda_t^A\right) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{68}$$

$$g_{c^B}^{\eta_t, z_t, \hat{b}_t, \text{ new}}\left(\hat{K}_t, \lambda_t^A\right) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{69}$$

Importantly, the parameters θ_i can differ not only across the exogenous shocks, but also across the wealth level. Figure 8 shows the household consumption of an advanced trader in aggregate shock z_2 and idiosyncratic shock η_1 for two different wealth levels along 300 realizations on the simulated path of the economy for the benchmark model. The top panel shows the consumption for a household without any financial wealth, $\hat{a}_t^{z_2} = 0$. The bottom panel shows the consumption for a rich household, who owns assets of value $\hat{a}_t^{z_2} = 26.194$, which corresponds to roughly 27 times the average annual labor earnings. First, we can observe that neither capital alone, nor the wealth share of the advanced traders alone would allow us to predict consumption accurately. Second, figure 8 illustrates that both, capital and the wealth share held by advanced traders, together with our simple functional form, allow us to obtain an excellent fit. Lastly, we also see that the relative importance of the two summary variables varies with the wealth of the household. For the poor household, capital is more important and would almost be sufficient to approximate their consumption. This is intuitive since the poor household depends on their wage and the income

[18]While this serves as an illustrative example, our algorithm does not require to forecast prices, but the households' consumption. In figure 8 we show that household consumption can be very accurately predicted as well.

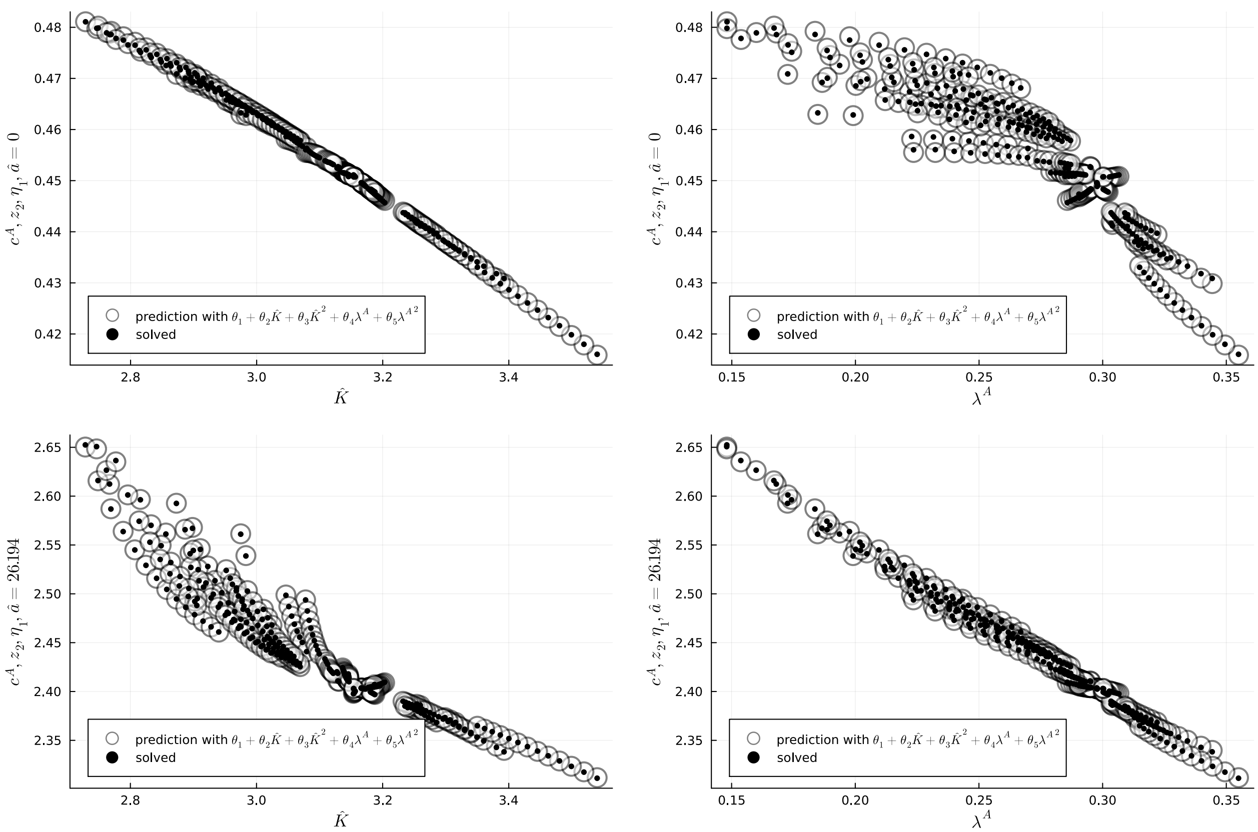

Figure 8: Consumption of an advanced trader in aggregate shock z_2 and idiosyncratic shock η_1 for two different wealth levels along 300 realizations on the simulated path of the economy for the benchmark model. The top panel shows the consumption for a household without any financial wealth, $\hat{a}_t^{z_2} = 0$. The bottom panel shows the consumption for a rich household, who owns assets of value $\hat{a}_t^{z_2} = 26.194$. The left panel shows the consumption level scattered against capital, and the right panel shows the consumption level scattered against the wealth share held by advanced traders.

stream from startup creation, which is decreasing in aggregate capital, and less on prices. For the rich household however, the wealth share held by advanced traders is a better prediction of their consumption than aggregate capital. This is intuitive since for the richest households, given their wealth, lower asset prices and higher asset returns are relatively more important than the income from labor and startup creation.

A.2 Calibration

For our calibration exercise, we need a mapping from the two free parameters, the time-preference parameter and the adjustment costs, to the moment of interest (see also Scheidegger and Bilionis, 2019; Catherine et al., 2022, for the use of surrogate models in economics). While our solution method is efficient and would hence allow us to solve the model on a dense grid for those two parameters, this is not necessary. Instead, use Gaussian Processes with a squared exponential kernel to approximate the quantities of interest in surrogate model.[19] A further advantage of using a surrogate model is that it smoothes out remaining fluctuations in mean quantities which arise from the fact that the mean quantities are obtained from Monte Carlo simulations. We initially solve the model on a coarse grid of parameter values for patience and adjustment costs.[20] Then we use the fitted surrogate model to guide us in choosing the parameter values for which we next solve the model. As more models are solved, update the surrogate model with the new information.

For example, to generate table 2, we solved the model for different parameter values all generating a market price of risk of about 45%. To generate figure 4, we solved the model for various exit rates, on top of various combinations of patience and adjustment costs. For a given firm exit rate, we first use the surrogate model for the mean interest rate to obtain the patience and adjustment costs parameters required to obtain an average interest rate of 0.8%. We then use the surrogate model for the expected log price earnings ratio to obtain an estimate for its value. To produce figure 5, we use a Gaussian process to estimate the mapping for the interest rate and the market price of risk to the log mean price of the infinitely lived console.

In the main text, figure 2 shows the market price of risk, the interest rate, as well as the objective function, which we aim to minimize with our calibration. Figure 3 shows how the standard deviation of aggregate consumption growth, the standard deviation of the interest rate, as well as the price-earnings ratio, depend on the time-preference parameter and the adjustment costs parameter. As discussed in section 3.2.2, these untargeted moments are roughly in line with what we measure in the data.

[19] See Rasmussen and Williams (2004) for a general introduction to Gaussian Processes and Scheidegger and Bilionis (2019) for applications in economics.

[20] We start out with $(\beta, \xi^{\text{adj}}) \in \{0.87, 0.93, 0.96\} \otimes \{2, 4, 6\}$.

B The calibration of the idiosyncratic income process (online only)

Storesletten et al. (2004) estimate a process for log earnings of the form

$$
\begin{aligned}
y_{it} &= g(x_{it}^h, Y_t) + u_{it}^h && (70)\\
u_{it} &= \alpha_i + z_{it}^h + \epsilon_{it} && (71)\\
z_{it}^h &= \rho z_{i,t-1}^{h-1} + \eta_{it} && (72)\\
\alpha_i &\sim \text{iid } N(0, \sigma_\alpha^2) && (73)\\
\epsilon_{it} &\sim \text{iid } N(0, \sigma_\epsilon^2) && (74)\\
\eta_{it} &\sim \text{iid } N(0, \sigma_t^2) && (75)\\
\sigma_t^2 &= \begin{cases} \sigma_E^2 & \text{in agg. expansions} \\ \sigma_C^2 & \text{in agg. contractions} \end{cases}. && (76)
\end{aligned}
$$

Storesletten et al. (2004) estimate $\sigma_\epsilon = 0.25$, $\rho = 0.95$ and frequency weighted average of σ_E and σ_C given by 0.17. We abstract from age and the CCV mechanism and focus on the non-permanent idiosyncratic component of log earnings, which we denote with x_{it}. We simulate a process for

$$
\begin{aligned}
x_{it} &= \epsilon_{it} + z_{it} && (77)\\
z_{it} &= \rho z_{it-1} + \eta_{it} && (78)\\
\epsilon_{it} &\sim \text{iid } N(0, \sigma_\epsilon^2) && (79)\\
\eta_{it} &\sim \text{iid } N(0, \sigma_\eta^2), && (80)
\end{aligned}
$$

where we take $\rho = 0.95$, $\sigma_\epsilon = 0.25$, and $\sigma_\eta = 0.17$ from Storesletten et al. (2004). We then fit an AR(1) process of the form

$$
\begin{aligned}
x_{it} &= \bar{x} + \rho_x x_{i,t-1} + \sigma_x \epsilon_{it} && (81)\\
\epsilon_t &\sim \text{iid } N(0, 1), && (82)
\end{aligned}
$$

and obtain $\rho_x = 0.785$, $\sigma_x = 0.372$, and $\bar{x} = -0.00$. We discretize the AR(1) process into a two-state Markov chain using Rouwenhorst (1995) algorithm. Next, we exponentiate the resulting state values and normalize them such that the average earnings are equal to 1. We obtain

$$
\begin{aligned}
X &= \begin{pmatrix} 0.463 \\ 1.537 \end{pmatrix} && (83)\\
\pi^X &= \begin{pmatrix} 0.892 & 0.108 \\ 0.108 & 0.892 \end{pmatrix} && (84)
\end{aligned}
$$

The resulting cross-sectional standard deviation of log earnings is 0.60 and matches the standard deviation of the process simulated in equations (70) - (76) and is in the ballpark of values typically used in the heterogeneous agents literature (see, *e.g.* Auclert et al., 2019).

References

Aguiar, M. A., Amador, M., and Arellano, C. (2021). Micro risks and pareto improving policies with low interest rates. Technical report, National Bureau of Economic Research.

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

Aiyagari, S. R. and Peled, D. (1991). Dominant root characterization of pareto optimality and the existence of optimal equilibria in stochastic overlapping generations models. *Journal of Economic Theory*, 54(1):69–83.

Amol, A. and Luttmer, E. G. (2022). Permanent primary deficits, idiosyncratic long-run risk, and growth. Technical report.

Arrow, K. J. (1974). *Essays in the theory of risk-bearing*, volume 121. North-Holland Amsterdam.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2019). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. Working Paper 26123, National Bureau of Economic Research.

Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.

Beeler, J. and Campbell, J. (2012). The long-run risks model and aggregate asset prices: An empirical assessment. *Critical Finance Review*, 1(1):141–182.

Blanchard, O. (2019). Public debt and low interest rates. *American Economic Review*, 109(4):1197–1229.

Bloise, G. and Reichlin, P. (2022). Low safe interest rates: A case for dynamic inefficiency? *Available at SSRN 4244152*.

Brumm, J., Feng, X., Kotlikoff, L. J., and Kubler, F. (2021). When interest rates go low, should public debt go high? Technical report, National Bureau of Economic Research.

Brumm, J. and Scheidegger, S. (2017). Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica*, 85(5):1575–1612.

Calvet, L. E., Campbell, J. Y., Gomes, F., and Sodini, P. (2021). The cross-section of household preferences. Technical report, National Bureau of Economic Research.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.

Catherine, S., Ebrahimian, M., Sraer, D., and Thesmar, D. (2022). Robustness checks in structural analysis. Technical report, National Bureau of Economic Research.

Chien, Y., Cole, H., and Lustig, H. (2011). A multiplier approach to understanding the macro implications of household finance. *The Review of Economic Studies*, 78(1):199–234.

Cochrane, J. (2009). *Asset pricing: Revised edition*. Princeton university press.

Constantinides, G. M. and Duffie, D. (1996). Asset pricing with heterogeneous consumers. *Journal of Political economy*, 104(2):219–240.

Crane, L. D., Decker, R. A., Flaaen, A., Hamins-Puertolas, A., and Kurz, C. (2022). Business exit during the covid-19 pandemic: Non-traditional measures in historical context. *Journal of macroeconomics*, 72:103419.

Druedahl, J. (2021). A guide on solving non-convex consumption-saving models. *Computational Economics*, 58(3):747–775.

Favilukis, J. (2013). Inequality, stock market participation, and the equity premium. *Journal of Financial Economics*, 107(3):740–759.

Fernández-Villaverde, J., Hurtado, S., and Nuno, G. (2019). Financial frictions and the wealth distribution. Technical report, National Bureau of Economic Research.

Guo, H. (2004). Limited stock market participation and asset prices in a dynamic economy. *Journal of Financial and Quantitative Analysis*, pages 495–516.

Guvenen, F. (2009). A parsimonious macroeconomic model for asset pricing. *Econometrica*, 77(6):1711–1750.

Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of political economy*, 99(2):225–262.

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

Kocherlakota, N. R. (2022). Infinite debt rollover in stochastic economies. Technical report, National Bureau of Economic Research.

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

Kubler, F. and Scheidegger, S. (2018). Self-justified equilibria: Existence and computation.

Lo, A. W. (2002). The statistics of sharpe ratios. *Financial analysts journal*, 58(4):36–52.

Lustig, H. and Verdelhan, A. (2012). Business cycle variation in the risk-return trade-off. *Journal of Monetary Economics*, 59:S35–S49.

Luttmer, E. G. (2011). On the mechanics of firm growth. *The Review of Economic Studies*, 78(3):1042–1068.

Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.

Mian, A., Straub, L., and Sufi, A. (2021). A goldilocks theory of fiscal policy. *NBER Working Paper*, (29351).

Modigliani, F. and Miller, M. H. (1958). The cost of capital, corporation finance and the theory of investment. *The American economic review*, 48(3):261–297.

Parker, J. A. and Vissing-Jorgensen, A. (2009). Who bears aggregate fluctuations and how? *American Economic Review*, 99(2):399–405.

Rasmussen, C. E. and Williams, C. K. (2004). Gaussian processes in machine learning. *Lecture notes in computer science*, 3176:63–71.

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. *Frontiers of business cycle research*, 1:294–330.

Scheidegger, S. and Bilionis, I. (2019). Machine learning for high-dimensional dynamic stochastic economies. *Journal of Computational Science*, 33:68 – 82.

Shleifer, A. and Vishny, R. (2011). Fire sales in finance and macroeconomics. *Journal of economic perspectives*, 25(1):29–48.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of political Economy*, 112(3):695–717.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of political Economy*, 110(4):825–853.

Vissing-Jørgensen, A. and Attanasio, O. P. (2003). Stock-market participation, intertemporal substitution, and risk-aversion. *American Economic Review*, 93(2):383–391.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of monetary economics*, 24(3):401–421.

Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41.