NBER WORKING PAPER SERIES

ASSET PRICING IN A LOW RATE ENVIRONMENT

Marlon Azinovic
Harold L. Cole
Felix Kubler

Working Paper 31832
http://www.nber.org/papers/w31832

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2023

We are grateful for the helpful comments we received from Andrew Abel, Nir Jaimovich, Urban Jermann and Narayana Kocherlakota, as well as the participants at the Wharton Brown Bag Seminar, the Bristol Macroeconomics Workshop, EEA-ESEM Conference, and Rochester's Stockman Memorial Workshop. 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". The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Asset Pricing in a Low Rate Environment
Marlon Azinovic, Harold L. Cole, and Felix Kubler
NBER Working Paper No. 31832
November 2023
JEL No. C6,E22,E43,E44,G12

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 as low as 2-3.5 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 the wealth distribution and asset prices in economies with very low real rates. We also 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 3.5 percent below the average growth rate.

Marlon Azinovic
University of Pennsylvania
azinovic@sas.upenn.edu

Harold L. Cole
Economics Department
University of Pennsylvania
3718 Locust Walk
160 McNeil Building
Philadelphia, PA 19104
and NBER
colehl@sas.upenn.edu

Felix Kubler
University of Zurich
Plattenstrasse 32
CH-8032 Zurich
Switzerland
and Swiss Financial Institute
fkubler@gmail.com

1 Introduction

We develop a tractable model of a production economy that matches key statistics of both asset prices and business cycle dynamics to examine outcomes in low interest rate environments. Our model 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 predict the impact of lower interest rates on the prices of long-lived 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% (see Mehra and Prescott (1985) and Beeler and Campbell (2012)). The growth rate of per capita output has been roughly 2%. 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%. This is while the real 10-year rate in the figure fell from around 4% in 1990 to averaging less than 0.5% between 2011-2021.

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).

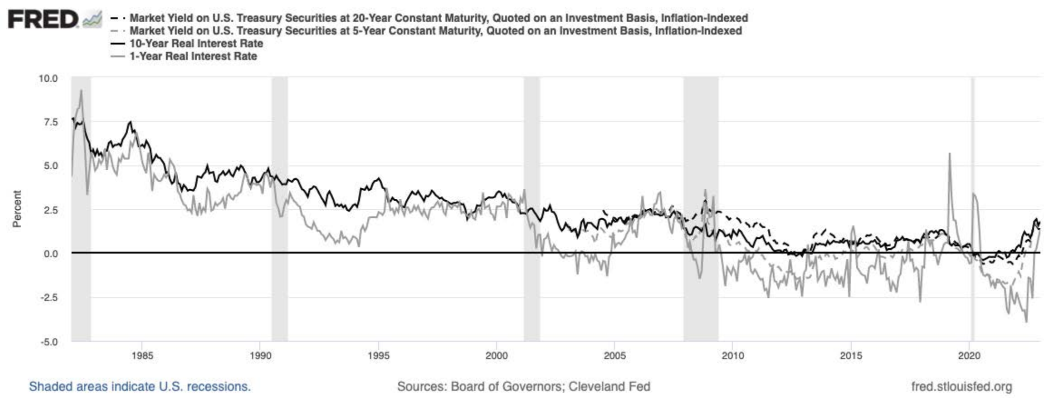

Figure 1: Real interest rates

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 5.5 and do not include a habit factor that can lead to much higher local degrees of aversion to risk. We consider this an important component of generating a reasonable pricing kernel 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 firms' market price can become very large. To match observed price-earnings 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 well-documented that every year a substantial fraction of firms and establishments in the US "die", while surviving firms, on average, grow in size, despite the fact that new firms enter. We calibrate our model to match these facts.

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 non-linearities and movements in the wealth distribution. Second, to solve for equilibrium prices of the Arrow securities, we need to employ approximation methods for the consumption functions with 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 our model's equilibrium, in which we alternate between simulating the economy and updating the policy, value, and forecasting functions. 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 aggregate state is sufficient to accurately forecast the next period's aggregate state for each exogenous aggregate shock. As is standard in Euler equation methods, we use the approximated policy functions (along with a forecast of next period's advanced trader wealth share) to determine the policies of the firm and households in the 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 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.[1]

An important aspect of our approach is that we can compute equilibria for our model even when (Arrow-Debreu) 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.

We examine the wealth distributions and asset prices in various different calibrations

[1] This means we approximate the households' policies by a different simple function for different wealth levels of the households.

resulting in different very low real rates. In our benchmark calibration, we set the real interest rate to 0.8% with a growth rate of 2%. We match a market price of risk of 50% as well as realistic moments for consumption- and output growth and realistic price-earnings ratios. We also calibrate the interest rate to be -1.5%. At such very low rates, the price-earnings ratio becomes slightly too high but remains in a reasonable range, due to firms' exit. The value of future labor income, however, explodes. We find that for very low real rates, the wealth inequality between advanced traders and bond-only traders increases substantially.

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)).

As an application, we, therefore, investigate whether a console 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 -1.5%), while the value of the console remains finite.

The possibility of infinite debt rollover is a sufficient condition for Pareto-improving government debt. In our model, it is not a necessary condition because of financial frictions. Obviously, government debt might help to ease borrowing constraints (in exchange for money today, agents pledge future labor income in the form of lump-sum taxes). However, in our calibration government debt always makes all agents (ex ante) worse off. It crowds out private savings and implies lower wages because of a lower average capital stock.

In our calibration, we assume a deterministic growth trend and transitory shocks to TFP. One can think of other calibrations where low interest rates necessarily imply that debt rollover must be possible. The purpose of our exercise is to provide one plausible calibration and show that very low interest rates are not inconsistent with finite values of long-lived assets.

In asset pricing, it is typically assumed that growth rates themselves are stochastic. Alvarez and Jermann (2005) point out that some asset pricing facts, in particular, the yield curve for government bonds strongly suggest that shocks must have a very large permanent component. We conduct a sensitivity analysis where we show that a relatively small permanent component to tfp shocks suffices to obtain a realistic yield curve in our model.

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 various calibrations with very low interest rates. In Section 5 we examine the possibility of debt rollover and discuss the effects of government debt on welfare. In Section 6 we examine the role of permanent tfp shocks for asset pricing in our model, and Section 7 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

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

where $\xi(z_t)$ is the stochastic productivity level, $A_t = (1+g)^t$ grows deterministically at a fixed rate g, k_t^i and l_t^i denote the amount of capital and labor firm i chooses in period t.

Every period starts with a unit measure of firms, of which the fraction Γ survive to produce and the fraction $1-\Gamma$ die. If firms exit, they do not pay dividends, and their capital evaporates.[2] Because death is random, the amount of capital in surviving firms is ΓK_t where K_t is the total amount of capital at the beginning of the period.

An equal measure of new firms replaces the dying firms. The new firms do not produce in the current period, and the initial amount of capital invested in them at the end of the period is $\bar{i}_t$. 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. Since initial capital does not suffer adjustment costs, there can be rents from firm creation.

The total measure of firms at the beginning of the next period is 1 by construction. Except for their capital stock, both new and surviving firms are identical starting next period, and both new and surviving firms in period t are subject to the survival shock at the beginning of the next period. 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 problem of surviving firms is standard. For firm i, 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

[2] 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).

period conditional on surviving. Firms producing at time t pay dividends

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

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

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

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

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

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

$$\omega_t = \xi_t A_t (1-\alpha) \left(\frac{k_t^i}{A_t l_t^i} \right)^{\alpha} \Leftrightarrow \frac{k_t^i}{A_t l_t^i} = \left(\frac{\omega_t}{A_t \xi_t (1-\alpha)} \right)^{\frac{1}{\alpha}} =: \mathcal{K}_t, \tag{6}$$

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

$$w_t = A_t (1-\alpha) \xi_t \mathcal{K}_t^{\alpha} \tag{7}$$

$$r_t^K = r_t^{Ki} = \alpha \xi_t \mathcal{K}_t^{\alpha-1} \tag{8}$$

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

$$1 + \underbrace{\frac{\partial \psi(k_{t+1}^i, k_t^i)}{\partial k_{t+1}^i}}_{=2\xi^{\text{adj}}\left(\frac{k_{t+1}^i}{k_t^i} - (1-\delta+x^{\text{target}})\right)} = \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} \left(r_{t+1}^K + 1 - \delta - \underbrace{\frac{\partial \psi(k_{t+2}^i, k_{t+1}^i)}{\partial k_{t+1}^i}}_{-\xi^{\text{adj}}\left(\left(\frac{k_{t+2}^i}{k_{t+1}^i}\right)^2 - (1-\delta+x^{\text{target}})^2\right)} \right) \tag{9}$$

The right-hand side of the Euler equation reflects the firm's anticipated policy tomorrow conditional on survival. Because of the homogeneity of production and adjustment costs, each incumbent firm chooses the same growth rate, which, for a given period, t, we denote by $(g^k)_t^i := \frac{k_{t+1}^i}{k_t^i} = g_t^k$.

We define K_t to be the aggregate capital in period t *at the beginning of the period before the survival lottery.* There is a measure Γ of incumbent firms from last period and a measure $1-\Gamma$ of that were created last period. Let $\mathcal{P}_t$ denote the set of indexes for firms, which are producing in period t. If $\tilde{K}_{t+1}$ is the per-period-t-producing-firm amount of capital chosen in period t by the producing firms, then

$$\tilde{K}_{t+1} = \frac{1}{\Gamma} \int_{i \in \mathcal{P}_t} k^i_{t+1} di \tag{10}$$

Similarly, let $\tilde{I}_t$ to be the per-period-t-producing-firm investment level for producing firms in period t,

$$\tilde{I}_t = \frac{1}{\Gamma} \int_{j \in \mathcal{P}_t} i^j_t dj.$$

Hence

$$\Gamma \tilde{K}_{t+1} = (1-\delta)\Gamma K_t + \Gamma \tilde{I}_t \Leftrightarrow \tilde{K}_{t+1} = (1-\delta) K_t + \tilde{I}_t, \tag{11}$$

where $\tilde{I}_t$ is the per-producing-firm amount of investment in period t.

Capital is chosen at the end of the period. This leads us to set the initial capital in new firms to be the fraction s of the per capita amount of capital being chosen by incumbent firms in period t:

$$\bar{i}_t = s\tilde{K}_{t+1} > 0.$$

The evolution of the beginning of period $t+1$ capital stock includes the capital coming from incumbent firms at time t and new entrants at time t:

$$K_{t+1} = \underbrace{\Gamma \tilde{K}_{t+1}}_{\text{incumbents}} + \underbrace{(1-\Gamma)\bar{i}_t}_{\text{entrants}} = [\Gamma + (1-\Gamma)s]\, \tilde{K}_{t+1} \tag{12}$$

The amount of capital available for production is ΓK_{t+1} and K_{t+1} is the average amount of capital at each producing firm in period $t+1$. Aggregate investment can be defined in a similar fashion. The aggregate investment, which includes the investment by startups as well as incumbents, is

$$I_t = \Gamma \tilde{I}_t + (1-\Gamma)\bar{i}_t = \Gamma \tilde{I}_t + (1-\Gamma)s\tilde{K}_{t+1} \tag{13}$$

The growth rate of capital at surviving producing firms is given by

$$g^k = \frac{\tilde{K}_{t+1}}{K_t} = \frac{1}{\Gamma + (1-\Gamma)s}\frac{K_{t+1}}{K_t}. \tag{14}$$

Aggregate output is given by

$$Y_t = \xi\left(z_t\right)\left(\Gamma K_t\right)^{\alpha}\left(A_t L_t\right)^{1-\alpha}, \tag{15}$$

which takes into account the amount of capital that survives and produces in period t. As a result, the wage is given by

$$\omega_t = \xi_t A_t(1-\alpha)\left(\frac{k_t^i}{A_t l_t^i}\right)^{\alpha} = \xi_t A_t(1-\alpha)\left(\frac{\Gamma K_t}{A_t L_t}\right)^{\alpha} \tag{16}$$

And, the aggregate dividends paid by producing firms are given by

$$D_t = Y_t - \omega_t L_t - \Gamma\left[\tilde{I}_t - K_t \xi^{\text{adj}}\left(g_t^k - (1-\delta + x^{\text{target}})\right)^2\right]. \tag{17}$$

Because all firm's invest and produce proportionate to their beginning of period capital stock, the value of a producing firm is simply equal to their capital times the value of a firm with a unit of capital; $\tilde{v}_t := v_t^i/k_t^i = v_t^j/k_t^j\ \forall i,j \in \mathcal{P}_t$. Therefore the total value of producing firms at time t is given by

$$V_t^{\text{firm}} = \Gamma K_t \tilde{v}_t = \int_{j \in \mathcal{P}_t} v_t^j dj = \int_{j \in \mathcal{P}_t} \left[\begin{array}{c} y_t^j - \omega_t l_t^j - i_t^j - k_t^j \ \xi^{\text{adj}} \left(g_t^k - (1 - \delta + x^{\text{target}}) \right)^2 \\ + \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^j \end{array} \right] dj \tag{18}$$

This condition can be restated in terms of the Bellman equation for a unit-sized firm is given by

$$\tilde{v}_t = \left[\ \frac{D_t}{\Gamma K_t} + \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} \tilde{v}_{t+1} g^k \ \right] \tag{19}$$

where $\tilde{v}_{t+1}$ is the anticipated value of a unit-sized surviving firm in the next period, and $D_t/(\Gamma K_t)$ is the dividend paid per-unit-of-capital by each producing firm.

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

$$\bar{v}_t := \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^{t\text{-entrant}*} = \Gamma \bar{i}_t \sum_{z_{t+1}} p_t^{z_{t+1}} \tilde{v}_{t+1} \tag{20}$$

where $v_{t+1}^{t\text{-entrant}*}$ denotes the value of an idiosyncratic firm in period $t+1$ with capital $k_{t+1}^j = \bar{i}_t$. The last expression follows from the fact that the value of an idiosyncratic firm is proportional to its capital. The net value of creating a new firm is simply $\bar{v}_t - \bar{i}_t$.

In terms of equating asset demand with supply, an important element is the value of firms producing next period, which includes the value of incumbent firms plus new entrants,

$$V_{t+1}^{\text{firm}} = \tilde{v}_{t+1} \Gamma \left[\Gamma \ \tilde{K}_{t+1} + (1 - \Gamma) \ \bar{i}_t \right] = \tilde{v}_{t+1} \Gamma K_{t+1} \tag{21}$$

How the aggregate value of currently producing firms evolves is instructive as to the differing roles of the firm death rate Γ and the relative size of new entrants s. The aggregate producing firm's Bellman equation is given by

$$V_t^{\text{firm}} = D_t + \frac{\Gamma}{\Gamma + (1-\Gamma)s} \sum_{z_{t+1}} p_t^{z_{t+1}} V_{t+1}^{\text{firm}}. \tag{22}$$

Expression (22) shows that fixing $s > 0$, a decrease in the death rate means that less of the future value of firm's dividends is compounded into the current value of producing firms. But it also shows that as $s \to 0$ all of the future value firm dividends ends up being compounded. This is because the missing future value of firms is a product of two elements, how many there are and how valuable each new firm is in terms of both it's initial capital and subsequent rents from creating capital.

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 the *beginning of period* aggregate capital in steady state must satisfy

$$\hat{K}_{t+1} = \underbrace{\Gamma \frac{1+g^k}{1+g} \hat{K}_t}_{\text{period t incumbents}} + \underbrace{(1-\Gamma)s \frac{1+g^k}{1+g} \hat{K}_t}_{\text{period t startups}} = [\Gamma + (1-\Gamma)s] \frac{1+g^k}{1+g} \hat{K}_t. \tag{23}$$

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

$$1 = [\Gamma + (1-\Gamma)s]\,\frac{1+g^k}{1+g} \implies 1+g^k = \frac{1+g}{\Gamma + (1-\Gamma)s} \tag{24}$$

As one can see from inspection, $g^k > g$ to the extent that $s < 1$. The share of employment in new firms is given by $(1-\Gamma)s/\left[(1-\Gamma)s + \Gamma\right]$, while the share of employment last period in firms who exit this period is given by $(1-\Gamma)$. 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.

Matching the share of the present value of firm rents coming from "existing firms" as opposed to future firms[3] is of fundamental importance to our exercise. Since all firms in our model grow at the same rate, this depends fundamentally on the share of capital in future exiting firms.[4] 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

[3]There are rents associated with capital because of convex adjustment costs, and also because new firms are started with capital that was not subject to this cost.

[4]In the data young firms grow faster due in large part to selection (see for example Rossi-Hansberg and Wright (2007)). We abstract from this feature of the data since our calculations primarily involve the long-run growth rate.

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.[5] 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.

$$U((c_t)_{t=0}^{\infty}) = E_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \tag{25}$$

where instantaneous utility is given by

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma} \tag{26}$$

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

where γ denotes the coefficient of relative risk aversion and β denotes 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 start-up, 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 simply means a cash transfer of value $\bar{v}_t - \bar{i}_t$ to the agent. This value is mostly positive, but may also become negative, in particular in economic downturns. We assume the $(1-\Gamma)/\Gamma$ start-ups per period to be equally distributed across the advanced traders. Each trader hence starts $(1-\Gamma)/(\Gamma\mu)$ start-ups, where μ denotes the measure of advanced traders. We denote the cash flow received from a start-up opportunity as $\zeta(z^t) := \bar{v}_t - \bar{i}_t$. The Bellman equation for advanced traders is given by

$$\begin{aligned} V(z^t, \eta_t, a_{t-1}^{z_t}) = &\max_{\{a_t^{z_{t+1}} \geq \underline{a}\}_{z_{t+1} \in \mathcal{Z}}} u(c_t) \\ &+ \beta \sum_{\tilde{z}_{t+1}} \sum_{\tilde{\eta}_{t+1}} \pi_z(\tilde{z}_{t+1}|z_t) \pi_\eta(\tilde{\eta}_{t+1}|\eta_t) V((z^t, \tilde{z}_{t+1}), \tilde{\eta}_{t+1}, a_t^{\tilde{z}_{t+1}}), \end{aligned} \tag{27}$$

where

$$c_t = l^A(\eta_t)\omega(z^t) + \frac{1-\Gamma}{\Gamma\mu}\zeta(z^t) + a_{t-1}^{z_t} - \sum_{z_{t+1} \in \mathcal{Z}} a_t^{z_{t+1}} p^{z_{t+1}}(z^t). \tag{28}$$

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

$$p^{z_{t+1}}(z^t) = p(z^{t+1}|z^t), \text{ for } z^{t+1} = (z^t, z_{t+1}) \succ z^t. \tag{29}$$

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

$$s_t^B := (z^t, \eta_t, b_{t-1}). \tag{30}$$

The Bellman equation of non-participants is given by

$$V(z^t, \eta_t, b_{t-1}) = \max_{b_t \geq \underline{b}} u(c_t) + \beta \sum_{\tilde{z}_{t+1} \in \mathcal{Z}} \sum_{\tilde{\eta}_{t+1} \in \mathcal{N}} \pi_z(\tilde{z}_{t+1}|z_t)\pi_\eta(\tilde{\eta}_{t+1}|\eta_t)V((z^t, \tilde{z}_{t+1}), \eta_{t+1}, b_t), \tag{31}$$

where

$$c_t = l^B(\eta_t)\omega(z^t) + b_{t-1} - b_t p^b(z^t). \tag{32}$$

Here, $p^b(z^t)$ denotes the price of a bond, which is given by

$$p^b(z^t) = \sum_{z^{t+1} \succ z^t} p(z^{t+1}|z^t). \tag{33}$$

2.3 Market clearing

There are $Z + 1$ relevant market-clearing conditions in this economy. One for labor and one for each aggregate state-contingent 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

$$L_t^{\text{households}} := L_t^A + L_t^B, \text{ where} \tag{34}$$

$$L_t^A := \mu \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{a}}^{\infty} l^A(\eta_t) \rho_t^A(\eta_t, x) \mathrm{d}x \right) \tag{35}$$

$$L_t^B := (1-\mu) \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{b}}^{\infty} l^B(\eta_t) \rho_t^B(\eta_t, x) \mathrm{d}x \right) \tag{36}$$

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

$$L_t^{\text{firm}} := \left(\frac{\omega_t}{(1-\alpha) A_t^{1-\alpha} \xi_t K_t^{\alpha}} \right)^{-\frac{1}{\alpha}}. \tag{37}$$

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

$$\omega_t = (1-\alpha) \xi_t K_t^{\alpha} L_t^{-\alpha}. \tag{38}$$

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

$$w_t^{\text{households}} := w_t^A + w_t^B, \text{ where} \tag{39}$$

$$w_t^A := \mu \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{a}}^{\infty} x \rho_t^A(\eta_t, x) \mathrm{d}x \right) \tag{40}$$

$$w_t^B := (1-\mu) \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{b}}^{\infty} x \rho_t^B(\eta_t, x) \mathrm{d}x \right). \tag{41}$$

Financial markets clear when

$$V_t^{\text{firm}} = w_t^{\text{households}}. \tag{42}$$

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 3.5%. 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%, while the establishment death rate is much higher at roughly 4.5%. So our death rate is the average of these two values. Concerning the appropriate size of entrants, the employment share of new entrants is reported to be 1% in the first year,[6] while Luttmer (2011) reports that firms' employment growth is roughly 1% on average. The latter suggests that $\bar{i}_t/\hat{K}_{t+1}$ is approximately 75%, given the firm death rate.

As mentioned above, we assume that the capital of exiting firms has zero value. We can generalize this by assuming that the value is positive but small relative to the value of capital in surviving firms. Clearly, our effects on value-earnings ratios become smaller the larger the resale value of capital for existing firms. We provide some further discussions in Appendix C.

We choose $x^{\text{target}} = \delta + g$ so that no adjustment costs have to be paid if $g_t^K = \frac{k_{t+1}^*}{k_t} = (1 - \delta + x^{\text{target}}) = 1 + g$. 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 5.5. This is well within the range considered realistic by Mehra and Prescott (1985)[7]. The time preference parameter, β, is a parameter used for matching observed prices.

We model shocks to total factor productivity as a 3-state discretized AR(1) process

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

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

for deviations from the deterministic trend[8]. For the AR(1) process we have

$$\log(\xi_t) = \rho^{\text{tfp}} \log(\xi_{t-1}) + \sigma^{\text{tfp}} \epsilon_t^{\xi}, \tag{43}$$

where $\epsilon_t^{\xi} \sim \mathcal{N}(0,1)$. We choose the auto-correlation to be $\rho^{\text{tfp}} = 0.8145$, following 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%).[9] We discretize the AR(1) process using the method from Rouwenhorst (1995).

We focus on temporary shocks to TFP. There is an important and large debate on whether there is a large random walk component in GNP. Cochrane (1988) finds little long-term persistence in US GNP data. On the other hand, Alvarez and Jermann (2005) argue that to rationalize return properties of long-term bonds in models where the stochastic discount factor depends on aggregate consumption, innovations to aggregate consumption need to have permanent effects. Applying the Cochrane test to time series on TFP, we find little evidence for permanent shocks.

Because of our 3-state process, 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

[8]In Section 6 we discuss the possibility of permanent shocks to TFP growth.

[9]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.

idiosyncratic risk with the aggregate state of the economy. Details can be found in Appendix B. We obtain

$$X = \begin{pmatrix} 0.46 \\ 1.54 \end{pmatrix}, \quad \pi^X = \begin{pmatrix} 0.89 & 0.11 \\ 0.11 & 0.89 \end{pmatrix} \tag{44}$$

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., 2021).

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[10] 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

[10] 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.

of aggregate state-contingent 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

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

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 our benchmark calibration, we follow Cochrane (2009) and target a value of 50% for the market price of risk. In their segmented market exchange economy, Chien et al. (2011) target moments of equity returns and report an MPR of 45%. Several empirical studies report a Sharpe ratio in annual US equity returns that lies significantly above 50% (see, e.g., Lo (2002) or Lustig and Verdelhan (2012)). In Section 5 below, we explain why, for the exercise there, we want to employ a calibration with a relatively low (but still realistic) MPR.

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[11]

$$\ell(\beta, \xi^{\text{adj}}) := \left(\frac{1}{2} \left(\frac{\overline{\text{mpr}}(\beta, \xi^{\text{adj}}) - \text{mpr}^{\text{target}}}{\text{mpr}^{\text{target}}} \right)^2 + \frac{1}{2} \left(\frac{\overline{r}(\beta, \xi^{\text{adj}}) - r^{\text{target}}}{\max\{0.01, r^{\text{target}}\}} \right)^2 \right)^{\frac{1}{2}}. \tag{45}$$

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, (45), 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 $\beta = 0.9273$ and $\xi^{adj} = 4.514$. 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 output volatility, the observed average risk-free rate, and an MPR of 50%.

The first two panels of 2 show how the two targeted moments vary with different values of the parameters. The market price of risk in increasing in adjustment costs and (somewhat depending on the region) decreasing in the time preference parameter,

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

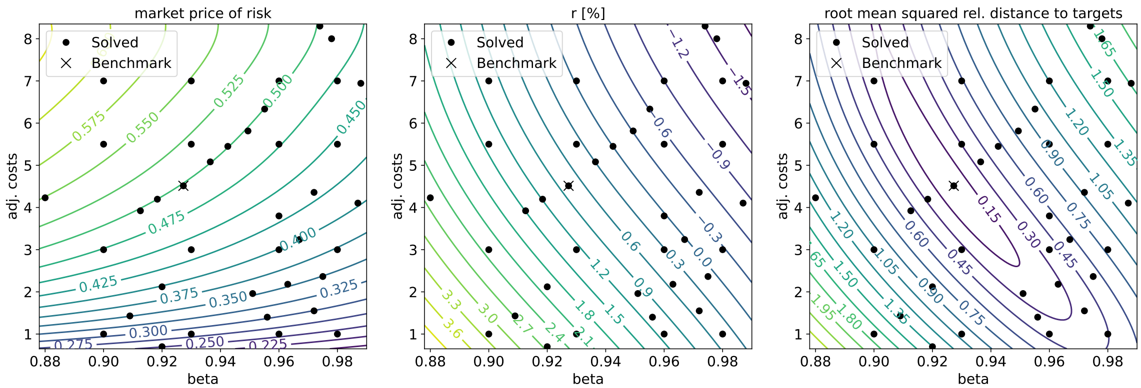

Figure 2: Dependence of the market price of risk (left panel), the interest rate (middle panel), and the root mean squared relative distance from the calibration targets (right panel) on the time-preference parameter β and the adjustment cost parameter ξ^{adj}.

β. The average real rate is decreasing in β and (somewhat depending on the region) increasing in the adjustment cost parameter. Note that an MPR of 55% or higher, together with a low real rate, is still attainable if one chooses higher values for the adjustment cost parameter and adjusts β.

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.[12] 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

[12]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.

	Model	Data	Rep. agent
r^{bond}	0.79%	0.8%	19%
MPR	50%	50%	12%
std output growth	2.6%	2.6%	2.6%
std r^{bond}	2.0%	1.8-2.9%	3.1%
std agg. consumption growth	2.0%	1.4-2.0%	2.2%
log V/E	2.8	2.8-3.0	1.8

Table 1: Key moments in model and data.

"Cyclically Adjusted Total Return Price Earnings Ratio P/E10".[13] Table 1 summarizes key statistics from our benchmark calibration of the model and compares them to the data and a representative agent specification.[14]

The first two targets are met almost exactly by construction. The volatility of the growth rates of real consumption per capita is 2% 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 mechanism already described in Chien et al. (2011). By construction,

[13] 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.

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

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 – we discuss the wealth distribution between and within classes of agents in detail in Section 4.2 below. 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. While this is not the focus of this paper, we also obtain a reasonable equity risk premium using standard assumptions on the leverage policy of firms. For example, with a mean debt-to-value ratio of 0.47 (implying a mean debt-to-capital ratio 0.52) we obtain that the average expected return on firms' equity is 4.9%, implying an equity risk premium of 4.1%.[15]

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 value of the loss function vary with the parameters β and ξ^{adj}. 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 below, our model setup allows us to match a wide range of possible interest rates while keeping the MPR constant and 50%. For this, three

[15] To obtain these numbers, we assume a constant level of (detrended) firm debt, such that the average debt-to-value ratio is 0.47.

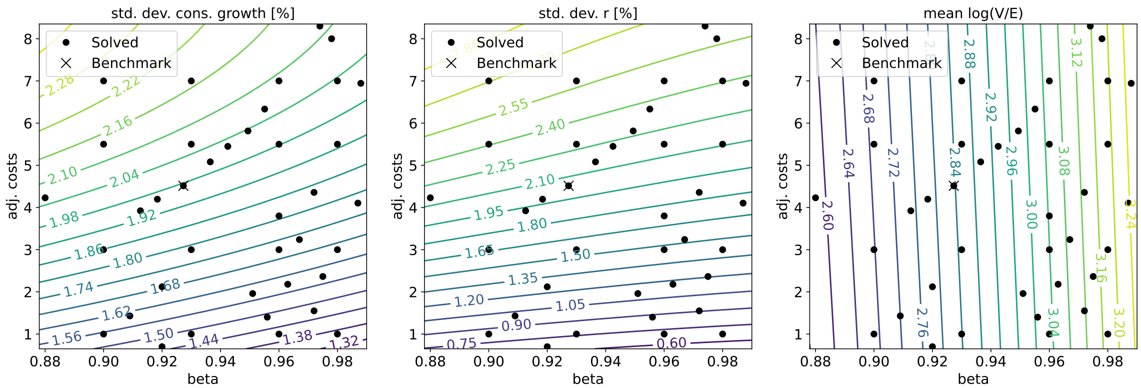

Figure 3: Standard deviation of consumption growth (left), standard deviation of the interest rate (middle), and the price-earnings ratio (right) for different values of the time-preference parameter (vertical axis), and the adjustment costs parameter (horizontal axis).

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.

3.4 Firms

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

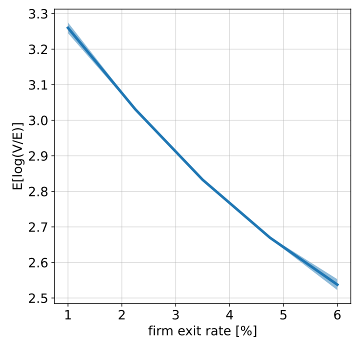

Figure 4: Mean log price-earnings ratio for different values of the firm exit rate. For a given firm exit rate, the adjustment costs and the patience are recalibrated to maintain an average market price of risk of 50% and an average interest rate of 0.8%. The shaded area shows the confidence interval of the prediction by the surrogate model.

recalibrate the adjustment costs and the patience to maintain a market price of risk of 50% 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.

4 Very low real rates

By varying the time-preference parameter, β as well as the adjustment cost parameter our model allows us to match a wide variety of (counter-factually low) interest rates with a realistic market price of risk and realistic volatilities of consumption- and output growth. This allows us to examine a range of economies where the average real rate range from -1.5% to +1.5% (with a +2% growth rate). As explained above, an average interest rate of 0.8% matches US averages well. An average real rate of -1.5% certainly seems to be a lower bound given the available data (although there have been historical episodes where the return on government bonds lay below -1.5%, this is certainly not a realistic average return).

Table 2 shows the resulting moments in our model for different interest rates, keeping the market price of risk at 50%. 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 and a lower average return to capital. Nevertheless, even for an average interest rate of -1.5%, dividends remain always positive. If we were to push the interest rate even lower, dividends would eventually become negative – the value of the firm would remain finite.

Surprisingly, the table shows our model matches output and consumption volatility for the entire range of real interest rates. The volatility of the real rate increases as the average real rate decreases as higher adjustment costs are necessary to obtain an MPR of 50% with a high value for β. The volatility of the return to capital also rises as a result.

β	ξ^{adj}	MPR	mean r^{bond}	std r^{bond}	mean $r^{\text{Firm, }K}$	std $r^{\text{Firm, }K}$	std Y_{t+1}/Y_t	std C_{t+1}/C_t
0.9740	8.300	49.9%	-1.55%	2.6%	1.5%	7.2%	2.6%	2.1%
0.9600	7.000	50.4%	-0.89%	2.5%	1.9%	6.7%	2.6%	2.1%
0.9551	6.333	49.9%	-0.60%	2.4%	2.0%	6.5%	2.6%	2.1%
0.9494	5.814	49.8%	-0.37%	2.3%	2.2%	6.0%	2.6%	2.0%
0.9425	5.449	50.1%	-0.06%	2.2%	2.4%	5.7%	2.6%	2.0%
0.9365	5.082	50.1%	0.33%	2.2%	2.6%	5.5%	2.6%	2.0%
0.9273	4.514	49.8%	0.79%	2.1%	2.9%	5.2%	2.6%	2.0%
0.9126	3.922	49.8%	1.47%	1.9%	3.4%	4.8%	2.6%	2.0%

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

We now want to examine the effects of such very low rates on the wealth distribution and on the prices of long-lived assets. In future tables will consider the exact same calibrations as in Table 2, and we will reference the different cases simply by the associated real rates, not repeating the adjustment cost and time-preference parameters that were used.

4.1 Asset prices

Decreasing interest rates will increase the price of long-lived assets. Shares in the firms are long-lived assets in our economy, but since we model firms' exit, they are not infinitely lived assets. In our model, we do not allow for the trade in infinitely lived assets, but we can obviously price such assets given our (unique) stochastic discount factor.

For this, we consider a fictitious console, an asset that pays a $(1+g)^t$ risk-free at any time t in the future. Columns three and give of Table 2 show how the price of the firms and the price of the console increase as the real rate decreases. Surprisingly

$\mathrm{E}\left[r^{\text{bond}}\right]$	$\mathrm{E}\left[\log(V/E)\right]$	$\mathrm{E}\left[p^{\text{console}}\right]$	$\mathrm{E}\left[\mathrm{pdv}(Lw)\right]$	$\mathrm{E}\left[\mathrm{pdv}(V)\right]$	$\mathrm{E}\left[\mathrm{pdv}(V^{\text{start-ups}})\right]$
-1.55%	3.16	524	455	3.83	47.2
-0.89%	3.05	195	171	3.62	19.3
-0.60%	3.02	162	140	3.52	16.1
-0.37%	2.99	133	115	3.45	13.6
-0.06%	2.94	106	90.8	3.34	11.2
0.33%	2.89	90	77.3	3.27	9.89
0.79%	2.84	73	61.8	3.15	8.31
1.47%	2.75	55	46.3	2.99	6.72

Table 3: Mean interest rates (first column), mean log price earnings ratios (second column), mean price of the console (third column), mean present discounted value of all present and future labor income (fourth column), mean value of operating firms (fifth column), and mean present discounted value of all present and future startups for different combinations of time-preference and adjustment costs parameters (as shown in table 2).

the price of an infinitely lived asset remains finite even with a real rate of -1.5%. We discuss this in detail in Section 5 below. The value-earnings ratio remains finite by our construction but increases substantially.

Somewhat surprisingly, the increase in the average value of the firm is much less pronounced. It increases from 2.99 at an interest rate of 1.5% to 3.83 at an interest rate of -1.5%. The reason for this is that the dividends of the firm decrease as the interest rate decreases - in fact, for some calibrations with very low interest rates these dividends can become negative (although this is not the case for the calibrations listed in Table 3). Note that when the interest rate is very low, dividends paid out in the far future greatly impact prices today. For the calibration with an interest rate of -1.5%, 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 making the assumption of firms exiting crucial.

While the price of current firms remains relatively stable as the interest rate becomes very low, the present discounted value of all future start-ups increases substantially. The same is true for the present discounted value of agents' future labor income. Crucially, the value of future start-ups and future labor income do not show up in our flow budget constraints, hence these constraints contain objects whose value does not increase substantially as the interest rate falls.

4.2 The wealth distribution

For calibrations with very low risk-free rates we need high values for the time preference parameter, β, which changes from 0.91 to 0.97, hence for the low-interest rate case, traders are substantially more patient. However, in equilibrium, this has differential effects on the wealth of Arrow traders versus bond-only traders. Table 4 shows how the wealth distribution across types changes across the different calibrations leading to different risk-free rates. The average wealth share of the Arrow traders increases as the interest rate decreases (recall that Arrow traders receive 10% of labor income). At the same time the variability of their wealth share increases. The differences in wealth between the best and worst aggregate state are very large for the Arrow traders but do not change with the real interest rate.

$\mathrm{E}\left[r^{\text{bond}}\right]$	$\mathrm{E}\left[\lambda^A\right]$	$\mathrm{Std}\left(\lambda^A\right)$	$\frac{\mathrm{E}[\lambda^A \mid z_t=\text{good}]}{\mathrm{E}[\lambda^A \mid z_t=\text{bad}]}$	$\mathrm{Std}\left(\frac{C^A_{t+1}}{C^A_t}\right)$	$\mathrm{Std}\left(\frac{C^B_{t+1}}{C^B_t}\right)$	$\mathrm{E}\left[\mu(\hat{a}_t=0)\right]$	$\mathrm{E}\left[\mu(\hat{b}_t=0)\right]$
-1.55%	35.5%	12.9%	2.8	8.5%	1.0%	1.1%	10.7%
-0.89%	33.4%	12.1%	2.8	8.6%	1.0%	1.3%	10.8%
-0.60%	32.0%	11.6%	2.8	8.7%	1.0%	1.4%	10.7%
-0.37%	31.0%	11.2%	2.7	8.5%	1.0%	1.5%	10.8%
-0.06%	30.5%	10.6%	2.7	8.3%	1.0%	1.7%	10.9%
0.33%	29.0%	10.6%	2.8	8.3%	1.0%	1.8%	10.9%
0.79%	27.6%	9.9%	2.7	8.3%	1.0%	2.0%	11.0%
1.47%	26.2%	9.0%	2.6	8.4%	1.0%	2.1%	11.3%

Table 4: Statistics on the wealth distribution for different combinations of time-preference and adjustment costs parameters (as shown in table 2). The first column shows the mean interest rate, the second column shows the mean wealth-share of advanced traders, and the third column shows it's standard deviation. The third column shows the ratio of the mean wealth share of advanced traders in the best relative to the worst aggregate shock. The fourth and fifth columns show the standard deviation of the growth of aggregate consumption for each of the two groups. The sixth and seventh columns show the fraction of constrained households for advanced traders (sixth column) and non-participants (seventh column).

Table 4 also shows the average fraction of agents within each type that are on their borrowing constraint. Clearly, a higher β implies that agents' borrowing constraints are less likely to be binding. On the other hand, the net present value of labor endowments exploding must mean that agents have a higher desire to borrow against the future and might be more likely to be constrained. The table shows that the effect is very different for Arrow traders compared to bond-only traders. The fraction of constrained bond-only traders stays roughly constant across our calibration, while the fraction of constrained Arrow traders falls by almost 50%.

The stark difference between the wealth of the two types can also be seen in Figure 5. The histogram of the average financial wealth of Arrow traders (by idiosyncratic shock) changes drastically as the interest rate changes from our benchmark of 0.8% to -1.5%. The wealth distribution of bond-traders only changes little.

These facts seem surprising at first. Although the interest rate decreases as β increases, the product $(1 + r)\beta$ actually increases (substantially) if we take r to be the average interest rate. Yet, the bond traders do not save (significantly) more. The

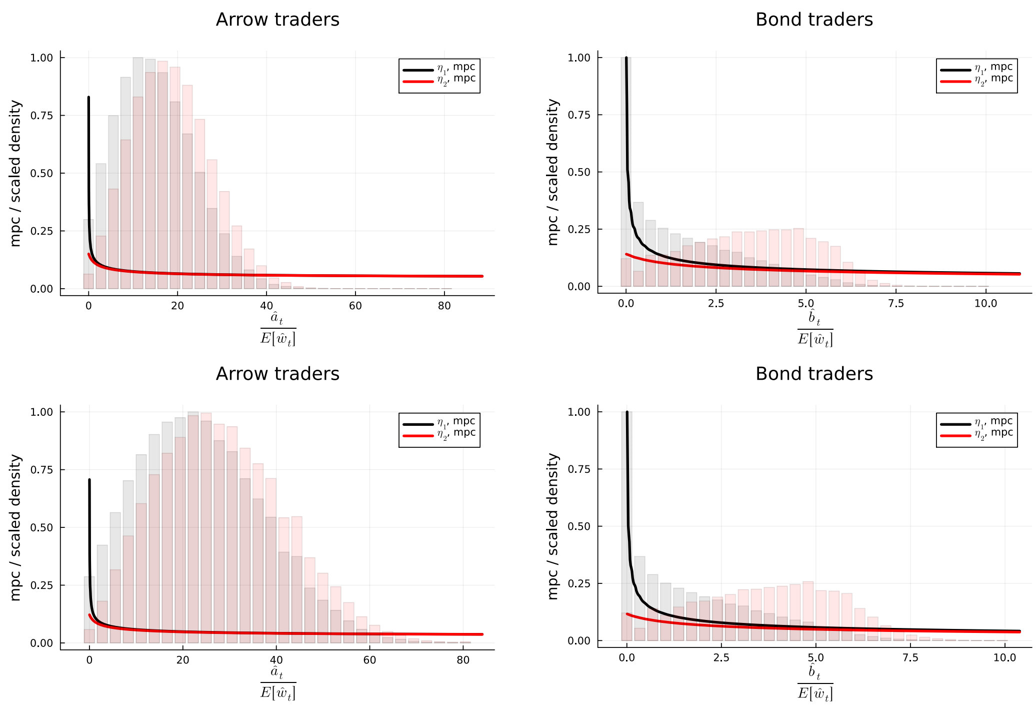

Figure 5: Histograms of the wealth distributions together with average MPC, by type and idiosyncratic shock. Upper panel for $r^{\text{bond}} = 0.79$, lower panel for $r^{\text{bond}} = -1.55$. In each panel, the histograms are normalized such that the highest histogram bin reaches one.

reason for this is that the interest rate volatility also increases as β increases. The correlation between the realized output growth rate and the current real interest rate lies around -0.35 in our calibration (and does not change much across the different cases). That is to say, when poor, the agents face a high interest rate, when rich the agents face a low interest rate. The increase in volatility of the interest rate implies that the bond agents do not increase savings substantially. The Arrow traders, on the other hand, save significantly more in the low interest rate calibration (as can be seen both in Figure 5 and Table 4). One important factor for this is the fact that the Arrow traders become much wealthier through the higher value of future start-ups discussed above. The other important factor is that the Arrow agents are, on average, much richer than the bond-only traders for all calibrations. Their MPC is lower and they save more out of additional income.

Finally, note from Table 4 that consumption growth volatility is very large for Arrow traders and does not change with the risk-free rate. As we mentioned above, this large consumption volatility is the main driver of the large market price of risk.

5 The fiscal costs of government debt

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. As we will review below, without structural assumptions on the economy, for example, just from observed prices) statements about the possibility of debt-rollover are impossible. To make statements, one needs to make strong assumptions on the underlying economy (i.e., specify endowments, preferences, technology, and a stochastic process for TFP). Clearly, there could be alternative specifications that also match our asset pricing facts and that produce the opposite conclusion.

A perhaps more important question is whether government debt can be Pareto-improving in our setting (see, e.g. Aguiar et al. (2021), Amol and Luttmer (2022) or Brumm et al. (2021)). Because of financial frictions, the fact that debt rollover is impossible does not imply anything about the welfare consequences of debt. We conduct a simple experiment where the government issues government debt and households can decide when to pay taxes to ensure that the debt is repaid in finite time. In our experiments, government debt is always detrimental to welfare because it induces crowding out of private savings and investment.

5.1 Debt rollover

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 (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 fraction of aggregate consumption for the infinite future, otherwise the value of those assets would be infinite. 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, we chose to introduce firm death in our economic model.

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

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

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.

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. It is then a quantitative question of how low interest rates can become with prices still being summable.

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 use the price of the (fictitious) console discussed above. Note that there are always cases where it is not decidable whether prices are summable - one cannot ensure 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.

In the last column of Table 2, $\mathrm{E}(p^{\text{console}})$ denotes the expected price for a console with payout $(1+g)^t$ in every period. We first note that for our benchmark calibration, prices are summable (and therefore, debt rollover is impossible). As Table 2 shows, as long as the MPR is held fixed at 50%, prices remain summable (i.e. debt rollover remains impossible) even when average interest rates are 3.5% 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[16] high MPR turns out to be crucial for this result. For example, with a coefficient of relative risk aversion of 5.5, $\beta = 0.972$, and a relatively low adjustment cost parameter of 1.55, we obtain a market price of risk of 29.9%. The average real rate is 0.47% and its standard deviation is 0.9%. In this economy, prices are not summable, making debt rollover possible.

To understand better the interacting effects of the average interest rate and the market price of risk we construct a surrogate model that maps average interest rates and the average market price of risk into the price of the console. 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 console price that we do not consider here. Figure 6 shows the price of the console for different combinations of the interest rate and the MPR. For lower interest rates computing the console'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 console is increasing faster than exponentially as the mean interest rate decreases. To see this,

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

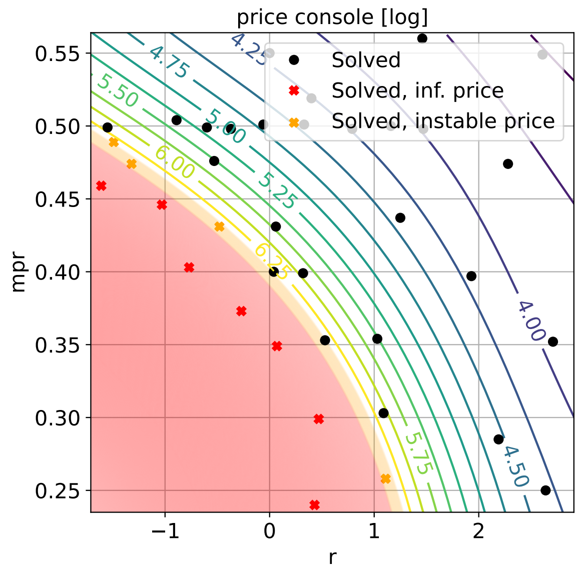

Figure 6: Mean price of a console, with a payout equal to $(1+g)^t$ in logs, plotted against the mean interest rate and the mean market price of risk.

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% 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.

5.2 Welfare effects of debt

Although in our calibration infinite rollover of debt is impossible and therefore, there must be a fiscal cost of debt, it is not clear what the welfare consequences of government debt are. The financial frictions imply that the competitive equilibria is suboptimal and there are no theoretical arguments to rule out the possibility that government debt could be Pareto improving. There are two key frictions in the model: (i) a lack of insurance for idiosyncratic risk for all households and aggregate risk for bond-only traders, and (ii) the borrowing constraint.

State contingent government debt could help with (i) but we focus here on the ability of a simple government policy intervention that can relax the borrowing constraint (from zero) through increasing the amount of risk free assets in the model. To do so, consider a government policy intervention of the following form. In the first period of the model, $t = 0$, the government gives each household some units of a government bond which we denote by b_0^g. Assume that thereafter, the government increases its bond issue by an amount so that the total for any household is

$$b_t^g = (1+g)b_{t-1}^g$$

thus keeping the growth deflated size of the debt held by a household constant. The government pays for its debt issuances by levying on the households lump-sum taxation in the future whose present-value is equal to the value of the issued bonds. By construction this is a purely Ricardian policy. Assume that the household can treat the government bonds as collateral and therefore short the government bond. Assume also that each household is free to devise their own optimal state-contingent repayment plan.

For the bond-only traders, this would alter their intertemporal budget constraint as follows.

$$\sum_{z^t} p(z^{t+k})c_t(z^t) = \sum_{z^t} p(z^{t+k}) \left[l^B(\eta_t)\omega(z^t) + b_{t-1} - b_t p^b(z^t) + \Delta b_t^g - T(z^t) \right] \quad (46)$$

where the last two terms involve the new debt transfer in period t, Δb_t^g, and the state-contingent repayment $T(z^t)$. Clearly for these two terms to net to zero, we simply need the original present-value budget constraint to hold. However, the ability to borrow against one's holding of government debt lowers the borrowing constraint by its level. So long as

$$E_t\left[\lim_{k\to\infty} p(z^{t+k}|z^t)(1+g)^t\right] = 0$$

with probability 1, the requirement that they stay weakly above the bond threshold $-b_0^g$ implicitly implies that there are repaying the government's bond transfer. This is equivalent to the requirement that the price of the government's growth consul stay finite. A similar argument holds for the advanced-traders. Therefore, if debt rollover is impossible, the above-described debt policy has the same effect as lowering the borrowing constraint of the household in each period from 0 to $-b_0^g$.

Clearly, a relaxation of their borrowing constraints will potentially increase the welfare of the bond traders. However, there is also a price-effect. The sign and magnitude of the overall effect is then a quantitative question. Here we present the results from lowering the borrowing constraint gradually to -.20 which can be roughly thought of in terms of percent mean consumption since this value is roughly one. We consider three calibrations for the average risk-free rate for the zero-borrowing case. As above, in the benchmark case the average interest rate is around 0.6 percent, we also consider a zero interest rate case as well as a case with an interest rate of 1.5 percent.

Because we change the preference parameter β to lower our interest rates, we cannot compare results across different calibrations, but only with respect to the impact of the lower borrowing constraint given a calibration. We are interested in how the change in the borrowing constraint impacts on asset prices, and through these prices, the level of capital. In addition, we are interested in welfare. Since preferences are given by $\mathrm{E}\left[(\right]1-\beta)\sum_t \beta^t u(c_t)$, we report the relative level of *normalized* certainty equivalent consumption level in the table. In table 5 we present some results for the relaxing the borrowing constraint for three calibrations, the benchmark economy and both a higher and lower interest rate economy.

Table 5: Impact of Government Policy to Lower Borrowing Constraint

Benchmark model								
$\underline{b}$	r	$\hat{K}$	$\hat{Y}$	λ	MPR	rel. $c^{*,A}$	rel. $c^{*,B}$	$\log(V/E)$
0.00	0.785%	2.838	1.395	27.7%	50.0%	1	1	2.837
-0.05	0.917%	2.806	1.390	27.8%	49.6%	0.9972	0.9992	2.819
-0.08	0.991%	2.790	1.387	28.0%	49.3%	0.9959	0.9989	2.810
-0.10	1.044%	2.777	1.385	28.0%	49.2%	0.9943	0.9985	2.802
-0.15	1.176%	2.747	1.380	28.1%	48.8%	0.9913	0.9974	2.786
-0.20	1.294%	2.724	1.376	28.2%	48.3%	0.9861	0.9966	2.770
Low interest rate								
$\underline{b}$	r	$\hat{K}$	$\hat{Y}$	λ	MPR	rel. $c^{*,A}$	rel. $c^{*,B}$	$\log(V/E)$
0.00	0.03%	2.980	1.418	30.0%	50.3%	1	1	2.932
-0.05	0.16%	2.950	1.413	30.1%	49.8%	0.996	1.000	2.914
-0.08	0.23%	2.932	1.410	30.2%	49.6%	0.994	1.000	2.904
-0.10	0.29%	2.917	1.408	30.2%	49.4%	0.992	1.000	2.895
-0.15	0.42%	2.885	1.402	30.3%	49.0%	0.991	0.999	2.878
-0.20	0.55%	2.853	1.398	30.4%	48.6%	0.988	0.998	2.860
Higher interest rate								
$\underline{b}$	r	$\hat{K}$	$\hat{Y}$	λ	MPR	rel. $c^{*,A}$	rel. $c^{*,B}$	$\log(V/E)$
0.00	1.47%	2.703	1.373	26.1%	50.0%	1	1	2.751
-0.05	1.61%	2.674	1.368	26.3%	49.6%	0.997	0.998	2.735
-0.08	1.69%	2.655	1.365	26.3%	49.4%	0.994	0.997	2.724
-0.10	1.74%	2.643	1.363	26.4%	49.2%	0.994	0.997	2.718
-0.15	1.88%	2.611	1.357	26.5%	48.9%	0.991	0.995	2.701
-0.20	2.01%	2.583	1.353	26.7%	48.5%	0.989	0.992	2.685

As can be seen from the table, government debt is never Pareto-improving. On the contrary, it typically makes all agents worse off (in the low-interest case, the welfare effects for the bond-traders are so small that they become insignificant). The table also shows the effects of government debt on the average capital stock. Government debt crowds out private savings and hence decreases the capital stock. Since in our calibrations the economy is *dynamically efficient* (meaning that a reduction in capital reduces aggregate consumption), average welfare decreases for both agents. It is well established in the literature (see, e.g., Bloise and Reichlin (2022)) that the possibility of debt rollover is a necessary but by far not sufficient condition for dynamic inefficiency. The table also shows that the risk-free rate increases substantially and the market price of risk decreases.

6 Transitory versus Permanent Shocks

As noted by Alvarez and Jermann (2005) the pricing of long-term risk-free bonds is very sensitive to the stochastic process for consumption in a representative agent model. They show that if consumption follows an AR1 process around a deterministic trend, then the one-period holding return of a long-term risk-free bond is in fact very risky. They show that this seems inconsistent with the observed term structure of US government bonds, which, on average seems to flatten out after 20 years. If, on the other hand, it is assumed that consumption follows a stochastic growth process, the resulting term structure seems much more consistent with US data. This fact obviously puts in question the validity of our calibration. Our results on the possibility of debt rollover crucially depend on the shocks being transitory.

There is a large literature that examines whether there is a unit root in the time series of aggregate consumption or GNP. Cochrane (1988) develops a measure of the persistence of fluctuations in GNP based on the variance of its long differences. He shows that this measure finds little long-term persistence in GNP. On the other hand, Cochrane (1994) points out that there is a significant unit root component in consumption. Chernov et al. (2021) show that consumption dynamics is best explained with a switching model with both transitory and aggregate states.

6.1 Evidence of a Unit Root in TFP, Output and Consumption

Cochrane (1988) examines the variance of the log of a time series, x_t with difference k; $var(x_t - x_{t-k})/k$; this measure should settle down to the variance of the unit root innovation with large enough k. The first panel of Figure 7 depicts Cochrane's measure

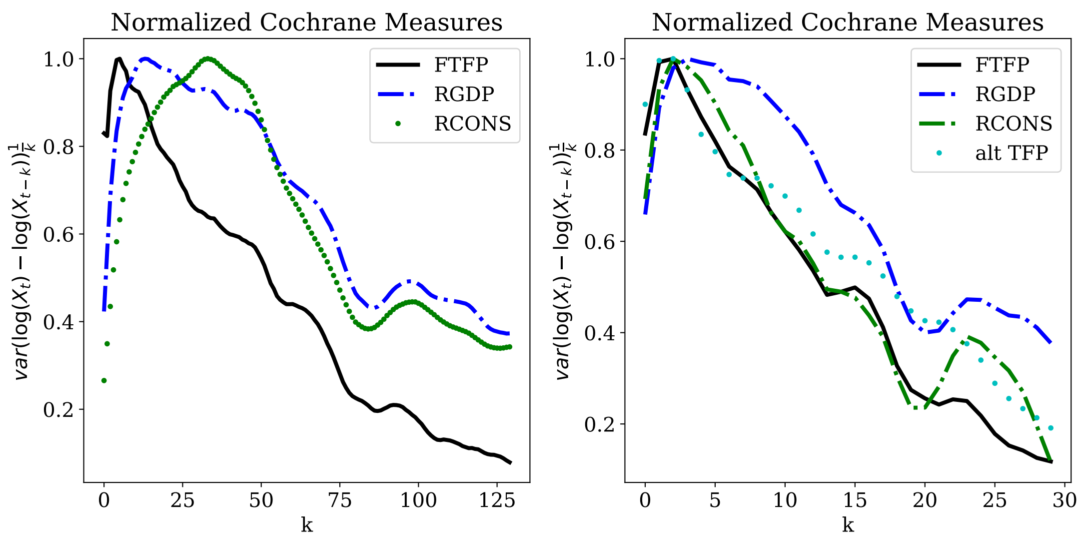

(a) Quarterly Data **(b)** Annual Data

Figure 7: Cochrane's Measure for TFP, GDP, and Consumption

for quarterly data 1960:Q1-2019:Q4 on TFP along with real GDP and consumption. The TFP measure is from Fernald (2014). As one can see from the figure, the evidence for a unit root in TFP appears more limited than for output or consumption. Because our model has an annual frequency, we have also plotted the same series on an annual basis in the second panel. In addition, we have added an alternative aggregate TFP series from FRED.[17] Overall, Figure 7 suggests that there may be a small unit root in consumption and output (though the measure is trending down) but at most a tiny

[17] University of Groningen and University of California, Davis, Total Factor Productivity at Constant National Prices for United States [RTFPNAUSA632NRUG], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/RTFPNAUSA632NRUG, September 22, 2023.

one in TFP. We will return to this evaluation once we incorporate a unit root into our benchmark model.

6.1.1 Evidence from the Real Yield Curve

Alvarez and Jermann (2005) demonstrate that asset pricing data points to a large unit root in consumption. The yield curve on pure discount bonds, especially at very long horizons, is very informative as to the size of the unit root in the representative agent's consumption. They focus on the implications of nominal bond yields for the real term structure – in Appendix D we review their argument in some detail.

In what follows, we focus on inflation-indexed bonds and ask whether our model can replicate the observed yield curve. Unfortunately, TIPS (Treasury inflation index securities) only became available starting in 2003 (and only for the 5 and 10-year bonds and even later for the longer term ones). We plot the yields on TIPS in figure 8. As one can plainly see, during this period the yield curve was almost always positively sloped, especially at the longer horizons. If we compare the yields during the period in

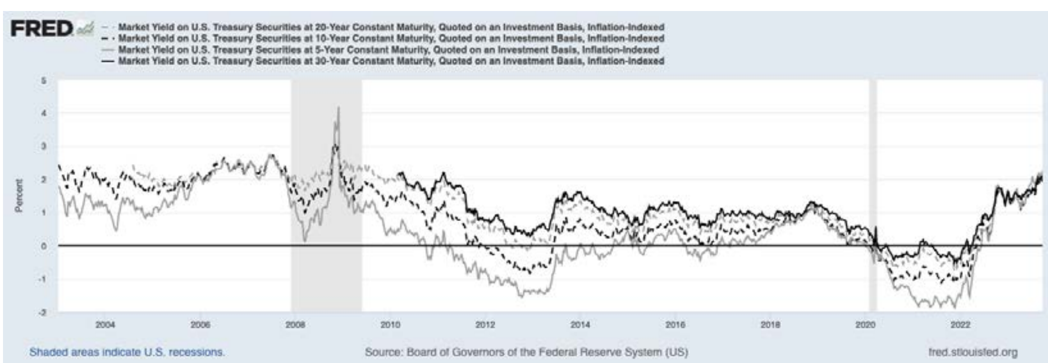

Figure 8: TIPS Yields

which we also have the longer bond and examine the fraction of the weeks in which the next longer bond yield was higher, the yield on a weekly yield on a 10-year bond was above that on a 5-year 87% of the time, on a 20-year vs. a 10-year 99% of the time, and on a 30-year vs. a 20-year 100% of the time. Over the period from 2004-2023, the average yield on the 5-, 10-, and 20- year bonds were 1.73, 1.99, and 2.19 respectively. Over the period from 2010 onwards, when we also have 30 year bonds, the respective averages were 1.38, 1.77, 2.03 and 2.14. The fact that the average yield curve for TIPS is upward sloping is also discussed in Chernov et al. (2021).[18]

Let us mention several cautionary notes about the TIPS evidence. First, this data has a limited time scale. Second, there are concerns about the pricing of TIPS: e.g. Haubrich et al. (2012) use inflation swap rates and inflation surveys, and find that there was substantial under-pricing of TIPS during their early years (prior to 2004) and again during the financial crisis of 2008-9. Third, during this period QE2 was widely thought to have lowered yields on medium and longer-term Treasuries (see, e.g. Vissing-Jorgensen and Krishnamurthy (2011)).

6.2 Quantitative Assessment with the Model

To examine the impact of adding in a stochastic growth component, we modified our model as follows. To keep the same three state Markov chain for the exogenous aggregate state, we add a growth component which is determined by the realization of the aggregate temporary shock z_t. In particular, we assume that if $z_t = z_j$ then the growth rate between t and $t+1$ is given by g_j, and the good (bad) temporary productivity state is associated with a higher (lower) long-run growth rate shock. Because our Markov chain is symmetric, each state is ex-ante equiprobable. Hence, we simply consider a shift factor for the growth rate $g \in [1.02 + x, 1.02, 1.02 - x]$ for $x \geq 1$. We consider two cases: small shocks with $x = 0.002$ and large shocks with $x = 0.006$. Note that

[18]The TIPS yields also include a coupon which distorts their comparison to the yields on a pure discount bond. The minimum yield on a TIPS bond is 0.125%, and the average yield over the whole period across the different maturities was 0.79%. We computed the pure discount bond yield when one adjusts for the average coupon rate and found this correction had a negligible effect.

we are adding volatility with these experiments since the transitory shock process is unchanged. Since we do not re-calibrate the model to match moments, the following results should be thought of as only indicative.

In Figure 9 we plot Cochrane's measure for our three cases, which we can compare to the results in Figure 7. Note that because the variance of the growth rate coming from the our specification of the growth process is $2 \times x^2/3$, in all of our cases the long-run level of the measure is very close to zero. In the short-run the benchmark and small growth shock cases reasonably closely match the data. In the large shock case, all three measures (TFP, Y, C) are quite high at the 30 year level, suggesting that the unit root here is too large relative to the data. Overall this figure highlights both how minor tweaks can generate measures in the data that suggest a unit root at modest horizons, and also the problematic nature of these measures; something that is well-known in the unit root literature (see Phillips and Xiao (1998) discussion of finite sample properties of unit root tests).

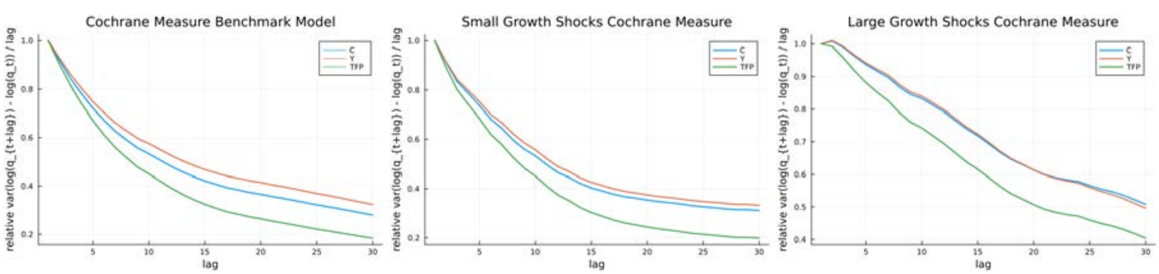

(a) Benchmark Model **(b)** Small growth shocks Model **(c)** Large growth shocks Model

Figure 9: Cochrane's Measure for Our Model

In Table 6 we report some selected results, and Figure 10 plots the yield curve for all three cases. One can see that the introduction of a unit root had the anticipated effect of flattening the yield curve and thus lowering the term risk premia. It also raised the one-period interest rate. We also report average yields from the data for TIPS for two different sample periods.

Focusing first on the yield curve data, because the level of interest rates shifts around, the proper measure is probably the relative yield curve (i.e. the yield relative to the shortest rate). In the data the relative yield curve in average yields for the longer period in which we do not have the 30-year bond is $[1.0, 1.15, 1.26]$ and for the shorter period with 30-year bonds it is $[1.0, 1.28, 1.47, 1.55]$. This data indicates that the yield curve, even in average terms, moves around, and also gives us a rough notion of the relative step size. In the model data the relative yield curve for the benchmark model values are $[1.0, 1.32, 1.61, 1.77]$, for the small growth shock version $[1.0, 1.25, 1.49, 1.61]$, while for the large growth shock model $[1.0, 1.20, 1.39, 1.47]$. The small and the large growth shocks bracket the relative yield curve in the data, and both reasonably approximate it (given the data issues mentioned above).[19]

Version	r	MPR	5-year	10-year	20-year	30-year	$\frac{Var(gr\_C)}{Var(gr\_Y)}$	$\frac{Var(gr\_I)}{Var(gr\_Y)}$
Benchmark	0.86%	49.8%	1.68%	2.21%	2.71%	2.98%	0.60	1.99
Small Shocks	1.07%	49.5%	1.76%	2.20%	2.62%	2.83%	0.65	1.78
Large Shocks	1.17%	48.5%	1.62%	1.95%	2.25%	2.38%	0.80	1.32
TIPS data 2004-2023			1.73%	1.99%	2.19%	n/a		
TIPS data 2010-2023			1.38%	1.77%	2.03%	2.14%		

Table 6: The TIPS sample mean in the longer sample is from 7/30/2004 to 9/22/2023, while in the sample with the 30-year bond data, it is from 2/26/2010 on.

In addition to its impact on the yield curve, the introduction of growth shocks impacted the model's fit relative to the data in several ways. First, it both raised the interest rate and lowered the market price of risk, especially for the large shock growth shock case. Matching risk pricing without major tweaks to the model has long been recognized as a major challenge. Second, it affected the relative volatility of consumption and investment. In our benchmark model the relative variance of

[19] The differences between our conclusion that the yield data suggests a small permanent component and AJ's conclusion that it has a large component comes from the fact that our real yield curve is upward sloping all the way out to thirty years, while in their data it is not.

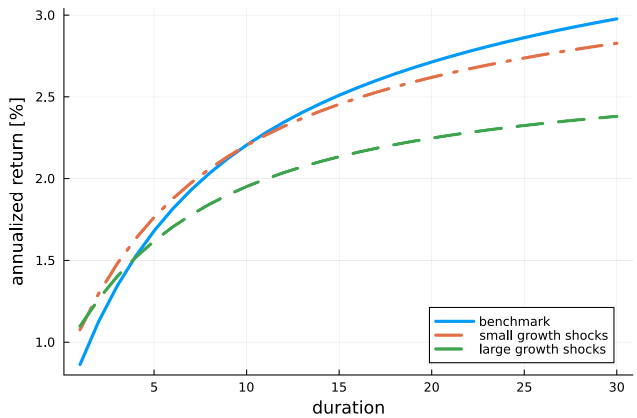

Figure 10: Yield Curve with growth shocks.

consumption-to-output, C/Y =0.60, while investment-to-output is I/Y = 1.99. By the time we get to the large shocks it has switched to C/Y = 0.80 and I/Y = 1.32. Given that investment volatility is low relative to the data in the benchmark model, this seems like an important trade-off with production based asset pricing models.[20] Since we simply added the long-run shocks to the model, this highlights the extent to which growth shocks directly dampen risk pricing and investment volatility.

7 Conclusion

The main departure of our paper from a standard asset pricing model with heterogeneous households that face idiosyncratic risk and with limited asset-market participation is incorporating firm production, exit, and entry. This allows us to examine economies with very low real rates. We need to consider three present-value bins: (i) labor income, (ii) net income from existing firms, and (iii) net income from firms entering at later dates. Because of household borrowing constraints, the operational budget constraint of the household only included current assets and current labor income.

[20]Since we do not have government spending and this component is not that volatile, hitting consumption volatility means necessarily being too low on investment volatility.

Hence, as long as the value of currently existing firms remains finite, the behavior of the model is surprisingly standard even with very low real interest rates.

We target different values for the average interest rate while holding fixed the market price of risk (which measures the extent of that which aggregate risk is priced). To achieve these targets, we adjust the preference discount factor and the adjustment costs on capital. The higher discounting induced by firm exit means that dividends stay positive even as the interest rate falls, and existing firm values do not explode. However, the horizon over which they must accurately be forecasted extends beyond 200 years making accurate numerical methods crucial for our analysis.

On the household side, while the average interest rate times the discount rate, $(1+r)\beta$, rises, the increase in the volatility of interest rates is sufficient to constrain the incentive to increase savings. As a result, the standard deviations of aggregate consumption and output growth are essentially unchanged. For our sophisticated traders, the standard deviation of consumption growth does rise along with their wealth share, but only modestly.

It is important to note here that because we go beyond the standard two-aggregate state Markov chain structure,[21] our advanced traders benefit more from their sophistication than in a model where they can only trade a bond and equity. Relaxing this restriction has significant implications regarding the exposure of households' wealth to specific aggregate shocks.

We considered the impact of the fall in interest rates on a growth console. Starting from the benchmark average real interest rate of 0.80%, reducing the real rate to -0.89% results in the real price of the console increasing by a factor of slightly more than 150%, while the present value of labor income grows by 177%. Despite this, the value of existing firms increases by only 15%. These results indicate both the

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

surprising extent to which present values stay finite in our model even at quite low rates and that the dampening effect of firm death is large. Unsurprisingly, we find that this dampening effect is quite sensitive to the firm exit rate. Lowering the exit rate from 3.5% to 1% results in clearly counterfactual values of the price-earnings ratio in our 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 (2019). 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,[22] would be problematic for two main reasons, which we illustrate in figure 11. Figure 11 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

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

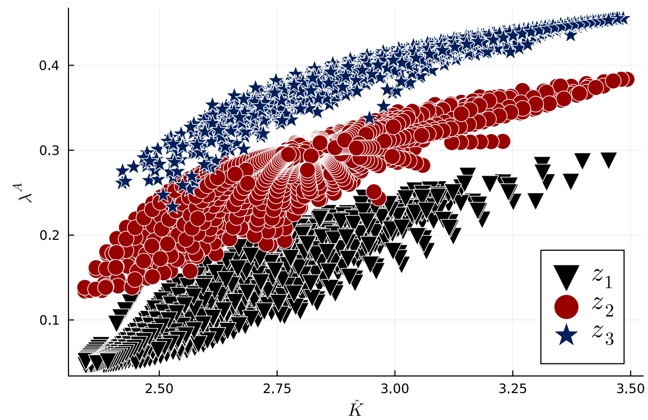

Figure 11: Realizations of the two endogenous aggregate variables, capital (horizontal axis) and the wealth share (vertical axis) of the advanced traders, for each of the exogenous aggregate shocks on a simulated path of 10000 periods for the benchmark model.

transitions for the endogenous aggregate variables are only observed on the simulated path, there is no obvious way how to extend a fitted forecasting 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); Auclert et al. (2021)). For a given value of asset holdings, which lie 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 (detrended) aggregate capital $\hat{K}_t$, 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

$$g_K\left(z_t, \hat{K}_t, \lambda_t^A\right) \approx \hat{K}_{t+1} \tag{47}$$

$$g_{V^{\text{firm}}}\left(z_t, \hat{K}_t, \lambda_t^A\right) \approx \hat{V}_t^{\text{firm}} \tag{48}$$

$$g_{\lambda^A}\left(z_t, \hat{K}_t, \lambda_t^A, z_{t+1}\right) \approx \lambda_{t+1|z_{t+1}}^A \tag{49}$$

$$g_{c^A}\left(z_t, \hat{K}_t, \lambda_t^A, \eta_t, \frac{\hat{a}_{t-1}^{z_t}}{1+g}\right) \approx \hat{c}_t^A \tag{50}$$

$$g_{c^B}\left(z_t, \hat{K}_t, \lambda_t^A, \eta_t, \frac{\hat{b}_{t-1}}{1+g}\right) \approx \hat{c}_t^B. \tag{51}$$

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 representative 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 (9). Where z_t, $\hat{K}_t$, and λ_t^A are given and $\hat{K}_{t+2}$ is obtained from the approximating function $\hat{K}_{t+2} = g_K\left(z_{t+1}, \hat{K}_{t+1}, \lambda_{t+1}^{A \text{ forecast}}\right)$. The wealth share of advanced traders in period $t+1$ is obtained by using the forecasting function $\lambda_{t+1}^{A \text{ forecast}} = g_{\lambda^A}\left(z_t, \hat{K}_t, \lambda_t^A, z_{t+1}\right)$.

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 $g_{c^A}^{z_t, \eta_t, \hat{a}_{t-1}^{z_t}/(1+g)}\left(\hat{K}_t, \lambda_t^A\right)$ for each combination of the discretized exogenous shocks z_t and η_t, and each asset holding on the pre-specified grid $\hat{a}_{t-1}^{z_t}/(1+g) \in \mathcal{A}^{\text{grid}}$, 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 $g_{c^A}^{z_t, \eta_t, \hat{a}_{t-1}^{z_t}/(1+g)}\left(\hat{K}_t, \lambda_t^A\right)$, $g_{c^B}^{z_t, \eta_t, \hat{b}_{t-1}^{z_t}/(1+g)}\left(\hat{K}_t, \lambda_t^A\right)$, 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 $\hat{b}_{t-1}^{z_t}/g$ on the asset grid $\mathcal{B}^{\text{grid}}$ with $\mu_t^{B,\eta}$ and we denote masses of advanced traders with idiosyncratic shock η and asset holdings $\hat{a}_{t-1}^{z_t}/(1+g)$ on the asset grid $\mathcal{A}^{\text{grid}}$ with $\mu_t^{A,\eta}$. 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:[23]

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 $\hat{b}_{t+1}^{z_t} \in \mathcal{B}^{\text{grid}}$

$$\hat{c}_{t+1}^{B,z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}} = g_{c^B}^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}\left(\hat{K}_{t+1}, \lambda_{t+1}^{A,\text{ forecast}}\right) \tag{52}$$

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

$$(V^{B'})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}} = u'(\hat{c}^{B,z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}})u'(1+g) \tag{53}$$

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 $\mathcal{B}^{\text{grid}}$.

$$(EV^{B'})_t^{\eta_t,\mathcal{B}^{\text{grid}}} = \sum_{z_{t+1},\eta_{t+1}} \pi^z(z_t, z_{t+1})\pi^\eta(\eta_t, \eta_{t+1})(V^{B'})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}} \tag{54}$$

4. Compute the consumption consistent with the households Euler equation and the

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

households savings choice lying on the grid $\mathcal{B}^{\text{grid}}$.

$$\mathcal{C}_t^{\eta_t,\mathcal{B}^{\text{grid}}} = (u')^{-1}\left(\frac{\beta}{p_t^{\text{bond}}}(EV^{B'})_t^{\eta_t,\mathcal{B}^{\text{grid}}}\right) \tag{55}$$

5. The obtained tuple, $\left(\mathcal{C}_t^{\eta_t,\mathcal{B}^{\text{grid}}}, \mathcal{B}^{\text{grid}}\right)$ 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 $(\hat{c}_t^{B,\eta_t}, \hat{b}_{t+1}^{\eta_t}) \in \left(\mathcal{C}_t^{\eta_t,\mathcal{B}^{\text{grid}}}, \mathcal{B}^{\text{grid}}\right)$

$$\hat{b}_t^{\eta_t} = \hat{c}_t^{B,\eta_t} + p_t^{\text{bond}}\hat{b}_{t+1}^{\eta_t}(1+g) - \hat{w}_t\eta_t \tag{56}$$

   The resulting tuple $\left(\hat{b}_t^{\eta_t}, \hat{b}_{t+1}^{\eta_t} \in \mathcal{B}^{\text{grid}}\right)$ 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-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

$$d_t^B = \sum_{\eta_t} \sum_{\hat{b}_t^{\eta_t} \in \mathcal{B}^{\text{grid}}} \mu^{B,\eta_t}(\hat{b}_t^{\eta_t}) \cdot \hat{b}_{t+1}^{\text{new},\eta_t}(\hat{b}_t^{\eta_t})(1+g) \tag{57}$$

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 $\hat{a}_{t+1}^{z_{t+1}} \in \mathcal{A}^{\text{grid}}$

$$\hat{c}_{t+1}^{A,z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} = g_{c^A}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}}\left(\hat{K}_{t+1}, \lambda_{t+1}^{A,\text{ forecast}}\right) \tag{58}$$

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

$$(V^{A'})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} = u'(\hat{c}^{A,z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}})u'(1+g) \tag{59}$$

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[24] 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}.

$$(EV^{A'})_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}} = \sum_{\eta_{t+1}} \pi^{\eta}(\eta_t, \eta_{t+1})(V^{A'})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} \tag{60}$$

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

[24] 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.

on the aggregate shock z_{t+1}.

$$\mathcal{C}_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}} = (u')^{-1}\left(\frac{\beta}{p_t^{z_{t+1}}}(EV^{A'})_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}}\right) \tag{61}$$

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-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

$$\hat{s}_t^{\text{common},\eta_t} = \sum_{z_{t+1}} p_t^{z_{t+1}} \hat{a}_{t+1}^{z_{t+1},\text{common},\eta_t}(1+g). \tag{62}$$

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

$$\hat{a}_t^{\text{common},\eta_t} = \hat{c}_t^{\text{common}} + \hat{s}_t^{\text{common},\eta_t} - \hat{w}_t\eta_t - \frac{1-\Gamma}{\Gamma\mu}\hat{\zeta}_t, \tag{63}$$

   where $\hat{\zeta}_t$ denotes the payoff of creating a start-up and the $\frac{1-\Gamma}{\Gamma\mu}$ denotes the start-ups 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

$$d_t^{A, z_{t+1}} = \sum_{\eta_t} \sum_{\hat{a}_t^{z_t, \eta_t} \in \mathcal{A}^{\text{grid}}} \mu^{A, \eta_t}(\hat{a}_t^{z_t, \eta_t}) \cdot \hat{a}_{t+1}^{\text{new}, z_{t+1} \eta_t}(\hat{a}_t^{z_t, \eta_t}) g \tag{64}$$

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

$$d_t^{A, z_{t+1}} + d_t^B = g_{V^{\text{firm}}}\left(z_{t+1}, \hat{K}_{t+1}, \lambda_{t+1}^{A \text{ forecast}}\right). \tag{65}$$

Together with the optimality condition for firm investment (*i.e.* equation (9)), this gives us four nonlinear equations to solve for the three prices and the next period's capital. Once the system of nonlinear equations is solved, we record 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

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

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 Analogously, 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

$$g_{\lambda^A}^{z_t,z_{t+1},\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\lambda_{t+1}^{A,z_{t+1},\text{pred}} + \alpha^{\text{update}}\lambda_{t+1}^{A,z_{t+1}}. \tag{67}$$

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.

$$g_{c^A}^{\eta_t, z_t, \hat{a}_t, \text{ new}} \left(\hat{K}_t, \lambda_t^A \right) \approx (1 - \alpha^{\text{update}}) \hat{c}_t^{A, \eta_t, z_t, \hat{a}_t, \text{pred}} + \alpha^{\text{update}} \hat{c}_t^{A, \eta_t, z_t, \hat{a}_t}, \tag{68}$$

$$g_{c^B}^{\eta_t, z_t, \hat{b}_t, \text{ new}} \left(\hat{K}_t, \lambda_t^A \right) \approx (1 - \alpha^{\text{update}}) \hat{c}_t^{B, \eta_t, z_t, \hat{b}_t, \text{pred}} + \alpha^{\text{update}} \hat{c}_t^{B, \eta_t, z_t, \hat{b}_t} \tag{69}$$

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 (22), 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

$$g_V^{z_t, \text{ new}} \left(\hat{K}_t, \lambda_t^A \right) \approx (1 - \alpha^{\text{update}}) \hat{V}_t^{\text{pred}} + \alpha^{\text{update}} \hat{V}_t^{\text{div}}. \tag{70}$$

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

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

$$g(\hat{K}, \lambda^A) = \theta_1 + \theta_2 \hat{K} + \theta_3 \hat{K}^2 + \theta_4 \lambda^A + \theta_5 (\lambda^A)^2 \tag{71}$$

and obtain the parameters $\{\theta_1, \theta_2, \theta_3, \theta_4, \theta_5\}$ by least-squares fitting the data.[25] 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.

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 5000 periods and the corresponding value predicted by our function approximators, which we obtained in the previous iteration. Table 7 shows that the remaining errors are low, with the 99th percentile well below 0.1%. Similarly, Table 8 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 and the exogenous shock 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 12 shows a scatter plot of capital against the state price for shock z_3, when the economy is in shock

[25] 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	0.07	0.01	0.04	0.01	0.02

Table 7: 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 5000 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 [%] $z_3 \to z_1$ $z_3 \to z_2$ $z_3 \to z_3$
			0.01	0.01	0.01
Mean	0.00	0.01	0.01	0.01	0.01
			0.01	0.01	0.01
			0.02	0.02	0.02
90th percentile	0.00	0.02	0.02	0.02	0.02
			0.02	0.02	0.01
			0.03	0.04	0.03
99th percentile	0.00	0.03	0.03	0.04	0.03
			0.03	0.04	0.03

Table 8: Statistics for the forecasting errors for aggregate capital and the wealth-share of the advanced traders on a simulated path of 5000 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%.

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 based on capital and the simple functional form we are using, provides a very good fit to the data.[26]

[26]While this serves as an illustrative example, our algorithm does not require to forecast prices, but the households' consumption. In figure 13 we show that household consumption can be very accurately predicted as well.

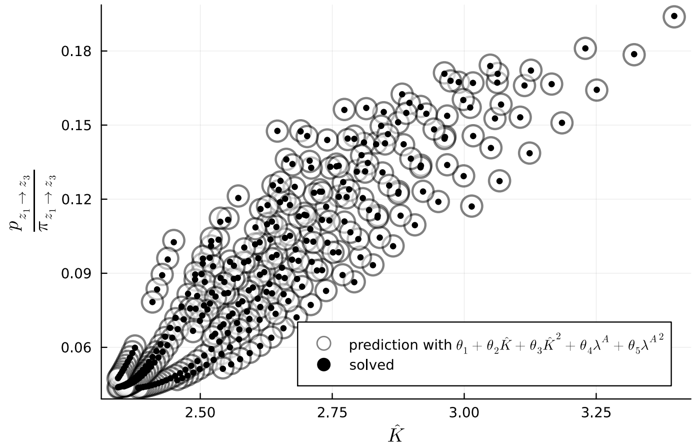

Figure 12: Realizations for the price of the aggregate-state contingent security for shock z_3, when the economy is in shock z_1, against the corresponding values of aggregate capital. The grey circles show the predictions of a fitted polynomial of the form $\theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2$.

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.

$$g_{c^A}^{\eta_t, z_t, \hat{a}_t, \text{ new}}\left(\hat{K}_t, \lambda_t^A\right) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{72}$$

$$g_{c^B}^{\eta_t, z_t, \hat{b}_t, \text{ new}}\left(\hat{K}_t, \lambda_t^A\right) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{73}$$

Importantly, the parameters θ_i can differ not only across the exogenous shocks, but also across the wealth level. Figure 13 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 $\hat{a}_t^{z_2} = 14.289$, which corresponds to roughly 15 times the average annual

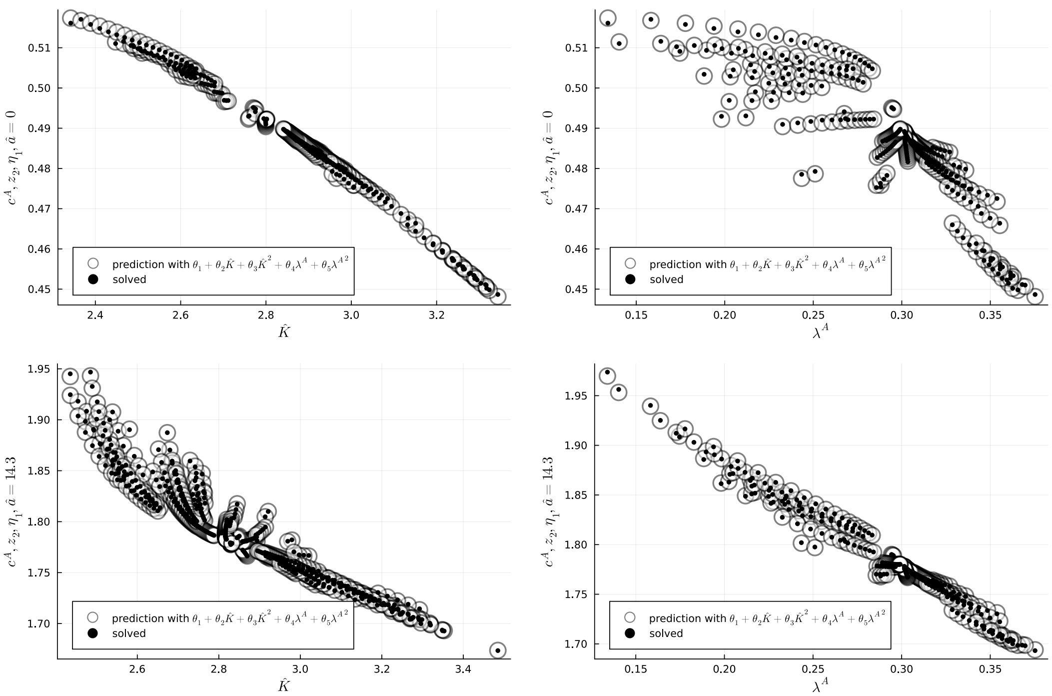

Figure 13: Consumption of an advanced trader in aggregate shock z_2 and idiosyncratic shock η_1 for two different wealth levels along 300 realizations on the simulated path of the economy for the benchmark model. The top panel shows the consumption for a household without any financial wealth, $\hat{a}_t^{z_2} = 0$. The bottom panel shows the consumption for a rich household, who owns assets of value $\hat{a}_t^{z_2} = 14.289$. 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.

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 13 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 stream from start-up 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 expected asset returns are relatively more important than the income from labor and start-up 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.[27] 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.[28] 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 50%. 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 6, we use a

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

[28] We start out with $(\beta, \xi^{\text{adj}}) \in \{0.93, 0.96, 0.98\} \otimes \{3, 5.5, 7\}$.

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.

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

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

$$y_{it} = g(x_{it}^h, Y_t) + u_{it}^h \tag{74}$$

$$u_{it} = \alpha_i + z_{it}^h + \epsilon_{it} \tag{75}$$

$$z_{it}^h = \rho z_{i,t-1}^{h-1} + \eta_{it} \tag{76}$$

$$\alpha_i \sim \text{iid } N(0, \sigma_\alpha^2) \tag{77}$$

$$\epsilon_{it} \sim \text{iid } N(0, \sigma_\epsilon^2) \tag{78}$$

$$\eta_{it} \sim \text{iid } N(0, \sigma_t^2) \tag{79}$$

$$\sigma_t^2 = \begin{cases} \sigma_E^2 & \text{in agg. expansions} \\ \sigma_C^2 & \text{in agg. contractions} \end{cases}. \tag{80}$$

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

$$x_{it} = \epsilon_{it} + z_{it} \tag{81}$$

$$z_{it} = \rho z_{it-1} + \eta_{it} \tag{82}$$

$$\epsilon_{it} \sim \text{iid } N(0, \sigma_\epsilon^2) \tag{83}$$

$$\eta_{it} \sim \text{iid } N(0, \sigma_\eta^2), \tag{84}$$

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

$$x_{it} = \bar{x} + \rho_x x_{i,t-1} + \sigma_x \epsilon_{it} \tag{85}$$

$$\epsilon_t \sim \text{iid } N(0, 1), \tag{86}$$

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

$$X = \begin{pmatrix} 0.463 \\ 1.537 \end{pmatrix} \tag{87}$$

$$\pi^X = \begin{pmatrix} 0.892 & 0.108 \\ 0.108 & 0.892 \end{pmatrix} \tag{88}$$

The resulting cross-sectional standard deviation of log earnings is 0.60 and matches the standard deviation of the process simulated in equations (74) - (80) and is in the ballpark of values typically used in the heterogeneous agents literature (see, *e.g.* Auclert et al., 2021).

C Impact of Firm Death in Deterministic Model (online only)

To get a sense of how firm survival rates Γ and initial firm size s interact, consider the following deterministic example. Assume that the growth rate is given by $A_{t+1} = (1+g)A_t$, which will imply that aggregate capital and output will grow at this rate in order to keep the marginal product of capital constant from (15). Hence, (14) implies that the growth rate of capital at the firm level is given by

$$g^k = \frac{1+g}{\Gamma + (1-\Gamma)s}.$$

This in turn implies that the value of a unit-sized firm is given by

$$\tilde{v}_t = d_t + \frac{g^k \Gamma}{1+r}\tilde{v}_{t+1} = d_t + D\tilde{v}_{t+1}, \text{ where } D := \frac{g^k \Gamma}{1+r}.$$

So, if the growth rate is 2% and the interest rate is 0.5%, we can construct the effective discount rate D on the future value of a unit sized firm, as shown in table 9. From these calculations, one can see that having a finite value for the unit-sized firm requires a fairly large initial firm size if the survival rate is 0.975 or higher; on the order of 0.75. With slightly lower survival rate, an effective discount $D < 1$ can be obtained, even when the initial firm size is 0.45.

So far, we have assumed that all of the capital in a dying firm is lost. If we instead assume that some fraction of it is converted into the final good (and therefore can be used either for investment in another firm or consumption), we get the following

Γ	s	g^k	D
0.975	1.00	1.020	0.990
0.975	0.75	1.026	0.996
0.975	0.50	1.033	1.002
0.965	1.00	1.020	0.979
0.965	0.75	1.029	0.988
0.965	0.50	1.038	0.997
0.965	0.45	1.040	0.999

Table 9: Effect of exit rate and relative startup size on the effective discount factor for firms in a deterministic model. The first columns shows the survival rate of firms, the second column shows the relative size of new firms, the third column shows the implied growth rate of capital at the firm level, and the last column shows the resulting effective firm discount factor.

expression for the value of a unit-sized firm, which shows that the recovery rate relative to the unit-sized-firm-value cannot be too large if we want to maintain that $D < 1$.

$$\tilde{v}_t = d_t + \frac{g^k}{1+r}\left[\Gamma \tilde{v}_{t+1} + (1-\Gamma)R\right] = d_t + \frac{g^k \tilde{v}_{t+1}}{1+r}\left[\Gamma + (1-\Gamma)\tilde{R}_{t+1}\right]$$
$$D := g^k \frac{\Gamma + (1-\Gamma)\tilde{R}_{t+1}}{1+r}, \quad \text{where } \tilde{R}_{t+1} = R/\tilde{v}_{t+1}$$

D Pricing long-term bonds

To explain the main insight from Alvarez and Jermann (2005), we follow their paper and consider the following simple example. If we denote the price of a bond which 1 unit in k periods by $p_k^b(z^t)$, then it is given by the natural extension of our one period price formula (33)

$$p_k^b(z^t) = \sum_{z^{t+k} \succ z^t} p(z^{t+k}|z^t). \tag{89}$$

The associated yield is simply $r_k(z^t) = \sqrt[k]{1/p_k^b(z^t)} - 1$. In the representative agent economy with CRRA preferences, the k-period ahead Arrow price is

$$p(z^{t+k}|z^t) = \beta^k \left(\frac{C(z^{t+k})}{C(z^t)}\right)^{-\sigma} Pr\{z^{t+k}|z^t\}$$

As in AJ's example, we focus on the holding period returns when the log of consumption follows a transitory and a random walk process. The one-period holding return on the bond is determined by the change in price, $p_{k-1}^b(z^{t+1})/p_k^b(z^t)$. If consumption follows an AR1 around a deterministic trend, and k is sufficiently large that the information about the future dies out, and this return is approximately,

$$\text{AR1:} \qquad \frac{p_{k-1}^b(z^{t+1})}{p_k^b(z^t)} \approx \beta\big(C(z^t)/C(z^{t+1})\big)^{-\sigma},$$

which is perfectly negatively correlated with the one-step ahead stochastic discount factor.

If instead consumption growth follows an i.i.d. process, then the expectation of $C(z^{t+k}) = E\{(1+g)^k\}C(z^t)$, and the condition expectation at time $t+1$ is $E\{(1+g)^{k-1}\}C(z^{t+1})$. In this case, the one-period holding return is given by

$$\text{Log RW:} \qquad \frac{p_{k-1}^b(z^{t+1})}{p_k^b(z^t)} = \frac{\beta^{k-1}E\left(\frac{C(z^{t+k})}{C(z^{t+1})}\right)^{-\sigma}}{\beta^k E\left(\frac{C(z^{t+k})}{C(z^t)}\right)^{-\sigma}} = \frac{E\left(\frac{C(z^{t+k})}{C(z^{t+1})}\right)^{-\sigma}}{\beta E\left(\frac{C(z^{t+k})}{C(z^{t+1})}\frac{C(z^{t+1})}{C(z^t)}\right)^{-\sigma}} = \frac{1}{\beta E\left(\frac{C(z^{t+1})}{C(z^t)}\right)^{-\sigma}}$$

which is just the inverse of the one-period risk-free rate. This is because the future scales with the realized growth rate in period $t+1$, and hence the expected return from $t+1$ to $t+k$ is unaffected.

This analysis directly implies that the yield curve on pure discount bonds, especially at very long horizons, is very informative as to the size of the unit root in the representative agent's consumption. Alvarez and Jermann (2005) focus on the implications of nominal bond yields for the real term structure. Stretching our model to assume it allowed for nominal prices, denoting the nominal price of the k-period ahead *nominal* Arrow price as $p^n(z^{t+k}|z^t)$, and assuming the absence of arbitrage, implies that the nominal price is related to the real price through the nominal price level $P(z^{t+k})$ as

$$p^n(z^{t+k}|z^t) = \frac{P(z^{t+k})}{P(z^t)} p(z^{t+k}|z^t).$$

From this it follows that the price of a nominal k-period ahead pure discount bond is given by

$$\begin{aligned} p^{n,b}(z^{t+k}|z^t) &= E\left[\frac{P(z^{t+k})}{P(z^t)} p(z^{t+k}|z^t)\right] \\ &= E\left[\frac{P(z^{t+k})}{P(z^t)}\right] E\left[p(z^{t+k}|z^t)\right] + Cov\left[\frac{P(z^{t+k})}{P(z^t)}, p(z^{t+k}|z^t)\right] \end{aligned}$$

From this expression one can see the nominal bond price includes the real bond price, expected inflation, and an inflation risk-premium that comes in the form of the covariance between the real Arrow bond price and inflation. In U.S. bond data, this covariance appears to have not only changed magnitudes over time but also changed sign (see Chen et al. (2016)), substantially complicating this relationship.

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.

Alvarez, F. and Jermann, U. J. (2005). Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6):1977–2016.

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. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.

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.

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.

Chen, A. Y., Engstrom, E., and Grishchenko, O. V. (2016). Has the inflation risk premium fallen? is it now negative? Technical report, Board of Governors of the Federal Reserve System (US).

Chernov, M., Lochstoer, L. A., and Song, D. (2021). The real channel for nominal bond-stock puzzles. 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.

Cochrane, J. H. (1988). How big is the random walk in gnp? *Journal of political economy*, 96(5):893–920.

Cochrane, J. H. (1994). Permanent and transitory components of gnp and stock prices. *The Quarterly Journal of Economics*, 109(1):241–265.

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.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco.

Fernández-Villaverde, J., Hurtado, S., and Nuño, G. (2019). Financial frictions and the wealth distribution. Working Paper 26302, 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.

Haubrich, J., Pennacchi, G., and Ritchken, P. (2012). Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *The Review of Financial Studies*, 25(5):1588–1629.

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. (2019). Self-justified equilibria: Existence and computation. *Available at SSRN 3494876*.

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.

Phillips, P. C. and Xiao, Z. (1998). A primer on unit root testing. *Journal of Economic Surveys*, 12(5):423–470.

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

Rossi-Hansberg, E. and Wright, M. L. J. (2007). Establishment size dynamics in the aggregate economy. *American Economic Review*, 97(5):1639–1666.

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.

Vissing-Jorgensen, A. and Krishnamurthy, A. (2011). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity*, 43(2):215–287.

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.