NBER WORKING PAPER SERIES

LOW RISK-FREE RATES AND INTERTEMPORAL ARBITRAGE

Marlon Azinovic-Yang
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, Revised November 2025

A previous version of the paper circulated under the title “Asset Pricing in a Low Rate Environment”. 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, the EEA-ESEM Conference and Rochester’s Stockman Memorial Workshop, the BI workshop on Production Based Asset Pricing, the University of Mannheim, Ohio State University, Oxford University and the Tinbergen Institute. Azinovic-Yang 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.

Low Risk-Free Rates and Intertemporal Arbitrage
Marlon Azinovic-Yang, Harold L. Cole, and Felix Kubler
NBER Working Paper No. 31832
November 2023, Revised November 2025
JEL No. C6, E22, E43, E44, G12

ABSTRACT

Is deficit finance free when real borrowing rates are routinely lower than growth rates? We study this question in a production-based asset-pricing model that features heterogeneous trading technologies as well as idiosyncratic and aggregate risk and can match the observed low average risk-free rate and the high market price of risk. We use our model to examine under which conditions realistic calibrations allow for (bounded) intertemporal arbitrage, that is to say, the possibility of infinite rollover of a short position. We give examples where this infinite rollover is possible. However, for our benchmark calibration that matches the high market price of risk observed in the data, rollover is impossible even if the average risk-free rate lies 3.5 percent below the average growth rate. The result is robust with respect to the introduction of permanent growth shocks.

Marlon Azinovic-Yang
University of North Carolina at Chapel Hill
marlonay@unc.edu

Harold L. Cole
University of Pennsylvania
Economics Department
and NBER
colehl@sas.upenn.edu

Felix Kubler
University of Zurich
and Swiss Financial Institute
fkubler@pm.me

1 Introduction

Starting in the 1980s, nominal interest rates in the U.S. went into a long period of decline. Rates, particularly short-term Treasury rates, were frequently below 1% between 2010-2021. 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)).[1] The growth rate of per capita output has been roughly 2%. 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 rollover of government debt is possible and whether deficit finance has no fiscal cost (see, e.g., Blanchard (2019), Kocherlakota (2023), Aguiar et al. (2024), Jiang et al. (2024)).

In this paper, we make five contributions to this literature. First, we deliver a tractable characterization of intertemporal arbitrage, linking rollover to the spectral radius of the growth-adjusted state-price matrix and, equivalently, to the price of a growth consol. In representative-agent, complete-market environments, this implies that the possibility of rollover depends exclusively on the discount factor, β (or on the growth rate and β in a model with growth), and is independent of aggregate uncertainty. We also extend the Cass-Zilcha criterion for dynamic efficiency to an economy with adjustment costs. Second, we build a tractable model of a production economy with heterogeneous agents and segmented markets that can match the observed high market price of risk and a very low average risk-free rate. Third, we evaluate the possibility of intertemporal arbitrage and dynamic inefficiency in our calibrated model. We show that infinite rollover fails, even with average interest rates well below the average growth rate. This result persists when considering a version of the model with both permanent and transitory shocks to productivity. Fourth, we map the average interest rate and the average market price of risk to the average consol price, characterizing the combinations under which rollover is possible. Fifth, we study policy and show that using government debt to relax borrowing constraints generally reduces average welfare due to capital being crowded out because the economy is dynamically efficient.

An empirical investigation into debt rollover and dynamic efficiency involves pro-

[1] Although real rates did increase after COVID in 2023, the average real rate is again trending down (see the Federal Reserve Bank of Cleveland, 1-Year Real Interest Rate [REAINTRATREARAT1YE], retrieved from FRED, Federal Reserve Bank of St. Louis).

jecting far into the future based on restricted present data. Given the notorious challenges associated with this task (see Kocherlakota (2024)), and considering that the average interest rate itself provides scant insight into the feasibility of perpetual debt rollover (as discussed in Bloise and Reichlin (2023), Kocherlakota (2023)), we employ a quantitative model-based approach to assess the effects of consistently low interest rates on dynamic efficiency and the feasibility of ongoing debt rollover.

We build a production-based asset pricing model with both idiosyncratic and aggregate risk, along with heterogeneous firms and households. The model includes adjustment costs on capital, firm death and entry,[2] along with borrowing constraints on households and heterogeneous access to financial markets along the lines of Chien et al. (2011): Some of our households can only invest in risk-free one period bonds (referred to as *bond traders* or *non-participants*) while others can invest in a full range of assets including aggregate-state-contingent securities (referred to as *advanced traders*).

Due to the presence of idiosyncratic labor income risk, all our households have precautionary motives to save and must contend with the borrowing constraint (as in Aiyagari (1994) and Huggett (1993)). On the household side this implies that even if the present-value of labor income becomes infinite, consumption stays finite.

On the firm side, adjustment costs to capital make the value of a firm meaningful, while incorporating firm death, beyond being a step towards greater realism, means that the value of firms being traded in asset markets does not include future entrants. This leads to a more realistic notion of firm value and also means that even if the present value of overall dividends becomes infinite, that from currently alive firms remains finite. As a result, equilibrium outcomes can be computed under very low interest rates while maintaining a plausible level of aggregate risk pricing.

We restrict ourselves to CRRA preferences with a degree of risk aversion of 5.5 and do not include a habit factor that can sharply increase local risk aversion. Despite this, our model is able to replicate our two key asset pricing facts: the high MPR and the low risk-free rate. Chien et al. (2011) showed that this was possible in an endowment model using a similar mechanism. The high market price of risk occurs because heterogeneous access to asset markets along with idiosyncratic risk leads to a wealth

[2] It is well documented (see, e.g. Crane et al. (2022)) that every year a substantial fraction of firms and establishments in the United States "die", while surviving firms, on average, grow in size, despite the fact that new firms enter. We calibrate our model to match these facts.

distribution for each type, which moves strongly in response to aggregate shocks. These large movements in the wealth distribution drive the high market price of risk. In addition, the combination of idiosyncratic risk and borrowing constraints can lead to a low risk-free rate in our model. Achieving this in a model with capital is more difficult than in pure exchange economies, since investment responds to intertemporal prices.

Solving our model is computationally challenging. While the firm-side satisfies a version of Gorman aggregation, the household-side features two types of heterogeneous agents. We solve the model using an extended version of the Krusell-Smith method. 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 currently alive 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 optimization problem still has a solution. However, computing an equilibrium of our model when the present value of future labor endowments is infinite is computationally challenging because the high prices of future consumption and endowments effectively lengthens pricing horizons. This reinforces the need for computational methods that are tailored precisely to our economic problem.

Our modeling setup also allows us to investigate whether a consol with a dividend growing at the long-run growth rate has a finite value. If this is the case, it is impossible to roll over debt infinitely. We find that as long as the parameterization of our model matches a realistically high market price of risk, average interest rates can become very small (with a growth rate of 2% we consider average interest rates down to -1.5%), while the value of the consol remains finite.

Associated with government debt is not only the issue of rollover, but also Pareto optimality. One can think of a government as being able to borrow against future tax payments and thereby using government debt to expand the borrowing capacity of households; effectively lowering their borrowing constraints. Government borrowing can also crowd out capital and, in some cases, make dynamically inefficient economies efficient as a result. Consistent with this, we examine the impact of relaxing the borrowing constraints of households on economic outcomes and welfare. Because our economy is dynamically efficient even at low rates, lowering agents' borrowing constraints via government policy causes crowding out and lowers the (average) welfare of both the advanced and the bond traders.

A fundamental calibration issue in any dynamic economic model of the economy is the extent to which the stochastic process for TFP includes a unit root. As a robustness check, we, therefore, examine the impact of including a unit root in our growth process on our results. To do this, we add an i.i.d. growth process and recalibrate the transitory shock process to TFP so that the one-step-ahead standard deviation (std) and autocorrelation are unchanged. When we incorporate a unit root with a 2% std into the TFP process, and recalibrate the model to match the interest rate, the MPR and the std of output to our benchmark values, we find that the std of the interest rate and consumption both rise, while the yield curve flattens but only modestly. The std of the interest rate remains within an empirically plausible range; however, the std of consumption is moderately outside of what we consider plausible. Despite this, our results for debt rollover and dynamic efficiency remain largely unchanged.

The paper is organized as follows. In Section 2 we briefly review the existing literature. In Section 3, we introduce the model and clarify the conditions for the possibility of infinite debt-rollover and for dynamic efficiency. Section 4 describes the calibration and key aspects of the quantitative implications of our model. 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. The Appendix contains proofs, additional material to our calibrations, and a detailed explanation of our computational approach.

2 Literature Review

A large recent literature asks whether low real interest rates relative to growth ($r < g$) permit perpetual debt rollover and whether deficit finance is fiscally costless. Key contributions include Blanchard (2019), Abel and Panageas (2025), Mian et al. (2021), Aguiar et al. (2024), Amol and Luttmer (2024), Brumm et al. (2024), and, on valuation, Jiang et al. (2024). Theoretically, Kocherlakota (2023) studies rollover with one-period debt under uncertainty, while Bloise and Reichlin (2023) relate rollover to Arrow prices.

We contribute a tractable Arrow-price characterization: rollover[3] is equivalent to the spectral radius of the growth-adjusted Arrow matrix exceeding one (equivalently, an infinite growth-consol price). In complete markets, the criterion depends only on the discount factor (or $((1+g)\beta)$) and is independent of aggregate uncertainty for a given discount factor.

Classic results show that government debt can raise welfare when economies with overlapping generations are dynamically inefficient (Samuelson, 1958; Cass, 1972; Abel et al., 1989). We extend the Cass-Zilcha test to include capital adjustment costs (Zilcha, 1990; Rangazas and Russell, 2005) and show that our calibrated economy is (typically) dynamically efficient even at very low rates.

We build on evidence that limited participation in the stock-market and heterogeneity are central to asset prices: participation is limited (Vissing-Jørgensen, 2002; Vissing-Jørgensen and Attanasio, 2003), high-consumption households bear more aggregate risk (Parker and Vissing-Jorgensen, 2009), and the required volatility of the SDF is high (Hansen and Jagannathan, 1991; Cochrane, 2009). In models, Guo (2004) and Chien et al. (2011) show that segmentation plus idiosyncratic risk can produce a low risk-free rate and a large countercyclical premium in endowment economies. Guvenen (2009) demonstrates related mechanisms in a production economy with two traders. Favilukis (2013) studies a production economy with overlapping generations, idiosyncratic risk, and fixed costs to participate in the stock market. He shows that loosening borrowing constraints, lower participation costs, and larger idiosyncratic risk can explain a decline in the interest rate and the equity premium. Our model differs along three margins central to the $r-g$ question: (i) two-type heterogeneity with uninsurable idiosyncratic risk (so wealth shares are shock-sensitive), (ii) complete markets for aggregate risk for advanced traders (so the Arrow matrix is well-defined in equilibrium), and (iii) a production block with adjustment costs and firm exit, which disciplines valuation ratios.

We incorporate firm exit and entry, which matters quantitatively for valuation ratios at low rates: with exit, traded value excludes future entrants, preventing mechanically explosive price earnings ratios in low-r environments (see also Rossi-Hansberg and Wright (2007); Luttmer (2011)). We also allow for capital adjustment costs, which tie investment and valuation tightly and are standard in quantitative

[3] We refer to the existence of a bounded Arrow-security Ponzi scheme as infinite debt rollover, which is a weaker condition for rollover than a one-period bond rollover.

production models.

Together, limited stock market participation with advanced traders and firm exit allow our calibrated model to match a low risk-free rate and a high market price of risk. Because of the presence of advanced traders, the prices of Arrow securities, which are crucial to the question of infinite debt rollover, are equilibrium outcomes of our model. Our paper is the first to examine when debt rollover is possible in a production based asset pricing model that matches the observed average market price of risk.

3 Model

We model a production economy with segmented financial markets and infinitely lived heterogeneous households that face aggregate and idiosyncratic risk and borrowing constraints. 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 of transitioning to the aggregate shock z_{t+1} when currently in the 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. We write $z^{t'} \succ z^t$ if node $z^{t'}$ succeeds z^t in the event tree.

3.1 Firms

We model a continuum of firms that take capital and labor as inputs 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 $\xi\left(z_t\right)$ 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 survives to produce and the fraction $1 - \Gamma$ dies. If firms exit, they do not pay dividends and their capital evaporates.[4] Because death is random, the amount of capital in the 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 chosen exogenously from the perspective of an individual entrant and is 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 in the next period, and both new and surviving firms in period t are subject to a survival shock at the beginning of the next period. Conditional on survival, new firms entering period t start producing in period $t+1$ with capital $k^i_{t+1} = \bar{i}_t$.

The problem of surviving firms is standard. For firm i, capital accumulates according to

$$k^i_{t+1} = k^i_t(1-\delta) + i^i_t \tag{2}$$

δ denotes the capital depreciation rate and k^i_{t+1} denotes the capital of firm i in the next period conditional on surviving. Firms producing at time t pay dividends

$$d^i_t = y^i_t - \omega_t l^i_t - i^i_t - \psi\left(k^i_{t+1}, k^i_t\right) \tag{3}$$

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

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

In this specification of the adjustment costs, the parameter ξ^{adj} denotes the level of the adjustment costs.

Firms have access to financial markets and trade in a complete set of aggregate-state contingent securities, *i.e.* Arrow securities for aggregate states. The price of an Arrow security that pays at node z^{t+1}, when traded at node z^t, is denoted by

[4]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).

$p(z^{t+1}|z^t)$. 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 in our model the Modigliani and Miller (1958) result holds, the firm's capital structure does not affect its value, and we therefore introduce the firm's problem without the consideration of financial policy. A firm's Bellman equation is given by

$$v^i(z^t) = \max_{k^i_{t+1}} d^i_t + \Gamma \sum_{z^{t+1} \succ z^t} p_t(z^{t+1}) v^i(z^{t+1}), \tag{5}$$

where k^i_{t+1} and v^i_{t+1} 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 implies 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 ratio and the same investment-to-capital ratio. Thus, there is a representative surviving firm and a representative new firm.

Aggregating the production sector, we obtain expressions for the aggregate value of producing firms V^{firm}_t, aggregate dividends D_t, the wage w_t, and aggregate profits from firm creation $\Pi^{\text{start-ups}}_t$. Furthermore, we obtain an aggregate Euler equation, which characterizes the optimal choice for aggregate capital K_{t+1}. We provide the expressions in Proposition 3.1 and give the derivations in Appendix A.1.

Proposition 3.1 *Let K_t denote the aggregate amount of capital in the beginning of period t, before the survival lottery. We have the following expressions for aggregate quantities output Y_t, wages w_t, marginal product of capital r^K_t, growth rate of capital in firms g^k_t, the firm's Euler equation (10), dividends D_t, the value of producing firms*

V_t^{firm}*, and the profits from startup creation* $\Pi^{startups}$*:*

$$Y_t = \xi_t(\Gamma K_t)^{\alpha}(A_t L_t)^{1-\alpha} \tag{6}$$

$$w_t = \xi_t A_t(1-\alpha)\left(\frac{\Gamma K_t}{A_t L_t}\right)^{\alpha} \tag{7}$$

$$r_t^K = \xi_t \alpha \left(\frac{\Gamma K_t}{A_t L_t}\right)^{\alpha-1} \tag{8}$$

$$g_t^k = \frac{1}{\Gamma + (1-\Gamma)s}\frac{K_{t+1}}{K_t} \tag{9}$$

$$1 = \frac{\Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} \left(r_{t+1}^K + 1 - \delta + \xi^{adj}\left((g_{t+1}^k)^2 - (1+g)^2\right)\right)}{1 + 2\xi^{adj}(g_t^k - (1+g))} \tag{10}$$

$$D_t = Y_t - \omega_t L_t - \Gamma\left[\left(g_t^k - (1-\delta)\right)K_t + K_t \xi^{adj}\left(g_t^k - (1+g)\right)^2\right] \tag{11}$$

$$V_t^{firm} = D_t + \frac{\Gamma}{\Gamma + (1-\Gamma)s}\sum_{z_{t+1}} p_t^{z_{t+1}} V_{t+1}^{firm} \tag{12}$$

$$\Pi^{startups} = \left(\frac{(1-\Gamma)s}{\Gamma + (1-\Gamma)s}\right)\left(\left(\sum_{z_{t+1}} p_t^{z_{t+1}} V_{t+1}^{firm}\right) - K_{t+1}\right). \tag{13}$$

Expression (12) shows that our model of firm exit and entry amounts to an additional discount of $\Gamma/(\Gamma + (1-\Gamma)s)$ in the aggregate Bellman equation. Fixing $0 < s \leq 1$, a decrease in the survival rate rate Γ means that less of the future value of firm's dividends is compounded into the current value of producing firms. It also shows that as $s \to 0$ all future firm dividends end up compounded. This is because the "missing" value of future firms is a product of two elements, how many there are and how valuable each new firms is in terms of both its initial capital and subsequent rents from creating capital. Put differently: $s > 0$ is required such that a positive fraction of tomorrow's stock market is not part of the stock market value today.

In Appendix B we provide additional details on the life-cycle of micro-level firms, which is active under the hood of our aggregation results.

3.2 Households

We model a unit measure of households. The households are of two types ex-ante and there is ex-post heterogeneity within type due to idiosyncratic productivity risk. The types of ex-ante heterogeneous households are based on Chien et al. (2011): a measure $0 < \mu \leq 1$ of households are *advanced traders* and a measure $1-\mu$ of households are *non-participants.* Advanced traders can trade the full set of aggregate-state contingent securities, while the non-participants only trade the risk-free one-period bond. Both types 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 of transitioning to an idiosyncratic shock η_{t+1} when currently in an idiosyncratic shock η_t. Each 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.

3.2.1 Preferences

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

$$U((c_t)_{t=0}^{\infty}) = E_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \text{ with } u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma}, \tag{14}$$

where γ denotes the coefficient of relative risk aversion and β denotes the time-preference parameter.

3.2.2 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 $z^{t+1} \succ z^t$ by $a_t^{z_{t+1}} = a^{z^{t+1}}(x^t)$. 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, in particular in economic downturns, may also become negative. We assume that the $(1-\Gamma)$ start-ups per period are equally distributed across advanced

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

traders. Thus, each trader starts $(1-\Gamma)/\mu$ start-ups, where μ denotes the measure of advanced traders. The Bellman equation for advanced traders is given by

$$\begin{aligned} V(z^t, \eta_t, a_{t-1}^{z_t}) = & \max_{\{a_t^{z_{t+1}} \geq \underline{a}\}_{z_{t+1} \in \mathcal{Z}}} u(c_t) \\ & + \beta \sum_{\tilde{z}_{t+1}} \sum_{\tilde{\eta}_{t+1}} \pi_z(\tilde{z}_{t+1}|z_t) \pi_\eta(\tilde{\eta}_{t+1}|\eta_t) V((z^t, \tilde{z}_{t+1}), \tilde{\eta}_{t+1}, a_t^{\tilde{z}_{t+1}}), \end{aligned} \tag{15}$$

where

$$c_t = l^A(\eta_t)\omega(z^t) + \frac{1-\Gamma}{\mu}\Pi^{startups}(z^t) + a_{t-1}^{z_t} - \sum_{z_{t+1} \in \mathcal{Z}} a_t^{z_{t+1}} p^{z_{t+1}}(z^t). \tag{16}$$

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

$$p^{z_{t+1}}(z^t) = p(z^{t+1}|z^t), \text{ for } z^{t+1} = (z^t, z_{t+1}) \succ z^t. \tag{17}$$

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

$$s_t^B := (z^t, \eta_t, b_{t-1}). \tag{18}$$

The Bellman equation of non-participants is given by

$$V(z^t, \eta_t, b_{t-1}) = \max_{b_t \geq \underline{b}} u(c_t) + \beta \sum_{\tilde{z}_{t+1} \in \mathcal{Z}} \sum_{\tilde{\eta}_{t+1} \in \mathcal{N}} \pi_z(\tilde{z}_{t+1}|z_t) \pi_\eta(\tilde{\eta}_{t+1}|\eta_t) V((z^t, \tilde{z}_{t+1}), \eta_{t+1}, b_t), \tag{19}$$

where

$$c_t = l^B(\eta_t)\omega(z^t) + b_{t-1} - b_t p^b(z^t). \tag{20}$$

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

$$p^b(z^t) = \sum_{z^{t+1} \succ z^t} p(z^{t+1}|z^t). \tag{21}$$

3.3 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

$$L_t^{\text{households}} := L_t^A + L_t^B, \text{ where} \tag{22}$$

$$L_t^A := \mu \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{a}}^{\infty} l^A(\eta_t) \rho_t^A(\eta_t, a_{t-1}^{z_t}) \mathrm{d}a_{t-1}^{z_t} \right) \tag{23}$$

$$L_t^B := (1-\mu) \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{b}}^{\infty} l^B(\eta_t) \rho_t^B(\eta_t, b_{t-1}) \mathrm{d}b_{t-1} \right). \tag{24}$$

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

$$L_t^{\text{firm}} := \left(\frac{\omega_t}{(1-\alpha) A_t^{1-\alpha} \xi_t K_t^{\alpha}} \right)^{-\frac{1}{\alpha}}. \tag{25}$$

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

$$\omega_t = (1-\alpha) \xi_t K_t^{\alpha} L_t^{-\alpha}. \tag{26}$$

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

$$w_t^{\text{households}} := w_t^A + w_t^B, \text{ where} \tag{27}$$

$$w_t^A := \mu \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{a}}^{\infty} a_{t-1}^{z_t} \rho_t^A(\eta_t, a_{t-1}^{z_t}) \mathrm{d}a_{t-1}^{z_t} \right) \tag{28}$$

$$w_t^B := (1-\mu) \sum_{\eta_t \in \mathcal{N}} \left(\int_{\underline{b}}^{\infty} b_{t-1} \rho_t^B(\eta_t, b_{t-1}) \mathrm{d}b_{t-1} \right). \tag{29}$$

Financial markets clear when

$$V_t^{\text{firm}} = w_t^{\text{households}}. \tag{30}$$

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.

3.4 Intertemporal arbitrage and overaccumulation of capital

Throughout this subsection, we assume that there is an irreducible Markov chain (z_t) with finite support and that the Arrow prices are a time-invariant function of z_t. Obviously, in the equilibrium of our model, this assumption does not hold.[6] However, note that for any computed approximate equilibrium, the assumption of finite support always holds because machine-precision implies the finiteness of possible states. The assumption of irreducibility is without loss of generality since we make our statements for specific initial conditions, and in the absence of irreducibility, one can always redefine the state space and the initial conditions accordingly.

[6] Very little is known about the exact equilibrium (as Brumm et al. (2017) explain, even the existence of Markov equilibria is, in general, an unresolved issue). In order to make statements about intertemporal arbitrage for the exact equilibrium, one would need to assume that the linear pricing operator has a strictly positive eigenprocess (see, e.g. Bloise and Reichlin (2023)). While the existence of an eigenprocess can be established via the Krein-Rutman theorem (see, e.g. Deimling (2013) Theorem 19.2), sufficient conditions are generally very demanding (see, e.g. Christensen (2017)), and nothing is known about the existence of a positive eigenprocess in models with incomplete markets and a continuum of heterogeneous agents.

3.4.1 Intertemporal arbitrage and debt rollover

We can define a non-negative matrix of Arrow prices P. Note that the price today of one unit of normalized output in j periods is then $[(1+g)P]^j[1]$. It turns out that intertemporal arbitrage is possible if and only if this explodes as $j \to \infty$. The following theorem formalizes this idea.

Theorem 3.1 *The following three statements are equivalent.*

1. *The largest eigenvalue of $(1+g)P$ is greater than or equal to one.*
2. *For any $B > 0$ there exists a (Markovian) trading strategy $(\theta(z_t))$ with $\|\theta(z_t)\|_\infty \leq B$ for all z_t that satisfies $p(z_0) \cdot \theta(z_0) < 0$ and*

$$\theta_{z_t}(z_{t-1}) - p(z_t) \cdot \theta(z_t) \geq 0 \textit{ for all } z^t$$

3. *The price of a growth consol that, at any time t pays $(1+g)^t$ in all states is infinite.*

As in Aiyagari and Peled (1991), the theorem shows that the largest eigenvalue of the matrix of Arrow security prices determines whether intertemporal arbitrage is possible and whether infinitely lived assets can have a finite price. It is then obvious to ask how this eigenvalue relates to other prices or to the fundamentals of the economy. The following theorem shows that the risk-free rate contains little information about this eigenvalue.

Theorem 3.2 *In each state i the inverse of the interest rate is given by $\frac{1}{R_i} = \sum_j p_{ij}$. Let P be an irreducible, non-negative matrix. Let λ_P denote the largest eigenvalue of P. Then*

$$\min_i \frac{1}{R_i} \leq \lambda_P \leq \max_i \frac{1}{R_i}.$$

Moreover, for each $\epsilon > 0$, there exist (non-negative and irreducible) P and P' such that $\| \min_i \frac{1}{R_i} - \lambda_P \| < \epsilon$ and $\| \max_i \frac{1}{R_i} - \lambda_{P'} \| < \epsilon$.

The result clarifies that nothing about the possibility of intertemporal arbitrage can be inferred from the average risk-free rate. The absence of intertemporal arbitrage merely places restrictions on the largest and smallest possible realization of the risk-free rate.[7]

[7] To see this, note that the last statement of the lemma implies that the bounds in the first

Note that the existence of intertemporal arbitrage is weaker than the possibility of rolling over one-period risk-free bonds (as examined in Kocherlakota (2023), see also Bloise and Reichlin (2023)). While Parts 1 and 3 of the theorem obviously give necessary conditions for rollover of one-period debt, Part 2 of Theorem 3.1 shows the existence of a bounded Ponzi-scheme in Arrow securities.[8] It is easy to construct examples where this does not imply that there is a Ponzi-scheme in the one period risk-free bond: Suppose there are two shocks, $g = 0$ and $R_1 = 1.05, R_2 = 0.95$ and the largest eigenvalue of P is $\frac{1}{0.95}$. Intertemporal arbitrage is possible, but if shock 1 realizes sufficiently often, the debt position must become arbitrarily large. This raises the question of why the government cannot issue Arrow-securities instead of one-period debt. Angeletos (2002) argues that by issuing sufficiently many bonds of different maturity, the government can actually replicate any Arrow-security portfolio. With this, the trading strategy in Theorem 3.1 can be viewed as debt-rollover, just that it is not exclusively one-period debt that is issued by the government. In the following, we will simply refer to the intertemporal arbitrage as the possibility of debt rollover. Since our main finding is that intertemporal arbitrage (and therefore infinite debt rollover) is not possible under realistic calibration, it is useful to employ the most general definition of debt-rollover.

3.4.2 Intertemporal arbitrage and equilibrium asset pricing

In this paper, we investigate the restrictions on P imposed by equilibrium asset pricing models. As one would expect, these depend crucially on the details of the model. In the simplest version without borrowing constraints and with complete markets, individual consumption will be Markovian. An (w.l.o.g. representative) agent with a consumption process $c(z_t)$ and CRRA utility $u(c)$ will satisfy Euler equations of the form

$$p_{z_{t+1}}(z_t) = \beta\pi(z_{t+1}|z_t)u'(1+g)\frac{u'(c(z_{t+1})}{u'(c(z_t))} \text{ for all } z_t, z_{t+1}.$$

We have the following theorem.

Theorem 3.3 *Suppose there is a positive (Markovian) process $(m(z_t))$ and a $\tilde{\beta} > 0$*

statement are tight, in the sense that, for some P, the largest eigenvalue can be arbitrarily close to the largest or smallest inverse of the risk free rate.

[8]As we show in Appendix C, the trading strategy is given by the eigenvector associated with the largest eigenvalue and is generally not the unit vector.

with

$$p_{z_{t+1}}(z_t) = \tilde{\beta}\pi(z_{t+1}|z_t)\frac{m(z_{t+1})}{m(z_t)} \text{ for all } z_t, z_{t+1}$$

Then the largest eigenvalue of P is equal to $\tilde{\beta}$. Moreover, for any $0 < \tilde{\beta} < 1$ and any $0 < R < 1$ there exist $(m(z_t))$ such that the average real rate is equal to R.

The theorem shows that the simplest equilibrium asset pricing framework does not establish any link between the average risk-free rate and the possibility of intertemporal arbitrage. It also shows that the discount factor β is the crucial parameter that determines the largest eigenvalue of P. In particular, this implies that, keeping the discount factor fixed, increasing aggregate uncertainty has no influence whatsoever on the possibility intertemporal arbitrage. Increasing aggregate risk, however, will reduce the average interest rate in the model. So if we infer the critical patience parameter by matching the average interest rate, its value, and therefore the possibility of intertemporal arbitrage, will depend on the assumptions on aggregate uncertainty.

3.4.3 Overaccumulation of capital

As Bloise and Reichlin (2023) show, the possibility of intertemporal arbitrage is equivalent to Pareto inefficiency of the competitive equilibrium in models of overlapping generations without financial frictions. In our model, equilibrium is always inefficient, independently of the existence of an intertemporal arbitrage. Following Zilcha (1990), Bloise and Reichlin (2023) also point out that in these models another stronger form of inefficiency is possible. It is possible that aggregate consumption can be (weakly) increased at all future date-events by decreasing investment in the current period. Following Bloise and Reichlin (2023) we refer to this possibility as "capital overaccumulation".[9]

We now turn to address this question in our full model. For theoretical tractability, we again study an equilibrium with finite support in that the aggregate capital only takes finitely many values and only depends on the Markov process (z_t).[10]

[9] Zilcha (1990) refers to this as "Type I inefficiency", it is also sometimes referred to as "dynamic inefficiency" although this terminology is slightly confusing, since the concept is strictly stronger than Pareto inefficiency.

[10] The Markov state z_t now includes the stochastic shocks as well as the level of capital and a finite approximation to the wealth distribution. As mentioned before, the finite support assumption is necessarily justified for any numerically computed equilibrium that uses finite precision arithmetic, such as the standard double-precision number format.

We consider the thought experiment where at some time τ investment is infinitesimally reduced and at all future times, investment is reduced by the implied reduction in output. Since this output, is subject to adjustment costs when reinvested, it is useful to define the return in capital units. Define a matrix

$$\mathbf{R}(z_t, z_{t+1}) = \begin{cases} \Gamma\left(\frac{\xi(z_{t+1})\alpha k_{t+1}^{\alpha-1} - \frac{\partial\psi(k_{t+2},k_{t+1})}{\partial k_{t+1}}}{1+\frac{\partial\psi(k_{t+2},k_{t+1})}{\partial k_{t+2}}} + (1-\delta)\right) & \text{if } \pi(z_{t+1}|z_t) > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Following Bloise and Reichlin (2023) the following theorem gives a simple necessary condition for the overaccumulation of capital.

Theorem 3.4 *Overaccumulation of capital is not possible if for any z_t there are a τ and $z_{t+1}, \ldots z_{t+\tau}$ satisfying*

$$\Pi_{i=0}^{\tau-1}\mathbf{R}(z_{t+i}, z_{t+i+1}) > (1+g)^{\tau}. \tag{31}$$

Note that Condition (31) relates to Zilcha's necessary condition for dynamic efficiency in a stationary OLG model (see Zilcha (1990); Rangazas and Russell (2005)). Zilcha's condition is the following.

$$E\log(\frac{\mathbf{R}}{1+g}) > 0. \tag{32}$$

It is easy to see that under our assumption that capital follows an irreducible Markov chain, if (32) holds, then for any sufficiently long sample path $(\mathbf{R}_t)_{t=0}^{T}$ we must have

$$\frac{1}{T}\sum_{t=0}^{T}\log(\frac{\mathbf{R}_t}{1+g}) > 0$$

which implies (31).

4 Calibration and Model Performance

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.

4.1 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 a parameter that is chosen to match key asset pricing moments and is 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%. 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,[11] 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 of 3.5%.[12] 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 G.

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).[13] The time preference parameter, β, is a parameter used to match observed prices and discussed below.

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

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

[12]With $s = 0.75$ and $\Gamma = 0.965$, we obtain $\Gamma + (1 - \Gamma)s \approx 0.99$.

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

have

$$\log(\xi_t) = \rho^{\text{tfp}} \log(\xi_{t-1}) + \sigma^{\text{tfp}} \epsilon_t^{\xi}, \tag{33}$$

where $\epsilon_t^{\xi} \sim \mathcal{N}(0, 1)$. We choose the autocorrelation 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%).[14] We discretize the AR(1) process using the method from Rouwenhorst (1995).

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. In Section 6 below we discuss this issue in detail and examine the impact of adding a stochastic growth component to TFP.

Due to 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 in being able to trade a richer set of assets, as our advanced traders do. The richer set of assets then also ensures that the value of the firm is well defined.

In addition to 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 the comovement of idiosyncratic risk with the aggregate state of the economy. Details can be found in Appendix F. We obtain

$$X = \begin{pmatrix} 0.46 \\ 1.54 \end{pmatrix}, \quad \pi^X = \begin{pmatrix} 0.89 & 0.11 \\ 0.11 & 0.89 \end{pmatrix} \tag{34}$$

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 range of values that are typically used in the literature on heterogeneous agents (see, *e.g.* Auclert et al., 2021).

[14]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.

4.2 Matching Moments

We have two remaining parameters that 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[15] 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.

Our second target is the market price of risk (MPR). Since there is a complete set 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, at a given node z^t, the market price of risk is given by

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

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 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 is significantly higher than 50% (see, e.g., Lo (2002) or Lustig and Verdelhan (2012)).

[15] 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.

4.2.1 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.[16]

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[17]

$$\ell(\beta, \xi^{\text{adj}}) := \left(\frac{1}{2} \left(\frac{\overline{\text{mpr}}(\beta, \xi^{\text{adj}}) - \text{mpr}^{\text{target}}}{\text{mpr}^{\text{target}}} \right)^2 + \frac{1}{2} \left(\frac{\overline{r}(\beta, \xi^{\text{adj}}) - r^{\text{target}}}{\max\{0.01, r^{\text{target}}\}} \right)^2 \right)^{\frac{1}{2}}. \tag{35}$$

In appendix E, 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 1 shows the value of our objective function, (35), for different combinations of β and ξ^{adj}. For our benchmark model, we chose values that 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 1 show how the two target moments vary with different values of the parameters. The market price of risk increases in adjustment costs

[16] See Appendix E.2 for an introduction to surrogate models.

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

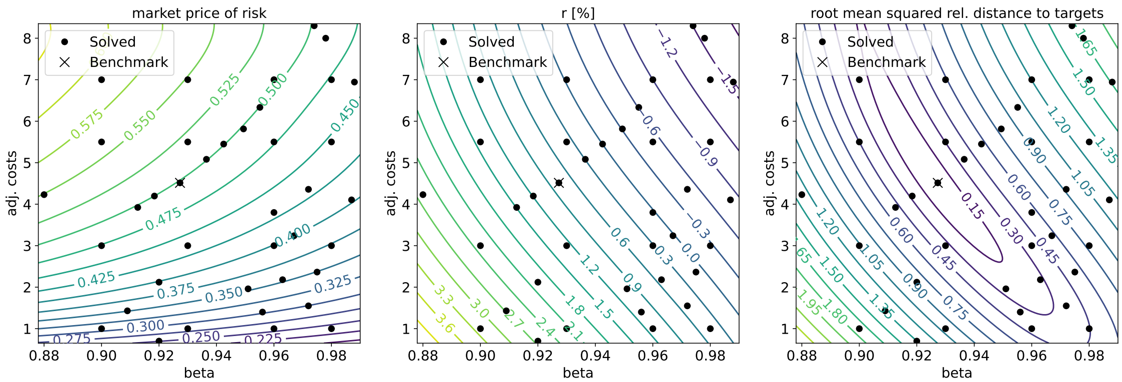

Figure 1: Dependence of the market price of risk (left panel), the interest rate (middle panel), and the root mean squared relative distance from the calibration targets (right panel) on the time-preference parameter β and the adjustment cost parameter ξ^{adj}.

and (depending on the region) decreases in the time preference parameter, β. The average real rate decreases in β and (somewhat depending on the region) increases 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.

4.3 Model performance

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 estimate the volatility of aggregate consumption growth to be 2.0% in the data.[18] The average of the log price-earnings ratio in the data lies somewhere between 2.8 and 2.99 depending on whether we consider the "Cyclically Adjusted Price Earnings Ratio P/E10" or if the "Cyclically Adjusted Total Return Price Earnings Ratio P/E10".[19]

[18]We take the series "A794RX0Q048SBEA" (Real personal consumption expenditures per capita) and the sum of "A796RX0Q048SBEA" (Real personal consumption expenditures per capita: Goods: Nondurable goods) and "A797RX0Q048SBEA" (Real personal consumption expenditures per capita: Services) from FRED between 1947 and 2008, compute the quarterly growth rates and aggregate them to a yearly frequency. We obtain a standard deviation of aggregate consumption growth of 2.0% using the first, and of 1.4% using the second measure.

[19]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,

	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.

Table 1 summarizes key statistics from our benchmark calibration of the model and compares them to the data and a representative agent specification.[20]

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 should be noted 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, the entire aggregate risk in asset returns is held by the advanced traders. Since these only constitute 10% of all agents in the economy, their individual consumption will be very volatile. In our benchmark calibration, the consumption of the group of advanced traders is roughly four times more volatile than aggregate consumption. The assets are priced based on the individual consumption of unconstrained advanced traders, which generates a high market price of risk in our model. The advanced traders, on

last accessed: January 2023) by the real earnings, take logs, and compute the average between 1947 and 2008. Earnings in the model are given by profits before investment $E_t = Y_t - \delta\Gamma K_t - w_t L_t - \Psi_t$, where Ψ_t denotes the aggregate capital adjustment costs.

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

average, hold more wealth and attain a higher consumption level. Although 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%.[21]

4.3.1 Mapping parameters to moments

As mentioned above, the first two panels in Figure 1 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 2 shows how the volatility of consumption growth and the volatility

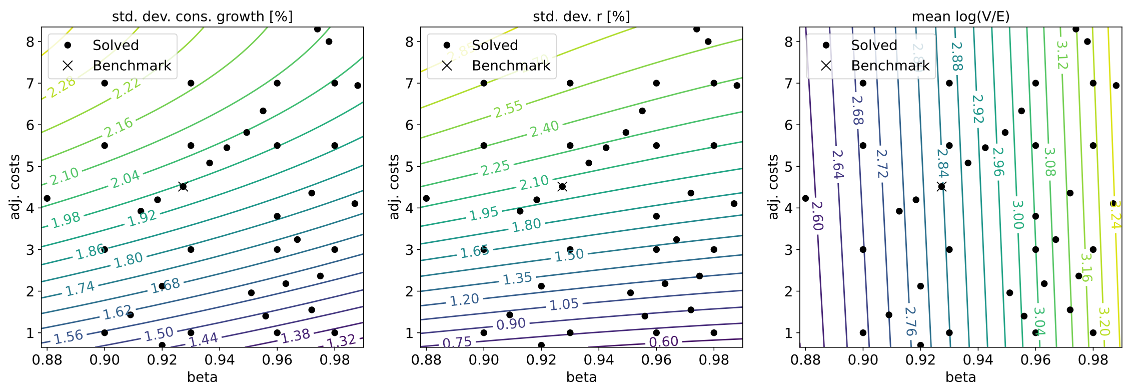

Figure 2: Standard deviation of consumption growth (left), standard deviation of the interest rate (middle), and the price-earnings ratio (right) for different values of the time-preference parameter (vertical axis), and the adjustment costs parameter (horizontal axis).

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 5 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 modeling ingredients play an important role, idiosyncratic income risk, borrowing

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

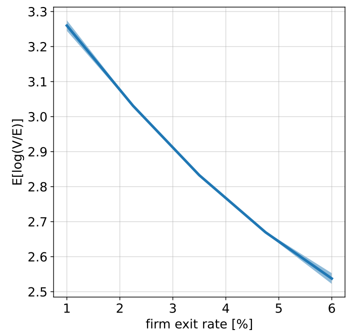

Figure 3: Mean log price-earnings ratio for different values of the firm exit rate. For a given firm exit rate, the adjustment costs and the patience are recalibrated to maintain an average market price of risk of 50% and an average interest rate of 0.8%. The shaded area shows the confidence interval of the prediction by the surrogate model.

constraints, and firms' death. Since Aiyagari (1994), the first two have become 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.

4.3.2 Firm value

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 E.2 for details). Figure 3 shows, based on the surrogate model,[22] how the mean price-earnings ratio in our model depends on the firm's exit rate. For each firm exit rate, we recalibrate the adjustment costs and patience parameter 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 3 shows, the price-earnings ratio in

[22] An introduction to surrogate models is provided in Appendix E.2.

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.

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 rollover of government debt is possible and whether deficit finance has no fiscal cost. 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.

It is a different question whether government debt can be Pareto-improving in our setting (see, e.g. Aguiar et al. (2024) or Brumm et al. (2024)). Due to 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 asymptotically. In our experiments, government debt is almost always detrimental to welfare because it induces crowding out of private savings and investment. We examine one case where debt rollover is possible and where welfare effects are ambiguous. However, in that case, the market price of risk is unrealistically low.

5.1 Debt Rollover

In a model without aggregate uncertainty, debt rollover is possible if and only if the risk-free rate is below the growth rate. As we argued above, with aggregate uncertainty matters can be very different. To better understand 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 consol.[23] We fix all parameters except for patience and adjustment costs at fixed values and compute equilibria for a variety of combinations of parameter. Infinite debt rollover is possible if and only if the value of a consol that pays $(1+g)^t$

[23] See Appendix E.2 for additional details on our surrogate models.

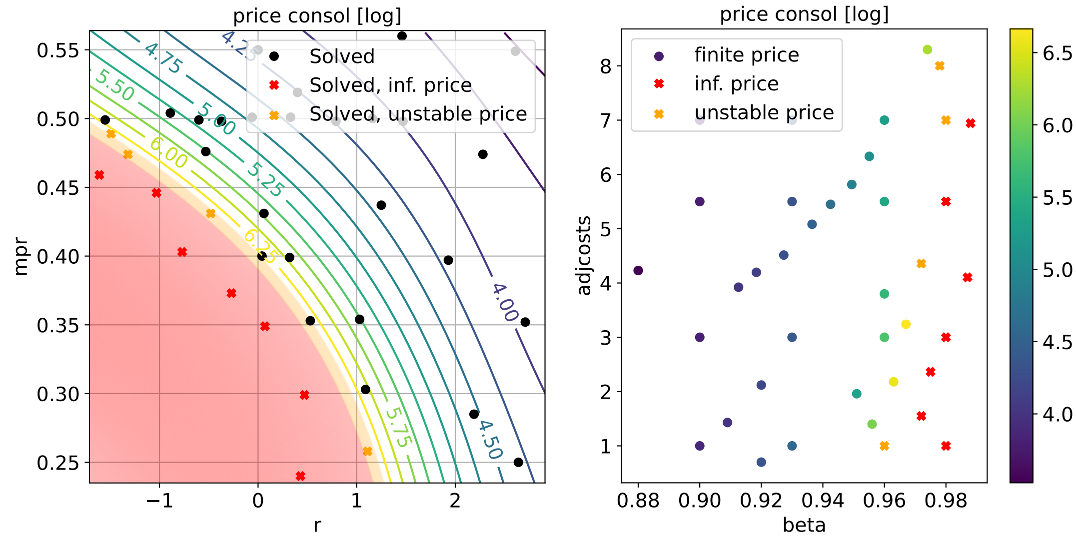

Figure 4: Mean price of a consol, in logs. The left panel shows the price plotted against the mean interest rate and the mean market price of risk. The right panel shows it plotted against the patience parameter and the adjustment cost parameter.

at any time t is infinite. Gaussian process regression is used to predict equilibrium behavior at all other parameters. The left panel in figure 4 shows the price of the consol for different combinations of the interest rate and the MPR.[24]

In regions where growth-adjusted Arrow prices are summable, the average price of the consol is increasing faster than exponentially as the mean interest rate decreases. To see this, observe that the distance between the contour lines is decreasing, despite the figure showing the log of the average price of the consol. A higher market price of risk lowers the price of the consol, and hence allows for summable prices with lower interest rates. A realistically[25] 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.

What are the assumptions on fundamentals that are crucial for the possibility of debt rollover? In the simple representative agent example in Section 3.4.1 the only

[24]For lower interest rates computing the consol price becomes numerically difficult before a clear divergence of the price can be observed. Hence, we mark the area in which we cannot observe either a clear divergence or convergence of the price with orange.

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

relevant parameters were the preference parameter β and the growth rates adjusted for risk aversion. In that example, debt-rollover is never possible in equilibrium. Here, borrowing constraints and the death of firms imply that there are specifications of the economy where debt rollover is possible. However, holding growth rates and risk-aversion constant, the value of the parameter β remains the crucial factor in determining whether debt rollover is possible. As the right panel in figure 4 shows, debt rollover is possible if β is greater than 0.975. Note that with our growth rate of 2 % and a risk-aversion parameter of 5.5 the utility maximization problem for each agent is well defined and equilibrium exists. The occasionally binding borrowing constraints imply that prices might not be summable even when the conditions derived in the representative agent example hold. Nevertheless, high values of the preference parameter β remain crucial for the possibility of debt rollover.

5.2 Welfare effects of relaxing borrowing constraints

Although in our benchmark 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 are suboptimal and there are no theoretical arguments to rule out the possibility that government debt could be Pareto improving. In particular, an active debt policy as, for example, in Bhandari et al. (2017), can be Pareto improving in the presence of borrowing constraints. On the other hand, even for the parameter specifications where debt rollover is possible, it is not obvious that government debt will always be Pareto improving. It is beyond the scope of the paper to solve for an "optimal" debt policy. Instead, we consider a government policy intervention that highlights the role government debt can play to alleviate borrowing constraints.

In the first period of the model, $t = 0$, the government gives each household some units of a one-period 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. In every period t, the government collects lump-sum taxes that amount to b_{t-1}^g. The only role this bond plays for the household is that it can be used as collateral for a short-position. Equivalently, the household can sell the bond, but still has to repay to the government next period.

Clearly, a relaxation of their borrowing constraints will potentially increase the agents' welfare independently of whether debt rollover is possible or not. However, there is also a price-effect. If there is no overaccumulation of capital, aggregate consumption will decrease as government debt increases. 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 percentage 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 0.8 percent, we also consider a zero interest rate case as well as a case with an average interest rate of -0.58 percent. In this last case, debt rollover is possible. In all four cases, our necessary condition for dynamic efficiency, Equation (32) is satisfied.

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 $U(\{c_t\}_{t=0}^{\infty}) = \mathrm{E}\left[\sum_t \beta^t u(c_t)\right]$, where $u(c) = \frac{c^{1-\ \gamma}}{1-\gamma}$, preferences are homothetic. To compare the utility levels for different the borrowing constraints, we compute the unconditional average lifetime utility in both specifications and then compute the factor $c^* = \left(\frac{U^{\text{lower}}}{U^{\text{baseline}}}\right)^{\frac{1}{1-\gamma}}$, such that $U^{\text{lower}}(\{c_t^{\text{lower}}\}_{t=0}^{\infty}) = U^{\text{baseline}}(\{c_t^{\text{baseline}} c^*\}_{t=0}^{\infty})$. The average is computed over the ergodic distribution of aggregate and idiosyncratic states. A factor $c^* > 1$ corresponds to a higher average utility with lower borrowing constraints, while $c^* < 1$ corresponds to a lower average utility in the model with loser borrowing contraints. In table 2 we present the effects of relaxing the borrowing constraint for three calibrations.

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. Absent price effects, relaxing occasionally-binding borrowing constraints obviously increases welfare. The values of $\mu^A(\underline{b})$ and $\mu^B(\underline{b})$ in Table 2 show the fraction of time the

Table 2: Impact of Government Policy to Lower Borrowing Constraints

Benchmark model, infinite debt rollover not possible										
$\underline{b}$	r	$\hat{K}$	MPR	rel. $c^{*,A}$	rel. $c^{*,B}$	$\log(V/E)$	λ^A	$\mu^A(\underline{b})$	$\mu^B(\underline{b})$	CZ
0.00	0.85%	2.829	49.8%	1	1	2.834	27.3%	2.1%	11.0%	6.4%
-0.10	1.03%	2.778	49.0%	0.994	0.999	2.801	27.9%	1.9%	10.7%	6.6%
-0.20	1.32%	2.712	48.2%	0.986	0.998	2.766	27.9%	1.9%	10.3%	6.8%
Low interest rate, infinite debt rollover not possible										
$\underline{b}$	r	$\hat{K}$	MPR	rel. $c^{*,A}$	rel. $c^{*,B}$	$\log(V/E)$	λ^A	$\mu^A(\underline{b})$	$\mu^B(\underline{b})$	CZ
0.00	0.02%	2.992	50.1%	1	1	2.937	30.1%	1.6%	10.9%	7.0%
-0.10	0.21%	2.930	49.2%	0.994	0.999	2.900	30.5%	1.5%	10.6%	7.1%
-0.20	0.60%	2.838	48.6%	0.987	0.998	2.853	29.9%	1.6%	10.2%	7.2%
Low interest rate and low MPR, infinite debt rollover possible										
$\underline{b}$	r	$\hat{K}$	MPR	rel. $c^{*,A}$	rel. $c^{*,B}$	$\log(V/E)$	λ^A	$\mu^A(\underline{b})$	$\mu^B(\underline{b})$	CZ
0.00	-0.58%	3.610	36.6%	1	1	3.259	22.4%	2.7%	8.5%	2.7%
-0.10	-0.45%	3.555	36.0%	0.999	1.001	3.223	22.9%	2.5%	8.3%	2.8%
-0.20	-0.27%	3.495	35.7%	0.997	1.003	3.188	23.0%	2.5%	8.1%	3.0%

borrowing constraint is binding for agents of type A and B . When debt rollover is possible, the Arrow traders have an arbitrage trade at their disposal. As borrowing constraints are relaxed the Arrow trader has a (limited) arbitrage strategy. Nevertheless, his equilibrium welfare decreases substantially in all three cases. The bond trader cannot make the same arbitrage trade, yet his welfare goes up slightly in the rollover-case and is very flat in the low interest rate case where rollover is impossible. The reason for this is that for this trader the borrowing constraint is binding much more often and a relaxation of that constraint allows him to smooth consumption after a series of negative idiosyncratic shocks.

Government debt crowds out private savings and hence decreases the capital stock. Since in our calibrations the economy is *dynamically efficient* in the sense that there is no overaccumulation of capital[26] (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 (2023)) that the possibility of debt rollover is a necessary but by far not sufficient condition for dynamic inefficiency.

[26] The column under CZ in Table 2 shows the estimated value of the lhs in Equation 32 – as we show above a non-positive value is necessary for capital overaccumulation.

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 argue 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, consumption is assumed to follow a stochastic growth process, the resulting term structure seems much more consistent with US data. On the other hand, Cochrane (1988) argues that the evidence of a unit root in the time-series of GDP is very mixed. Chernov et al. (2021) show that consumption dynamics is best explained with a switching model with both transitory and aggregate states. We discuss the empirical evidence for a unit root in consumption, TFP and output in Appendix D.

It seems difficult to argue that there is no transitory component in the stochastic process of TFP or that there is no permanent component in the process. As a robustness exercise, we examine the impact of adding a permanent stochastic component to the TFP process, and we modify our model as follows. We add i.i.d. growth shocks with standard deviation σ^{perm}, which we discretize into a two-state process. This gives us six exogenous aggregate shocks in total. We re-calibrate the AR(1) parameters of the transitory shock ρ^{trans} and σ^{trans} such that the standard deviation (2.7%) and auto-correlation (-0.10) of productivity growth remains the same. That is to say, as we increase σ^{perm}, we decrease ρ^{trans} and σ^{trans}. We consider $\sigma^{\text{perm}} = 2\%$. Table 3 shows the analogue of 1 for this case. Consumption volatility is too large, investment volatility too small. Increasing the size of the permanent component would make this even worse.

Figure 5 shows the effect of the permanent component on the yield curve. Together with the yield curve in the data for inflation indexed treasury securities with different maturities.[27] If we just add the growth-shock without recalibrating the model, the risk-free rate increases and the MPR decreases. Recalibrating the preference parameters leaves us with a yield curve for the permanent shock case that does not flatten

[27] We compute the mean annual yield from the data series `WFII5`, `DFII7`, `DFII10`, `DFII20`, and `DFII30` obtained from https://fred.stlouisfed.org on 09/09/2025.

	Model perm + trans	Model trans only	Data
r^{bond}	0.8%	0.8%	0.8%
MPR	50%	50 %	50%
std output growth	2.6%	2.6%	2.6%
std r^{bond}	2.3%	2.0%	1.8-2.9%
std agg. consumption growth	2.3%	2.0%	1.4-2.0%
log V/E	2.8	2.8	2.8-3.0

Table 3: Key moments in the model with permanent and transitory shocks, the model with only transitory shocks, and the data.

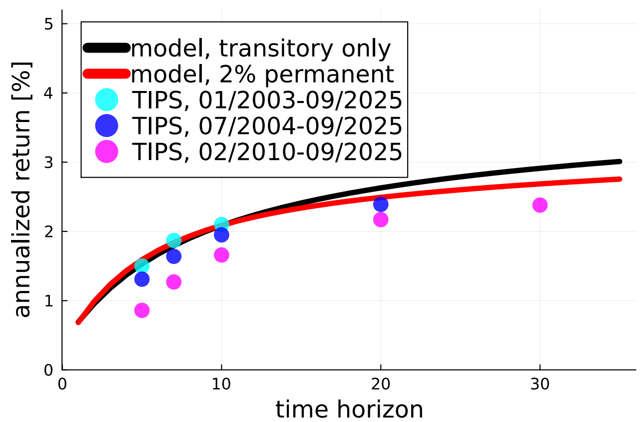

Figure 5: Yield Curve implied by the model with (red solid line) and without (black solid line) permanent productivity shocks. The dots show the annualized yield of inflation indexed treasury securities with different maturities. The model implied yield curve is based on 20000 simulated periods.

out much more than our original yield curve. To compare our real yield curve to the data, we also plotted various averages for Treasury inflation-indexed yields (TIPS). The real yield curve for TIPS does not appear to be flatter than our model's prediction. Unfortunately, TIPS only started being issued relatively recently, with 20 years of 5-20 year yields and only 10 years of 30-year yields. However, they do have the advantage of being a highly liquid explicit real bond (rather than the real return on a nominal bond). We discuss the evidence for a flattening out of the real yield curve in Appendix D.

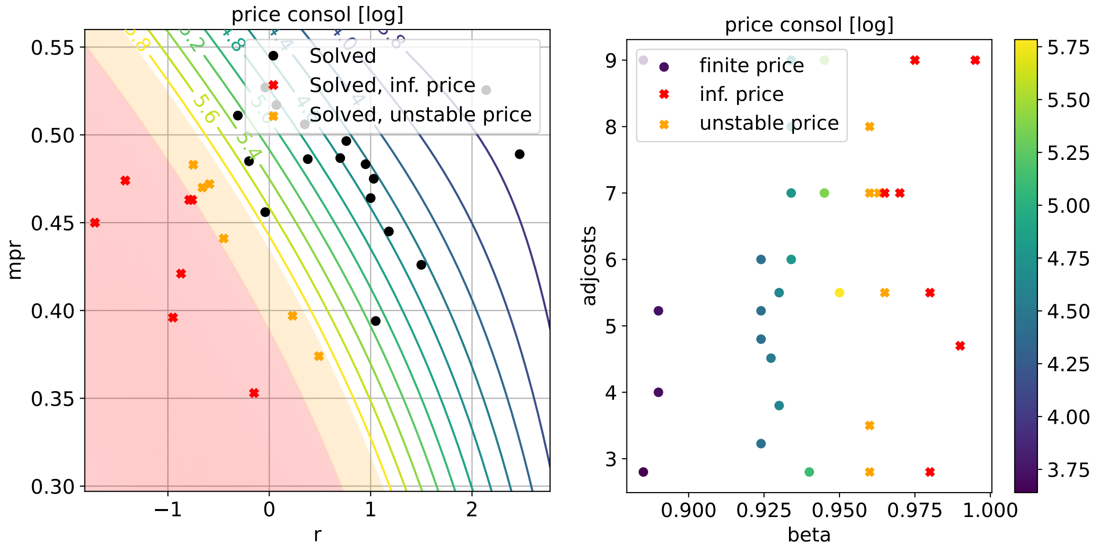

Figure 6: Mean price of a consol, in logs. The left panel shows the price plotted against the mean interest rate and the mean market price of risk. The right panel shows it plotted against the patience parameter and the adjustment cost parameter.

6.1 The possibility of debt-rollover

When considering permanent growth shocks, we examined debt-rollover for "growth consols"; bonds whose payout grows with the realized growth rate. The reason for this is that average growth rates imply little or nothing about relative levels in this context. As explained in Appendix C.2, analogues of Theorems 3.1, 3.2 and 3.3 can be shown to hold if one considers prices of "growth consol", i.e. a consol that pays $\Pi_{i=0}^{t}\gamma(z_i)$ units at node $z^t = (z_0, z_1, \ldots, z_t)$ (see Bloise and Reichlin (2023) for a similar approach).

As in Section 5, debt-rollover is impossible for our benchmark calibration with transitory and permanent shocks. As above, to better understand 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 consol. We fix all parameters except for the patience and the adjustment costs at fixed values; hence other moments, such as the standard deviation of the risk-free rate, also vary as the market price of risk varies. The left panel in Figure 6 shows the price of the growth consol for different combinations of the interest rate and the MPR.

Comparing the figure with the case of purely transitory shocks, we find that differences between these two calibrations are very small. It is again the case that for realistic values of risk-free rate and market price of risk, debt rollover is impossible. The presence of a permanent component in the TFP process appears to have little effect on the possibility of debt rollover. This is confirmed in the right panel of Figure 6 that shows the relation of β and the possibility of debt rollover. As in the representative agent example in Section 3.4.1, if the growth rates and risk aversion are fixed, the crucial parameter that determines whether debt rollover is possible is β. As in the case of purely transitory shocks, the borrowing constraints imply that debt-rollover is consistent with the existence of equilibrium, but it is inconsistent with a realistic MPR.

7 Conclusion

We examine the possibility of infinite debt-rollover in a heterogeneous agent stochastic production economy. We find that lowering the interest rate, while maintaining a realistically high market price of risk, does not lead to rollover being possible, even for very low interest rates. We show that the question of debt-rollover is simply a function of the discount rate in any representative agent economy, in which the outcomes can be represented as a finite Markov chain. This is true independently of the amount of aggregate risk. Aggregate risk however, lowers the average interest rate for a given discount factor. The answer to this question with heterogeneity is much less obvious. Despite this, we find a tight quantitative link between the discount rate and the rollover question. We conjecture that this arises because despite the presence of borrowing constraints, there is a positive measure of households who can trade aggregate contingent securities and who are off their borrowing constraints for long periods. Moreover, we find that even with very low interest rates, our economy is generally dynamically efficient. Because of this, lowering agents' borrowing constraints via government policy causes crowding out and lowers the welfare of both types of traders for calibrations with a realistic real interest rate.

Appendix

A Proofs

A.1 Aggregation of the firm side

In this section we derive the expressions for the aggregation of the firm side of our model, *i.e.* equations (6) - (13). To establish this formally, note that each firm i takes the wage as given and chooses l_t^i such that

$$\omega_t = \xi_t A_t(1-\alpha)\left(\frac{k_t^i}{A_t l_t^i}\right)^{\alpha} \Leftrightarrow \frac{k_t^i}{A_t l_t^i} = \left(\frac{\omega_t}{A_t\xi_t(1-\alpha)}\right)^{\frac{1}{\alpha}} =: \mathcal{K}_t, \tag{36}$$

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

$$w_t = A_t(1-\alpha)\xi_t\mathcal{K}_t^{\alpha} \tag{37}$$

$$r_t^K = r_t^{Ki} = \alpha\xi_t\mathcal{K}_t^{\alpha-1}. \tag{38}$$

All incumbent firms face the same first-order condition for optimality, which is given by

$$1 + \underbrace{\frac{\partial\psi(k_{t+1}^i, k_t^i)}{\partial k_{t+1}^i}}_{=2\xi^{\text{adj}}\left(\frac{k_{t+1}^i}{k_t^i}-(1+g)\right)} = \Gamma\sum_{z_{t+1}} p_t^{z_{t+1}}\left(r_{t+1}^K + 1 - \delta - \underbrace{\frac{\partial\psi(k_{t+2}^i, k_{t+1}^i)}{\partial k_{t+1}^i}}_{=-\xi^{\text{adj}}\left(\left(\frac{k_{t+2}^i}{k_{t+1}^i}\right)^2-(1+g)^2\right)}\right). \tag{39}$$

The right-hand side of the Euler equation reflects the firm's anticipated policy tomorrow conditional on survival. Due to the homogeneity of production and adjustment costs, each incumbent firm chooses the same growth rate, which, for a given period, t, we denote by $(g^k)_t^i := \frac{k_{t+1}^i}{k_t^i} = g_t^k$. The constant g_t^k establishes equation (10) as an aggregate Euler equation.

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

$$\tilde{K}_{t+1} = \frac{1}{\Gamma} \int_{i \in \mathcal{P}_t} k_{t+1}^i di \tag{40}$$

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

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

Hence

$$\Gamma \tilde{K}_{t+1} = (1-\delta)\Gamma K_t + \Gamma \tilde{I}_t \Leftrightarrow \tilde{K}_{t+1} = (1-\delta)K_t + \tilde{I}_t. \tag{41}$$

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:

$$K_{t+1} = \underbrace{\Gamma \tilde{K}_{t+1}}_{\text{incumbents}} + \underbrace{(1-\Gamma)\bar{i}_t}_{\text{entrants}} = [\Gamma + (1-\Gamma)s]\, \tilde{K}_{t+1} \tag{42}$$

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 $t+1$. We have that

$$g_t^k k_t^i = k_{t+1}^i \Rightarrow g_t^k \underbrace{\int_{i \in \mathcal{P}_t} k_t^i di}_{=\Gamma K_t} = \underbrace{\int_{i \in \mathcal{P}_t} k_{t+1}^i di}_{=\Gamma \tilde{K}_{t+1}}, \tag{43}$$

and, hence, we obtain

$$g_t^k = \frac{\tilde{K}_{t+1}}{K_t} = \frac{1}{\Gamma + (1-\Gamma)s} \frac{K_{t+1}}{K_t}. \tag{44}$$

This establishes equation (9).

Aggregate output is given by

$$Y_t = \int_{i \in \mathcal{P}_t} y_t^i di = \int_{i \in \mathcal{P}_t} \xi_t (k_t^i)^\alpha (A_t l_t^i)^{1-\alpha} di \tag{45}$$

$$= \int_{i \in \mathcal{P}_t} \xi_t (k_t^i)^\alpha (A_t l_t^i)^{1-\alpha} di = \int_{i \in \mathcal{P}_t} \xi_t \mathcal{K}_t^\alpha (A_t l_t^i) di \tag{46}$$

$$= \xi_t A_t L_t \mathcal{K}_t^\alpha = \xi_t A_t L_t \left(\frac{\Gamma K_t}{A_t L_t} \right)^\alpha \tag{47}$$

$$= \xi_t (\Gamma K_t)^\alpha (A_t L_t)^{1-\alpha}, \tag{48}$$

where we used $\mathcal{K}_t = \frac{\Gamma K_t}{A_t L_t}$. This establishes equation (6). Similarly, the wage and the marginal product of capital are given by

$$\omega_t = \xi_t A_t (1-\alpha) \mathcal{K}_t{}^\alpha = \xi_t A_t (1-\alpha) \left(\frac{\Gamma K_t}{A_t L_t} \right)^\alpha \tag{49}$$

$$r_t^K = \alpha \xi_t \mathcal{K}_t{}^{\alpha-1} = \xi_t \alpha \left(\frac{\Gamma K_t}{A_t L_t} \right)^{\alpha-1}. \tag{50}$$

This establishes equations (7) and (8).

Aggregate investment, which includes the investment by startups as well as incumbents, is

$$I_t = \Gamma \tilde{I}_t + (1-\Gamma) \bar{i}_t = \Gamma \tilde{I}_t + (1-\Gamma) s \tilde{K}_{t+1} \tag{51}$$

Aggregate dividends are given by

$$D_t = Y_t - \omega_t L_t - \Gamma \left[\tilde{I}_t + K_t \xi^{\text{adj}} \left(g_t^k - (1+g) \right)^2 \right]. \tag{52}$$

Using equations (41) and (44), we obtain

$$D_t = Y_t - \omega_t L_t - \Gamma \left[(g_t^k K_t - (1-\delta) K_t) + K_t \xi^{\text{adj}} \left(g_t^k - (1+g) \right)^2 \right], \tag{53}$$

and obtain the equation (11).

Because all firms invest and produce proportionate to their 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

$$V_t^{\text{firm}} = \Gamma K_t \tilde{v}_t = \int_{j\in\mathcal{P}_t} v_t^j dj = \int_{j\in\mathcal{P}_t} \left[\begin{array}{c} y_t^j - \omega_t l_t^j - i_t^j \\ -k_t^j \ \xi^{\text{adj}} \left(g_t^k - (1+g)\right)^2 \\ +\Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^j \end{array} \right] dj \tag{54}$$

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

$$\tilde{v}_t = \left[\ \frac{D_t}{\Gamma K_t} + \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} \tilde{v}_{t+1} g^k \ \right] \tag{55}$$

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 presence of convex adjustment costs means that there are rents that accrue to the firm and that each firm faces a non-trivial intertemporal optimization problem.

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

$$\bar{v}_t := \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^{t\text{-entrant}*} = \Gamma \bar{i}_t \sum_{z_{t+1}} p_t^{z_{t+1}} \tilde{v}_{t+1} \tag{56}$$

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,

$$V_{t+1}^{\text{firm}} = \tilde{v}_{t+1}\Gamma \left[\Gamma \ \tilde{K}_{t+1} + (1-\Gamma) \ \bar{i}_t\right] = \tilde{v}_{t+1}\Gamma K_{t+1} \tag{57}$$

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. Intgerating the right hand side of equation (54), we obtain the aggregate producing firm's Bellman equation is given by

$$V_t^{\text{firm}} = D_t + \frac{\Gamma}{\Gamma + (1-\Gamma)s} \sum_{z_{t+1}} p_t^{z_{t+1}} V_{t+1}^{\text{firm}}. \tag{58}$$

To see this, notice that

$$\int_{j \in \mathcal{P}_t} \left(\Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} v_{t+1}^j \right) dj = \Gamma \sum_{z_{t+1}} p_t^{z_{t+1}} \int_{j \in \mathcal{P}_t} v_{t+1}^j dj \tag{59}$$

and that

$$V_{t+1}^{\text{firm}} = (\Gamma + (1-\Gamma)s) \left(\int_{j \in \mathcal{P}_t} v_{t+1}^j dj \right). \tag{60}$$

Equation (60) comes from the fact that the aggregate value of the firms producing in period $t+1$ consists of two parts: the firms that have been producing in period t and the firms that were start-ups in period t. At the same time, we know that the size of the startups is a fraction s of the size of average incumbent. This establishes equation (12).

B 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. 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, which we denote as $\hat{K}_t := K_t/(1+g)^t$, in steady state must satisfy

$$\hat{K}_{t+1} = \underbrace{\Gamma \frac{1+g^k}{1+g} \hat{K}_t}_{\text{period t incumbents}} + \underbrace{(1-\Gamma)s \frac{1+g^k}{1+g} \hat{K}_t}_{\text{period t startups}} = [\Gamma + (1-\Gamma)s] \frac{1+g^k}{1+g} \hat{K}_t. \tag{61}$$

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

$$1 = [\Gamma + (1-\Gamma)s]\,\frac{1+g^k}{1+g} \implies 1+g^k = \frac{1+g}{\Gamma + (1-\Gamma)s} \tag{62}$$

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[28] is of fundamental importance for our exercise. Since all firms in our model grow at the same rate, this fundamentally depends on the share of capital in future exiting firms.[29] 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.

C Intertemporal arbitrage

The following lemma plays a crucial role throughout this subsection. The first part is known as Perron-Frobenius theorem, and Part 4 is sometimes referred to as the Collatz-Wielandt formula. For further details and proofs, see, e.g. Meyer (2023).

Lemma C.1 *Suppose A is a non-negative $n \times n$ matrix that is regular in the sense that for sufficiently large T, A^T is strictly positive. Then the following is true*

[28]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.

[29]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.

1. *The largest eigenvalue of A is real and positive and the associated eigenvector is strictly positive.*

2. *Let λ be the largest eigenvalue of A. Then*

$$\max_{i=1,\dots,n} \sum_{j=1}^{n} a_{ij} \geq \lambda \geq \min_{i=1,\dots,n} \sum_{j=1}^{n} a_{ij}$$

3. *The geometric sum $\sum_{i=0}^{t} A^t$ converges as $t \to \infty$ if and only if the largest eigenvalue of A is less than one.*

4. *For $x \in \mathbb{R}^n_{++}$ define $g(x) = \min_i \frac{[Ax]_i}{x_i}$. Then $g(.)$ attains its maximum at the eigenvector associated with the largest eigenvalue of A.*

C.1 Transitory shocks

Proof of Theorem 3.1

Since the underlying Markov chain is irreducible, the matrix P must be regular in that $P^\tau \gg 0$ for some τ. The equivalence of 1. and 3. follows directly from Lemma C.1. 1. directly implies 2. if we take θ to be the eigenvector associated with the largest eigenvalue of A, appropriately scaled so that its largest entry is below B. Part 4. of Lemma C.1 implies that if intertemporal arbitrage is only possible if the largest eigenvalue of P is at least one.

Proof of Theorem 3.2

The first part follows directly from Lemma C.1 Part 2. To see why the bounds are tight, define a matrix

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

It is easy to see that for $\epsilon = 0$, the largest solution to $\det(\lambda I - \tilde{P}_0) = 0$ must be $\lambda = \frac{1}{R_1}$. Moreover, for sufficiently small $\epsilon > 0$, the implicit function theorem can be applied, and it follows that the largest eigenvalue varies smoothly with ϵ.

Proof of Theorem 3.3

It is easy to see that β must be an eigenvalue of the matrix P, since

$$\det(\beta I - P) = \beta \det(\hat{\pi}),$$

where $\hat{\pi}_{ii} = (1 - \pi_{ii})$ and $\hat{\pi}_{ij} \cdot \hat{\pi}_{ji} = \pi_{ij}\pi_{ji}$ for all $i \neq j$. If there is a single agent who has a (strictly positive) consumption process $c \in \mathcal{L}_0$, and who has an expected utility function

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

there can be no equilibrium where the value of the agent's endowments is infinite; hence for an equilibrium to exist, the largest eigenvalue of P must be smaller than one. If (and only if) $\beta < 1$, an equilibrium always exists. Hence, debt rollover is possible if and only if $\beta \geq 1$.

On the other hand, it is easy to see that without restrictions on risk aversion and consumption volatility, the risk-free rate can be arbitrary. To see this, consider the 2-state case. Direct computation shows that the invariant distribution is given by

$$\pi^* = \left(\frac{\pi(1|2)}{\pi(2|1) + \pi(1|2)}, \frac{\pi(2|1)}{\pi(2|1) + \pi(1|2)} \right). \tag{63}$$

Let $x = \frac{u'(1)}{u'(2)}$, , i.e. the marginal rate of substitution in state 2 with respect to state 1. The average risk-free rate can then be written as

$$1 + \bar{r} = \frac{1}{\beta} \left(\pi_1^* \frac{1}{\pi(2|1)\frac{1}{x} + (1 - \pi(2|1)} + \pi_2^* \frac{1}{\pi(1|2)x + (1 - \pi(1|2))} \right)$$

Assuming an Inada condition and making consumption in state 1 arbitrarily small, we can assume x to be arbitrarily large and obtain

$$1 + \bar{r} \simeq \frac{1}{\beta} \left(\pi_1^* \frac{1}{1 - \pi(2|1)} \right),$$

which, by (63) can be made arbitrarily small with small enough $\pi(1|2)$.

C.2 Permanent growth shocks

Given growth rates $\gamma \in \mathbb{R}^Z_{++}$, we can repeat the above arguments and prove analogues of our three theorems. For this define γ-Arrow prices to be the prices of Arrow securities that pay the growth rate γ_z when shock z realizes, i.e.

$$\tilde{p}(z^{t+1}|z^t) = p(z^{t+1}|z^t)\gamma(z_{t+1}).$$

The analogues of Theorems 3.1 and 3.2 follow directly: the spectral radius of $\tilde{P}$ (its largest eigenvalue) is greater than or equal to one if and only if the value of a risky consol whose payoffs grow at rates $\gamma(z_t)$ is infinite. This is the case if and only if there is an intertemporal arbitrage.

For the analogue of Theorem 3.3 assume

$$p(z^{t+1}|z^t) = \beta\pi(z^{t+1}|z^t)\frac{u'(c(z^{t+1}))}{u'(c(z^t))}, \tag{64}$$

for some concave and increasing function $u(.)$ and some $\beta > 0$. We now consider the other extreme in that there are no temporary shocks and that $\frac{c(z^{t+1})}{c(z^t)} = \gamma(z_{t+1})$ for some stochastic (gross) growth rates $\gamma \in \mathbb{R}^Z_{++}$. We denote the invariant distribution of the Markov chain by π^*, and assume that the average growth rate is fixed to be $(1+g)$, i.e.

$$\sum_z \pi^*(z)\gamma_z = 1 + g.$$

To obtain stationary prices in (64), we assume CRRA utility with coefficient of relative risk-aversion σ and obtain $\tilde{p}(z^{t+1}|z^t) = \tilde{p}(z_{t+1}|z_t) = \beta\pi(z_{t+1}|z_t)\gamma(z_{t+1})^{1-\sigma}$.

With CRRA, define the growth-adjusted Arrow matrix

$$\tilde{P} \;:=\; \big(\tilde{p}(z'|z)\big)_{z,z'\in Z} \;=\; \Pi\,\mathrm{diag}\big(\beta\,\gamma^{1-\sigma}\big).$$

We now state the precise link between $\rho(\tilde{P})$ and utility finiteness.

Lemma C.2 *Utility is finite if and only if $\rho(\tilde{P}) < 1$, where $\rho(\tilde{P})$ denotes the spectral radius of $\tilde{P}$. In this case the Euler conditions are necessary and sufficient (the TVC holds).*

Proof of Lemma C.2: Let $c_0 > 0$ and $c_t = c_0 \prod_{j=1}^{t} \gamma(z_j)$. Under CRRA,

$$\mathrm{E}\left[u(c_t)\right] = \frac{c_0^{1-\sigma}}{1-\sigma}\, \pi_0^{\top}\big(\Pi\, \mathrm{diag}(\gamma^{\,1-\sigma})\big)^{t}\mathbf{1}.$$

Hence

$$\sum_{t\geq 0} \beta^t \mathrm{E}\left[u(c_t)\right] \;\propto\; \pi_0^{\top}\Big(\sum_{t\geq 0}\big(\Pi\, \mathrm{diag}(\beta\gamma^{\,1-\sigma})\big)^{t}\Big)\mathbf{1} \;=\; \pi_0^{\top}\Big(\sum_{t\geq 0}\tilde{P}^{\,t}\Big)\mathbf{1},$$

which converges if and only if $\rho(\tilde{P}) < 1$ by C.1 (part 3). With $\rho(\tilde{P}) < 1$, the present value is finite, so the TVC holds and the Euler equations are sufficient. This concludes the proof of Lemma C.2.

From Lemma C.1 it follows directly that

$$\beta \max_s \sum_z \pi(z|s)\gamma_z^{1-\sigma} < 1 \tag{65}$$

is a sufficient condition for debt-rollover to be impossible. In the special case of i.i.d. shocks, when the rows of Π are identical, condition (65) is also necessary.

As in the case of transitory shocks, the average interest rate can again be arbitrarily small, independently of the average growth rate. To see this, observe that in the i.i.d. case, (65) is a condition on the return on an asset that is a claim to risky consumption. The equity premium is the gap between that and the risk-free rate, $1+\bar{r} = \sum_z \pi_z^* \frac{1}{\beta \sum_s \pi(s|z)\gamma_s^{-\sigma}}$, and can be arbitrarily large if shocks become arbitrarily large.

To verify this claim, consider again the simple case of two shocks, and suppose the shocks are i.i.d. with probabilities π^*. Consider $Z = \{1, 2\}$ with $0 < \gamma_1 < 1 < \gamma_2$ and choose $\pi_1^* = \gamma_1$ and $\pi_2^* = 1 - \gamma_1$. Fix the mean growth by choosing γ_2 so that $\gamma_1\gamma_1 + (1-\gamma_1)\gamma_2 = 1+g$; as $\gamma_1 \to 0$, $\gamma_2 \to 1+g$ and remains bounded. For $\sigma > 2$, set

$$\beta \;=\; \tfrac{1}{4}\,\gamma_1^{\sigma-2}.$$

Each row sum of $\tilde{P} = \Pi\, \mathrm{diag}(\beta\gamma^{\,1-\sigma})$ equals

$$\beta\Big[\gamma_1\,\gamma_1^{1-\sigma} + (1-\gamma_1)\,\gamma_2^{1-\sigma}\Big] = \tfrac{1}{4} \;+\; \tfrac{1}{4}\,(1-\gamma_1)\,\gamma_1^{\sigma-2}\,\gamma_2^{1-\sigma} \;<\; 1$$

for γ_1 small, so condition (65) holds and rollover is impossible. At the same time,

$$R(2) \;=\; \left[\beta \sum_s \pi(s|2)\gamma_s^{-\sigma}\right]^{-1} \approx \left[\beta\, \pi(1|2)\, \gamma_1^{-\sigma}\right]^{-1} = \left[\tfrac{1}{4}\gamma_1^{\sigma-2}\cdot\gamma_1\cdot\gamma_1^{-\sigma}\right]^{-1} = 4\,\gamma_1,$$

so

$$1+\bar{r} \approx \pi_2^*\, R(2) \;=\; 4\,(1-\gamma_1)\,\gamma_1$$

which can be made arbitrarily small.

D Is there a unit root in TFP?

D.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

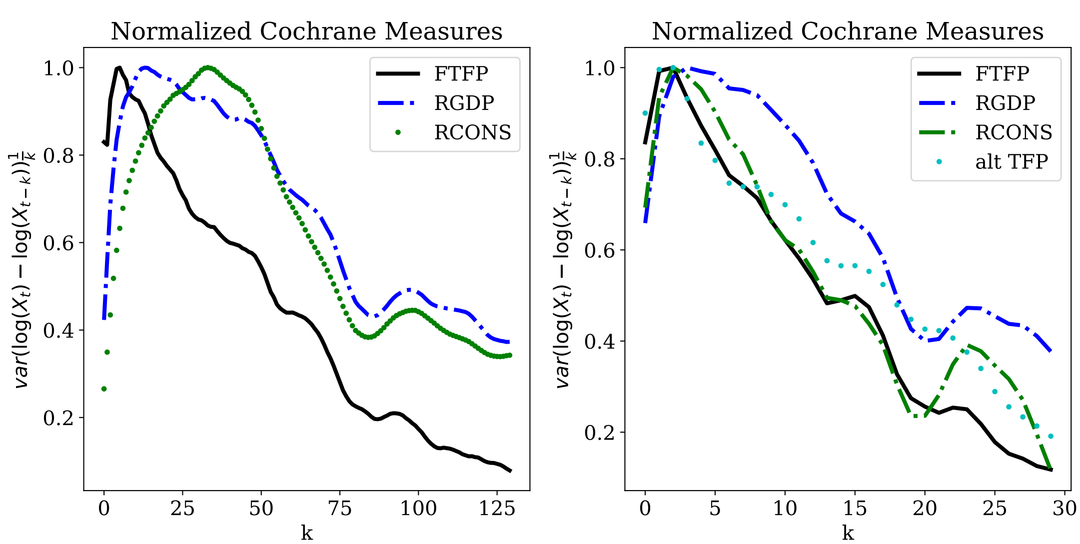

(a) Quarterly Data **(b)** Annual Data

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.[30] 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 one in TFP.

D.2 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 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

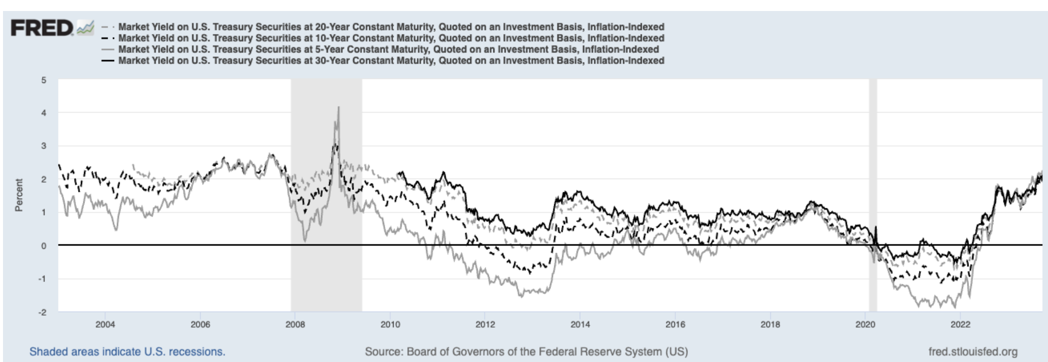

Figure 8: TIPS Yields

the period in which we also have the longer bond, 83.6% of the weeks the yields were strictly upward sloping ($30 > 20 > 10 > 7 > 5$), and 100% of the weeks the 30-year

[30]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.

yield was above the 20-year. Moreover, the average gap in the yields was sizeable even at the long end, with the 30-year less the 20-year gap equal to 0.21 percentage points. The fact that the average yield curve for TIPS is upward sloping is also discussed in Chernov et al. (2021).[31]

Table 4: Average Yields and Yield Curve Inequalities

Average Yields					
	5-year	7-year	10-year	20-year	30-year
2003–2025	1.50	1.87	2.10		
2004–2025	1.31	1.64	1.94	2.39	
2010–2025	0.86	1.27	1.66	2.17	2.38
	0.41	0.39	0.52	0.21	
Yield Curve Inequalities					
	5<7	7<10	10<20	20<30	ALL
2010–2025	0.8422	0.8533	0.9901	1.0000	0.8360

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)).

[31] Source is Board of Governors of the Federal Reserve System (US), Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity, Quoted on an Investment Basis, Inflation-Indexed [WFII10], retrieved from FRED, Federal Reserve Bank of St. Louis. Our data is weekly and runs from 2003-01-03 to 2025-09-05. 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.

E Computational Method

E.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[32] or any approximation method based on a hypercubic domain,[33] would be problematic for two main reasons, which we illustrate in figure 9. Figure 9 shows realizations of the two endogenous aggregate variables, capital (horizontal axis) and the wealth share (vertical axis) of the advanced traders, for each of the exogenous aggregate shocks. The first problem is that the computational effort spent to solve the model for capital wealth-share combinations in the upper left and lower right corners would be completely wasted, since the model never reaches those states. The second problem is that, since the correct shock-contingent transitions for the endogenous aggregate variables are only observed on the simulated path, there is no obvious way how to extend a fitted forecasting 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

[32] The classical Krusell and Smith (1998) approach has only a single summary statistic for the endogenous wealth distribution, so the extrapolation problem is much less severe.

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

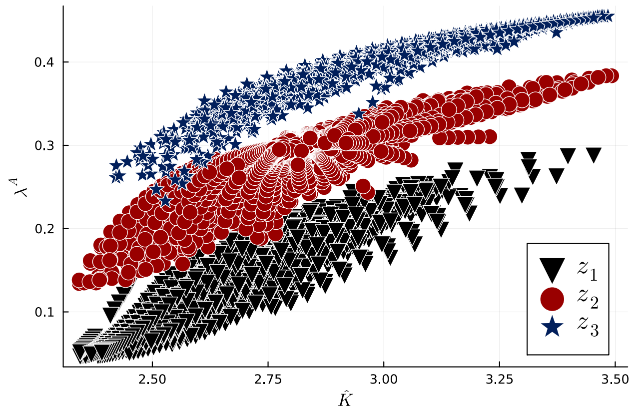

Figure 9: Realizations of the two endogenous aggregate variables, capital (horizontal axis) and the wealth share (vertical axis) of the advanced traders, for each of the exogenous aggregate shocks on a simulated path of 10000 periods for the benchmark model.

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.

E.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

$$g_K\left(z_t, \hat{K}_t, \lambda_t^A\right) \approx \hat{K}_{t+1} \tag{66}$$

$$g_{V^{\text{firm}}}\left(z_t, \hat{K}_t, \lambda_t^A\right) \approx \hat{V}_t^{\text{firm}} \tag{67}$$

$$g_{\lambda^A}\left(z_t, \hat{K}_t, \lambda_t^A, z_{t+1}\right) \approx \lambda_{t+1|z_{t+1}}^A \tag{68}$$

$$g_{c^A}\left(z_t, \hat{K}_t, \lambda_t^A, \eta_t, \frac{\hat{a}_{t-1}^{z_t}}{1+g}\right) \approx \hat{c}_t^A \tag{69}$$

$$g_{c^B}\left(z_t, \hat{K}_t, \lambda_t^A, \eta_t, \frac{\hat{b}_{t-1}}{1+g}\right) \approx \hat{c}_t^B. \tag{70}$$

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 (39). 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:[34]

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

$$\hat{c}_{t+1}^{B,z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}=g_{c^B}^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}\left(\hat{K}_{t+1},\lambda_{t+1}^{A,\text{ forecast}}\right)\tag{71}$$

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

$$(V^{B\prime})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}=u'(\hat{c}^{B,z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}})u'(1+g)\tag{72}$$

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

$$(EV^{B\prime})_t^{\eta_t,\mathcal{B}^{\text{grid}}}=\sum_{z_{t+1},\eta_{t+1}}\pi^z(z_t,z_{t+1})\pi^\eta(\eta_t,\eta_{t+1})(V^{B\prime})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{B}^{\text{grid}}}\tag{73}$$

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

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

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

$$\mathcal{C}_t^{\eta_t,\mathcal{B}^{\text{grid}}} = (u')^{-1}\left(\frac{\beta}{p_t^{\text{bond}}}(EV^{B\prime})_t^{\eta_t,\mathcal{B}^{\text{grid}}}\right) \tag{74}$$

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

$$\hat{b}_t^{\eta_t} = \hat{c}_t^{B,\eta_t} + p_t^{\text{bond}}\hat{b}_{t+1}^{\eta_t}(1+g) - \hat{w}_t\eta_t \tag{75}$$

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

$$d_t^B = \sum_{\eta_t}\sum_{\hat{b}_t^{\eta_t}\in\mathcal{B}^{\text{grid}}} \mu^{B,\eta_t}(\hat{b}_t^{\eta_t}) \cdot \hat{b}_{t+1}^{\text{new},\eta_t}(\hat{b}_t^{\eta_t})(1+g) \tag{76}$$

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

$$\hat{c}_{t+1}^{A,z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} = g_{c^A}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}}\left(\hat{K}_{t+1}, \lambda_{t+1}^{A,\text{ forecast}}\right) \tag{77}$$

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

$$(V^{A\prime})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} = u'(\hat{c}^{A,z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}})u'(1+g) \tag{78}$$

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[35] 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}.

$$(EV^{A\prime})_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}} = \sum_{\eta_{t+1}} \pi^\eta(\eta_t,\eta_{t+1})(V^{A\prime})_{t+1}^{z_{t+1},\eta_{t+1},\mathcal{A}^{\text{grid}}} \tag{79}$$

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

$$\mathcal{C}_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}} = (u')^{-1}\left(\frac{\beta}{p_t^{z_{t+1}}}(EV^{A\prime})_t^{\eta_t,z_{t+1},\mathcal{A}^{\text{grid}}}\right) \tag{80}$$

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

[35] 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.

$\hat{c}_t^{\text{common}}$. 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

$$\hat{s}_t^{\text{common},\eta_t} = \sum_{z_{t+1}} p_t^{z_{t+1}} \hat{a}_{t+1}^{z_{t+1},\text{common},\eta_t}(1+g). \tag{81}$$

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

$$\hat{a}_t^{\text{common},\eta_t} = \hat{c}_t^{\text{common}} + \hat{s}_t^{\text{common},\eta_t} - \hat{w}_t\eta_t - \frac{1-\Gamma}{\Gamma\mu}\hat{\zeta}_t, \tag{82}$$

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

$$d_t^{A,z_{t+1}} = \sum_{\eta_t} \sum_{\hat{a}_t^{z_t,\eta_t}\in\mathcal{A}^{\text{grid}}} \mu^{A,\eta_t}(\hat{a}_t^{z_t,\eta_t}) \cdot \hat{a}_{t+1}^{\text{new},z_{t+1}\eta_t}(\hat{a}_t^{z_t,\eta_t})(1+g) \tag{83}$$

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

$$d_t^{A,z_{t+1}} + d_t^B = g_{V^{\text{firm}}}\left(z_{t+1}, \hat{K}_{t+1}, \lambda_{t+1}^{A \text{ forecast}}\right)(1+g). \tag{84}$$

Together with the optimality condition for firm investment (*i.e.* equation (39)), 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

$$g_K^{z_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\hat{K}_{t+1}^{\text{pred}} + \alpha^{\text{update}}\hat{K}_{t+1}. \tag{85}$$

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

$$g_{\lambda^A}^{z_t,z_{t+1},\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\lambda_{t+1}^{A,z_{t+1},\text{pred}} + \alpha^{\text{update}}\lambda_{t+1}^{A,z_{t+1}}. \tag{86}$$

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.

$$g_{c^A}^{\eta_t, z_t, \hat{a}_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\hat{c}_t^{A,\eta_t,z_t,\hat{a}_t,\text{pred}} + \alpha^{\text{update}}\hat{c}_t^{A,\eta_t,z_t,\hat{a}_t}, \tag{87}$$

$$g_{c^B}^{\eta_t, z_t, \hat{b}_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\hat{c}_t^{B,\eta_t,z_t,\hat{b}_t,\text{pred}} + \alpha^{\text{update}}\hat{c}_t^{B,\eta_t,z_t,\hat{b}_t} \tag{88}$$

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 (58), 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

$$g_V^{z_t,\text{ new}}\left(\hat{K}_t, \lambda_t^A\right) \approx (1-\alpha^{\text{update}})\hat{V}_t^{\text{pred}} + \alpha^{\text{update}}\hat{V}_t^{\text{div}}. \tag{89}$$

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

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

$$g(\hat{K}, \lambda^A) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{90}$$

and obtain the parameters $\{\theta_1, \theta_2, \theta_3, \theta_4, \theta_5\}$ by least-squares fitting the data.[36] 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.

[36] 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 5: 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%.

	$\hat{K}$ [%]	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 6: 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%.

E.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 5 shows that the remaining errors are low, with the 99th percentile well below 0.1%. Similarly, Table 6 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.

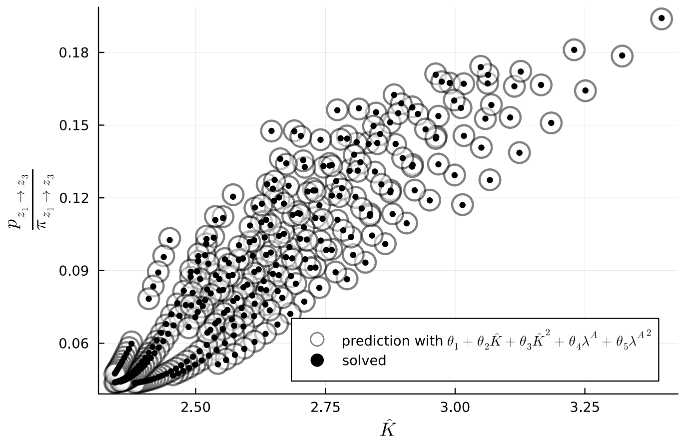

Figure 10: Realizations for the price of the aggregate-state contingent security for shock z_3, when the economy is in shock z_1, against the corresponding values of aggregate capital. The grey circles show the predictions of a fitted polynomial of the form $\theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2$.

E.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 10 shows a scatter plot of capital against the state price for shock z_3, when the economy is in shock z_1 for 300 realizations along the simulated path in the benchmark model, normalized by the transition probability $\pi_{z_1 \to z_3}$. The circles show a prediction of the price when fitting a polynomial of the form $\theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2$, where $\hat{K}$ denotes aggregate capital and λ^A denotes the wealth share of advanced traders. First, we can see that aggregate capital alone is not enough to predict the price accurately. Second, we can see that the prediction based on capital and the simple functional form we are using, provides a very good fit to the data.[37]

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

[37] While this serves as an illustrative example, our algorithm does not require to forecast prices, but the households' consumption. In figure 11 we show that household consumption can be very accurately predicted as well.

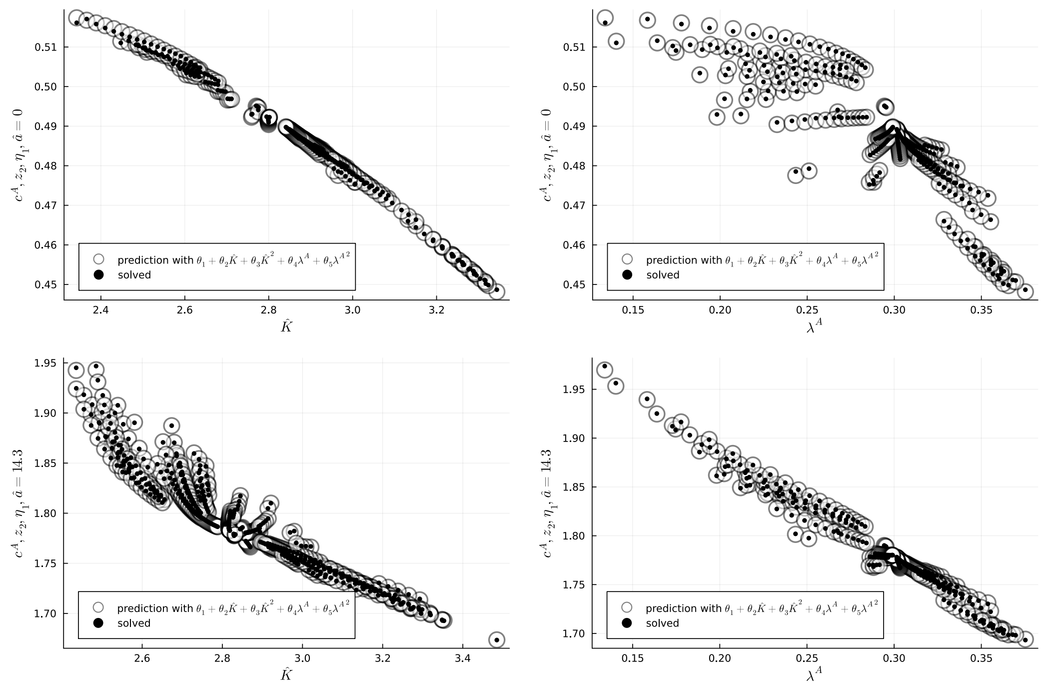

Figure 11: Consumption of an advanced trader in aggregate shock z_2 and idiosyncratic shock η_1 for two different wealth levels along 300 realizations on the simulated path of the economy for the benchmark model. The top panel shows the consumption for a household without any financial wealth, $\hat{a}_t^{z_2} = 0$. 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.

for each of the functions.

$$g_{c^A}^{\eta_t, z_t, \hat{a}_t, \text{ new}}\left(\hat{K}_t, \lambda_t^A\right) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{91}$$

$$g_{c^B}^{\eta_t, z_t, \hat{b}_t, \text{ new}}\left(\hat{K}_t, \lambda_t^A\right) = \theta_1 + \theta_2\hat{K} + \theta_3\hat{K}^2 + \theta_4\lambda^A + \theta_5(\lambda^A)^2 \tag{92}$$

Importantly, the parameters θ_i can differ not only across the exogenous shocks, but also across the wealth level. Figure 11 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 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 11 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.

E.2 Details on surrogate models and 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. To obtain such a mapping we use a surrogate model. We proceed to provide a short introduction to surrogate models in section E.2.1 (see also Scheidegger and Bilionis, 2019; Catherine et al., 2022, for the use of surrogate models in economics) and then provide additional details on how we use them for our calibration in section E.2.2.

E.2.1 Surrogate models

Suppose we have a function $f : x \in \mathbb{R}^N \to y = f(x) \in \mathbb{R}$. Further, suppose that the function is unknown, but we can evaluate the function at any point x_i to obtain $y_i = f(x_i)$ at high computational cost. In our application, x_i will be a vector of model parameters $x_i = [\beta, \xi^{\text{adj}}] \in \mathbb{R}^2$, and the function $y_i = f(x_i)$ a moment of interest implied by the model for the specified parameters, such as the average interest rate. Obtaining y_i for a specific x_i is costly because it requires computing the equilibrium and then performing a long stochastic simulation of the model to compute the implied average interest rate.

The idea of a surrogate model is to approximate the function $f(x_i)$ by a surrogate function $\hat{f}(x)$, such that: first, the surrogate is a good approximation of the actual function $\hat{f}(x) \approx f(x)$ and second, the surrogate is computationally cheap to evaluate.

Because the function f is expensive to evaluate, we want to be able to obtain a good surrogate model from a small number of data points $\mathcal{D}_{N_{\text{data}}} = \{(x_i, y_i)\}_{i=1}^{N_{\text{data}}}$. In practice, Gaussian Processes, a non parametric machine learning method to obtain a function approximation $\hat{f}(x)$ from a data set $\mathcal{D}_{N_{\text{data}}}$, provide us with a good way to obtain surrogate models.[38] What makes Gaussian processes desirable is that they can used with unstructured data,[39] and that they are suitable for giving good approximations with comparatively few data points.

E.2.2 Calibration

We use Gaussian Processes with a squared exponential kernel to approximate the quantities of interest in surrogate model. 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.[40] 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 Figure 3, 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% and an average market price of risk of 0.50. We then use the surrogate model for the expected log price earnings ratio to obtain an estimate for its value. To produce figure 4, we use a Gaussian process to estimate the mapping for the interest rate and the market price of risk to the log mean price of the infinitely lived consol.

In the main text, Figure 1 shows the market price of risk, the interest rate, as well as the objective function, which we aim to minimize with our calibration. Figure 2 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

[38] See, Rasmussen and Williams (2004) for a general introduction to Gaussian Processes and Scheidegger and Bilionis (2019) or Kübler et al. (2025) for applications in economics.

[39] This is in difference to grid-based approximations methods such as sparse grid (see, Krueger and Kubler, 2004) or adaptive sparse grids(see, Brumm and Scheidegger, 2017).

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

4, these untargeted moments are roughly in line with what we measure in the data.

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

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

$$y_{it} = g(x_{it}^h, Y_t) + u_{it}^h \tag{93}$$
$$u_{it} = \alpha_i + z_{it}^h + \epsilon_{it} \tag{94}$$
$$z_{it}^h = \rho z_{i,t-1}^{h-1} + \eta_{it} \tag{95}$$
$$\alpha_i \sim \text{i.i.d. } N(0, \sigma_\alpha^2) \tag{96}$$
$$\epsilon_{it} \sim \text{i.i.d. } N(0, \sigma_\epsilon^2) \tag{97}$$
$$\eta_{it} \sim \text{i.i.d. } N(0, \sigma_t^2) \tag{98}$$
$$\sigma_t^2 = \begin{cases} \sigma_E^2 & \text{in agg. expansions} \\ \sigma_C^2 & \text{in agg. contractions} \end{cases}. \tag{99}$$

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

$$x_{it} = \epsilon_{it} + z_{it} \tag{100}$$
$$z_{it} = \rho z_{it-1} + \eta_{it} \tag{101}$$
$$\epsilon_{it} \sim \text{i.i.d. } N(0, \sigma_\epsilon^2) \tag{102}$$
$$\eta_{it} \sim \text{i.i.d. } N(0, \sigma_\eta^2), \tag{103}$$

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

$$x_{it} = \bar{x} + \rho_x x_{i,t-1} + \sigma_x \epsilon_{it} \tag{104}$$
$$\epsilon_t \sim \text{i.i.d. } N(0, 1), \tag{105}$$

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

$$X = \begin{pmatrix} 0.463 \\ 1.537 \end{pmatrix} \tag{106}$$

$$\pi^X = \begin{pmatrix} 0.892 & 0.108 \\ 0.108 & 0.892 \end{pmatrix} \tag{107}$$

The resulting cross-sectional standard deviation of log earnings is 0.60 and matches the standard deviation of the process simulated in equations (93) - (99) and is in the ballpark of values typically used in the heterogeneous agents literature (see, *e.g.* Auclert et al., 2021).

G 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 (48). Hence, (44) 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 7. 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.

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 7: 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.

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 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 $D < 1$.

$$\tilde{v}_t = d_t + \frac{g^k}{1+r}\left[\Gamma\tilde{v}_{t+1} + (1-\Gamma)R\right] = d_t + \frac{g^k\tilde{v}_{t+1}}{1+r}\left[\Gamma + (1-\Gamma)\tilde{R}_{t+1}\right]$$

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

References

Abel, A. B., Mankiw, N. G., Summers, L. H., and Zeckhauser, R. J. (1989). Assessing dynamic efficiency: Theory and evidence. *The Review of Economic Studies*, 56(1):1–19.

Abel, A. B. and Panageas, S. (2025). Running primary deficits forever in a dynamically efficient economy: Feasibility and optimality. *Econometrica*, 93(5):1601–1633.

Aguiar, M., Amador, M., and Arellano, C. (2024). Micro risks and (robust) pareto-improving policies. *American Economic Review*, 114(11):3669–3713.

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. J. (2024). Unique implementation of permanent primary deficits? Technical report, Federal Reserve Bank of Minneapolis.

Angeletos, G.-M. (2002). Fiscal policy with noncontingent debt and the optimal maturity structure. *The Quarterly Journal of Economics*, 117(3):1105–1131.

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.

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.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2017). Public debt in economies with heterogeneous agents. *Journal of Monetary Economics*, 91:39–51.

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

Bloise, G. and Reichlin, P. (2023). Low safe interest rates: A case for dynamic inefficiency? *Review of Economic Dynamics*, 51:633–656.

Brumm, J., Feng, X., Kotlikoff, L., and Kubler, F. (2024). When interest rates go low, should public debt go high? *American Economic Journal: Macroeconomics*, 16(4):432–469.

Brumm, J., Kryczka, D., and Kubler, F. (2017). Recursive equilibria in dynamic economies with stochastic production. *Econometrica*, 85(5):1467–1499.

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.

Cass, D. (1972). On capital overaccumulation in the aggregative, neoclassical model of economic growth: A complete characterization. *Journal of Economic Theory*, 4(2):200–223.

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

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.

Christensen, T. M. (2017). Nonparametric stochastic discount factor decomposition. *Econometrica*, 85(5):1501–1536.

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.

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.

Deimling, K. (2013). *Nonlinear functional analysis*. Springer Science & Business Media.

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.

Jiang, Z., Lustig, H., Van Nieuwerburgh, S., and Xiaolan, M. Z. (2024). The us public debt valuation puzzle. *Econometrica*, 92(4):1309–1347.

Kocherlakota, N. R. (2023). Infinite debt rollover in stochastic economies. *Econometrica*, 91(5):1629–1658.

Kocherlakota, N. R. (2024). Difficulties in testing for capital overaccumulation. *Quantitative Economics*, 15(1):89–114.

Krueger, D. and Kubler, F. (2004). Computing equilibrium in olg models with stochastic production. *Journal of Economic Dynamics and Control*, 28(7):1411–1436.

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*.

Kübler, F., Scheidegger, S., and Surbek, O. (2025). Using machine learning to compute constrained optimal carbon tax rules. *arXiv preprint arXiv:2507.01704*.

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.

Meyer, C. D. (2023). *Matrix analysis and applied linear algebra*. SIAM.

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.

Rangazas, P. and Russell, S. (2005). The zilcha criterion for dynamic inefficiency. *Economic Theory*, 26(3):701–716.

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.

Samuelson, P. A. (1958). An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of political economy*, 66(6):467–482.

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.

Zilcha, I. (1990). Dynamic efficiency in overlapping generations models with stochastic production. *Journal of Economic Theory*, 52(2):364–379.