

Heterogeneity and Asset Prices: A Different Approach

Nicolae Gârleanu

UC Berkeley-Haas and NBER

Stavros Panageas

UCLA Anderson School of Management and NBER*

June 2018

Abstract

We develop a macro-asset pricing framework that links volatile asset prices and high risk premiums to non-volatile, but persistent, movements in the cross-sectional income and consumption distributions. We propose a novel empirical approach to infer low frequency, time-series movements in the marginal agent's consumption over a long period of time by utilizing cross-sectional (rather than time-series) information on the consumption of different cohorts. In a calibration we use these inferred low frequency components to assess the model's ability to account quantitatively for the stylized asset pricing facts (high market price of risk, equity premium, volatility, return predictability, etc.). We also show how imperfect risk sharing suffices to allow anticipations of future movements in discount rates to become a self-fulfilling source of asset price volatility.

Keywords: asset pricing, heterogeneity, imperfect risk sharing, overlapping generations, cohort studies

JEL Classification: G01, G12

*We are grateful for comments and suggestions from Valentin Haddad, Jon Ingersoll, Alberto Bisin, Pietro Veronesi, as well as participants at WFA, NBER Conference on multiple Equilibria, Princeton, Minneapolis FED, Yale SOM, UT Austin, Boston College, Chicago FED, NYU Stern, University of Notre Dame, Kellogg, Chicago Booth, Berkeley Haas-Stanford GSB joint seminar, UCLA Macro-Finance lunch. An earlier version of this paper was circulated under the title "What to Expect when Everyone is Expecting: Self-Fulfilling Expectations and Asset-Pricing Puzzles".

1 Introduction

We construct equilibria of continuous-time overlapping-generations economies with imperfect risk sharing that can jointly account for volatile asset-price fluctuations and high risk premiums in an economy with deterministic aggregate growth and non-volatile changes in the cross-sectional consumption and wealth distributions. Unlike the seminal paper of Constantinides and Duffie (1996), where the risk premiums are driven by i.i.d. shocks to the cross-sectional variance of consumption changes, in our construction all the moments of the cross-sectional distribution of consumption evolve in a non-volatile, indeed locally deterministic¹ and persistent manner — features that seem empirically attractive.

Because of imperfect risk sharing, the stochastic discount factor does not coincide with aggregate per-capita consumption growth. Rather, the identity and weights of agents who are marginal for asset pricing keeps changing. We develop a novel empirical methodology to infer the consumption growth of the “representative, marginal agent” in the absence of perfect risk sharing. By utilizing a minimal part of the model structure, the methodology is flexible enough to account in a comprehensive manner for a multitude of factors (different cohort sizes, age-dependent life-cycle effects, demographic shifts in the population, different cohort productivities, etc.). Equally importantly, the methodology utilizes a few repeated cross sections of data to infer a long time series path of the marginal agent’s consumption growth. We show that the marginal agent’s consumption growth in the absence of perfect risk sharing exhibits substantially stronger persistence and long-run variance than a representative agent model would suggest.

We next describe the theoretical structure and the empirical methodology in greater detail.

The framework is a continuous-time, overlapping generations economy. Agents arrive continuously and are either “workers” or “entrepreneurs.” The difference lies exclusively in the stochastic properties of their endowments: Workers obtain a stochastic path of wages over

¹Throughout the paper “locally deterministic” refers to a time-differentiable process. By definition, such a process has no diffusion component, but a possibly stochastic drift process. This means that the volatility of ever shorter increments disappears in the limit.

their life-times, while entrepreneurs come endowed with the shares of a newly created firm that produces a stochastic stream of dividends. To illustrate the difference with the standard representative agent models, all shocks in this economy are re-distributional: They affect the fraction of dividend (resp. labor) income that is obtained by firms (resp. workers) born at different times, while *aggregate* labor and dividend income and aggregate consumption are *deterministic* process all growing at the same constant rate. Moreover, to distinguish our setup from the one employed by Constantinides and Duffie (1996), the shares of labor or dividend income accruing to a given cohort of investors are locally deterministic processes, i.e., their volatility vanishes as the time interval becomes small; however, these processes are random over the long run.

In this setup we introduce a) imperfect inter-cohort risk sharing and b) recursive preferences with preferences for early resolution of uncertainty. We utilize this framework to perform two exercises, a theoretical and an empirical.

Our theoretical exercise parallels the exercise in Constantinides and Duffie (1996). Specifically, we show the following “possibility” result: Share processes exist that support (essentially) any given stationary processes for both the market price of risk and the price-dividend ratio. The main difference between the construction in Constantinides and Duffie (1996) and our paper pertains to the time-series implications for inequality measures, such as the cross-sectional variance of consumption. Constantinides and Duffie (1996) rely on heterogeneous period-by-period changes in individual consumption-growth dispersion, which lead to period-by-period movements in inequality; by contrast, we rely on dispersion and uncertainty in the integrated consumption growth experienced by cohorts born at different times. As a result in our approach inequality measures are substantially less volatile on a period-by-period basis, but quite persistent. An important byproduct of our possibility result is that it makes the model quite tractable. We show how to judiciously postulate a functional form for the stochastic processes followed by the dividend and labor shares, so that we can obtain a simple closed form solution for asset prices.

For the empirical exercise, we infer the persistent component of the marginal agent’s consumption process that is driven by the re-distributional shocks, arising either by different

cohort sizes or productivities. The novel aspect of this technique is that the time series properties of the persistent component are inferred from cross-sectional rather than time-series data. Specifically, the procedure involves estimating a regression whereby cross-sectional consumption is regressed on time, age, and cohort dummies. Using the information in age and cohort effects, we can re-construct a time series path of the consumption growth that the marginal agents (i.e., the surviving agents between times t and $t + 1$) must have experienced over the sample. Since this technique does not rely on time-series filtering methods, the estimation error disappears as the number of observations inside each cohort tends to infinity.

To illustrate the quantitative implications of the model, we calibrate it, document its quantitative properties, and compare them to the data. We show that the model produces realistic risk premiums, return predictability, interest rate levels, and volatility, even though consumption and dividends have small annual volatility. We also illustrate that the cross sectional consumption variance has negligible volatility but follows a near-unit-root process, consistent with the data.

We conclude the paper with a theoretical result that is meant to illustrate some of the conceptual differences between representative-agent and heterogeneous-agent modeling of long-run risks. Specifically, we show that in our model long-run fluctuations in individual consumption growth may be caused by, rather than be causal for, asset-price fluctuations. To this end we employ a variant of our model in which new company creation is endogenous and reacts to current prices. We show that in such a model all uncertainty can arise endogenously, due to coordinated shifts in agents' anticipations of future discount rates, which manifest themselves in asset-price fluctuations. These asset-price fluctuations lead to endogenous consumption dynamics that justify the assumed anticipations, making the assumed expectation of long run consumption changes self-confirming. Importantly, this result does not rely on bubbles or money, the usual suspects for indeterminacy in overlapping generations models and may obtain in an economy in which the average interest rate is higher than the growth rate.

1.1 Relation to the literature

The theoretical exercise of the paper is similar to Constantinides and Duffie (1996). Like Constantinides and Duffie (1996), we show how to construct a process for inequality that can help address asset pricing puzzles. However, our approach does not rely on volatile year-over-year changes in dispersion, but rather on low frequency movements in cross-sectional income and consumption dispersion.

The model we employ is a stochastic, endowment version of the Blanchard (1985) perpetual youth model. The attractiveness of this framework is that it allows a meaningful discussion of imperfect risk sharing across cohorts, but sidesteps the technical complications of more conventional overlapping generations models, where the wealth distribution across different cohorts becomes a state variable.² By sidestepping this complication, we can tractably model a continuum of co-existing cohorts and continuous trading, while retaining tractability.

Gârleanu et al. (2012) also use a perpetual youth model with stochastic fluctuations in the labor income and profit shares of different cohorts. A shortcoming of Gârleanu et al. (2012) is that there is no time-variation in the Sharpe ratio, equity premium, return volatility etc. Moreover, the continuous-time limit of Gârleanu et al. (2012) would imply that there is no finite-rate specification for the arrival of new wealth.³ This feature would imply volatile short run changes in consumption cohort effects, a difficulty that we overcome in the present model. In addition, we present a novel technique to empirically identify the impact of redistributive shocks on the consumption growth of marginal agents using exclusively cross-sectional data.

The paper also relates to the literature on long run risks, which was initiated by Bansal and Yaron (2004). The point we make in this paper is complementary to Bansal and Yaron (2004). We show that risk sharing imperfections may imply uncertain long run growth rates for the marginal agent, even though aggregate consumption and dividends are both

²Indicative examples of such papers are Constantinides et al. (2002), Gomes and Michaelides (2005), Storesletten et al. (2007), and Piazzesi and Schneider (2009). The literature on demographic shocks to asset prices, which we don't attempt to summarize here, is also (remotely) related to the present paper.

³The Gamma process, which that paper employs, is a discontinuous process in its continuous time limit.

deterministic.⁴

There is a voluminous literature that utilizes time-, age- and cohort- decompositions in repeated cross-sections. We do not attempt to summarize this literature here. One common issue in that literature is that linear trends in time-, age-, and cohort- effects cannot be separated. We contribute to that literature by showing that despite this issue, it is possible to identify the consumption evolution of the “marginal agent” over a sample length that is equal to the number of cohort dummies, rather than time-dummies. To the best of our knowledge, this result is new, and its applicability extends beyond the specific asset pricing application that we consider in this paper.

The last section of our paper relates to the literature on indeterminacy in overlapping generations models. Even though there are numerous ways to obtain indeterminacy in macro models, a distinguishing feature of our approach is that it does not require bubbles, money, or increasing returns to scale. Rather, the mechanism relies on anticipations of future discount rates, whose impact on cross-sectional inequality renders them self-sustaining.

2 Model

We present the baseline model in two steps. In a first step we lay out the assumptions of the model assuming that agents have expected utility, logarithmic preferences. In a second step we extend the analysis to recursive preferences.

2.1 Consumers

Time is continuous. Each agent faces a constant hazard rate of death $\lambda > 0$ throughout her life, so that a fraction λ of the population perishes at each instant. A new cohort of mass λ is born per unit of time, so that the total population remains at $\lambda \int_{-\infty}^t e^{-\lambda(t-s)} ds = 1$.

Consumers maximize the utility they derive from their stream of consumption. In Section 3.3, where we derive the main result, the preferences take the form of recursive utility

⁴A similar intuition is present in Kogan et al. (2016). The focus of our paper is different, however: We utilize a simpler model, and focus on obtaining an in-depth understanding of both the theoretical and empirical connections between the lack of inter-cohort risk sharing and asset pricing.

with unitary intertemporal elasticity of substitution (IES). In Sections 3.1–3.2 we illustrate our approach in the special case of logarithmic utility, i.e., consumers maximize

$$E_s \left[\int_s^\infty e^{-\rho(t-s)} \log(c_{t,s}) dt \right], \quad (1)$$

where s is the time of their birth and t is calendar time.

Consumers have no bequest (or gift) motives for simplicity.

2.2 Endowments

Following a long tradition in asset pricing, we consider an endowment economy. The total endowment of the economy is denoted by Y_t and evolves exogenously according to

$$\frac{\dot{Y}_t}{Y_t} \equiv g, \quad (2)$$

where $g > 0$. We intentionally model the aggregate endowment as a deterministic, constant-growth process in order to isolate the effect of redistribution shocks.⁵ In this section we specify how this aggregate endowment accrues to the agents populating the economy.

At birth, agents are of two types, to which we refer as “entrepreneurs” and “workers.”

Entrepreneurs are a fraction ε of the population. Per unit of time a mass $\lambda\varepsilon$ of entrepreneurs is born. An entrepreneur born at time s introduces a new firm into the market, whose time $t > s$ dividends are equal to $\frac{D_{t,s}}{\lambda\varepsilon}$, where $D_{t,s}$ is the total dividend stream accruing to all firms born at time s and is given by

$$D_{t,s} = \alpha Y_t \eta_s^d e^{-\int_s^t \eta_u^d du} \quad (3)$$

for all $t \geq s$. The term $\alpha \in (0, 1)$ is a constant, while $\eta_t^d \geq 0$ is assumed to follow a non-negative diffusion

$$d\eta_t^d = \mu_t^d dt + \sigma_t^d dB_t.$$

⁵Extending the analysis to allow for a more general aggregate endowment process is straightforward.




Figure 1: Illustration of $D_{t,s}$ for a constant η_t^d .

In equation (3) we can interpret α as the fraction of output that is paid out as dividends, and $\eta_s^d e^{-\int_s^t \eta_u^d du} \geq 0$ as the fraction of dividends accruing to firms of vintage s , since $\int_{-\infty}^t \eta_s^d e^{-\int_s^t \eta_u^d du} ds = 1$ for any path of η_t^d .⁶ Accordingly, aggregating across firms of all vintages gives

$$D_t^A \equiv \int_{-\infty}^t D_{t,s} ds = \alpha Y_t \int_{-\infty}^t \eta_s^d e^{-\int_s^t \eta_u^d du} ds = \alpha Y_t. \quad (4)$$

Figure 2.2 illustrates the time-paths of dividends for firms of different vintages in the simple case where $\eta_t^d = \eta^d$ is a constant. The figure shows that firms belonging to any given cohort s account for a smaller and smaller fraction of aggregate dividends as time t goes by. This is an empirically motivated feature of the model.

We note that even though we have specified the allocation of dividends as a primitive of the model, such a specification would arise naturally in any standard creative destruction

⁶Specifically, this statement holds true for paths of η_t^d satisfying $\int_{-\infty}^t \eta_s^d ds = \infty$. For the type of stochastic processes that we shall consider for η_t^d this property will hold almost surely.

model where new firms embody ideas for producing consumption goods that are rivalrous to previous firms.

We next turn to workers. The specification of workers' endowments mirrors the one for dividends and is a simple extension of the specification in Blanchard (1985). Specifically, per unit of time a mass $(1 - \varepsilon)\lambda$ is born. Accordingly, the time- t density of surviving workers who were born at time s is given by $l_{t,s} = \lambda(1 - \varepsilon)e^{-\lambda(t-s)}$. The time- t endowment $w_{t,s}$ of a worker who was born at time $s \leq t$ is given by

$$w_{t,s} \equiv \frac{(1 - \alpha) Y_t \eta_s^l e^{-\int_s^t \eta_u^l du}}{l_{t,s}}, \quad (5)$$

where $\eta_t^l \geq 0$ is assumed to follow a non-negative diffusion

$$d\eta_t^l = \mu_t^l dt + \sigma_t^l dB_t.$$

As with dividend income, the term $\eta_s^l e^{-\int_s^t \eta_u^l du}$ can be interpreted as the share of aggregate earnings that accrues to the cohort of workers born at time s . Repeating the observations we made earlier, we have

$$\int_{-\infty}^t w_{t,s} l_{t,s} ds = (1 - \alpha) Y_t.$$

2.3 Markets

Markets are dynamically complete. Investors can trade in instantaneously maturing riskless bonds in zero net supply, which pay an interest rate r_t . Consumers can also trade claims on all existing firms (normalized to unit supply). Finally, investors can access a market for annuities through competitive insurance companies as in Blanchard (1985), allowing them to receive an income stream of $\lambda W_{t,s}$ per unit of time, where $W_{t,s}$ is the consumer's financial wealth. In exchange, the insurance company collects the agent's financial wealth when she dies. Entering such a contract is optimal for all agents, given the absence of bequest motives.

A worker's dynamic budget constraint is given by

$$dW_{t,s} = (r_t + \lambda) W_{t,s} dt + w_{t,s} dt + \theta_{t,s} (dP_t + D_t^A dt - r_t P_t dt), \quad (6)$$

where P_t is the value of the market portfolio at time t and $\theta_{t,s}$ is the number of shares of the market portfolio. Specification (6) assumes that the consumer trades only in shares of the market portfolio, rather than individual firms. This is without loss of generality in our setup, since all existing firms have identical dividend growth rates, and hence any firm (and any portfolio of firms) must have the same return, otherwise there would be an arbitrage. So, in order to save notation, we simply assume that the consumer trades shares of the market index.

For a worker, $W_{t,t} = 0$. An entrepreneur's dynamic budget constraint is identical, except that the term $w_{t,s}$ is replaced by zero and the initial wealth $W_{t,t}$ is given by the value of the firm that the entrepreneur creates.

2.4 Equilibrium

The equilibrium definition is standard. We let $c_{t,s}^e$ (resp. $c_{t,s}^w$) denote the time- t consumption of an entrepreneur (resp. worker) born at time s and $\theta_{t,s}^e$ ($\theta_{t,s}^w$) her holding of stock. With $c_{t,s} = \varepsilon c_{t,s}^e + (1 - \varepsilon) c_{t,s}^w$ the per-capita consumption of cohort s and, similarly, $\theta_{t,s} = \varepsilon \theta_{t,s}^e + (1 - \varepsilon) \theta_{t,s}^w$, we look for consumption processes, asset allocations $\theta_{t,s}$, asset prices $P_{t,s}$, and an interest rate r_t such that a) consumers maximize (1) subject to (6), b) the goods market clears, i.e., $\lambda \int_{-\infty}^t e^{-\lambda(t-s)} c_{t,s} ds = Y_t$, and c) assets markets clear, i.e., $\int_{-\infty}^t \lambda e^{-\lambda(t-s)} \theta_{t,s} ds = 1$ and $\int_{-\infty}^t \lambda e^{-\lambda(t-s)} (W_{t,s} - \theta_{t,s} P_t) ds = 0$.

Once we have determined the market-clearing price of the market index P_t , it is straightforward to determine the price of any individual firm by the absence of arbitrage. Since all existing stocks experience the same dividend growth rates, their price-to-dividend ratios

must be the same and therefore⁷

$$P_{t,s} = P_t \frac{D_{t,s}}{D_t^A} = P_t \eta_s^d e^{-\int_s^t \eta_u^d du}.$$

3 Solution and Analysis

This section contains the theoretical results of the paper. This section mirrors the spirit of the exercise performed in Constantinides and Duffie (1996). Similar to Constantinides and Duffie (1996), we establish the existence of share processes η_t^l and η_t^d that can support given processes for the asset pricing quantities as equilibrium outcomes.

The section is divided into four subsections. In subsection 3.1 we derive, under the logarithmic-preference assumption, a key relation linking the processes η_t^d and η_t^l to the dynamics of the price-dividend ratio, which we denote by q . We use this relation in subsection 3.2 to establish the existence of processes η_t^d and η_t^l that can support any given process for q as an equilibrium outcome. In subsection 3.3 we enrich the setup to allow for recursive preferences and show how to obtain any (joint) dynamics for the price-dividend ratio and the market price of risk (Sharpe ratio) in equilibrium.

Besides providing a comprehensive mapping from assumptions on η_t^l and η_t^d to the equilibrium processes for the price-dividend ratio and the market price of risk, the propositions of this section have a practical implication: They can help determine the functional forms to assume for the diffusions η_t^d and η_t^l to ensure a given (closed-form) expression for the price-dividend ratio and the Sharpe ratio. We illustrate this statement with two examples.

The last subsection (3.4) concludes with a discussion of the implications of the model for the joint dynamics of inequality and asset prices and highlights the differences of our framework from the literature.

⁷We note in passing that the market-clearing condition $\int_{-\infty}^t \lambda e^{-\lambda(t-s)} \theta_{t,s} ds = 1$ directly implies that the price $P_{t,s'}$ clears the market for stocks of vintage s' for any s' . Indeed, by holding the market portfolio a consumer born at time s allocates a weight equal to $w_{t,s'} = \eta_{s'}^d e^{-\int_s^t \eta_u^d du}$ to stocks of vintage s' , and hence holds $\frac{w_{t,s'} W_{t,s}}{P_{t,s'}}$ shares of firm s' . Aggregating across consumers and noting that $w_{t,s'} = \frac{P_{t,s'}}{P_t}$ implies that $\int_{-\infty}^t \lambda e^{-\lambda(t-s)} \left(\frac{w_{t,s'} W_{t,s}}{P_{t,s'}} \right) ds = \frac{\int_{-\infty}^t \lambda e^{-\lambda(t-s)} W_{t,s} ds}{P_t} = 1$.

3.1 Logarithmic utility

We start by conjecturing that in this economy investors' consumption processes are locally deterministic and since agents have expected utility preferences, there are no risk premia and the equilibrium stochastic discount factor m_t follows the dynamics

$$\frac{dm_t}{m_t} = -r_t dt,$$

for an interest rate process that will be determined in equilibrium. (In subsection 3.3 we allow agents to have recursive preferences, and show that the stochastic discount factor exhibits a positive market price of risk). We employ the following definition.

Definition 1 Let $q_{t,s}^d$ denote the ratio of the present value of the dividend stream $D_{u,s}$ to the current dividend:

$$q_{t,s}^d \equiv \frac{E_t \int_t^\infty \frac{m_u}{m_t} D_{u,s} du}{D_{t,s}}. \quad (7)$$

Similarly, we define the respective valuation ratio for earnings, $q_{t,s}^l$:

$$q_{t,s}^l \equiv \frac{E_t \int_t^\infty e^{-\lambda(u-t)} \frac{m_u}{m_t} w_{u,s} du}{w_{t,s}}. \quad (8)$$

Remark 1 Both $q_{t,s}^d$ and $q_{t,s}^l$ are independent of s , since $\frac{D_{u,s}}{D_{t,s}}$ and $\frac{w_{u,s}^l}{w_{t,s}^l}$ are not functions of s . Accordingly, we shall write q_t instead of $q_{t,s}$.

We will refer to q_t^d as the price-to-dividend ratio. We note the following simplification due to unitary IES.

Lemma 1 Let $\beta \equiv \rho + \lambda$. In any (bubble-free) equilibrium,

$$\alpha q_t^d + (1 - \alpha) q_t^l = \frac{1}{\beta}. \quad (9)$$

Equation (9) is intuitive. It states that the sum of the present values of all dividend income accruing to existing firms ($q_t^d \alpha Y_t$) and all earnings accruing to existing agents

$(q_t^l (1 - \alpha) Y_t)$ equals the present value of the aggregate consumption of existing agents $(\frac{C_t}{\beta})$.

By Lemma 1, q_t^l can be expressed as a simple (affine) function of q_t^d . Therefore, from now on we shall concentrate our efforts on determining q_t^d and we'll simplify notation by writing q_t instead of q_t^d .

The goal of this section is to solve for the equilibrium quantities r_t and, especially, q_t as functions of the input variables η_t^d and η_t^l . The result to take away is Lemma 4, which we use in the next subsection to provide an answer to the main question of the paper, namely the construction of a mapping from the dynamics of q_t to those of η_t^d and η_t^l . The remainder of the subsection works towards establishing Lemma 4.

Applying Ito's Lemma to (7) yields the drift of the diffusion process q_t as⁸

$$\mu_{q,t} \equiv (r_t - g + \eta_t^d) q_t - 1. \quad (10)$$

Equation (10) is an arbitrage relation between stocks and bonds. After some re-arranging, it states that the expected percentage capital gain on stocks, $\frac{\mu_q}{q}$, plus the dividend yield, $\frac{1}{q}$, minus the depreciation (or appreciation) rate $\eta^d - g$, should equal the interest rate r .

The drift of q_t depends on the equilibrium interest rate r_t . To determine this equilibrium interest rate, we proceed in three steps.

First, the Euler equation for agents with log preferences implies that the consumption dynamics of any agent are given by⁹

$$\frac{\dot{c}_{t,s}}{c_{t,s}} = -(\rho - r_t), \quad (11)$$

which also implies that $\frac{\dot{c}_{t,s}}{c_{t,s}}$ is independent of s .

Second, as we show in the appendix, the market clearing condition for aggregate con-

⁸To see this, note that (7) implies that $m_t q_{t,s}^d D_{t,s} + \int_{-\infty}^t m_u D_{u,s} du$ must be a martingale, and hence the drift of this expression must be zero.

⁹For a derivation of the Euler equation in our perpetual youth model we refer to Gârleanu and Panageas (2015).

sumption implies that the consumption growth of an existing agent equals

$$\frac{\dot{c}_{t,s}}{c_{t,s}} = g + \lambda - \lambda \frac{c_{t,t}}{Y_t}, \quad (12)$$

which is very intuitive: the growth in consumption to an existing agent consists of the growth in aggregate consumption (g), plus the consumption share that perishing agents do not consume (λ), minus the consumption shares accruing to newly born agents ($\lambda c_{t,t}/Y_t$).

Finally, the intertemporal budget constraint at the time of a consumer's birth leads to:

Lemma 2 *Let $\varphi_t \equiv \eta_t^d - \eta_t^l$. The arriving agents' consumption is given by*

$$\frac{c_{t,t}}{Y_t} = \frac{\beta}{\lambda} ((1 - \alpha) \eta_t^l q_t^l + \alpha \eta_t^d q_t^d) \quad (13)$$

$$= \frac{1}{\lambda} (\eta_t^l + \alpha \beta q_t \varphi_t). \quad (14)$$

Equation (13) states that the per-capita consumption of an arriving cohort of agents is given by the consumption-to-wealth ratio for an investor with unit elasticity of substitution, β , multiplied by the sum of the per-capita value of new firms, $\alpha \eta_t^d q_t^d Y_t$, and the cohort's present value of labor income at birth, $(1 - \alpha) \eta_t^l q_t^l Y_t$. Equation (14) follows from Lemma 1.

Combining equations (11), (12), and (14) leads to

Lemma 3 *The equilibrium interest rate is given by*

$$r_t = \beta + g - \eta_t^l - \alpha \beta \varphi_t q_t. \quad (15)$$

Unlike in a representative-agent economy, where the interest rate would be constant and equal to $\rho + g$, in our economy the interest rate is stochastic, since φ_t , η_t^l , and q_t are all random. This is due to the fact that even though aggregate consumption is deterministic, any given agent's consumption has random drift.

Having solved for the equilibrium interest rate, we can now substitute (15) into (10) to obtain the following important result.

Lemma 4 *The drift of q is given by*

$$\mu_{q,t} = (\beta + \varphi_t) q_t - \beta\alpha\varphi_t q_t^2 - 1. \quad (16)$$

Equation (16) is central for our purposes, since it encapsulates all the equilibrium requirements that our model places on the drift of the price-dividend ratio. Indeed, it is the main equation that we use to prove the main results of the paper.

3.2 Obtaining a process for q_t as an equilibrium outcome

In this section we ask whether, taking two functions f and σ as given, one can specify a Markovian diffusion for $\varphi_t = \eta_t^d - \eta_t^l$ such that the equilibrium process for q_t is given by

$$dq_t = f(q_t) dt + \sigma(q_t) dB_t. \quad (17)$$

We leave some technical details for the appendix, and present here the main argument, followed by an illustrative example and a general proposition.

Any process φ_t that supports (17) as an equilibrium must satisfy equation (16) with $\mu_{q,t} = f(q_t)$. We use this equation to define a function φ that relates φ_t to q_t :

$$\varphi_t = \varphi(q_t) = \frac{1 - \beta q_t + f(q_t)}{q_t(1 - \beta\alpha q_t)}. \quad (18)$$

We assume that φ thus defined is a strictly decreasing function, so that its inverse exists. (We note that simple differentiation of (18) shows that φ decreases for $q \leq \frac{1}{2} \frac{1}{\alpha\beta}$. Hence, setting $q^{\max} < \frac{1}{2} \frac{1}{\alpha\beta}$, or choosing a function f that has a sufficiently negative first derivative, will ensure that φ is strictly decreasing.)

The dynamics of the process φ_t are easily written as an application of Ito's Lemma.

Important, the argument q_t equals $\varphi^{-1}(\varphi_t)$, so that φ_t follows as Markov process:

$$d\varphi_t = \underbrace{\varphi'(\varphi^{-1}(\varphi_t)) \left(f(\varphi^{-1}(\varphi_t)) + \frac{1}{2} \sigma^2(\varphi^{-1}(\varphi_t)) \frac{\varphi''(\varphi^{-1}(\varphi_t))}{\varphi'(\varphi^{-1}(\varphi_t))} \right)}_{\equiv \mu_\varphi(\varphi_t)} dt + \underbrace{\varphi'(\varphi^{-1}(\varphi_t)) \sigma(\varphi^{-1}(\varphi_t))}_{\equiv \sigma_\varphi(\varphi_t)} dB_t. \quad (19)$$

Equation (19) provides the answer to the question that we posed at the outset. Specifically, if we started out with the primitive assumption that φ_t follows the Markov diffusion

$$d\varphi_t = \mu_\varphi(\varphi_t)dt + \sigma_\varphi(\varphi_t)dB_t, \quad (20)$$

with $\mu_\varphi(\varphi_t)$ and $\sigma_\varphi(\varphi_t)$ defined in equation (19), then — by construction — the equilibrium dynamics of the price-dividend ratio are given by (17).

Before stating a formal general result, we illustrate the above ideas with a concrete example that allows closed-form expressions for the functions φ and φ^{-1} .

Example 1 *Suppose that the targeted valuation ratio is given by $q_t = a_1 + a_2x_t$, where*

$$dx_t = (-v_1x_t + v_2(1 - x_t)) dt - \sigma_x \sqrt{x_t(1 - x_t)} dB_t \quad (21)$$

and a_1 , a_2 , v_1 , v_2 , and σ_x are positive constants. It is known (see, e.g., Karlin and Taylor (1981), p. 221) that x_t has support in $[0, 1]$ as long as v_1 and v_2 are larger than $\frac{\sigma_x^2}{2}$.

In this example, equation (18) becomes

$$\varphi(q_t) = \frac{1 - \beta q_t + a_2 v_2 - (v_1 + v_2)(q_t - a_1)}{q_t(1 - \beta \alpha q)} = \frac{1 - \beta(a_0 + a_1 x_t) + a_2 v_2 - a_2(v_1 + v_2)x_t}{(a_1 + a_2 x_t)(1 - \beta \alpha(a_1 + a_2 x_t))}.$$

The second equality above provides the dynamics of φ_t that support the price dividend ratio $q_t = a_1 + a_2x_t$ as an equilibrium function.

The price-dividend ratio (i.e., the inverse function $q_t = \varphi^{-1}(\varphi_t)$) can be computed from

the first equality as

$$\varphi^{-1}(\varphi_t) = \frac{1}{2\beta\alpha} \left(\frac{\varphi_t + \beta + v_1 + v_2}{\varphi_t} - \sqrt{\left(\frac{\varphi_t + \beta + v_1 + v_2}{\varphi_t} \right)^2 - \frac{4\beta\alpha}{\varphi_t} (1 + a_2v_2 + \alpha_1(v_1 + v_2))} \right).$$

Using the above expression for $\varphi^{-1}(\varphi_t)$ inside (19) allows us to express φ_t as a (Markovian) stochastic differential equation for φ_t that supports $q_t = a_1 + a_2x_t$ as an equilibrium process for the price-dividend ratio.

The following proposition provides the general result.

Proposition 1 *Suppose that technical Assumption 1 in the appendix is satisfied, and that $\varphi(q)$ in equation (18) is decreasing. Then the equilibrium stochastic process for q_t is given by (17) if, and only if, φ_t follows the (Markovian) dynamics (19). Moreover, q_t is stationary and takes values in $[q^{\min}, q^{\max}]$.*

Proposition 1 states that one can support a wide range of diffusions for q_t as an equilibrium outcome, even though aggregate consumption and dividends are both deterministic. A technical assumption is offered in the appendix to ensure a stationary distribution for q_t supported by $[q^{\min}, q^{\max}]$.

We conclude with two remarks. First, the process φ_t that supports a given equilibrium stochastic process for q_t is unique. Second, the process q_t only determines $\varphi_t = \eta_t^d - \eta_t^l$. The individual processes η_t^d and η_t^l can be chosen freely as long as their difference obeys the dynamics (19). For instance, one choice is to set $\eta_t^l = \eta^l = \varphi(q^{\max})$ and $\eta_t^d = \eta^l + \varphi_t$.

3.3 Recursive preferences and risk premiums

The central purpose of the previous section is to illustrate our approach to reverse engineering the redistribution processes η^d and η^l to obtained desired asset-pricing implications. In the context of the expected-utility model, these implications are subject to an important limitation: Any agent's consumption is locally deterministic and so is their marginal utility, and therefore the market price of risk in this economy is zero.

Since the market price of risk is an equilibrium quantity of central importance to the determination of asset prices, in this section we allow for recursive preferences and show how to support any given dynamics for the price-dividend ratio and the market price of risk jointly. The construction of the appropriate processes η_t^d and η_t^l is conceptually quite similar to the construction in the previous section. Hence, in order to avoid repetition, we simply state the main results and refer the reader to the appendix for the derivations.

We start by introducing the more general preferences needed for a non-zero risk premium. Specifically, we continue to assume that investors have unit intertemporal elasticity of substitution, but allow for a risk aversion higher than one. In mathematical terms, taking into account the consumer's death probability, her instantaneous utility flow is given by the aggregator

$$f(c_{t,s}, V_{t,s}) = \beta\gamma V_{t,s} (\log(c_{t,s}) - \gamma^{-1} \log(\gamma V_{t,s})), \quad (22)$$

in that

$$V_{t,s} = E_t \left[\int_t^\infty f(c_{u,s}, V_{u,s}) du \right]. \quad (23)$$

Here, $V_{t,s}$ is a consumer's value function and $\gamma < 0$ is a parameter that controls the risk aversion of the investor. Utilities of this form are introduced and discussed extensively in Duffie and Epstein (1992). They correspond to the continuous-time limit of Epstein-Zin-Weil utilities with unit elasticity of substitution.

Since markets are dynamically complete for existing agents, the hazard rate of death is constant, and their preferences continue to be homothetic, it follows that $\frac{\dot{c}_{t,s}}{c_{t,s}}$ is independent of the cohort s to which the consumer belongs. Accordingly, equation (12) continues to hold and so do Lemmas 1 and 2, which follow from agents' budget constraints. Therefore, combining (12) with Lemma 2 yields

$$\frac{\dot{c}_{t,s}}{c_{t,s}} = \lambda + g - \nu_t \quad (24)$$

where ν_t , the fraction of consumption accruing to new agents, is defined as

$$\nu_t \equiv (1 - \alpha\beta q_t)\eta_t^l + \alpha\beta q_t\eta_t^d = \eta_t^l + \alpha\beta q_t\varphi_t, \quad (25)$$

with φ_t defined in Lemma 2. By Lemma 1, it follows that $\alpha\beta q_t \in [0, 1]$ and hence ν_t is a convex combination of η_t^l and η_t^d . Since $\frac{\dot{c}_{t,s}}{c_{t,s}}$ is independent of s , we shall henceforth write $\frac{\dot{c}_t}{c_t}$.

The only object that changes when agents have recursive preferences is the stochastic discount factor, described by the following result.

Lemma 5 *Define the process Z_t as the solution to the backward stochastic differential equation*

$$Z_t \equiv E_t \int_t^\infty e^{-\beta(u-t)} \left(\gamma \left(\frac{\dot{c}_u}{c_u} \right) + \frac{1}{2} \sigma_{Z,u}^2 \right) du, \quad (26)$$

where $\sigma_{Z,t}$ is the volatility of Z_t . Then the stochastic discount factor evolves according to

$$\frac{dm_t}{m_t} = -r_t dt - \kappa_t dB_t,$$

where r_t , the interest rate in this economy, continues to be given by equation (15), while κ_t is the market price of risk in this economy and is given by $\kappa_t = -\sigma_{Z,t}$.

Recursive preferences imply a risk premium. Intuitively, a risk premium arises because investors worry not only about the immediate impact of redistribution risks, but also about the long run impact of these risks on their consumption. This long run impact is captured by the definition of Z_t and the magnitude of the market price of risk (or Sharpe ratio) κ_t reflects the volatility of Z_t .

We next ask a question similar to the one we asked in the previous subsection. Is it possible to choose diffusion processes for η_t^l and η_t^d to support given stock-market dynamics (q) and given dynamics of the Sharpe ratio (κ)?

To provide an answer to this question, we proceed as in the previous section. Specifically, we fix functions f_Z , σ_Z , f_q , and σ_q and intervals $[Z^{\min}, Z^{\max}]$ and $[q^{\min}, q^{\max}]$ and try to determine a (vector) Markov process (η_t^l, η_t^d) such that the equilibrium process Z_t — to target

a particular Sharpe ratio κ , all we need is that the process Z_t have volatility $\sigma_Z(Z_t) = -\kappa_t$ — has support in $[Z^{\min}, Z^{\max}]$ and follows the dynamics

$$dZ_t = f_Z(Z_t) + \sigma_Z(Z_t) dB_t, \quad (27)$$

while the process for q_t has support in $[q^{\min}, q^{\max}]$ and follows the dynamics

$$dq_t = f_q(q_t) + \sigma_q(q_t) dB_t. \quad (28)$$

As we show in the appendix, the equilibrium dynamics of Z_t and q_t obey equations (27) and (28) when and only when the functions f_Z and f_q satisfy the relations

$$f_Z(Z_t) = \beta Z_t + \gamma \nu_t - \frac{1}{2} \sigma_Z^2(Z_t) - \gamma(\lambda + g) \quad (29)$$

$$f_q(q_t) = (\beta + \varphi_t) q_t - \beta \alpha \varphi_t q_t^2 - 1 - \sigma_Z(Z_t) \sigma_q(q_t). \quad (30)$$

Comparing the right-hand sides of (30) and (16) shows that the two expressions are identical, except for the last term in equation (30) which captures the presence of an equity premium.

Solving for ν from equation (29) and for φ from equation (30) leads to

$$\nu(Z) = \frac{1}{\gamma} \left(f_Z(Z) + \frac{1}{2} \sigma_Z^2(Z) - \beta Z \right) + \lambda + g, \quad (31)$$

$$\varphi(q, Z) = \frac{1 - \beta q + f_q(q) + \sigma_Z(Z) \sigma_q(q)}{q(1 - \beta \alpha q)}. \quad (32)$$

Once again, we wish to be able to invert this mapping, which we can do under the conditions given in the following lemma.

Lemma 6 *Suppose that $\frac{d\nu}{dZ} > 0$ for all $Z \in [Z^{\min}, Z^{\max}]$ and also $\frac{\partial \varphi}{\partial q} < 0$ for any $Z \in [Z^{\min}, Z^{\max}]$ and $q \in [q^{\min}, q^{\max}]$. Then the mapping (31)–(32) is invertible.*

Given invertibility, we obtain, from Ito's Lemma, two (jointly Markov) diffusion processes for ν_t and φ_t that support (29) and (30) as equilibrium outcomes. The values of η_t^d and η_t^l

follow easily as solutions to the linear 2×2 system $\varphi_t = \eta_t^d - \eta_t^l$ and equation (25):¹⁰

$$\eta_t^d = \nu_t + (1 - \alpha\beta q_t) \varphi_t \quad (33)$$

$$\eta_t^l = \nu_t - \alpha\beta q_t \varphi_t. \quad (34)$$

We record the formal result:

Proposition 2 *Consider intervals $[q^{\min}, q^{\max}] \subset \left(0, \frac{1}{\alpha\beta}\right)$ and $[Z^{\min}, Z^{\max}]$ and continuous functions f_Z, f_q, σ_Z , and σ_q such that the assumptions of Lemma 6 hold. Then if (and only if) ν_t follows the (Markovian) dynamics (79) in the Appendix and the analogous ones for φ_t , the equilibrium stochastic process for Z_t and q_t are given by the diffusions (27) with support $[Z^{\min}, Z^{\max}]$ and (28) with support $[q^{\min}, q^{\max}]$.*

The main goal of Proposition 2 is to provide an explicit mapping between assumptions on the share processes η_t^d and η_t^l and the resulting diffusion processes for the Sharpe ratio and the price-to-dividend ratio. The restrictions placed on these latter two processes by the assumptions of Lemma 6 are rather mild and in practical applications amount to simple parametric restrictions, as the next example illustrates.

Example 2 *Suppose that x_t follows the process (21) and that we wish to obtain $Z_t = b_1 + b_2 x_t$ and $q_t = a_1 + a_2 x_t$ as equilibrium outcomes with $b_1 = \frac{\gamma}{\beta}(\lambda + g)$ and some constants $a_1 > 0$, $a_2 > 0$, and $b_2 < 0$.*

In that case equation (31) implies

$$\nu_t = \frac{1}{\gamma} \left(b_2 v_2 - (v_1 + v_2) (Z_t - b_1) + \frac{\sigma_x^2}{2} (Z_t - b_1) (b_2 + b_1 - Z_t) - \beta (Z_t - b_1) \right) \quad (35)$$

$$= \frac{b_2}{\gamma} \left(v_2 - (v_1 + v_2) x_t + \frac{b_2 \sigma_x^2}{2} x_t (1 - x_t) - \beta x_t \right). \quad (36)$$

¹⁰A possible issue with (33) and (34) is that η_t^d and η_t^l could be negative. In that case redefine ν_t as $\nu_t + L$, where the constant L is large enough to ensure that the resulting values for η_t^d and η_t^l are both non-negative. (Such a value always exists, because ν_t and φ_t are continuous functions of Z_t, q_t , which take values on a bounded set, so that η_t^d and η_t^l are bounded below). Since the volatility of Z_t is unaffected, simply increasing η_t^d and η_t^l by L supports an equilibrium with the same dynamics for q_t and the Sharpe ratio.

We require

$$v_1 + v_2 + \frac{b_2 \sigma_x^2}{2} + \beta > 0$$

so that the right hand side of (35) decreases in Z_t for all $Z_t \geq b_1 + b_2$, so that the inverse function $Z_t = \nu^{-1}(\nu_t)$ exists, and ν_t can be written as a (Markovian) stochastic differential equation.¹¹ With this specification the Sharpe ratio is given by $|b_2| \sqrt{x_t(1-x_t)}$.

The dynamics of φ_t that generate $q_t = a_1 + a_2 x_t$ follow from equation (32), namely

$$\varphi_t = \frac{1 - \beta q_t + a_2 v_2 - (v_1 + v_2)(q_t - a_1) + \frac{b_2}{a_2} \sigma_x^2 (q_t - a_1)(a_1 + a_2 - q_t)}{q_t(1 - \beta \alpha q_t)} \quad (37)$$

$$= \frac{1 - \beta(a_1 + a_2 x_t) + a_2 v_2 - a_2(v_1 + v_2)x_t + a_2 b_2 \sigma_x^2 x_t(1-x_t)}{(a_1 + a_2 x_t)(1 - \beta \alpha(a_1 + a_2 x_t))} \quad (38)$$

We assume that parameters are such that the right-hand side of (37) is decreasing in q_t , so that the relation between φ_t and q_t is invertible. Accordingly, φ_t can be written as a (Markovian) stochastic differential equation.¹² We note here that $a_2 b_2 < 0$ means that the Sharpe ratio is negatively related to the price-dividend ratio.

As a final remark, we note that we have assumed throughout that Z_t and q_t^d (and by implication η_t^d and η_t^l) are driven by the same Brownian motion. It is straightforward to extend the analysis to allow Z_t and q_t^d to be driven by separate Brownian motions with some correlation coefficient ρ .¹³

¹¹Its inverse is given by

$$Z_t = \nu^{-1}(\nu_t) = b_1 + \frac{\frac{\sigma_x^2}{2} b_2 - v_1 - v_2 - \beta + \sqrt{\left(v_1 + v_2 + \beta - \frac{\sigma_x^2}{2} b_2\right)^2 + 2\sigma_x^2 \left(b_2 v_2 - \frac{\gamma}{\beta} \nu_t\right)}}{\sigma_x^2}.$$

Using the same steps as the ones we used to arrive at (19), one can derive the stochastic differential equation for $d\nu_t$.

¹²A sufficient condition is that $v_1 + v_2 + b_2 \sigma_x^2 + \beta > 0$ and $a_1 + a_2 < \frac{1}{2\alpha\beta}$.

¹³The only modification required in that case is that the term $\sigma_Z \sigma_q$ in equations (30) and (32) needs to be changed to $\rho \sigma_Z \sigma_q$. As a result, equation (37) in Example 2 would feature Z on the right-hand side, as well, so that q_t follows as a function of both φ and Z , and therefore φ and ν .

3.4 Relation to the literature

Proposition 2 is a “possibility” result, similar to the one provided in Constantinides and Duffie (1996). It shows that the model is –at least in principle– able to produce a wide range of dynamics for the price-dividend and the Sharpe ratio despite constant consumption and dividend growth.

It is an empirical matter to establish whether the share processes η_t^l and η_t^d that are required to reproduce the empirically observed asset pricing moments are empirically plausible or not. We address this question in the next section.

Before proceeding, we conclude this section with two observations about how this model differs from Constantinides and Duffie (1996). One obvious difference is that we do not require independent innovations to the stochastic discount factor, and show how to accommodate instead a Markovian structure.

The more important difference between the two models (and indeed to many other models of heterogeneous agents) pertains to the dynamic behavior of inequality. To see this, it is useful to define the cross-sectional variance of log consumption as

$$G_t = \lambda \int_{-\infty}^t e^{-\lambda(t-s)} \left(\log(c_{t,s}) - \lambda \int_{-\infty}^t e^{-\lambda(t-u)} \log(c_{t,u}) du \right)^2 ds.$$

Time-differentiating G_t we obtain the following dynamics

$$dG_t = -\lambda G_t dt + \lambda \left(\log(c_{t,t}) - \lambda \int_{-\infty}^t e^{-\lambda(t-u)} \log(c_{t,u}) du \right)^2 dt \quad (39)$$

An immediate implication of the above equation is that G_t is a locally deterministic process, i.e. has no diffusion component. Accordingly, when integrated over short periods of time (say a quarter or a year), this process will appear substantially less volatile than asset returns and will exhibit an essentially zero correlation with asset returns. Inspection of the first term on the right hand side of (39) shows that the process G_t is quite persistent, since it mean-reverts at the rate λ , the rate of population death. By contrast, the mean reversion of the price-to-dividend ratio need not be λ , that is the price-dividend ratio may exhibit faster

mean reversion than the cross-sectional variance of log consumption. This is an attractive feature of this model, since measures of inequality in the data tend to exhibit small year-over-year changes, but near unit root persistence. We also note that this dynamic behavior of inequality is qualitatively different from other models of heterogeneous agents as well (such as Gârleanu and Panageas, 2015): In these models there is strong contemporaneous correlation between asset returns and cross sectional inequality, which is absent here.

4 Empirical Implications

The discussion so far was focused on the theoretical possibility of supporting some given dynamics for the price-to-dividend ratio and the market price of risk as equilibrium outcomes. In the process, we derived a practical method for determining the functional forms for the dynamics of the share processes that result in closed form dynamics for the price-dividend ratio and the Sharpe ratio.

In this section we show how to measure these share processes in the data and illustrate some basic qualitative implications of the model. Specifically in section 4.1 we show how to measure η_t^d using stock market data and in section 4.2 we show how to measure ν_t (which is a convex combination of η_t^d and η_t^l) using cross-sectional consumption data.

In section 5 we calibrate the dynamics of the processes η_t^d and ν_t so as to match their empirical counterparts and examine the resulting asset price dynamics.

4.1 The measurement of η_t^d

A straightforward empirical proxy for η_t^d implied by the model is the ratio of the market value of additions to the market index to the total market value of the index:

$$\eta_t^d = \frac{D_{t,t}}{D_t} = \frac{\overbrace{P_{t,t}}^{\text{Market cap of new additions}}}{\underbrace{P_t}_{\text{Total market cap}}} . \quad (40)$$




Figure 2: Total real logarithm of S&P 500 dividends per share, real log-aggregate consumption and real log-aggregate dividends. The CPI is used as a deflator for all series. The line “Index Dividends + New Cap” is equal to real log-dividends per share plus the cumulative (log) gross growth in the shares of the index that are due to the addition of new firms. Sources: R. Shiller’s website, FRED, Personal Dividend Income series, and CRSPSift.

This is the measure that we use for our calibration exercise. (We provide further details in section 5).

As a parenthetical remark, we find it interesting to illustrate an implication of the model in the data. Specifically, we note that according to the model η_t^d is also the discrepancy between aggregate dividend growth and the dividend growth rate of the market portfolio:

$$\eta_t^d = \underbrace{\frac{\dot{D}_t^A}{D_t^A}}_{\text{Aggregate dividend growth}} - \underbrace{\int_{-\infty}^t \pi_s \frac{\dot{D}_{t,s}}{D_{t,s}} ds}_{\text{Dividend growth of a self-financing strategy}}, \quad (41)$$

where π_s is the set of market weights (or any weights integrating to one for that matter).

Figure 2 illustrates equation (41) in the data. The solid line in the figure depicts the (log) dividends of the S&P 500 since its inception. The figure also depicts the National Income and Product Account (NIPA) aggregate consumption series and NIPA aggregate dividend series.

The figure suggests that aggregate dividends and aggregate consumption are co-trending, while the index dividends appear follow a markedly slower growth path. If, however, we add the percentage increase of the market capitalization of the index that is due to additions every year (the right hand side of equation (40)) to the series labeled “index” dividends, we obtain a series that is very close to the aggregate NIPA dividend series, consistent with the implication of equations (40) and (41).¹⁴

4.2 Cohort effects and the measurement of ν_t

The origin of a positive market price of risk in this model is that a given individual “marginal agent’s” consumption growth ($\frac{\dot{c}_{t,s}}{c_{t,s}}$) differs from aggregate (per capita) consumption growth.¹⁵ To summarize, the main difference between the two concepts is that the first concept follows the consumption growth of a given cohort, whereas the second concept aggregates everyone’s consumption at each point in time and follows the consumption growth of that aggregate. In this subsection we wish to i) provide a way to measure the consumption growth of the marginal agent and examine its statistical properties, ii) show that the discrepancy between marginal and per capita consumption growth is largely independent from fluctuations in aggregate per capita consumption growth, and iii) to provide support for the relation between the behaviour of the real (expected) interest rate and marginal agent consumption growth.

We start with some notation. Specifically, we define aggregate consumption as

$$C_t^A = \int_{-\infty}^t \Lambda(t, s) c_{t,s} ds, \quad (42)$$

where $\Lambda(t, s)$ is the surviving time- t population measure that was born at time s . In the model $\Lambda(t, s) = \lambda e^{-\lambda(t-s)}$ in order to simplify the model. However, for measurement purposes we will allow a general $\Lambda(t, s)$.

Letting $\omega_{t,s} = \frac{\Lambda(t,s)c_{t,s}}{C_t^A}$ denote the time- t consumption share of cohort s , time-differentiating

¹⁴Some discrepancies arise because aggregate dividends is based on NIPA data, which include listed and non-listed companies.

¹⁵Inside the model aggregate and per capita consumption growth are identical, since population is constant.

(42), and re-arranging gives

$$\int_{-\infty}^t \omega(t, s) \frac{dc_{t,s}}{c_{t,s}} ds = \frac{dC_t^A}{C_t^A} - \int_{-\infty}^t \omega(t, s) \frac{d\Lambda(t, s)}{\Lambda(t, s)} ds - \omega(t, t) dt. \quad (43)$$

We note that for measurement purposes we will allow aggregate consumption (and marginal agent consumption) to follow a general stochastic process, which may not be time-differentiable (such as a diffusion). (This is reflected in the notation dC_t^A rather than $\dot{C}_t^A dt$ and $dc_{t,s}$ rather than $\dot{c}_{t,s} dt$.)

As mentioned above, the main insights of the paper rely on the fact that that marginal-agent consumption growth $\int_{-\infty}^t \omega(t, s) \frac{dc_{t,s}}{c_{t,s}} ds$ differs from aggregate consumption growth $\frac{dC_t^A}{C_t^A}$. In the remainder of this section, we wish to use only a minimal part of the model structure in order to measure this discrepancy in the data. The main challenge with this measurement is that while there is long time-series (NIPA) data on aggregate consumption growth $\frac{dC_t^A}{C_t^A}$ and (life-table) data on $\Lambda(t, s)$, there is no long time series data on $c_{t,s}$. (Cross sectional consumption data starts from the eighties onwards).

In order to overcome this problem we use the structure of the model. The Euler equation $\frac{\dot{c}_{t,s}}{c_{t,s}} = r_t - \rho$ implies that consumption can be written in the form

$$\log c_{t,s} = \underbrace{A_s}_{\text{“cohort effect”}} + \underbrace{L_t}_{\text{“time effect”}} + \underbrace{G_{t-s}}_{\text{“age effect”}}. \quad (44)$$

Inside the model, these components equal to $L_t = \int^t r_u du$, $A_s = \log c_{s,s} - L_s$, $G_{t-s} = -\rho(t-s)$. However, a time-, age-, and cohort- decomposition of the form (44) is not special to this model. It would readily extend to a wide class of models that imply risk sharing across the cohorts of agents that are marginal at each point in time.

Equation (44) makes it possible to use exclusively cross sectional data to reconstruct a path of the discrepancy between “marginal-agent consumption growth” $\int_{-\infty}^t \omega(t, s) \frac{dc_{t,s}}{c_{t,s}} ds$ and aggregate consumption growth. The following Lemma shows how:

Lemma 7 *Assume that there exist processes A_s, L_t and a function of age G_{t-s} so that $\log c_{t,s}$*

is given by (44). Define

$$F_t \equiv \int_{-\infty}^t \Lambda(t, s) e^{A_s + G_{t-s}} ds. \quad (45)$$

Then we have the following relations

$$\int_{-\infty}^t \omega(t, s) \frac{dc_{t,s}}{c_{t,s}} ds - \int_{-\infty}^t \omega(t, s) dG_{t,s} = dL_t = \frac{dC_t^A}{C_t^A} - \frac{dF_t}{F_t} \quad (46)$$

Corollary 1 If $dG_{t-s} = -\rho dt$ then

$$\int_{-\infty}^t \omega(t, s) \frac{dc_{t,s}}{c_{t,s}} ds = dL_t - \rho dt$$

Lemma 7 suggests an indirect approach of inferring the time effects L_t , which capture the common, age independent variation in the consumption growth of all marginal agents: Fix two times t and $t-1$ (say this year and last year) and suppose that the econometrician has information about the consumption of cohorts born at times $s = t, t-1, ..t-T, ..etc.$ At first pass, it would seem impossible to measure a time series of the time-effects L_t ; it would seem that for such a task the econometrician would need to have as many cross sections as the time dummies that she wishes to estimate. Lemma 7 circumvents this problem by using estimates on a sequence of cohort dummies A_s and age effects G_{t-s} , $s = t, t-1, ..t-T, ..etc.$ (which can be estimated with as many as two cross sections) to infer a whole time series path of L_t . Specifically, with estimates of the cohort and age dummies, one can easily compute a time series of $F_t, F_{t-1}, F_{t-2}, \dots$ by using the discrete time approximation

$$\begin{aligned} F_t &= \sum_{s=-\infty \dots t} \Lambda(t, s) e^{A_s + G_{t-s}}, \\ F_{t-1} &= \sum_{s=-\infty \dots t-1} \Lambda(t-1, s) e^{A_s + G_{t-1-s}}, \dots, etc. \end{aligned}$$

Assuming that the econometrician also has access to a time series of aggregate consumption data, it becomes now possible to utilize the second equality in (46) to reconstruct the path of L_t .

We conclude with two remarks.

First, we note that the quantity $\frac{dF_t}{F_t}$ can only be identified up to an additive constant. The reason is that time-, age-, and cohort- effects can only be identified up to a linear term in the data, i.e., the data do not allow one to separate the model of equation (44) from the alternative model

$$\log c_{t,s} = \underbrace{A_s + \chi s}_{\text{modified cohort effect}} + \underbrace{L_t - \chi t}_{\text{modified time effect}} + \underbrace{G_{t-s} + \chi(t-s)}_{\text{modified age effect}},$$

where χ is an arbitrary constant. Because of this reason $\frac{dF_t}{F_t}$ in equation (83) and dL_t can only be identified up to an additive constant.¹⁶ This is not a major obstacle for our purposes, since we are interesting in the time-variation rather than the level of dL_t .

Second, we note that inside our model the variation in the time effects dL_t would coincide with marginal agent consumption growth $\int_{-\infty}^t \omega(t,s) \frac{dc_{t,s}}{c_{t,s}} ds$, up to an additive constant (see corollary 1). When we estimate the model in the data, however, we do not impose that age effects are simply given by a linear trend; we estimate them without restriction. Lemma 7 shows how to infer dL_t for arbitrary estimates of age effects. One issue that arises when age effects are not simply a linear function of age is that each cohort will experience both a (stochastic) common variation in its consumption but also a (deterministic) age variation. Hence forth we will focus exclusively on dL_t as our measure of marginal agent consumption growth, since in our model variations in the investment opportunity set are related to the common, stochastic changes in consumption growth.¹⁷ We note in passing that for the results

¹⁶As we show in the appendix $\frac{dF_t}{F_t} = \int_{-\infty}^t \omega_{t,s} \left(\frac{d\Lambda(t,s)}{\Lambda(t,s)} + dG_{t-s} \right) ds + \omega_{t,t} dt$, and since dG_{t-s} is available only up to an additive constant, $\frac{dF_t}{F_t}$ is only available up to an additive constant.

¹⁷If we were to extend our model by introducing an age-specific subjective discount factor ρ_{t-s} we would simply modify the Euler relation to $\frac{\dot{c}_{t,s}}{c_{t,s}} = r - \rho_{t-s}$. Upon integrating we obtain the following time-, age-, and cohort- decomposition:

$$\log c_{t,s} = \underbrace{\log c_{s,s} - M_s}_{\text{Cohort Effect } A_s} + \underbrace{M_t}_{\text{Time effect } L_t} - \underbrace{\int_0^{t-s} \rho_u du}_{\text{Age Effect } G_{t-s}}, \quad (47)$$

where $M_t = \int^t r_u du$. The introduction of age effects does not change the relation between L_t and the (integrated) interest rate M_t .

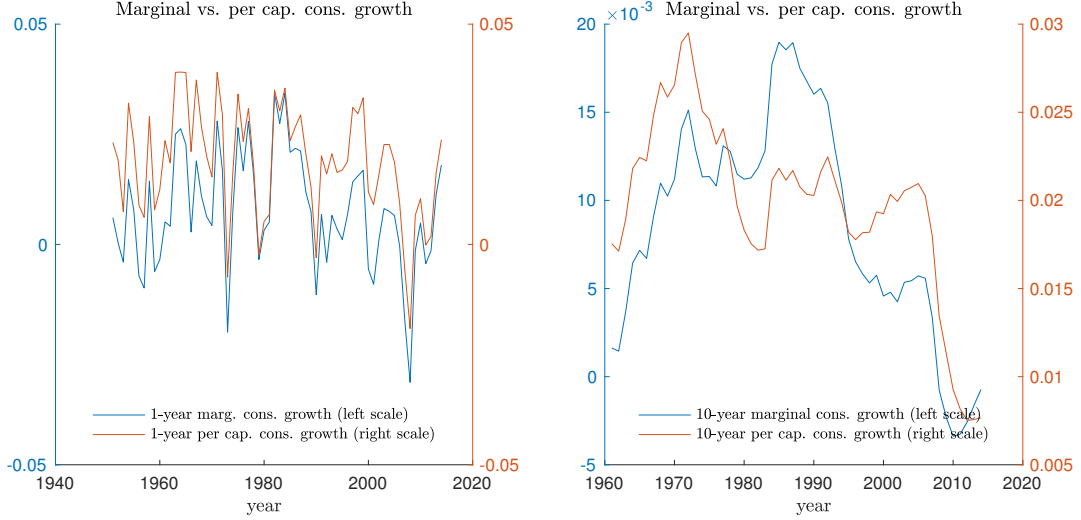


Figure 3: Left plot: Yearly per capita consumption growth of the marginal agent $\Delta L_t = \Delta \log C_t^A - \Delta \log F_t$ (left scale) and yearly per capita consumption growth $\Delta \log C_t^A - \Delta \log N_t$, (right scale). Right plot: Same plot, but for 10-year moving averages of consumption growth.

we present next, it makes essentially no difference whether we add $\int_{-\infty}^t \omega(t, s) dG_{t,s}$ to dL_t or not.

We next describe briefly the data and how we implement our measurement procedure. We relegate a more detailed description of the data, the sample choice, and the estimation procedure to the appendix.

To compute dL_t from equation (46) we need data on aggregate consumption growth $\frac{dC_t^A}{C_t^A}$, population of different cohorts at different points in time $\Lambda(t, s)$, and finally cross sectional consumption data to estimate the cohort effects A_s and the age effects G_{t-s} .

Our measure of aggregate consumption growth is from the National Income and Product Accounts (NIPA). We use the consumption of services and non-durables, deflated by the respective deflators. We use annual data (since 1929), as our cohort and age effects can only be measured at annual frequency. For most of our illustrative results we focus on the post-1950 part of the sample, simply because consumption data tend to be less volatile post world war II. However, for the econometric results of table 1, we use the full sample.

The demographic life table data $\Lambda(t, s)$, is from the census. For the cross sectional consumption data we use the Consumer Expenditure survey (CEX). We used the data processed

and compiled by John Sabelhaus (as available on the NBER website) from 1980 to 2003 and then extended that data to the latest available cross section (2016). As is common in the literature, we included only households that completed all the four surveys. We chose to include incomplete income respondents. (However, whether we include or drop incomplete income respondents does not matter for the results we report). For our measure of age we used the age of the reference person at the initial CEX interview, and defined “cohort s ” as the set of people who were 20 years old at time s . For our expenditure measure, we made similar choices as Sabelhaus. To cross-validate our consumption measures, we used the time period from 1996 to 2003, where the Sabelhaus dataset and our data overlap. The estimated cohort and age effects we obtained for either sample were essentially the same. The appendix contains further details on the computations.

Figure 3 plots the estimated consumption growth of marginal agents $\Delta L_t = \Delta \log C_t^A - \Delta \log F_t$, where Δ is the first difference operator at annual frequency. To compare, the figure also plots per capita consumption growth $\Delta \log C_t^A - \Delta \log N_t$, where N_t is the US population at time t . Clearly, the two series look quite similar at annual frequencies, since in the short run the movements in the two series are dominated by their common component, namely aggregate consumption growth $\Delta \log C_t^A$. However, the two time series look noticeably different when we time aggregate them over 10-year-long moving average growth rates.

To understand these time series properties, Figure 4 plots the difference between marginal and per-capita consumption growth $\Delta \log N_t - \Delta \log F_t$, which corresponds to $-\nu_t$ in our model. It is evident that this time series has small year-over-year volatility, but is quite persistent. This is in contrast to the common component of the two time series (aggregate consumption growth $\Delta \log C_t^A$), which is comparatively more volatile on a year-over-year basis, but far less persistent. By averaging over longer periods, the less persistent aggregate consumption growth fades in importance compared to the less volatile, but more persistent difference between the two series $\Delta \log N_t - \Delta \log F_t$.

Table 1 provides a formal econometric framework to model the joint time series properties of i) per capita consumption growth and ii) the difference between marginal and per capita

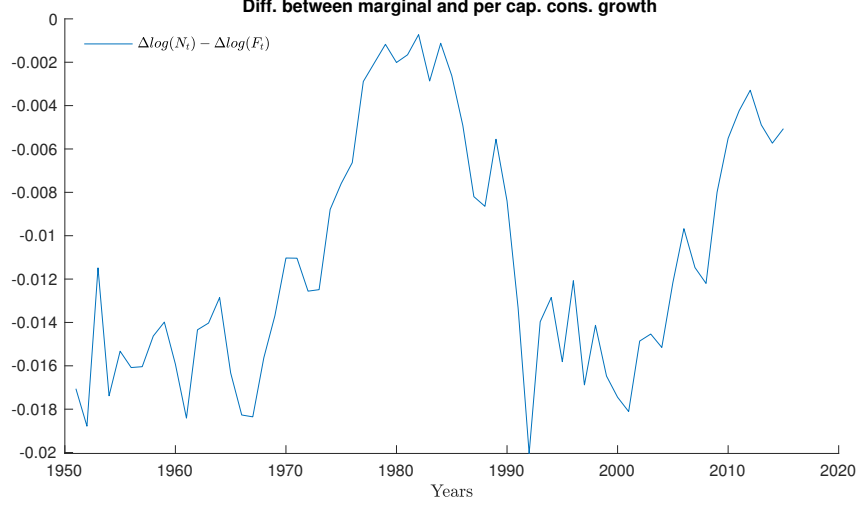


Figure 4: Difference between marginal agent consumption growth and per capita consumption growth.

	$\Delta \log C_t^A - \Delta \log N_t$	$\Delta \log N_t - \Delta \log F_t$
Lag $\Delta \log C_t^A - \Delta \log N_t$	0.4449 (0.0905)	-0.0207 (0.2867)
Lag $\Delta \log N_t - \Delta \log F_t$	0.0147 (0.0152)	0.8834 (0.0480)
R^2	0.22	0.80
σ_ϵ	0.0182	0.0030
N(obs)	87	87

Table 1: Bivariate Vector Autoregression of i) per capita consumption growth ($\Delta \log C_t^A - \Delta \log N_t$) and ii) the difference between marginal and per capita consumption growth ($\Delta \log N_t - \Delta \log F_t$).

consumption growth rate as a bivariate first order vector autoregression

$$\begin{bmatrix} \Delta \log C_t^A - \Delta \log N_t \\ \Delta \log N_t - \Delta \log F_t \end{bmatrix} = \begin{bmatrix} \Delta \log C_{t-1}^A - \Delta \log N_{t-1} \\ \Delta \log N_{t-1} - \Delta \log F_{t-1} \end{bmatrix} B + \begin{bmatrix} \epsilon_t \\ u_t \end{bmatrix}.$$

Using the obtained estimates for B and the covariance matrix of the residuals Σ , we compute the long run covariance matrix of the two time series

$$\Omega = [I + B + B^2 + \dots] \Sigma [I + B + B^2 + \dots]' = \frac{1}{100} \times \begin{bmatrix} 0.1082 & -0.0077 \\ -0.0077 & 0.0646 \end{bmatrix}$$




Figure 5: Short term interest rate, inflation, ex-ante inflation expectations. Sources: Interest rates: Robert E. Shiller, Online long term data on stock and bond returns, CPI inflation: Bureau of Economic Analysis, Inflation expectations: Philadelphia FED.

There are two conclusions we wish to draw from the equation immediately above. First, the fact that the off diagonal elements of Ω are essentially zero implies that the two time series are practically uncorrelated. Hence, the re-distributional fluctuations in $\Delta \log N_t - \Delta \log F_t$ (which correspond to ν_t in our model) are distinct from long run consumption growth fluctuations in per capita consumption growth. Phrased simply, the redistributional risks that arise from imperfect risk sharing are a separate source of long run consumption uncertainty, and fluctuations in aggregate growth don't offset them, which is a feature of our model.

Second, adding together per capita consumption growth $\Delta \log C_t^A - \Delta \log N_t$ and $\Delta \log N_t - \Delta \log F_t$ to arrive at a marginal agent's consumption growth implies that the log run variance of consumption growth of the marginal agent $[1, 1] \Omega [1, 1]'$ is about 50% higher than the long run variance that would be implied by per capita consumption growth. (This number remains roughly unchanged if we only consider post 1950s data).

Having obtained a measure of $\nu_t = \Delta \log N_t - \Delta \log F_t$, and its time series properties,

we have the empirical targets for the calibration exercise that follows.

We conclude this section by presenting some evidence that is not directly useful for the calibration, but relates to a testable implications of the model, namely the connection between our measure of marginal agent consumption growth and the expected real interest rate.

To measure the expected real interest rate, we use the short term nominal interest rate from Robert Shiller’s online data set minus the (ex ante) expected inflation rate as formed in December of the preceding year (Source: Philadelphia FED Inflation expectations survey). Figure 5 plots the nominal interest rate, the ex ante inflation expectation for the respective year, and the realized CPI inflation for the year.

The top left plot of figure 6 plots the expected real interest rate and marginal agent consumption growth for the respective year. Clearly the two series differ, because in reality there are shocks to aggregate consumption that our model abstracts from. However, we would expect that if the real interest rate reflects the expected (rather than the realized) marginal agent consumption growth, then we should find that the co-movement between the two series rises as we time aggregate the two series over longer horizons. Indeed, this is what the top right plot of Figure 6 shows. 10-year moving averages of expected consumption growth and real expected interest rates exhibit very similar fluctuations (the correlation coefficient is around 0.9). The bottom two plots show that, by comparison, this co-movement is weaker for per capita consumption growth.

Figure 7 shows that the steady increase of the R^2 when regressing marginal agent consumption growth on the expected real interest rate is unlikely to be the result of randomness, by performing a bootstrap exercise enforcing the null hypothesis that the two series are uncorrelated.¹⁸

In the next section we evaluate the model’s quantitative ability to reproduce stylized asset-pricing facts.

¹⁸Specifically, we draw 10,000 random time series of marginal agent consumption growth with replacement, time-aggregate both the real interest rate and the marginal agent consumption growth for each sample and compute the R^2 for each of those 10,000 samples.

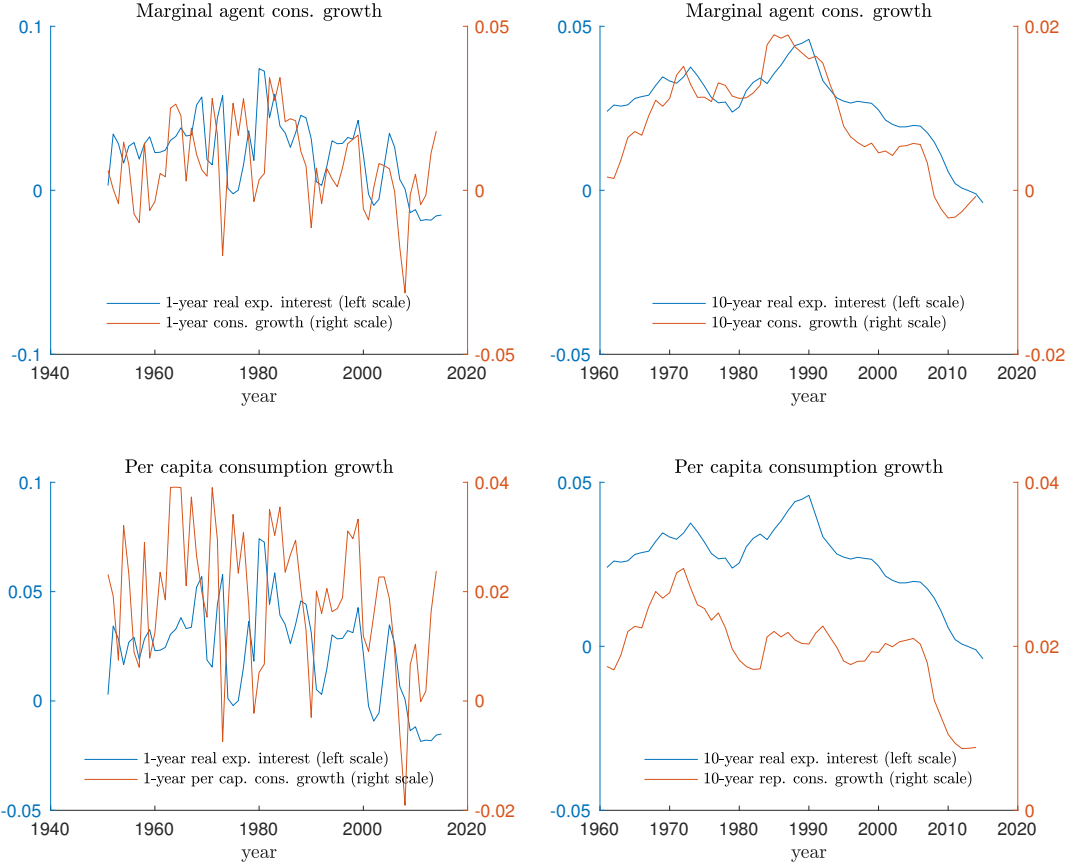


Figure 6: Top left plot: Expected real interest rate at the beginning of each year and marginal agent consumption growth over the year. Top right plot: 10-year moving averages yearly marginal agent consumption growth and 10-year moving average of real expected real rate. Bottom left and right plots: Identical to the top plots, except that marginal agent consumption growth is replaced with per capita consumption growth.

5 Calibration

The exercise we perform is straightforward. First, we choose functional forms for the dynamics of η_t^l and η_t^d . These functional forms are chosen judiciously to support closed-form solutions for the dynamics of the price-dividend ratio q_t^d and the Sharpe ratio κ_t . Second, we choose the parameters governing the dynamics of η_t^l and η_t^d to match the empirical moments of η_t^d and ν_t in the data. (Note that by equations (34) and (33) there is a one-to-one correspondence between the pair η_t^l, η_t^d and η_t^d, ν_t). Then we examine the resulting moments




Figure 7: Top left plot: R^2 of Regression of 1,2,3...,12 year moving average of marginal consumption growth on 1,2,3...,12 moving average of the real expected interest rate. The solid line refers to the data, the dotted line refers to the 95% confidence bands obtained by drawing 10,000 random time series of marginal agent consumption growth with replacement from the data, time-aggregating both the real interest rate and the marginal agent consumption growth for each sample, computing the R^2 for each of those 10,000 samples, and reporting the top 95-th percentile of R^2 .

of asset-price dynamics and compare them to the data.

Specifically, we employ a functional form specification similar to Example 2. Using the definition of x_t in equation (21), we specify ν_t as in equation (36). Then by construction, the equilibrium Sharpe ratio is given by $|b_2|\sqrt{x_t(1-x_t)}$. The only modification to example 2 is that we wish that the logarithm (rather than the level of q_t) be linear in x_t , so that $\log(q_t^d) = a_1 + a_2x_t$, and accordingly Ito's Lemma implies that the drift of q_t^d be given by

$$\frac{f_q(q_t^d)}{q_t^d} = v_1(a_1 - \log(q_t^d)) + v_2(a_2 - \log(q_t^d) + a_1) + \frac{\sigma_x^2}{2}(\log(q_t^d) - a_1)(a_2 - \log(q_t^d) + a_1). \quad (48)$$

Plugging this expression into (32) provides the dynamics for φ_t . The pair (φ_t, ν_t) is (joint) Markov for the parameters we choose in the calibration. These parameters are listed in Table 2.

We fix preference parameters to $\beta = 0.03$ (sum of discount- and death- rates) and $\gamma = -8$, which implies a risk aversion of $1 - \gamma = 9$. We set α to a level that reflects the share of capital income in output (0.3). The aggregate growth rate is set to 0.025, in line with historical

ϕ_1	0.012	β	0.03	γ	8
ϕ_2	0.078	α	0.3	g	0.025
σ_x	0.12	a_1	1.1	a_2	2.25
b_2	7				

Table 2: Parameters used in the model calibration.

	Data	Model
Median arrival rate of new firms	2.2%	2.44% (1.45%)
Standard deviation of the arrival rate of new firms	1.6%	2.00% (0.92%)
Autocorrelation of the arrival rate of new firms	0.77	0.89 (0.0532)
Median value of ν_t		3.08% (1.45%)
Standard deviation of of ν_t	0.55 %	0.58% (0.21%)
Autocorrelation of imputed ν_t	0.89	0.89 (0.0514)

Table 3: Targeted moments: Model and Data. We simulate 1000 paths of similar length as the data, and compute each of the six moments for every path. We then report the mean and standard deviation (across the 1000) paths for each moment. The term “arrival rate of new firms” refers to the ratio of the value of the market value of additions to the market portfolio to the total value of the market portfolio.

data. We choose the six parameters v_1 , v_2 , σ_x , b_2 , a_1 , and a_2 to (approximately) match six moments, namely the mean, standard deviation, and autocorrelation of the inferred values of ν_t and η_t^d in the data.

Table 3 shows that these parameter choices allow us to plausibly reproduce the targeted empirical moments within our model. To account for estimation error, we do not only report average values of the targeted moments within our model, but also the standard deviation for the model-implied values, when we simulate our model over similar sample lengths to the data. As can be seen, the moments in the data are within two standard deviations of their simulated means inside the model. Figure 8 provides an alternative, graphical illustration of Table 3 by comparing the empirical and the simulated distributions of ν_t and φ_t .

Having determined the parameters to match the moments of the redistribution processes,




Figure 8: Left plot: Histogram of (de-meaned) inferred values of ν_t and kernel density of the respective quantity inside the model. Right plot: Histogram of the market capitalization of new index additions over the existing market capitalization of the index (de-meaned) and kernel density of the respective quantity inside the model. To obtain the model-implied quantities, we simulate 1000 paths of length identical to the length of the data sample, and de-mean the simulated data separately on each sample path, to account for sampling error in the computation of the means. We then compare the distribution of the de-meaned simulated data to the de-meaned empirical data.

we next examine what these parameter choices imply for asset pricing moments. Table 4 provides a comparison between the model-implied unconditional moments and the respective moments in the data. In reporting the results we follow the approach of Barro (2006) to relate the results of our model (which produces implications for an all-equity financed firm) to the data (where equity is levered). Specifically, we use the well known Modigliani-Miller formula relating the returns of levered equity to those of unlevered equity, along with the historically observed debt-to-equity ratio, to report model-implied levered returns. (Specifically, we set the ratio of levered to unlevered equity returns to be equal to 1.7, as in Barro (2006).)

Inspection of Table 4 shows that the model accounts for a sizable fraction of all asset pricing moments. To put these numbers in the proper relation to the literature, it is worth highlighting that aggregate consumption and dividend growth are constant in this model.

	Data	Model
Aggregate consumption growth rate	2.3%	2.3%
Standard deviation of consumption growth rate	3.3	0
Sharpe ratio	0.29	0.26
Stock market volatility	18.2	14.32%
Equity premium	5.2%	4.13%
Average interest rate	2.8%	2.37%
Standard deviation of real interest rate	0.92%	0.72%
Average (log) PD ratio	2.9	3.05
Standard deviation of (log) PD ratio	0.27	0.21
Autocorrelation of (log) PD ratio	0.89	0.92

Table 4: Unconditional moments for the data and the model. The data for the average equity premium, the volatility of returns, and the level of the interest rate are from the long historical sample available from the website of R. Shiller (<http://www.econ.yale.edu/?shiller/data/chapt26.xls>). The volatility of the real rate is inferred from the yields of 5-year constant maturity TIPS.

The numbers should therefore be interpreted as the asset-pricing moments that would obtain in an economy where one abstracts from all aggregate sources of uncertainty and examines the impact of the redistribution processes in isolation.

Table 4 only pertains to unconditional moments. To evaluate the model's ability to account for variations in conditional moments we turn to Figure 9 and Table 5. Figure 9 plots the equity premium, market price of risk, interest rate, and stock-return volatility as a function of the log-price-earnings ratio. The Sharpe ratio and the equity premium are both declining functions of the log-price-earnings ratio. This counter-cyclicality is responsible for the model's ability to reproduce the predictability relations documented in Table 5. This table reports results of simulated predictability regressions inside the model and compares the results with the data. The main takeaway of the table is that the model-implied predictability is close to the respective time-variation in the data.

The model implies only modest annual volatility of the cross sectional variance of log consumption, and it does not require that the log price-dividend ratio and the cross-sectional variance of log consumption have similar persistence. Figure 10 helps visualize these points. The figure plots a simulated path of the cross sectional variance of log consumption over a sample that is similar in length to the post-war sample. It also plots the log price-dividend




Figure 9: Calibration results. Equity premium, market price of risk (Sharpe ratio), interest rate, and stock return volatility for the baseline parametrization. We plot each variable against the price-to-earnings ratio $\log(q_t^d)$. The range of values of $\log(q_t^d)$ correspond to the interval between the bottom 1% and the top 99% percentiles of the stationary distribution of $\log(q_t^d)$.

ratio (on the right scale). There are a few patterns that emerge from this graph, which are consistent with the data. The year-over-year changes in inequality (as measured by the cross-sectional variance of log consumption) are negligible compared to the time variation of the log-price dividend ratio. The two time-series operate at different frequencies, with inequality being substantially more persistent — essentially close to a random walk.

These patterns are very much consistent with the data. In the data, consumption inequality only changes by a few basis points on a yearly basis.¹⁹ Moreover, in the data inequality measures behave like unit root processes, consistent with Figure 10. These features distinguish this model from Constantinides and Duffie (1996), where the mechanism

¹⁹For instance, Krueger and Perri report that the cross sectional variance of log consumption changed from about 0.18 to about 0.23 from 1980 to 2003.

Year	β (Data)	β (Model)	R^2 (Data)	R^2 (Model)
1	-0.130	-0.092	0.040	0.038
		[-0.235 -0.003]		[0.000 0.134]
3	-0.350	-0.251	0.090	0.098
		[-0.573 -0.031]		[0.002 0.255]
5	-0.600	-0.376	0.180	0.144
		[-0.820 -0.089]		[0.011 0.350]
7	-0.750	-0.476	0.230	0.180
		[-0.893 -0.070]		[0.006 0.395]

Table 5: Long-horizon regressions of excess returns on the log P/D ratio. The simulated data are based on 1000 independent simulations of 100-year long samples. For each of these 100-year long simulated samples, we run predictive regressions of the form $\log R_{t \rightarrow t+h}^e = \alpha + \beta \log(P_t/D_t)$, where $\log R_{t \rightarrow t+h}^e$ denotes the time-t gross excess return over the next h years. We report the mean values for the coefficient β and the R^2 of these regressions, along with the respective [0.025, 0.975] percentiles.

becomes quantitatively powerful for large year-over-year variation in cross-sectional inequality. These features also distinguish this paper from models that rely on (preference, belief, or investment-opportunity) heterogeneity, which tend to imply that the price dividend ratio has the same persistence as consumption inequality. Both in the data and in our model, the autocorrelation coefficient for the price-dividend ratio is around 0.9, while the autocorrelation coefficient for the cross-sectional variance of log consumption is essentially one.

6 Endogenous displacement and intrinsic shocks

In this section we show that when the displacement process is endogenized, the resulting model can feature (real) equilibrium indeterminacy. Accordingly, shocks to consumers' expectations about future discount rates can be self-propagating and become a source of asset price fluctuations, even if there are no extrinsic shocks to the economy. Importantly, *real* indeterminacy arises despite the absence of “bubbles” in this model, a result that is of independent theoretical interest.

Since the goal of this section is mostly illustrative, we assume that investors have expected utility preferences rather than recursive preferences, as in the baseline model of Section 2.

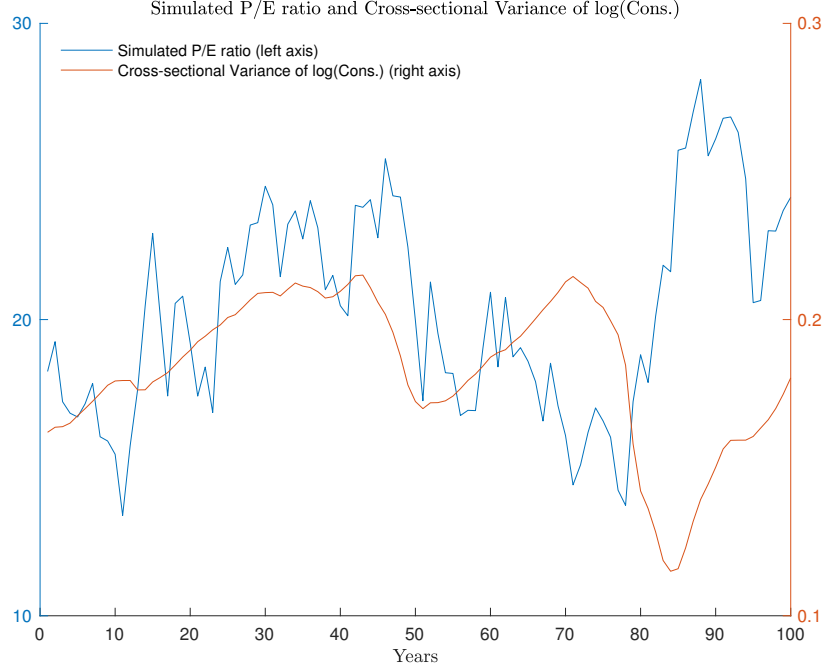


Figure 10: An indicative, model-implied path of the Price-Earnings ratio (left scale) and the cross-sectional standard deviation of log consumption (right scale).

Moreover, we assume that $\eta_t^l = \eta^l$ is constant. The only departure from the baseline model is that the displacement process is endogenous. Specifically, at the time of their birth only a fraction $\bar{\varepsilon} < 1$ of arriving agents have the ability to become entrepreneurs. These potential entrepreneurs, which we index by $i \in [0, \bar{\varepsilon}]$, are faced with two choices at birth: a) the “safe” choice of introducing a company that produces dividends

$$D_{t,s}^{(i)} = \psi \alpha Y_t e^{-\int_s^t \eta_u^d du} \quad (49)$$

for all $t \geq s$, where $\psi > 0$, and b) the “risky” choice of introducing a company that is successful with probability $\pi \in (0, 1)$, as reflected in its dividend process

$$D_{t,s}^{(i)} = \begin{cases} \xi^i \alpha Y_t e^{-\int_s^t \eta_u^d du} & \text{with probability } \pi > 0 \\ 0 & \text{with probability } 1 - \pi \end{cases} \quad (50)$$

for times $t \geq s$. The quantity ξ^i is entrepreneur-specific and known to the entrepreneur before she makes her choice. Without loss of generality, we assume that $\xi^i : [0, \bar{\varepsilon}] \rightarrow R^+$ is a

decreasing function.

If the firm ends up being worthless, then the entrepreneur must become a worker for the remainder of her life, thus financing a positive consumption stream. The choice between the safe and the risky options is made once, at birth, and the uncertainty associated with the risky choice is resolved immediately and publicly after the entrepreneur makes the risky choice.

To make matters interesting, we assume $\pi\xi(\bar{\varepsilon}) > \psi$. This assumption implies that a risk-neutral entrepreneur would always prefer to make the risky choice, since the expected dividends of the risky choice exceed the ones of the safe choice even for the entrepreneur with the least productive risky tree $\xi(\bar{\varepsilon})$. With risk-averse entrepreneurs, however, there is a meaningful tradeoff.

The rest of the assumptions of the model remain the same. In particular, to keep the aggregate dividend a constant function of aggregate consumption, we continue to assume

$$\frac{D_{t,t}}{\alpha Y_t} = \eta_t^d. \quad (51)$$

The only difference is that now η_t^d is endogenous and depends on the prevailing valuation ratios. Specifically, if a measure $\zeta_t \leq \bar{\varepsilon}$ of entrepreneurs chooses the risky choice, then aggregating gives

$$D_{t,t} = \int_0^{\bar{\varepsilon}} D_{t,t}^{(i)} di = \alpha Y_t \left(\pi \int_0^{\zeta_t} \xi^i di + (\bar{\varepsilon} - \zeta_t) \psi \right). \quad (52)$$

Combining (51) with (52) gives

$$\eta_t^d = \left(\pi \int_0^{\zeta_t} \xi^i di + (\bar{\varepsilon} - \zeta_t) \psi \right). \quad (53)$$

The next lemma is key for our purposes.

Lemma 8 *The measure ζ_t of agents choosing the risky options is a decreasing function of $\frac{q_t}{q_t'}$. Accordingly, by equation (53) $\eta_t^d = \eta(q_t)$ with $\eta' < 0$.*

Lemma 8 is intuitive: A smaller ratio of $\frac{q_t}{q_t'}$ makes an entrepreneur more willing to take




Figure 11: An illustration of Proposition 3.

risk, since — even if the tree turns out to be worthless — the present value of her earnings will now be larger. By Lemma 1, $\frac{q_t}{q_t^d} = (1 - \alpha) \frac{q_t}{\frac{1}{\beta} - \alpha q_t}$ is a monotonically increasing function of q_t , and hence Lemma 8 follows.

With Lemma 8 in hand, and assuming that there are no extrinsic shocks to the economy, we can proceed to construct a deterministic equilibrium. Specifically, repeating identical steps to Section 3, the dynamics of q_t continue to be given by

$$\dot{q}_t^d = \mu(q_t) \equiv (\beta + \bar{\eta}(q_t)) q_t - \beta \alpha \varphi(q_t) \times (q_t)^2 - 1, \quad (54)$$

with the only difference that now $\bar{\eta}(q_t) = \eta^d(q_t) - \eta^l$ is endogenous to q_t rather than being an exogenous process. We obtain the following proposition.

Proposition 3 *For any three real numbers $0 < q_1 < q_2 < q_3 < \frac{1}{\alpha\beta}$, there exist parameters under which q_i , $i \in \{1, 2, 3\}$, are roots of $\mu(q)$ with $\mu'(q_1) > 0$, $\mu'(q_2) < 0$, and $\mu'(q_3) > 0$.*

Figure 11 illustrates Proposition 3. The figure shows that \dot{q}_t is positive between q_1 and q_2 and negative between q_2 and q_3 . An immediate implication is that the dynamic system (54) has a stable steady state. Any initial value $q_{t_0} \in (q_1, q_3)$ is associated with a different

equilibrium transition path to the steady state q_2 . Interestingly, all of these paths constitute different, perfect-foresight equilibria and the economic structure cannot rule out any of them.

Indeterminacy arises because different expectations about future discount rates become self-fulfilling. Say for instance that everyone becomes convinced that interest rates will be lower on the transition path than at the steady state. This redistributes consumption to the arriving cohorts of agents by raising the present value of their human capital and increasing the creation of new firms, which appropriate profits from existing firms. Hence, low discount rates redistribute wealth from the old to the young, and reduce the consumption growth of existing cohorts who are marginal in asset markets, thus confirming the anticipation of low rates.

An immediate implication of the multiplicity identified above is the potential for so-called “sunspot” equilibria, i.e., stochastic equilibria in which uncertainty is not about preferences, endowments, etc., but rather reflects random fluctuations in agents’ (self-fulfilling) perceptions about the equilibrium paths that the economy will follow. To construct such equilibria, we introduce a standard brownian motion B_t . This Brownian motion reflects random “noise” that is extrinsic to the economy; however, everyone understands (and knows that everyone else also understands) that this noise acts as a coordination device for investor expectations.

The next proposition shows the existence of equilibria, whereby investor perceptions that the noise B_t is useful in coordinating expectations ends up becoming self-fulfilling, in the sense that it affects both asset price dynamics and equilibrium consumption allocations.

Proposition 4 *For q_1 and q_3 as in Proposition 3, take an interval $[q^{min}, q^{max}] \subset [q_1, q_3]$ with $\mu(q^{min}) > 0$ and $\mu(q^{max}) < 0$. Then choose a bounded function $\sigma_q : [q^{min}, q^{max}] \rightarrow R^+$ with the properties $\sigma_q(q^{min}) = \sigma_q(q^{max}) = 0$ and*

$$\lim_{q \rightarrow q^{max}} \frac{\sigma_q^2(q)}{q^{max} - q} = v_1 < 2 |\mu(q^{max})|, \quad \lim_{q \rightarrow q^{min}} \frac{\sigma_q^2(q)}{q - q^{min}} = v_2 < 2 |\mu(q^{min})|. \quad (55)$$

Then there exists an equilibrium whereby the equilibrium stochastic process for q_t is given by

the diffusion

$$dq_t = \mu(q_t) dt + \sigma_q(q_t) dB_t. \tag{56}$$

In such an equilibrium q_t possesses a stationary distribution, q_t^l continues to be given by (9), and r_t continues to satisfy (15).

Proposition 4 is reminiscent of Proposition 1. Indeed, similar to Section 3, if we wished to support a given drift function for $\mu(q_t)$ as an equilibrium outcome, we could just make assumptions on the distribution of the risky trees ξ_i so that the resulting function $\varphi(q_t)$ is given by (18). The extrinsic-uncertainty (shocks are technological) and the intrinsic-uncertainty (the shocks affect belief formation, with the displacement process being endogenous) versions of the model would be observationally equivalent.

7 Conclusion

In this paper we propose a simple mechanism to relate low frequency movements in inequality with volatile asset price movements. We exploit the structure of an overlapping generations economy, which allows different cohorts of agents to experience different (and random) consumption growth paths over their lifetimes, even though aggregate consumption evolves deterministically. Combining this observation with recursive preferences, we prove a possibility result (for the model's ability to match asset pricing phenomena) that is similar in spirit to Constantinides and Duffie (1996). However, we use a different set of assumptions on agents' endowments, which imply non-volatile, but persistent, inequality.

We also develop an empirical strategy to infer the persistent components of consumption growth by utilizing a time, age, and cohort decomposition of cross-sectional consumption data. Since it does not require time-series information, this technique can be implemented using readily available data sources, such as the CEX.

Finally, we show the theoretical possibility that volatility of asset prices can become a self-fulfilling prophecy, if one extends the model to allow for endogenous firm creation. In this

version of the model uncertain long-run variations in consumption growth of the marginal agent may be jointly “caused”, rather than be causal, for fluctuations in asset prices. This feature of the model helps illustrate a conceptual difference with representative agent models, where only extrinsic uncertainty can affect asset prices.

References

- Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59(4), 1481–1509.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics* 121(3), 823–66.
- Blanchard, O. J. (1985, April 1985). Debt, deficits, and finite horizons. *Journal of Political Economy* 93(2), 223–247.
- Constantinides, G., J. B. Donaldson, and R. Mehra (2002). Junior can't borrow: A new perspective on the equity premium puzzle. *The Quarterly Journal of Economics* 117(1), 269–296.
- Constantinides, G. and D. Duffie (1996). Asset pricing with heterogeneous consumers. *Journal of Political Economy* 104, 219–240.
- Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica* 60(2), 353–394.
- Gârleanu, N., L. Kogan, and S. Panageas (2012). The demographics of innovation and asset returns. *Journal of Financial Economics* 105, 491–510.
- Gârleanu, N. and S. Panageas (2015). Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy* 123, 670–685.
- Gomes, F. and A. Michaelides (2005). Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance* 60(2), 869–904.
- Karlin, S. and H. Taylor (1981). *A Second Course in Stochastic Processes*. Number v. 2. Academic Press.
- Kogan, L., D. Papanikolaou, and N. Stoffman (2016). Winners and losers: Creative destruction and the stock market. Unpublished working paper. MIT.

Piazzesi, M. and M. Schneider (2009). Inflation and the price of real assets. Federal Reserve Bank of Minneapolis, Staff Report: 423.

Storesletten, K., C. I. Telmer, and A. Yaron (2007). Asset pricing with idiosyncratic risk and overlapping generations. *Review of Economic Dynamics* 10(4), 519–548.

A Proofs

Proof of Lemma 1. The absence of bubbles together with the assumption of a unit elasticity of substitution implies that aggregate consumption is given by $C_t = \beta (\bar{W}_t + \bar{H}_t)$, where

$$\bar{W}_t = \int_{-\infty}^t q_{t,s}^d D_{t,s} ds = \alpha q_t^d Y_t \quad (57)$$

is the present value of all dividends to be paid by existing firms. Similarly, the total value of all human capital of existing agents is

$$\bar{H}_t = \int_{-\infty}^t q_{t,s}^l l_{t,s} w_{t,s} ds = (1 - \alpha) q_t^l Y_t. \quad (58)$$

Combining goods market clearing ($C_t = Y_t$) with (57) and (58) and re-arranging leads to (9). ■

Proof of Lemma 2. The present value of all newly-born workers' wages is given by $q_t^l (1 - \alpha) \eta^l Y_t$, while the present value of all newly created firms is $\alpha \eta^d q_t^d Y_t$. The sum of these quantities gives the total wealth of newly born agents. Given that the consumption-to-wealth ratio for investors is β , the per-capita consumption of the newly born, as a proportion of total consumption, follows as given by (13). ■

Proof of Lemma 3. The only step of the proof not made completely explicit in the proof is the one yielding equation (12). To show this relation, time-differentiate aggregate consumption $C_t = \lambda \int_{-\infty}^t e^{-\lambda(t-s)} c_{t,s} ds$ to get

$$\dot{C}_t = -\lambda C_t + \lambda \int_{-\infty}^t e^{-\lambda(t-s)} \dot{c}_{t,s} + \lambda c_{t,t} = -\lambda C_t + \dot{c}_{t,s} + \lambda c_{t,t}. \quad (59)$$

■

Proof of Lemma 4. Contained in the text. ■

The following technical restrictions on f and σ ensure existence of stationary q .

Assumption 1 *The functions f and σ are Lipschitz continuous on the bounded interval $[q^{min}, q^{max}] \subset (0, \frac{1}{\alpha\beta})$. Moreover, f is twice differentiable, monotonically decreasing, and satisfies $f(q^{min}) > 0$ and $f(q^{max}) < 0$. Finally, $\sigma(q) \geq 0$, $\sigma(q^{min}) = \sigma(q^{max}) = 0$, and*

$$\lim_{q \rightarrow q^{max}} \frac{\sigma^2(q)}{q^{max} - q} < 2|f(q^{max})|, \quad \lim_{q \rightarrow q^{min}} \frac{\sigma^2(q)}{q - q^{min}} < 2|f(q^{min})|. \quad (60)$$

Proof of Proposition 1. Let's start with a solution \hat{q} to the stochastic differential equation (17). By construction, the process $\varphi(\hat{q}_t)$ solves the SDE (19). Since φ_t is bounded, we can construct two positive processes η_t^l and η_t^d such that $\eta_t^d - \eta_t^l = \varphi_t$.

Posit that $q_t^d = q_t = \hat{q}_t$ and q_t^l as given by Lemma 1. Further, conjecture the interest rate r_t as given by equation (15).

Given the dynamics of q_t , the definition of $\varphi(\cdot)$, and the definition of r_t , pricing equation (10) is satisfied. Further, using also the definition of q_t^l , which implies $\alpha dq_t^d + (1-\alpha)dq_t^l = 0$, we obtain the analogous pricing equation for q_t^l :

$$\mathbb{E}[dq_t^l] = (r_t - g + \eta_t^d) q_t^l - 1. \quad (61)$$

Agents' consumption optimality require $c_{t,t} = \beta W_{t,t}$, yielding Lemma 2, as well as the Euler equation (11). Starting with equation (11), then applying Lemma 2 and equation (15)

in succession, we obtain

$$C_t = \int_{-\infty}^t \lambda e^{-\lambda(t-s)} e^{\int_s^t (r_u - \rho) du} c_{s,s} ds \quad (62)$$

$$= \int_{-\infty}^t e^{\int_s^t (r_u - \beta) du} (\eta_t^l + \alpha \beta q_t^l \varphi_t) Y_s ds \quad (63)$$

$$= \int_{-\infty}^t e^{\int_s^t (r_u - \beta - g) du} (\beta + g - r_s) Y_t ds \quad (64)$$

$$= Y_t, \quad (65)$$

given that r is bounded above away from $\beta + g$. Proposed consumption processes are therefore optimal and clear the consumption market, given the interest rate.

Finally, with q_t^d and q_t^l the valuation ratios, the total wealth in the economy is

$$\frac{1}{\beta} C_t = \frac{1}{\beta} Y_t = \alpha q_t^d Y_t + (1 - \alpha) q_t^l Y_t. \quad (66)$$

Given that the newly born agents' wealth equals $(1 - \alpha) \eta^l q_t^l + \alpha \eta_t^d q_t$, the change in wealth of all agents alive at time t equals, with \bar{t} fixed equal to t ,

$$\beta^{-1} C_t (r_t - \beta) dt = \beta^{-1} g Y_t dt - ((1 - \alpha) \eta^l q_t^l + \alpha \eta_t^d q_t) dt \quad (67)$$

$$= d \left(q_t^d \int_{-\infty}^{\bar{t}} D_{t,s} ds \right) \Big|_{\bar{t}=t} + d \left(q_t^l \int_{-\infty}^{\bar{t}} l_{t,s} w_{t,s} ds \right) \Big|_{\bar{t}=t} + Y_t dt - C_t dt, \quad (68)$$

which states that the (representative) agent alive at t invests her entire wealth by buying all available shares, and does not use the bond market — thus, asset markets clear. Equation (68) holds because $(1 - \alpha) q_t^l + \alpha q_t^d = \beta^{-1}$ and because of the dynamics of $D_{t,s}$ and $w_{t,s}$, i.e., equations (3) and (5).

Uniqueness of the process φ_t is a direct consequence of the analysis in the text, in particular equation (16).

We end the proof with a technical detail — a sketch of an argument that shows that q_t is stationary. We make use of results in Karlin and Taylor (1981). Specifically, we start by

defining

$$s(q) \equiv \exp \left\{ - \int^q \frac{2A(\xi)}{\sigma^2(\xi)} d\xi \right\},$$

noting that by assumption (60) there exists $\bar{v} > 1$ such that, for ε small enough and $q \in (q^{\max} - \varepsilon, q^{\max})$ we have

$$\frac{s(q)}{s(q^{\max} - \varepsilon)} = \exp \left\{ - \int_{q^{\max} - \varepsilon}^q \frac{2A(\xi)}{\sigma^2(\xi)} d\xi \right\} < \exp \left\{ - \int_{q^{\max} - \varepsilon}^q \frac{\bar{v}}{q^{\max} - \xi} d\xi \right\} = \left(\frac{q}{q^{\max} - \varepsilon} \right)^{-\bar{v}}.$$

Hence, for q “close” to q^{\max} the function $s(q)$ (and accordingly the speed measure $S(q) = \int^q s(\eta) d\eta$) behaves as in Example 5 on page 221 in Karlin and Taylor (1981). (A similar argument applies to the boundary $q = q^{\min}$). It then follows that the boundaries q^{\min} and q^{\max} are entrance boundaries whenever condition (60) holds and a stationary distribution exists. ■

Proof of Lemma 5. The fact that m_t is a (spanned) stochastic discount factor (SDF) means

$$d \log(m_t) = -r_t dt - \frac{\kappa_t^2}{2} dt - \kappa_t dB_t, \tag{69}$$

where κ_t is the market price of risk (the maximal Sharpe ratio). In the special case when preferences are specified by (22), and given the existence of annuities, the dynamics of the process $\log(m_t)$ are

$$d \log(m_t) = \beta (\gamma \log(c_t) - \log(\gamma V_t)) dt - \rho dt + d \log(\gamma V_t) - d \log(c_t). \tag{70}$$

An agent’s value function V is homogeneous of degree γ in the her total wealth W , which is the sum of her financial wealth and the present value of her future earnings. We

consequently write

$$V_t(W) = \frac{W^\gamma}{\gamma} e^{\tilde{Z}_t} \quad (71)$$

for an appropriate process \tilde{Z} . Furthermore, from the envelope condition we have

$$\frac{\gamma}{W} V_t = \frac{\partial V_t}{\partial W} = f_c = \frac{\beta \gamma V_t}{c},$$

giving $c = \beta W$.

For any $s < t$, the definition of V_t implies

$$V_t + \int_s^t \beta \gamma V_u \left(\log(c_u) - \frac{1}{\gamma} \log(\gamma V_u) \right) du = E_t \int_s^\infty \beta \gamma V_u \left(\log(c_u) - \frac{1}{\gamma} \log(\gamma V_u) \right) du.$$

Since the right-hand side is a martingale, the drift of the left-hand side equals zero, implying

$$dV_t = - \left(\beta \gamma V_t (\log(c_u) - \gamma^{-1} \log(\gamma V_u)) \right) dt + \sigma_V dB_t$$

and therefore

$$d \log(\gamma V_t) = \left(\frac{\mu_V}{V} - \frac{1}{2} \frac{\sigma_V^2}{V^2} \right) dt + \frac{\sigma_V}{V} dB_t \quad (72)$$

$$= -\beta \gamma (\log(c) - \gamma^{-1} \log(\gamma V)) dt - \frac{1}{2} \frac{\sigma_V^2}{V^2} dt + \frac{\sigma_V}{V} dB_t. \quad (73)$$

Plugging this last formula in (70), we obtain

$$d \log(m_t) = -\beta dt - \frac{1}{2} \frac{\sigma_V^2}{V^2} dt - d \log(c_t) + \frac{\sigma_V}{V} dB_t.$$

Comparison with (69), along with the fact that aggregate consumption growth is determin-

istic and as a result individual consumption growth needs to be locally deterministic implies

$$\frac{\sigma_V}{V} = -\kappa_t \tag{74}$$

$$\dot{c}_t = (r_t - \rho)c_t. \tag{75}$$

Consumption c_t is therefore locally deterministic, and so is $W_t = \beta^{-1}c_t$, which, upon using equation (71), leads to

$$\frac{\sigma_V}{V} = \sigma_{\tilde{Z}} = -\kappa_t.$$

Note that the dynamics of consumption depend only on q_t as long as r_t has this property, giving the value function \tilde{Z} as a function of q_t , as well. We go further and compute the function \tilde{Z} to derive $\kappa = -\sigma_{\tilde{Z}} = -\tilde{Z}_q \sigma_q$. Specifically, from the martingale condition,

$$\begin{aligned} \frac{\mu_V}{V} &= -\beta\gamma (\log(c) - \gamma^{-1} \log(\gamma V)) \\ &= -\beta\gamma (\log(\beta) - \gamma^{-1} \tilde{Z}). \end{aligned} \tag{76}$$

Using (71) and the fact that $c_t = \beta W_t$ implies that

$$d \log(\gamma V_t) = d\tilde{Z}_t + \gamma d \log c_t. \tag{77}$$

Since $\log c_t$ is locally deterministic, it follows that $\frac{\sigma_V}{V} = \sigma_{\tilde{Z}}$.

Combining (77), (73), and (76), and letting $Z_t = \tilde{Z}_t - \gamma \log(\beta)$, we obtain

$$dZ_t = -\gamma d \log c_t - \beta Z dt - \frac{1}{2} \sigma_Z^2 dt + \sigma_Z dB_t. \tag{78}$$

Integrating, and noting that σ_Z is bounded, gives equation (26). ■

Proof of Lemma 6. Since the right hand side of (31) depends only on Z_t , it is immediate

that strict monotonicity is equivalent to invertibility. Fixing Z_t and therefore ν_t , $\frac{\partial \varphi(q_t, Z_t)}{\partial q_t} < 0$ implies that there is a unique $q_t = \varphi^{-1}(\varphi_t, \nu_t)$. ■

Proof of Proposition 2. The proof of the proposition follows the same logic as that of Proposition 1. In the interest of completeness, we start by invoking Ito's Lemma to write down the SDE for ν :

$$d\nu_t = \nu'(\nu^{-1}(\nu_t)) \left(f_Z(\nu^{-1}(\nu_t)) + \frac{\sigma_Z^2(\nu^{-1}(\nu_t)) \nu''(\nu^{-1}(\nu_t))}{2 \nu'(\nu^{-1}(\nu_t))} \right) dt + \nu'(\nu^{-1}(\nu_t)) \sigma_Z(\nu^{-1}(\nu_t)) dB_t \quad (79)$$

Similarly, one can write the dynamics of

$$\varphi_t = \varphi(q_t, Z_t) \quad (80)$$

based on the dynamics of q_t and Z_t , and then plug in $q_t = \varphi^{-1}(\varphi_t, \nu_t)$ and $Z_t = \nu^{-1}(\nu_t)$.

The existence of the inverse functions ν^{-1} and φ^{-1} is ensured by Lemma 6. To avoid repetition, we only justify two key statements in the text, namely equations (29) and (30).

As before, the definition of q_t implies

$$m_t q_t D_{t,s} + \int_s^t m_t D_{t,s} = E_t \int_s^\infty m_u D_{u,s} du \quad (81)$$

is a martingale. Using Ito's Lemma and $\kappa_t = -\sigma_{Z,t}$ yields equation (30).

From equation (78), and using equation (24), the drift of Z_t equals

$$\beta Z_t - \gamma \frac{\dot{c}_t}{c_t} - \frac{1}{2} \sigma_Z^2(Z_t) = \beta Z_t - \gamma(\lambda + \rho - \nu_t) - \frac{1}{2} \sigma_Z^2(Z_t), \quad (82)$$

which is equated to $f_Z(Z_t)$ to yield equation (29). ■

Proof of Lemma 7. Re-writing $c_{t,s}$ as

$$c_{t,s} = e^{A_s + L_t + G_{t-s}},$$

computing $\frac{dc_{t,s}}{c_{t,s}}$, multiplying by $\omega(t, s)$ and integrating gives the first equality in (46). To obtain the second equality note that

$$\begin{aligned} \frac{dF_t}{F_t} &= \int_{-\infty}^t \left(\frac{d\Lambda(t, s)}{\Lambda(t, s)} + dG_{t-s} \right) \left(\frac{\Lambda(t, s) e^{A_s + G_{t-s}}}{\int_{-\infty}^t \Lambda(t, s) e^{A_s + G_{t-s}} ds} \right) ds + \frac{\Lambda(t, t) e^{A_t + G_0}}{\int_{-\infty}^t \Lambda(t, s) e^{A_s + G_{t-s}} ds} dt \\ &= \int_{-\infty}^t \omega_{t,s} \left(\frac{d\Lambda(t, s)}{\Lambda(t, s)} + dG_{t-s} \right) ds + \omega_{t,t} dt, \end{aligned} \quad (83)$$

where the second line of the above equation follows from

$$\omega_{t,s} = \frac{\Lambda(t, s) c_{t,s}}{C_t} = \frac{e^{A_s + L_t + G_{t-s}}}{\int_{-\infty}^t \Lambda(t, s) e^{A_s + L_t + G_{t-s}} ds} = \frac{e^{A_s + G_{t-s}}}{\int_{-\infty}^t \Lambda(t, s) e^{A_s + G_{t-s}} ds}.$$

Using (83) inside (43) leads to the second equality in (46). ■

Proof of Lemma 8. Since the value function of a newly-born person is logarithmic in wealth, the index $i \in [0, \bar{\varepsilon}]$ of the entrepreneur who is indifferent between the risky and the riskless choice is given by

$$\pi \log [\xi(i) q_t \alpha Y_t] + (1 - \pi) \log \left[q_t^l \frac{\eta^l (1 - \alpha)}{1 - \bar{\varepsilon} + (1 - \pi) i} Y_t \right] = \log [\psi q_t \alpha Y_t]. \quad (84)$$

The left hand side gives the value function of trying the risky choice which succeeds with probability π and fails with probability $(1 - \pi)$, in which case the entrepreneur shares the labor income accruing to her cohort. The right hand side is the (certain) payoff of the riskless choice. Simplifying and re-arranging gives

$$(1 - \pi) \log \left[\frac{q_t}{q_t^l} \right] = \pi \log [\xi(i) \alpha] + \log \left[\frac{\eta^l (1 - \alpha)}{1 - \bar{\varepsilon} + (1 - \pi) i} \right] - \log [\psi \alpha] \quad (85)$$

The right hand side of the above equation is a decreasing function of i . Hence the value of i that makes the above equation hold is decreasing in $\frac{q_t}{q_t^l}$. An implication of Lemma 1 is that

$\frac{q_t}{q_t}$ is increasing in q_t , which concludes the proof. ■

Proof of Proposition 3. Let $B(q; \bar{\eta}) \equiv (\beta + \bar{\eta})q - \beta\alpha\bar{\eta}q^2 - 1$ and define the function

$$\bar{\eta}^*(q) = \frac{1 - \beta q}{q(1 - \alpha\beta q)} \quad (86)$$

for $q \in (0, \frac{1}{\alpha\beta})$. By construction, $B(q; \bar{\eta}^*(q)) = 0$. It is easy to verify (e.g., by direct differentiation) that $\bar{\eta}^*$ decreases strictly. Note also that $\frac{\partial}{\partial \bar{\eta}} B(q; \bar{\eta}) = q(1 - \alpha\beta q) > 0$.

Let $\bar{\eta}(q)$ be continuously differentiable and decreasing with the following properties: (i) $\bar{\eta}(q_i) = \bar{\eta}^*(q_i)$ for $i \in \{1, 2, 3\}$; (ii) $\bar{\eta}'(q_i) > \bar{\eta}^{*'}(q_i)$, $i \in \{1, 3\}$; (iii) $\bar{\eta}'(q_2) < \bar{\eta}^{*'}(q_2)$. Given that $\frac{\partial}{\partial \bar{\eta}} B > 0$, these properties ensure that the proposition holds. (Note that properties (ii) and (iii) require that $\bar{\eta}$ be flatter than $\bar{\eta}^*$ around the extreme \bar{q}_i , and steeper around q_2 .)

Turning now to the model primitives under which $\bar{\eta}$ takes the form chosen above, we note that equation (85) can be solved uniquely for $\xi(\zeta(q))$ as a function of q and $\zeta(q)$ (over an appropriate range). We also have

$$\bar{\eta}(q) = \pi \int_0^{\zeta(q)} \xi^i di + (\bar{\varepsilon} - \zeta(q))\psi - (\eta^l + g), \quad (87)$$

which implies

$$\bar{\eta}'(q) = (\pi\xi(\zeta(q)) - \psi)\zeta'(q), \quad (88)$$

or

$$\zeta'(q) = \frac{\bar{\eta}'(q)}{\pi\xi(\zeta(q)) - \psi}. \quad (89)$$

Given $\xi(\zeta(q))$ from (85), this is a first-order ODE in $\zeta(q)$.

We wish that a decreasing solution exists on $[q_1, q_3]$ with image in $[0, \bar{\varepsilon}]$. To ensure the existence of such a solution, we can build one as follows under appropriate parameter choices.

Let $\zeta(q_3) = 0$, and pick $\eta^l + g$ so that (87) is satisfied. Note now that, from (85), for η^l low

enough ξ can be bounded below uniformly by an arbitrarily large value. Consequently the value of ζ , obtained by solving (integrating) equation (89), can be kept as small as desired. ■

Proof of Proposition 4. Equations (11) and (15) continue to be valid irrespective of whether q_t is stochastic or not. So do Lemmas 2, 1, and 8. Accordingly q_t^l and r_t continue to be given by (9) and (15) respectively. Applying Ito's Lemma to (7), while substituting r_t from (15) implies that the drift of q_t in any stochastic equilibrium must necessarily be given by $A(q_t)$. Moreover, if the dynamics of q_t are given by (56), the Feynman-Kac formula implies that q_t satisfies (7). In a nutshell, if agents perceive that the dynamics of asset prices are given by (56), then the resulting optimal dynamics of consumption will be such that the market-clearing interest rate will be given by (15) and the equilibrium (arbitrage-free) price of each firm will indeed be given by $q_t D_{t,s}$. ■