Survival and long-run dynamics with heterogeneous beliefs under recursive preferences*

Jaroslav Borovička

University of Chicago

Federal Reserve Bank of Chicago

borovicka@uchicago.edu

January 2, 2012

Abstract

I study the long-run behavior of a two-agent economy where agents differ in their beliefs and are endowed with homothetic recursive preferences of the Duffie-Epstein-Zin type. When preferences are separable, the economy is dominated in the long run by the agent whose beliefs are relatively more accurate, a result consistent with the market selection hypothesis. However, recursive preference specifications lead to equilibria in which both agents survive, or to ones where either agent can dominate the economy with a strictly positive probability. In this respect, the market selection hypothesis is not robust to deviations from separability. I derive analytical conditions for the existence of nondegenerate long-run equilibria, and show that these equilibria exist for plausible parameterizations when risk aversion is larger than the inverse of the intertemporal elasticity of substitution. These results provide a justification for models that combine belief heterogeneity and recursive preferences.

*For the latest version of the paper, check http://home.uchicago.edu/∼borovicka/research.html. I am indebted to Lars Peter Hansen for his advice and continuous support. I appreciate helpful comments from Fernando Alvarez, Gadi Barlevy, Katarína Borovičková, Hui Chen, Valentin Haddad, Narayana Kocherlakota, Alan Moreira, Stavros Panageas, Christopher Phelan, Thomas Sargent, Harald Uhlig, Yuichiro Waki, Mark Westerfield, and the participants of the Economic Dynamics working group at the University of Chicago. The views expressed herein are those of the author and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System. All typos and errors are mine.

1 Introduction

The market selection hypothesis first articulated by Alchian (1950) and Friedman (1953) is one of the supporting arguments for the plausibility of the rational expectations theory. The hypothesis states that agents who systematically evaluate the distributions of future quantities incorrectly (and are therefore called 'irrational') lose wealth on average, and will ultimately be driven out of the market. Thus, in a long-run equilibrium, the dynamics of the economy are only determined by the behavior of the rational agents whose beliefs about the future are in line with the true probability distributions.

However, rationality does not guarantee survival, nor do deviations from rational preferences imply extinction. Ultimately, survival in a market is driven by the consumption-saving decision and willingness to take risky positions with high expected return vis-à-vis market prices. In this sense, rationality may *facilitate* survival if it prevents overconsumption and/or leads the agent to take appropriate bets. On the other hand, specific forms of irrationality may, at least in theory, provide even stronger incentives for survival in the long run, despite not being optimal in the rational sense. Since individual decisions depend on equilibrium prices, survival analysis only makes sense in the context of a fully specified model, including preferences and belief formation of the market participants and their trading opportunities.

Survival of agents with incorrect beliefs has been studied extensively in complete market models populated by agents endowed with separable preferences. The existing literature on market survival has demonstrated how differences in beliefs can be counteracted by differences in preferences. Yan (2008) analyzes a model with constant relative risk aversion (CRRA) preferences and constructs a quantity called the survival index that aggregates the role of intertemporal elasticity of substitution (IES), time preference, and belief distortion into a single number that determines survival. When preferences are identical across agents, then only agents whose beliefs are closest to the truth will survive in the long run.

This insight is central for the understanding of the survival mechanism and can be rephrased as follows. If rich and poor agents are alike, in the sense that rich agents behave as scaled versions of poor ones, then agents with relatively more incorrect beliefs cannot survive in the long run. This is precisely correct for the homothetic CRRA preferences. In their lucid analysis, Kogan, Ross, Wang, and Westerfield (2009) show that this statement is also true for a class of preferences with bounded relative risk aversion, i.e., preferences that are in some norm uniformly 'close' to the homothetic CRRA case.

This paper shows that support in favor of the market survival hypothesis weakens considerably once the assumption of separability in preferences is relaxed. In order to focus solely on the impact of belief heterogeneity, I endow agents with identical homothetic recursive preferences axiomatized by Kreps and Porteus (1978), and developed by Epstein and Zin (1989) and Weil (1990) in discrete time, and by Duffie and Epstein (1992b) in continuous time. These preferences allow one to disentagle the risk aversion with respect to intratemporal gambles from the intertemporal elasticity of substitution, and include the CRRA utility as a special case. Thanks to the additional degree of flexibility, this class of preferences is widely used in the asset pricing literature to provide a better fit of the constructed models to empirically observed patterns in asset returns. Homotheticity assures that survival results are not driven by exogenous differences in the local properties of the utility functions.

The decoupling of risk aversion and IES proves to have a crucial impact on survival. Agents in the economy can accumulate wealth, and thus avoid extinction, in two ways — by holding portfolios with high expected logarithmic returns, and by choosing a high saving rate. The risk aversion parameter governs the portfolio allocation decision of the agents and equilibrium risk premia associated with risky assets. These two quantities jointly determine the difference in expected returns on agents' portfolios. The IES parameter then drives the difference in saving rates as a function of the difference in perceived expected returns on individual portfolios.

To shed more light on the survival mechanism, consider an endowment economy as in this paper populated by agents who are optimistic or pessimistic about the growth rate of the stochastic aggregate endowment. More optimistic agents hold a larger share of their wealth in the risky claim on aggregate endowment. An increase in the level of risk aversion in the economy has two effects. First, it increases risk premia, which increases disproportionately more the returns on portfolios of the more optimistic agents with their large positions in the risky claim. Second, the higher risk aversion reduces betting — the portfolio positions of different agents become more alike, which compresses differences in expected returns. Overall, the increase in risk aversion tends to increase the wedge between the expected logarithmic returns on the portfolios of optimistic and pessimistic agents, improving the survival chances of the more optimistic agents. When risk aversion is sufficiently high, optimistic agents unambiguously earn a higher expected logarithmic return on their portfolios than their pessimistic counterparts.

On the other hand, an increase in IES facilitates, through the consumption-saving decision, the survival of agents who are relatively more optimistic about the expected return on their portfolio. When IES is low, agents prefer flat consumption profiles over time. Agents who expect, *under their subjective beliefs*, high returns on their portfolio will therefore have a higher consumption rate than agents with lower perceived expected portfolio returns. As IES increases, the difference in consumption rates changes sign (agents will have identical constant consumption rates when IES is equal to one). With a high IES, agents who perceive high expected returns on their portfolios also have high saving rates. Both effects aid the survival of optimistic agents.

The consumption-saving choice may also facilitate the survival of pessimistic agents. An agent who is pessimistic about the growth rate of aggregate endowment would be optimistic about the growth rate of a portfolio involving a short position in the claim on aggregate endowment. When the short position or the degree of pessimism are sufficiently large, the *perceived* expected return on this portfolio will exceed the perceived expected returns on other agents' portfolios. With a high IES, the pessimistic agent will prevent her extinction by choosing a high saving rate.

Under CRRA preferences, risk aversion and IES are inversely related, and the two effects offset each other. Increasing the risk aversion increases relatively more the expected returns on portfolios held by optimistic agents, but the associated decrease in IES makes them, at the same time, consume more out of their wealth, relative to agents with correct beliefs. Pessimistic agents, on the other hand, are paying more to insure against the states with low consumption growth as risk aversion increases. The relatively lower willingness to consume out of wealth linked to the associated decrease in IES in conjunction with lower perceived expected returns is not strong enough to compensate for the cost.

The survival mechanism that supports the coexistence of agents with heterogeneous beliefs in the long run does not hinge on the stochastic nature of the aggregate endowment. Important is the existence of a betting device over which the agents disagree. If agents write state-contingent contracts on the outcomes of the betting device, they will perceive different expected returns on their portfolios, and the consumption-saving mechanism remains operational. When IES is sufficiently high, both agents will survive in the long run. Even with a constant aggregate endowment, the economy will exhibit stochastic dynamics of wealth shares and asset prices.

I analyze these mechanisms in a two-agent, continuous-time endowment economy with complete markets and an aggregate endowment process modeled as a geometric Brownian motion. The continuous-time, Brownian information framework is not critical for the qualitative results but offers analytical tractability which allows sharp closed-form characterization of the results.

I find that both agents survive and a nondegenerate equilibrium exists in the long run for wide regions of the parameter space when risk aversion is larger than the inverse of IES. An optimistic agent will dominate the economy in the long run when risk aversion is sufficiently large. A pessimistic agent survives in the long run when IES is sufficiently high and risk aversion is not excessive. In the opposite case, when risk aversion is sufficiently lower than the inverse of IES, only one agent survives in the long run, but the surviving agent can be either of the two agents with a strictly positive probability.

The market selection hypothesis is thus not robust to departures from separable preferences. Survival crucially depends on the interaction between risk attitudes that drive the portfolio selection decision and the IES that influences the consumption-saving decision. Recursive preferences provide an additional degree of freedom compared to the separable case that allows one to separate these two effects.

Given the homotheticity of preferences and complete markets, the survival results are driven purely by prices endogenously determined in general equilibrium. Crucially, non-degenerate long-run equilibria arise for preference parameterizations that are considered plausible in the asset pricing literature, which lends support to asset pricing models that combine heterogeneous beliefs and recursive preferences.

1.1 Methodology and literature overview

The modern approach in the market survival literature[1] originates from the work of De Long, Shleifer, Summers, and Waldmann (1991), who study wealth accumulation in a partial equilibrium setup with exogenously specified returns and find that irrational noise traders can outgrow their rational counterparts and dominate the market. Similarly, Blume and Easley (1992) look at the survival problem from the vantage point of exogenously specified saving rules, albeit in a general equilibrium setting.

Subsequent research has shown that taking into account general equilibrium effects and intertemporal optimization eliminates much of the support for survival of agents with incor-

[1]Modeling of economies populated by agents endowed with heterogeneous beliefs constitutes a quickly growing branch of literature, and a thorough overview of the literature is beyond the scope of this paper. Here, I primarily focus on the intersection of this literature with the analysis of recursive nonseparable preferences. Bhamra and Uppal (2009) provide a more general survey that also focuses on asset pricing implications of belief and preference heterogeneity.

I also omit the discussion of evolutionary literature which predominantly focuses on the analysis of the interaction between agents with exogenously specified portfolio rules and price dynamics. The survival mechanism in this paper critically hinges on the interaction the of endogenous consumption-saving decision and portfolio allocation vis-à-vis general equilibrium prices driven by the dynamics of the wealth shares, and is thus only loosely related. See Hommes (2006) for a survey of the evolutionary literature, and Evstigneev, Hens, and Schenk-Hoppé (2006) for an analysis of portfolio rule selection.

rect beliefs that models with ad hoc price dynamics produce. Sandroni (2000) and Blume and Easley (2006) base their survival results on the evolution of relative entropy as a measure of disparity between subjective beliefs and the true probability distribution. In their models, aggregate endowment is bounded from above and away from zero. As a result, changes in the curvature of the utility function are immaterial for survival when mild regularity conditions are satisfied. Controlling for pure time preference, the long-run fate of economic agents is determined solely by belief characteristics, and only agents whose beliefs are in a specific sense asymptotically 'closest' to the truth can survive.

With unbounded aggregate endowment, local properties of the utility function become an additional survival factor. Even if preferences are identical across agents, the local curvature of the utility function at low and high levels of consumption can be sufficiently different to outweigh the divergence in beliefs, and lead to survival of agents with relatively more incorrect beliefs. Kogan, Ross, Wang, and Westerfield (2009) show that a sufficient condition to prevent this outcome is the boundedness of the relative risk aversion coefficient. This condition can be interpreted as a bound on deviations of the utility function from homotheticity.

Importantly, survival analysis under separable preferences corresponds to analyzing a sequence of time- and state-indexed static problems that are only interlinked through the initial marginal utility of wealth, which is largely innocuous for the long-run characterization of the economy. The survival literature frequently exploits martingale methods to characterize the long-run divergence of subjective beliefs and marginal utilities of consumption.

Nonseparability of preferences breaks this straightforward link, and I therefore develop a different method that is more suitable for this environment. I utilize the planner's problem derived in Dumas, Uppal, and Wang (2000) and extend it to include heterogeneity in beliefs. The solution of the planner's problem involves endogenously determined processes that can be interpreted as stochastic Pareto weights. The analysis under separable preferences reflects the purely *intratemporal* tradeoff in the allocation of consumption vis-à-vis changes in the local curvature of the period utility function. The nonseparable nature of recursive preferences introduces an additional *intertemporal* component captured in the dynamics of the Pareto weights.

The analysis of market survival then corresponds to investigating the long-run behavior of scaled Pareto weights. I present tight sufficient conditions for the existence of nondegenerate long-run equilibria and for dominance and extinction. While the full model requires a numerical solution, I show that the behavior at the boundaries, which is essential for survival analysis, can be established analytically. I thus provide closed-form solutions for the regions of the parameter space in which the survival conditions are satisfied.

The method utilizes asymptotic properties of a differential equation for the planner's problem to characterize the asset price dynamics at the boundaries in the decentralized equilibrium. The resulting conditions from the planner's problem translate naturally into conditions on the relative logarithmic growth rates in agents' wealth.

The applicability of the derived solution method is not limited to fixed distortions. I discuss how to extend the procedure to include learning and robust preferences of Anderson, Hansen, and Sargent (2003). Explicit solutions of these problems are left for future work.

The approach based on the characterization of the behavior of the endogenously determined Pareto weights is closely linked to the literature on endogenous discounting, initiated by Koopmans (1960) and Uzawa (1968), and to models of heterogeneous agent economies under recursive preferences, studied by Lucas and Stokey (1984) and Epstein (1987) under certainty and by Kan (1995) under uncertainty. The survival conditions derived in this paper resemble a sufficient condition for the existence of a stable interior steady state in Lucas and Stokey (1984), called increasing marginal impatience. This condition postulates that agents discount future less as they become poorer. I show that my analysis crucially depends on a similar quantity that I call *relative patience*. The key difference lies in the determination of the two quantities. While Lucas and Stokey require that the time preference exogenously encoded in the utility specification changes with the level of consumption, in this paper the variation in relative patience arises endogenously as an equilibrium outcome driven by belief differences.

Anderson (2005) studies Pareto optimal allocations under heterogeneous recursive preferences in a discrete-time setup using similar methods but he does not consider survival under belief heterogeneity. Mazoy (2005) discusses long-run consumption dynamics when agents differ in their IES. Colacito and Croce (2010) prove the existence of nondegenerate long-run equilibria in a two-good economy when agents are endowed with risk-sensitive preferences and differ in the preferences over the two goods. However, none of these papers treats systematically the case of belief heterogeneity. This work aims at filling this gap.

The paper is organized as follows. Section 2 outlines the economic environment, provides a theoretical exposition to recursive preferences, and derives the planner's problem that is central to the analysis. Section 3 presents the survival results. I provide in analytical form tight sufficient conditions for survival and extinction and discuss the economic interpretation of the results. This analytical part is followed by numerical analysis of consumption and price dynamics for economies with nondegenerate long-run equilibria in Section 4. Section 5 summarizes the findings and outlines extensions of the developed framework involving learning and endogenously determined belief distortions derived, for instance, from robust preferences. The Appendix contains proofs omitted from the main text. Further material that provides more details and extends the analysis is available in the online appendix.[2]

2 Optimal allocations under heterogeneous beliefs

I analyze the dynamics of equilibrium allocations in a continuous-time endowment economy populated by two types of infinitely-lived agents endowed with identical recursive preferences. I call an economy where both agents have strictly positive wealth shares a heterogeneous economy. A homogeneous economy is populated by a single agent only. The term 'agent' refers to an infinitesimal competitive representative of the particular type.

Agents differ in their subjective beliefs about the distribution of future quantities but are firm believers in their probability models and 'agree to disagree' about their beliefs as in Morris (1995). Since they do not interpret their belief differences as a result of information asymmetries, there is no strategic trading behavior.

Without introducing any specific market structure, I assume that markets are dynamically complete in the sense of Harrison and Kreps (1979). This allows me to sidestep the problem of directly calculating the equilibrium by considering a planner's problem. The discussion of market survival then amounts to the analysis of the dynamics of Pareto weights associated with this planner's problem. Optimal allocations and continuation values generate a valid stochastic discount factor and a replicating trading strategy for the decentralized equilibrium.

In this section, I specify agents' preferences and belief distortions, and lay out the planner's problem. I utilize the framework introduced by Dumas, Uppal, and Wang (2000), and exploit the observation that belief heterogeneity can be analyzed in their framework without increasing the degree of complexity of the problem. The method then leads to a Hamilton-Jacobi-Bellman equation for the planner's value function.

2.1 Information structure and beliefs

The stochastic structure of the economy is given by a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ with an augmented filtration defined by a family of σ-algebras $\{\mathcal{F}_t\}, \ t \geq 0$ generated by a

[2] http://home.uchicago.edu/~borovicka/files/research/heterogeneous_beliefs_online_appendix.pdf

univariate Brownian motion W. Given the continuous-time nature of the problem, equalities are meant in the appropriate almost-sure sense. I also assume that all processes, in particular belief distortions and permissible trading strategies, satisfy regularity conditions like square integrability over finite horizons, so that stochastic integrals are well defined and pathological cases are avoided. Under the parameter restrictions below, constructed equilibria satisfy these assumptions.

The scalar aggregate endowment process Y satisfies

$$\frac{dY_t}{Y_t} = \mu_y dt + \sigma_y dW_t, \qquad Y_0 > 0 \tag{1}$$

with constant parameters μ_y and σ_y.

Agents of type $n \in \{1, 2\}$ are endowed with identical preferences but differ in their subjective probability measures Q^n that they use to assign probabilities to future events. I assume that measures Q^n and P are equivalent for all finite-horizon events, so that there exists the Radon-Nikodým derivative

$$\left(\frac{dQ^n}{dP}\right)_t \doteq M_t^n = \exp\left(-\frac{1}{2}\int_0^t |u_s^n|^2 \, ds + \int_0^t u_s^n dW_s\right), \tag{2}$$

where u^n is an adapted process. The martingale M^n measures the disparity between the subjective and true probability measures and is commonly called the belief ratio. While a likelihood evaluation of the past observed data reveals that the view of an agent with distorted beliefs becomes less and less likely to be correct as time passes, absolute continuity of the measure Q^n with respect to P over finite horizons implies that he cannot refute his view of the world as impossible in finite time. The main results of the paper are developed using a constant u^n, but the computational strategy allows me to incorporate more general distortion processes, which I discuss in the concluding remarks.

The Girsanov theorem implies that agent n, whose deviation from rational beliefs is described by M^n, views the evolution of the Brownian motion W as distorted by a drift component u^n, i.e., $dW_t = u_t^n dt + dW_t^n$, where W^n is a Brownian motion under Q^n. Consequently, the aggregate endowment is perceived to contain an additional drift component $u^n \sigma_y$, and u^n can be interpreted as a degree of optimism or pessimism about Y. When $\sigma_y = 0$, this distinction loses its meaning but the survival problem is still nondegenerate, as long as the agents can contract upon the realizations of the process W.

2.2 Recursive utility

Agents endowed with separable preferences reduce intertemporal compound lotteries (different payoff streams allocated over time) to atemporal simple lotteries that resolve uncertainty at a single point in time. In the Arrow-Debreu world with separable preferences, once trading of state-contingent securities for all future periods is completed at time 0, uncertainty about the realized path of the economy can be resolved immediately without any consequences for the ex-ante preference ranking of the outcomes by the agents.

Kreps and Porteus (1978) relaxed the separability assumption by axiomatizing discrete-time preferences where temporal resolution of uncertainty matters and preferences are not separable over time. While intratemporal lotteries in the Kreps-Porteus axiomatization still satisfy the von Neumann-Morgenstern expected utility axioms, intertemporal lotteries cannot in general be reduced to atemporal ones. The work by Epstein and Zin (1989, 1991) extended the results of Kreps and Porteus (1978), and initiated the widespread use of recursive preferences in the asset pricing literature. Duffie and Epstein (1992a,b) formulated the continuous-time counterpart of the recursion.[3]

I utilize a characterization based on the more general variational utility approach studied by Geoffard (1996) in the deterministic case and El Karoui, Peng, and Quenez (1997) in a stochastic environment.[4] They show that recursive preferences can be represented as a solution to the maximization problem

$$\lambda_t^n V_t^n = \sup_{\nu^n} E_t^{Q^n} \left[\int_t^\infty \lambda_s^n F\left(C_s^n, \nu_s^n\right) ds \right] \tag{3}$$

subject to

$$\frac{d\lambda_t^n}{\lambda_t^n} = -\nu_t^n dt, \quad t \geq 0; \quad \lambda_0^n = 1, \tag{4}$$

where ν^n is called the discount rate process, and λ^n the discount factor process. The felicity function $F(C, \nu)$ encodes the contribution of the consumption stream C to present utility. This representation closely links recursive preferences to the literature on endogenous discounting, initiated by Koopmans (1960) and Uzawa (1968).

[3] Duffie and Epstein (1992b) provide sufficient conditions for the existence of the recursive utility process for the infinite-horizon case but these are too strict for the preference specification considered in this paper. However, the Markov structure of the problem allows me to utilize existence results derived Duffie and Lions (1992). Schroder and Skiadas (1999) establish conditions under which the continuation value is concave, and provide further technical details. Skiadas (1997) shows a representation theorem for the discrete time version of recursive preferences with subjective beliefs.

[4] Hansen (2004) offers a tractable summary of the link between the recursive and variational utility. Interested readers may refer to the online appendix for a more detailed discussion.

For the case of the Duffie-Epstein-Zin preferences, the felicity function is given by

$$F(C,\nu) = \beta \frac{C^\gamma}{\gamma} \left(\frac{\gamma - \rho \frac{\nu}{\beta}}{\gamma - \rho} \right)^{1-\frac{\gamma}{\rho}},$$

with parameters satisfying $\gamma, \rho < 1$, and $\beta > 0$. Preferences specified by this felicity function[5] are homothetic and exhibit a constant relative risk aversion with respect to intratemporal wealth gambles $\alpha = 1 - \gamma$ and (under intratemporal certainty) a constant intertemporal elasticity of substitution $\eta = \frac{1}{1-\rho}$. Parameter β is the time preference coefficient. Assumption 2 below restricts parameters to assure sufficient discounting for the continuation values to be finite in both homogeneous and heterogeneous economies. In the case when $\gamma = \rho$, the utility reduces to the separable CRRA utility with the coefficient of relative risk aversion α.

Formula (3), together with an application of the Girsanov theorem, suggests that it is advantageous to combine the contribution of the discount factor process λ^n and the martingale M^n that specifies the belief distortion in (2):

Definition 1 *A modified discount factor process $\bar{\lambda}^n$ is a discount factor process that incorporates the martingale M^n arising from the belief distortion, $\bar{\lambda}^n \doteq \lambda^n M^n$.*

Applying Itô's lemma to $\bar{\lambda}^n$ leads to a maximization problem under the true probability measure

$$\bar{\lambda}_t^n V_t^n = \sup_{\nu^n} E_t \left[\int_t^\infty \bar{\lambda}_s^n F(C_s^n, \nu_s^n)\, ds \right] \tag{5}$$

subject to

$$\frac{d\bar{\lambda}_t^n}{\bar{\lambda}_t^n} = -\nu_t^n dt + u_t^n dW_t, \;\; t \geq 0; \;\; \bar{\lambda}_0^n = 1. \tag{6}$$

The problem (5-6) indicates that $F(C,\nu)$ can be viewed as a generalization of the period utility function with a potentially stochastic rate of time preference ν that depends on the properties of the consumption process and thus arises endogenously in a market equilibrium. Moreover, belief distortions are now fully incorporated in the framework of Dumas, Uppal, and Wang (2000) — the only difference is that the modified discount factor process is not locally predictable.

The diffusion term $u_s^n dW_s$ has an intuitive interpretation. Consider an optimistic agent with $u^n > 0$. This agent's beliefs are distorted in that the mass of the distribution of dW_s is shifted to the right — the agent effectively overweighs good realizations of dW_s. Formula

[5] The cases of $\rho \to 0$ and $\gamma \to 0$ can be obtained as appropriate limits. The maximization problem (3) assumes that the felicity function is concave in its second argument. When it is convex, the formulation becomes a minimization problem.

(6) indicates that under the true probability measure, positive realizations of dW_s increase the term $d\bar{\lambda}_s^n/\bar{\lambda}_s^n$, which implies that the optimistic agent discounts positive realizations of dW_s less than negative ones.

From the perspective of the utility-maximizing agent, assigning a higher probability to an event and a lower discounting of the utility contribution of this event have the same effect. In fact, equation (3) suggests that we can understand the belief distortion as a preference shock and view $\bar{\lambda}^n F(C^n, \nu^n)$ as a state-dependent felicity function. However, interpreting the martingale M^n as a belief distortion is more appealing since it bears a clearer economic meaning, separating the structure of beliefs and preferences.

2.3 Planner's problem and optimal allocations

The problem of an individual agent (3–4) is homogeneous degree one in the modified discount factors and homogeneous degree γ in consumption. In the homogeneous economy, there exists a closed-form solution for the continuation value $V_t^n(Y) = \gamma^{-1} Y_t^\gamma \tilde{V}^n$ where

$$\tilde{V}^n = \left(\beta^{-1}\left[\beta - \rho\left(\mu_y + u^n\sigma_y - \frac{1}{2}(1-\gamma)\sigma_y^2\right)\right]\right)^{-\frac{\gamma}{\rho}} \tag{7}$$

with the associated discount rate

$$\nu^n = \frac{\beta}{\rho}\left(\gamma + (\rho-\gamma)\left(\tilde{V}^n\right)^{-\frac{\rho}{\gamma}}\right) = \beta + (\gamma-\rho)\left(\mu_y + u^n\sigma_y - \frac{1}{2}(1-\gamma)\sigma_y^2\right). \tag{8}$$

Assumption 2 *The parameters in the model satisfy the restrictions*

$$\beta > \max_n \rho\left(\mu_y + u^n\sigma_y - \frac{1}{2}(1-\gamma)\sigma_y^2\right), \tag{9}$$

$$\beta > \max_n \rho\left(\mu_c + u^{\sim n}\sigma_y - \frac{1}{2}(1-\gamma)\sigma_y^2\right) + \frac{\rho}{1-\rho}\left[\left(u^n - u^{\sim n}\right)\sigma_y + \frac{1}{2}\frac{\left(u^n - u^{\sim n}\right)^2}{1-\gamma}\right] \tag{10}$$

where $\sim n$ is the index of the agent other than n.

The first restriction is sufficient for the continuation values in the homogeneous economies to be well-defined. The second restriction, which may be, depending on the parameterization, somewhat tighter, is a sufficient condition assuring that the wealth-consumption ratio is asymptotically well-behaved in the survival proofs when the agent becomes infinitesimally small. Observe that both conditions are restrictions on the time-preference parameter of the agents and can always be jointly satisfied by making the agents sufficiently impatient.

Since survival results will not depend on β, Assumption 2 does not introduce substantial restrictions for the analysis of the problem.

In the heterogeneous economy, I can follow Dumas, Uppal, and Wang (2000) and introduce a fictitious planner who maximizes a weighted average of the continuation values of the two agents. Given a pair of strictly positive initial Pareto weights $\alpha = (\alpha^1, \alpha^2)$, the planner's time-0 objective function $J_0(\alpha)$ is the solution to the problem

$$J_0(\alpha) = \sup_{(C^1, C^2, \nu^1, \nu^2)} \sum_{n=1}^{2} E_0 \left(\int_0^{\infty} \bar{\lambda}_t^n F(C_t^n, \nu_t^n)\, dt \right) \tag{11}$$

subject to the law of motion for the modified discount factors,

$$\frac{d\bar{\lambda}_t^n}{\bar{\lambda}_t^n} = -\nu_t^n dt + u_t^n dW_t, \quad t \geq 0; \quad \bar{\lambda}_0^n = \alpha^n \tag{12}$$

for $n \in \{1, 2\}$, and the feasibility constraint $C^1 + C^2 \leq Y$.

The validity of this approach for a finite-horizon economy is discussed in Dumas, Uppal, and Wang (2000) and Schroder and Skiadas (1999). The infinite-horizon problem in (11-12) is a straightforward extension when individual continuation values are well-defined. The planner's objective function is bounded from above by the weighted average of continuation values from the homogeneous economies, $J_0(\alpha) \leq \alpha^1 V_0^n(Y) + \alpha^2 V_0^n(Y)$, and the supremum in (11) thus exists. Since the continuation values are concave, first-order conditions are sufficient for the supremum problem. The following Lemma describes the behavior of the objective function at the boundaries.

Lemma 3 *The objective function $J_0(\alpha)$ can be continuously extended at the boundaries as $\alpha^1 \searrow 0$ or $\alpha^2 \searrow 0$ by the continuation values calculated for the homogeneous economies, i.e., for $\alpha^2 > 0$*

$$J_0\left(0, \alpha^2\right) \doteq \lim_{\alpha^1 \searrow 0} J_0\left(\alpha^1, \alpha^2\right) = \alpha^2 V_0^2(Y) \tag{13}$$

and $\lim_{\alpha^1 \searrow 0} C^2(\alpha^1, \alpha^2) = Y$. The case $\alpha^2 \searrow 0$ is symmetric.

The planner's problem (11-12) suggests that we can interpret the modified discount factor processes $\bar{\lambda}^n$ as stochastic Pareto weights. Indeed, if $\bar{\lambda}_0^n = \alpha^n$ are the initial weights, then $\bar{\lambda}_t^n$ are the consistent state-dependent weights for the continuation problem of the planner at time t.[6,7]

The evolution of the weights involves the drift component ν^n and thus can only be determined in equilibrium unless agent n's preferences are separable, in which case $\nu^n = \beta$. The variation in Pareto weights arises from the interaction of two components in the model — the nonseparable preference structure and the belief distortion that drives the diffusion component in (12).[8]

Observe that the introduction of belief heterogeneity kept the structure of the problem unchanged. For instance, Dumas, Uppal, and Wang (2000) show that in a Markov environment, the discount factor processes λ^n serve as new state variables that allow a recursive formulation of the problem using the Hamilton-Jacobi-Bellman (HJB) equation. The same conclusion is true for the modified discount factor processes $\bar{\lambda}^n$, once belief heterogeneity is incorporated. Belief distortions thus do not introduce any additional state variables into the problem, as long as the distorting processes u^n are functions of the existing state variables.

2.4 Hamilton-Jacobi-Bellman equation

From now on, I assume that both agents have constant belief distortions u^n, a frequently considered case in the survival literature. Extensions involving endogenously determined distortion processes including learning dynamics are considered in Section 5.

The planner's problem has an appealing Markov structure. Homogeneity of the planner's

[6]Similar techniques, which extend the formulation of the representative agent provided by Negishi (1960) to representations with nonconstant Pareto weights, can be used to study models with incomplete markets where changes in the Pareto weights reflect the tightness of the binding constraints. See Cuoco and He (2001) for a general approach in discrete time and Basak and Cuoco (1998) for a model with restricted stock market participation in continuous time.

[7]Jouini and Napp (2007) approach the problem from a different angle to show that a planner's problem formulation with constant Pareto weights is in general not feasible under heterogeneous beliefs. Given an equilibrium with heterogeneous beliefs, they define a hypothetical representative agent with a utility function constructed as a weighted average of individual utility functions, with weights given by the inverses of marginal utilities of wealth. The implied consensus belief of the representative agent that would replicate the equilibrium allocation is not a proper belief but can be decomposed into the product of a proper belief and a discount factor. This discount factor would mimic the dynamics of the Pareto shares in problem (11–12).

[8]The belief heterogeneity introduces an additional risk component arising through the stochastic reweighing of wealth shares. The diffusion component in the weight dynamics will have a direct impact on local risk prices. Notice that other sources of heterogeneity, including differences in risk aversion and IES parameters, do not lead to a diffusion component in the dynamics of the Pareto weights (12), and reweighting therefore has no impact on local risk prices in these cases. Different types of market participation constraints as in Basak and Cuoco (1998) may also introduce a diffusion term into the Pareto weight dynamics.

problem (11-12) in $\left(\bar{\lambda}^1, \bar{\lambda}^2\right)$ suggests a transformation of variables

$$\theta^1 = \bar{\lambda}^1 \left(\bar{\lambda}^1 + \bar{\lambda}^2\right)^{-1} \qquad \theta^2 = \bar{\lambda}^1 + \bar{\lambda}^2. \tag{14}$$

The single state variable θ^1 represents the Pareto share of agent 1. The dynamics of θ^1 are central to the study of survival in this paper. Obviously, θ^1 is bounded between zero and one. It will become clear that for strictly positive initial weights, the boundaries are unattainable, so that θ^1 evolves on the open interval $(0,1)$. Since the objective function of the planner is also homogeneous degree γ in Y, the planner's problem can be characterized as a solution to an ordinary differential equation with a single state variable θ^1.

Proposition 4 *The objective function for the planner's problem (11-12) is*

$$J_0(\alpha) = \left(\alpha^1 + \alpha^2\right) \gamma^{-1} Y_0^{\gamma} \tilde{J}\left(\alpha^1 / \left(\alpha^1 + \alpha^2\right)\right),$$

where $\tilde{J}\left(\theta^1\right)$ is the solution to the nonlinear ordinary differential equation

$$\begin{aligned} 0 \;=\; & \theta^1 \frac{\beta}{\rho} \left(\zeta^1\right)^{\rho} \left(\tilde{J}^1\right)^{1-\frac{\rho}{\gamma}} + \left(1-\theta^1\right) \frac{\beta}{\rho} \left(1-\zeta^1\right)^{\rho} \left(\tilde{J}^2\right)^{1-\frac{\rho}{\gamma}} + \\ & + \left(-\frac{\beta}{\rho} + \mu_y + \left(\theta^1 u^1 + \left(1-\theta^1\right) u^2\right) \sigma_y + \frac{1}{2}\left(\gamma - 1\right) \sigma_y^2\right) \tilde{J} + \\ & + \theta^1 \left(1-\theta^1\right) \left(u^1 - u^2\right) \sigma_y \tilde{J}_{\theta^1} + \frac{1}{2}\frac{1}{\gamma} \left(1-\theta^1\right)^2 \left(\theta^1\right)^2 \left(u^1 - u^2\right)^2 \tilde{J}_{\theta^1\theta^1} \end{aligned} \tag{15}$$

with boundary conditions $\tilde{J}(0) = \tilde{V}^2$ and $\tilde{J}(1) = \tilde{V}^1$, where $\tilde{V}^n$ are defined in (7). The functions $\tilde{J}^n\left(\theta^1\right)$ are the continuation values of the two agents scaled by $\gamma^{-1} Y^{\gamma}$,

$$\begin{aligned} \tilde{J}^1\left(\theta^1\right) &\doteq \tilde{J}\left(\theta^1\right) + \left(1-\theta^1\right) \tilde{J}_{\theta^1}\left(\theta^1\right) \\ \tilde{J}^2\left(\theta^1\right) &\doteq \tilde{J}\left(\theta^1\right) - \theta^1 \tilde{J}_{\theta^1}\left(\theta^1\right). \end{aligned} \tag{16}$$

and the consumption share ζ^1 is given by

$$\zeta^1\left(\theta^1\right) = \frac{\left(\theta^1\right)^{\frac{1}{1-\rho}} \left[\tilde{J}^1\left(\theta^1\right)\right]^{\frac{1-\rho/\gamma}{1-\rho}}}{\left(\theta^1\right)^{\frac{1}{1-\rho}} \left[\tilde{J}^1\left(\theta^1\right)\right]^{\frac{1-\rho/\gamma}{1-\rho}} + \left(1-\theta^1\right)^{\frac{1}{1-\rho}} \left[\tilde{J}^2\left(\theta^1\right)\right]^{\frac{1-\rho/\gamma}{1-\rho}}}. \tag{17}$$

Unfortunately, equation (15) does not have a general closed-form solution. However, the Pareto share θ^1 of agent 1 remains the only state variable. This considerably simplifies numerical solutions, and, more importantly, allows one to formulate the survival problem in terms of the boundary behavior of a scalar Itô process. Despite the nonexistence of a closed-form solution for $\tilde{J}(\theta^1)$, this boundary behavior can be characterized analytically by studying the limiting behavior of the objective function.

Equation (15) is not specific to the planner's problem (11-12). For instance, Gârleanu and Panageas (2010) use the martingale approach to directly analyze the equilibrium in an economy with agents endowed with heterogeneous recursive preferences, and show that they can derive their asset pricing formulas in closed form up to the solution of a nonlinear ODE that has the same structure as (15), which they have to solve for numerically. The analytical characterization of the boundary behavior of the ODE derived in this paper is thus applicable to a wider class of recursive utility models, and can aid numerical calculations which are often unstable in the neighborhood of the boundaries in this type of problems.

3 Survival

Survival chances of agents with distorted beliefs have been studied extensively under separable utility. Kogan, Ross, Wang, and Westerfield (2009) show a tight link between the behavior of the belief ratio, consumption shares, and the risk aversion coefficient as a measure of curvature of the utility function. To provide a simple illustration, consider a period utility function $U(C)$ and the corresponding Euler equation that prices a payoff Z_{t+s} at time t

$$P_t^z = E_t^{Q_n}\left[e^{-\beta s}\frac{U'\left(C_{t+s}^n\right)}{U'\left(C_t^n\right)}Z_{t+s}\right] = E_t\left[e^{-\beta s}\frac{U'\left(C_{t+s}^n\right)}{U'\left(C_t^n\right)}\frac{M_{t+s}^n}{M_t^n}Z_{t+s}\right].$$

Since prices are observed in equilibrium, agents have to agree on them. When markets are complete, the objects

$$\frac{U'\left(C_{t+s}^n\right)}{U'\left(C_t^n\right)}\frac{M_{t+s}^n}{M_t^n}$$

have to be equalized across agents n, and deviations in beliefs have to be offset by reciprocal deviations in marginal utilities. Survival analysis thus corresponds to analyzing a sequence of state- and time-indexed static problems that are interlinked only by the initial relative marginal utilities of wealth of the two agents, whose choice is largely innocuous for the long-run results. If agent 1 has a constant belief distortion $u^1 \neq 0$ and agent 2 is rational, then M^1 is a strictly positive supermartingale and $\lim_{s\to\infty} M_{t+s}^1 = 0$ (P-a.s.), and thus $\lim_{s\to\infty} U'\left(C_{t+s}^1\right)/U'\left(C_{t+s}^2\right) = +\infty$ (P-a.s.). For a class of utility functions that includes the

CRRA utility (the special case when $\gamma = \rho$ in this paper), this implies $\lim_{s\to\infty} \zeta^1_{t+s}/\zeta^2_{t+s} = 0$ (P-a.s.).

When preferences are not separable, this straightforward link breaks down because marginal utilities also depend on continuation values and the stochastic discount factor involves the evolution of the endogenously determined discount rate process ν^n between t and $t+s$. Since these continuation values and discount rate processes are not available in closed form, they have to in general be solved for numerically.

I show in this section that in order to evaluate the survival chances of individual agents, a complete solution for the consumption allocation, continuation values, and the implied discount rate processes is not necessary. In fact, it is sufficient to characterize the wealth dynamics in the limiting cases when the wealth share of one of the agents becomes negligible, and this limiting behavior can be solved for in closed form. This characterization of survival requires taking an approach that is different from the majority of the literature, which typically analyzes the global properties of relative entropy as a measure of disparity between subjective beliefs and the true probability distribution, and its convergence as $t \nearrow \infty$.

Instead, I derive the local dynamics of the Pareto share θ^1 and rely on its ergodic properties, which allow me to investigate the existence of a unique stationary distribution for θ^1 that is closely related to survival. The derived sufficient conditions are tightly linked to the behavior of the difference of endogenous discount rates of the two agents. In a decentralized economy, these *relative patience* conditions can be reinterpreted in terms of the difference in expected logarithmic growth rates of individual wealth.

Since the analyzed model includes growing and decaying economies, I am interested in a measure of relative survival. The following definition distinguishes between survival along individual paths and almost-sure survival.

Definition 5 *Agent 1 becomes extinct along the path* $\omega \in \Omega$ *if* $\lim_{t\to\infty} \theta^1_t(\omega) = 0$. *Otherwise, agent 1 survives along the path* ω. *Agent 1 dominates in the long run along the path* ω *if* $\lim_{t\to\infty} \theta^1_t(\omega) = 1$.

Agent 1 becomes extinct (under measure P*) if* $\lim_{t\to\infty} \theta^1_t = 0$, *$P$-a.s. Agent 1 survives if* $\limsup_{t\to\infty} \theta^1_t > 0$, *$P$-a.s. Agent 1 dominates in the long run if* $\lim_{t\to\infty} \theta^1_t = 1$, *$P$-a.s.*

Kogan, Ross, Wang, and Westerfield (2009) or Yan (2008) use the consumption share ζ^1 as a measure of survival. Since the consumption share (17) is continuous and strictly increasing in θ^1 and the limits are $\lim_{\theta^1\searrow 0} \zeta^1(\theta^1) = 0$ and $\lim_{\theta^1\nearrow 1} \zeta^1(\theta^1) = 1$, the two measures are equivalent in this setting.

3.1 Dynamics of the Pareto share and long-run distributions

Recall the dynamics of the modified discount factor processes $\bar{\lambda}^n$ in (12). An application of Itô's lemma to $\theta^1 = \bar{\lambda}^1 / \left(\bar{\lambda}^1 + \bar{\lambda}^2\right)$ yields

$$\begin{aligned} \frac{d\theta_t^1}{\theta_t^1} &= \left(1-\theta_t^1\right)\left[\nu_t^2 - \nu_t^1 + \left(\theta_t^1 u^1 + \left(1-\theta_t^1\right) u^2\right)\left(u^2 - u^1\right)\right] dt + \\ &\quad + \left(1-\theta_t^1\right)\left(u^1 - u^2\right) dW_t. \end{aligned} \tag{18}$$

Both heterogeneous beliefs and heterogeneous recursive preferences lead to nonconstant dynamics of the Pareto share, although with different implications. Under nonseparability, preference heterogeneity induces a smooth evolution of the Pareto weights, while belief heterogeneity leads to dynamics with a nonzero volatility term. Identical belief distortions ($u^1 = u^2$) under separable preferences with identical time preference coefficients or under identical recursive preferences imply a constant Pareto share $\theta_t^1 \equiv \alpha^1 / \left(\alpha^1 + \alpha^2\right)$. In what follows, I abstract from this situation, and assume $u^1 \neq u^2$.

Under nonseparable preferences, the discount rates are determined endogenously in the model as a solution to problem (11–12) and are given by

$$\nu^n\left(\theta^1\right) = \frac{\beta}{\rho}\left(\gamma + \left(\rho - \gamma\right)\left(\frac{\zeta^n\left(\theta^1\right)}{\tilde{J}^n\left(\theta^1\right)^{1/\gamma}}\right)^{\rho}\right). \tag{19}$$

The discount rates ν^n are twice continuously differentiable functions of the state variable θ^1, and thus θ^1 is an Itô process on the open interval $(0,1)$ with continuous drift and volatility coefficients.[9] Intuitively, one would expect a stationary distribution for θ^1 to exist if the process exhibits sufficient pull toward the center of the interval when close to the boundaries. This is formalized in the following Proposition:

Proposition 6 *Define the following 'repealing' conditions (i) and (ii), and their 'attracting' counterparts (i') and (ii').*

(i) $\lim_{\theta^1 \searrow 0}\left[\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)\right] > \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]$ **(i')** $<$

(ii) $\lim_{\theta^1 \nearrow 1}\left[\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)\right] < \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]$ **(ii')** $>$

Then the following statements are true:

(a) *If conditions (i) and (ii) hold, then both agents survive under P.*

[9] The unattainability of the boundaries follows from the proof of Proposition 6.

(b) *If conditions (i) and (ii') hold, then agent 1 dominates in the long run under P*

(c) *If conditions (i') and (ii) hold, then agent 2 dominates in the long run under P.*

(d) *If conditions (i') and (ii') hold, then there exist sets $S^1, S^2 \subset \Omega$ which satisfy*

$$S^1 \cap S^2 = \emptyset, \quad P\left(S^1\right) \neq 0 \neq P\left(S^2\right), \quad \textit{and} \quad P\left(S^1 \cup S^2\right) = 1$$

such that agent 1 dominates in the long run along each path $\omega \in S^1$ and agent 2 dominates in the long run along each path $\omega \in S^2$.

The conditions are also the least tight bounds of this type.

Given the dynamics of the Pareto share (18), conditions (i) and (ii) are jointly sufficient for the existence of a unique stationary density $q(\theta^1)$. The proof of Proposition 6 is based on the classification of boundary behavior of diffusion processes, discussed in Karlin and Taylor (1981). The four 'attracting' and 'repealing' conditions are only sufficient and their combinations stated in Proposition 6 are not exhaustive. There are delicate cases involving equalities, which are however only of limited importance in the analysis below.

I call the difference in the discount rates $\nu^2(\theta^1) - \nu^1(\theta^1)$ *relative patience* because it captures the difference in discounting of future felicity in the variational utility specification (3) between the two agents. Conditions in Proposition 6 have an intuitive interpretation. Survival condition (i) states that agent 1 survives under the true probability measure even in cases when his beliefs are more distorted, $|u^1| > |u^2|$, as long as his relative patience becomes sufficiently high to overcome the distortion when his Pareto share vanishes.

Lucas and Stokey (1984) impose a similar condition called increasing marginal impatience that is sufficient to guarantee the existence of a nondegenerate steady state as an exogenous restriction on the preference specification. This condition requires the preferences in their framework to be nonhomothetic, and rich agents must discount future more than poor ones. In this model, preferences are homothetic, and variation in relative patience arises purely as a response to the market interaction of the two agents endowed with heterogeneous beliefs. The discount rate ν^n encodes not only a pure time preference but also an interaction through the dynamics of the optimal consumption stream.

3.2 CRRA preferences

The framework introduced in this paper includes as a special case the separable constant relative risk aversion preferences when $\gamma = \rho$. Yan (2008) and Kogan, Ross, Wang, and

Westerfield (2009) show that in the economy presented in this paper, the agent whose beliefs are less distorted dominates in the long run under measure P. The conditions in Proposition 6 confirm these results as follows:

Corollary 7 *Under separable CRRA preferences ($\gamma = \rho$), agent n dominates in the long run under measure P if and only if $|u^n| < |u^{\sim n}|$. Agent n survives under P if and only if the inequality is non-strict. Further, agent n survives under measure Q^n and dominates in the long run under Q^n if and only if $u^n \neq u^{\sim n}$.*

Under separable CRRA preferences, the dynamics of the Pareto share (18) do not depend on the characteristics of the endowment process. The survival result in Corollary 7 thus extends to an arbitrary adapted aggregate endowment process Y that satisfies elementary integrability conditions, including a constant one, as long as the two agents can write contracts on the realizations of the Brownian motion W. It is a special case of the analysis in Kogan, Ross, Wang, and Westerfield (2009), who show that this survival result holds, under mild conditions, for any separable preferences with bounded relative risk aversion. In this sense, the separable environment generates a robust result about the extinction of agents whose beliefs are relatively inaccurate.

A specific situation in Corollary 7 arises when $u^n = -u^{\sim n}$. The proof of the corollary shows that although none of the agents becomes extinct, a nondegenerate long-run distribution for θ^1 does not exist.

3.3 The nonseparable case

When preferences are not separable, consumption choices across periods are interlinked through the endogenously determined discount rate processes ν^n, which opens another channel for intertemporal tradeoff and thus potential survival. This endogenous discounting is reflected in the evolution of the Pareto share θ^1. In this section, I derive closed-form formulas for the boundary behavior of ν^n, and evaluate analytically the region in the parameter space in which the conditions of Proposition 6 hold.

The proof strategy in this section relies on a decentralization argument and utilizes the asymptotic properties of the differential equation (15) for the planner's continuation value. The economy is driven by a single Brownian shock, and two suitably chosen assets that can be continuously traded are therefore sufficient to complete the markets in the sense of Harrison and Kreps (1979). Let the two traded assets be an infinitesimal risk-free bond in zero net supply that yields a risk-free rate $r_t = r(\theta_t^1)$ and a claim on the aggregate endowment with price $\Xi_t = Y_t \xi\left(\theta_t^1\right)$, where $\xi\left(\theta^1\right)$ is the aggregate wealth-consumption ratio. Individual wealth levels are denoted $\Xi_t^n = Y_t \zeta^n\left(\theta_t^1\right) \xi^n\left(\theta_t^1\right)$, where $\xi^n\left(\theta^1\right)$ are the individual wealth-consumption ratios.

The results reveal that as the Pareto share of one of the agents converges to zero, the infinitesimal returns associated with the two assets converge to those which prevail in a homogeneous economy populated by the agent with the large Pareto share. Solving the individual optimization problems yields the required limits for the discount rates ν^n. The limiting problems correspond to the analysis of homogeneous economies for which analytical solutions exist, and thus the limits for ν^n are also available in closed form.

The proof also shows that the conditions on the limiting behavior of the discount rates in Proposition 6 that assure the existence of a nondegenerate long-run equilibrium can be directly restated as conditions on the limiting expected growth rates of the logarithm of individual wealth levels in a decentralized economy.

3.3.1 Equilibrium prices

Homotheticity of preferences implies that individual wealth-consumption ratios are given by

$$\xi^n\left(\theta^1\right) = \frac{1}{\beta}\left(\frac{\tilde{J}^n\left(\theta^1\right)^{1/\gamma}}{\zeta^n\left(\theta^1\right)}\right)^{\rho}. \tag{20}$$

I start by assuming that $\xi^n\left(\theta^1\right)$ are functions that are bounded and bounded away from zero. This, among other things, implies that the discount rate functions $\nu^n\left(\theta^1\right)$ in (19) are bounded and that the drift and volatility coefficients in the stochastic differential equation for θ^1, (18), are bounded as well. The assumption will ultimately be verified by a direct calculation of the limits of $\xi^n\left(\theta^1\right)$ as $\theta^1 \searrow 0$ or $\theta^1 \nearrow 1$.

Without loss of generality, it is sufficient to focus on the case $\theta^1 \searrow 0$. First notice some asymptotic results for the planner's continuation value $\tilde{J}\left(\theta^1\right)$.

Lemma 8 *The solution of the planner's problem implies that*

$$\lim_{\theta^1 \searrow 0} \theta^1 \tilde{J}_{\theta^1}\left(\theta^1\right) = \lim_{\theta^1 \searrow 0} \left(\theta^1\right)^2 \tilde{J}_{\theta^1\theta^1}\left(\theta^1\right) = \lim_{\theta^1 \searrow 0} \left(\theta^1\right)^3 \tilde{J}_{\theta^1\theta^1\theta^1}\left(\theta^1\right) = 0.$$

The Markov structure of the problem implies that the evolution of the continuation values and consumption shares can be written as

$$\frac{d\tilde{J}^n\left(\theta_t^1\right)}{\tilde{J}^n\left(\theta_t^1\right)} = \mu_{\tilde{J}^n}\left(\theta_t^1\right)dt + \sigma_{\tilde{J}^n}\left(\theta_t^1\right)dW_t \tag{21}$$

$$\frac{d\zeta^n\left(\theta_t^1\right)}{\zeta^n\left(\theta_t^1\right)} = \mu_{\zeta^n}\left(\theta_t^1\right)dt + \sigma_{\zeta^n}\left(\theta_t^1\right)dW_t, \tag{22}$$

where the drift and volatility coefficients are functions of θ^1, and the results from Lemma 8 allow the characterization of their limiting behavior.

Lemma 9 *The coefficients in equations (21–22) for agent 2 satisfy*

$$\lim_{\theta^1\searrow 0}\mu_{\tilde{J}^2}\left(\theta^1\right) = \lim_{\theta^1\searrow 0}\sigma_{\tilde{J}^2}\left(\theta^1\right) = \lim_{\theta^1\searrow 0}\mu_{\zeta^2}\left(\theta^1\right) = \lim_{\theta^1\searrow 0}\sigma_{\zeta^2}\left(\theta^1\right) = 0. \tag{23}$$

The result follows from an application of Itô's lemma to $\tilde{J}^2$ and ζ^2. Utilizing formulas (16) and (17), the coefficients will contain expressions for the value function $\tilde{J}\left(\theta^1\right)$ and its partial derivatives up to the third order, and all the expressions can be shown to converge to zero using Lemma 8.

Using the construction from Duffie and Epstein (1992a), the stochastic discount factor process for agent 2 under the subjective probability measure Q^2 is given by

$$S_t^2 = \exp\left(-\int_0^t \nu^2\left(\theta_s^1\right)ds\right)\left(\frac{Y_t}{Y_0}\right)^{\gamma-1}\left(\frac{\zeta^2\left(\theta_t^1\right)}{\zeta^2\left(\theta_0^1\right)}\right)^{\rho-1}\left(\frac{\tilde{J}^2\left(\theta_t^1\right)}{\tilde{J}^2\left(\theta_0^1\right)}\right)^{1-\frac{\rho}{\gamma}}. \tag{24}$$

Since $\lim_{\theta^1\searrow 0}\nu^2\left(\theta^1\right) = \nu^2$, which is given in (8), and Lemma 9 states that the local drift and volatility of the last two terms decline to zero as $\theta^1 \searrow 0$, the infinitesimal risk-free rate and the local price of risk converge to their homogeneous economy counterparts. Moreover, the price of aggregate endowment Ξ converges as well, and so does the local return on aggregate wealth. The following Proposition summarizes the limiting pricing implications.

Proposition 10 *As $\theta^1 \searrow 0$, the infinitesimal risk-free rate $r\left(\theta^1\right)$, the aggregate wealth-consumption ratio $\xi\left(\theta^1\right)$, and the drift and volatility coefficients of the aggregate wealth pro-*

cess $d\Xi_t/\Xi_t = \mu_\Xi(\theta^1_t)\,dt + \sigma_\Xi(\theta^1_t)\,dt$ *converge to their homogeneous economy counterparts:*

$$\begin{aligned}
\lim_{\theta^1 \searrow 0} r\left(\theta^1\right) &= r(0) = \beta + (1-\rho)\left(\mu_y + u^2\sigma_y\right) - \frac{1}{2}(2-\rho)(1-\gamma)\sigma_y^2, \\
\lim_{\theta^1 \searrow 0} \xi\left(\theta^1\right) &= \xi(0) = \left[\beta - \rho\left(\mu_y + u^2\sigma_y - \frac{1}{2}(1-\gamma)\sigma_y^2\right)\right]^{-1}, \\
\lim_{\theta^1 \searrow 0} \mu_\Xi\left(\theta^1\right) &= \mu_y, \quad \textit{and} \quad \lim_{\theta^1 \searrow 0} \sigma_\Xi\left(\theta^1\right) = \sigma_y.
\end{aligned}$$

Consequently, the infinitesimal return on the claim on aggregate wealth,

$$\left[\left[\xi\left(\theta^1_t\right)\right]^{-1} + \mu_\Xi\left(\theta^1_t\right)\right]dt + \sigma_\Xi\left(\theta^1_t\right)dW_t, \tag{25}$$

has coefficients that converge as well.

Notice that the convergence of the coefficients of the wealth process is not an immediate consequence of the convergence of the aggregate wealth-consumption ratio. It may be that the wealth-consumption ratio $\xi(\theta^1)$ converges as $\theta^1 \searrow 0$, yet its price dynamics are such that $\mu_\Xi(\theta^1)$ and $\sigma_\Xi(\theta^1)$ do not converge to μ_y and σ_y, respectively. The fact that this does not happen is closely linked to the dynamics of $\log\theta^1$. The bounded drift and volatility coefficient of $\log\theta^1$ assure that the local variation in $\xi(\theta^1)$ becomes irrelevant as $\log\theta^1 \searrow -\infty$.

The results in Proposition 10 are sufficient to proceed with the construction of the main result. As a side note, prices of finite-horizon risk-free claims and individual cash flows from the aggregate endowment converge as well:

Corollary 11 *For every fixed maturity t, prices of a zero-coupon bond and a claim to a payout from the aggregate endowment stream converge to their homogeneous economy counterparts as $\theta^1 \searrow 0$.*

3.3.2 Decision problem of an agent with negligible wealth

I have now established that the actual *general equilibrium* price dynamics that agent 1 with infinitesimal wealth takes as given when solving his portfolio problem are locally the same as those in an economy populated only by agent 2. However, the construction of the main result is not completed yet. The marginal utility of agent 1 is forward looking due to the nonseparable nature of the preferences, and will thus depend on agent's 1 continuation value. It remains to show that the continuation value of agent 1 converges as well.

Agent 1, whose wealth Ξ^1 is close to zero, solves

$$\bar{\lambda}_t^1 V_t^1 = \max_{C^1,\pi^1,\nu^1} E_t \left[\int_t^\infty \bar{\lambda}_s^1 F\left(C_s^1, \nu_s^1\right) ds \right]$$

subject to (6) and the budget constraint,

$$\frac{d\Xi_t^1}{\Xi_t^1} = \left[r\left(\theta_t^1\right) + \pi_t^1 \left(\left[\xi\left(\theta_t^1\right)\right]^{-1} + \mu_\Xi\left(\theta_t^1\right) - r\left(\theta_t^1\right) \right) - \frac{C_t^1}{\Xi_t^1} \right] dt + \pi_t^1 \sigma_\Xi\left(\theta_t^1\right) dW_t, \tag{26}$$

where π^1 is the portfolio share invested in the stock. The homogeneity of the problem motivates the guess

$$\gamma V_t^1 = \left(\Xi_t^1\right)^\gamma \hat{V}^1\left(\theta_t^1\right). \tag{27}$$

The drift and volatility coefficients depend explicitly on θ^1 because Ξ^1 and θ^1 are linked through

$$\Xi_t^1 = Y_t \zeta^1\left(\theta_t^1\right) \beta^{-\frac{1}{1-\rho}} \left[\hat{V}^1(\theta)\right]^{\frac{\rho}{\gamma}\frac{1}{1-\rho}}. \tag{28}$$

Recall that we are interested in the characterization of the limiting solution as $\theta^1 \searrow 0$. The associated HJB equation leads to a second-order ODE (omitting dependence on θ^1)

$$\begin{aligned} 0 \;=\; & \max_{C^1,\pi^1,\nu^1} \frac{1}{\rho}\beta^{\frac{1}{1-\rho}}\left(\hat{V}^1\right)^{1-\frac{\rho}{\gamma}\frac{1}{1-\rho}} + \hat{V}^1\left(-\frac{\beta}{\rho} + \mu_{\Xi^1} + u^1\sigma_{\Xi^1} - \frac{1}{2}(1-\gamma)\left(\sigma_{\Xi^1}\right)^2\right) + \\ & + \hat{V}_{\theta^1}^1 \theta^1 \left(\frac{1}{\gamma}\left(\mu_{\theta^1} + u^1\sigma_{\theta^1}\right) + \sigma_{\theta^1}\sigma_{\Xi^1}\right) + \hat{V}_{\theta^1\theta^1}^1 \left(\theta^1\right)^2 \frac{1}{2}\frac{1}{\gamma}\left(\sigma_{\theta^1}\right)^2, \end{aligned} \tag{29}$$

which yields the first-order conditions on C_t^1 and π_t^1:

$$\begin{aligned} \frac{C_t^1}{\Xi_t^1} &= \beta^{\frac{1}{1-\rho}}\left(\hat{V}^1\left(\theta_t^1\right)\right)^{-\frac{\rho}{\gamma}\frac{1}{1-\rho}} \\ \pi_t^1 &= \frac{\left[\xi\left(\theta_t^1\right)\right]^{-1} + \mu_\Xi\left(\theta_t^1\right) + u_1\sigma_\Xi\left(\theta_t^1\right) - r\left(\theta_t^1\right) + \frac{\theta \hat{V}_{\theta^1}^1\left(\theta^1\right)}{\hat{V}^1\left(\theta^1\right)}\sigma_{\theta^1}\left(\theta_t^1\right)\sigma_{\Xi^1}\left(\theta_t^1\right)}{(1-\gamma)\left(\sigma_\Xi\left(\theta_t^1\right)\right)^2}, \end{aligned} \tag{30}$$

where μ_{Ξ^1} and σ_{Ξ^1} are the drift and volatility coefficients on the right-hand side of (26), and μ_{θ^1} and σ_{θ^1} are the coefficients associated with the evolution of $d\theta_t^1/\theta_t^1$. Notice that the portfolio choice π^1 almost corresponds to the standard Merton (1971) result, except the last term in the numerator which explicitly takes into account the covariance between agent's 1 wealth and the evolution in the state variable θ^1 imposed by (28).

The solution of this equation determines the consumption-wealth ratio of agent 1 and, consequently, the evolution of his wealth. While a closed-form solution of this equation is not available, it is again possible to characterize the asymptotic behavior as $\theta^1 \searrow 0$.

Lemma 12 *The following results hold:*

$$\lim_{\theta^1 \searrow 0} \theta^1 \hat{V}^1_{\theta^1}\left(\theta^1\right) = \lim_{\theta^1 \searrow 0} \left(\theta^1\right)^2 \hat{V}^1_{\theta^1\theta^1}\left(\theta^1\right) = 0.$$

These results are similar to those in Lemma 8. They imply that the derivative terms in the ODE (29) vanish as $\theta^1 \searrow 0$, and we obtain the limit for $\hat{V}^1\left(\theta^1\right)$ and the evolution of Ξ^1 in closed form.

Proposition 13 *The consumption-wealth ratio of agent 1 converges to*

$$\begin{aligned} \lim_{\theta^1 \searrow 0} \beta^{\frac{1}{1-\rho}} \left(\hat{V}^1\left(\theta^1\right)\right)^{-\frac{\rho}{\gamma}\frac{1}{1-\rho}} &= \beta - \rho\left(\mu_y + u^2\sigma_y - \frac{1}{2}\left(1-\gamma\right)\left(\sigma_y\right)^2\right) - \\ &\quad -\frac{\rho}{1-\rho}\left[\left(u^1 - u^2\right)\sigma_y + \frac{1}{2}\frac{\left(u^1-u^2\right)^2}{1-\gamma}\right] \end{aligned} \tag{31}$$

and the wealth share invested into the claim on aggregate consumption to

$$\lim_{\theta^1 \searrow 0} \pi^1\left(\theta^1\right) = 1 + \frac{u^1 - u^2}{\left(1-\gamma\right)\sigma_y}. \tag{32}$$

It follows that the asymptotic coefficients for the evolution of agent's 1 wealth are

$$\begin{aligned} \lim_{\theta^1 \searrow 0} \mu_{\Xi^1}\left(\theta^1\right) &= \mu_y + \frac{1}{1-\rho}\left(u_1 - u_2\right)\sigma_y + \frac{1}{2}\frac{2-\rho}{1-\rho}\frac{\left(u^1-u^2\right)^2}{1-\gamma} - \frac{u^1\left(u^1-u^2\right)}{1-\gamma} \\ \lim_{\theta^1 \searrow 0} \sigma_{\Xi^1}\left(\theta^1\right) &= \sigma_y + \frac{u^1-u^2}{\left(1-\gamma\right)}. \end{aligned}$$

Naturally, the wealth evolution must track the evolution of the aggregate endowment when $u^1 = u^2$.

3.3.3 Limiting relative patience and relationship to wealth growth

Importantly, these results allow one to calculate the limiting discount rate $\nu^1\left(\theta^1\right)$ and state the main result of this section.

Proposition 14 *The expressions for the limiting behavior of the relative patience in Proposition 6 are*

$$\lim_{\theta^1 \searrow 0} \nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right) = \frac{\rho-\gamma}{1-\rho}\left[\left(u^1-u^2\right)\sigma_y + \frac{1}{2}\frac{\left(u^1-u^2\right)^2}{1-\gamma}\right], \tag{33}$$

$$\lim_{\theta^1 \nearrow 1} \nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right) = \frac{\rho-\gamma}{1-\rho}\left[\left(u^1-u^2\right)\sigma_y - \frac{1}{2}\frac{\left(u^1-u^2\right)^2}{1-\gamma}\right]. \tag{34}$$

Section 3.4 discusses which regions of the parameter space satisfy the individual survival and extinction conditions from Proposition 6. It remains for me to verify that the assumption about the boundedness of wealth consumption ratios indeed holds.

Corollary 15 *Under parameter restrictions in Assumption 2, the wealth-consumption ratios are bounded and bounded away from zero.*

Notice that while Assumption 2 imposes a restriction on the time preference parameter β of the agents, the survival conditions do not explicitly depend on β. The survival results thus always hold with the implicit assumption that time discounting is sufficiently large.

The construction of the main survival result utilized the link between the planner's problem and the competitive equilibrium. It turns out that relative patience conditions that assure survival can be restated as conditions on the relative growth rates of individual wealth.

Corollary 16 *The survival conditions in part a) of Proposition 6 are equivalent to:*

(i) $\lim_{\theta^1 \searrow 0} \mu_{\Xi^1}\left(\theta^1\right) - \frac{1}{2}\left[\sigma_{\Xi^1}\left(\theta^1\right)\right]^2 > \lim_{\theta^1 \searrow 0} \mu_{\Xi^2}\left(\theta^1\right) - \frac{1}{2}\left[\sigma_{\Xi^2}\left(\theta^1\right)\right]^2,$

(ii) $\lim_{\theta^1 \nearrow 1} \mu_{\Xi^1}\left(\theta^1\right) - \frac{1}{2}\left[\sigma_{\Xi^1}\left(\theta^1\right)\right]^2 < \lim_{\theta^1 \nearrow 1} \mu_{\Xi^2}\left(\theta^1\right) - \frac{1}{2}\left[\sigma_{\Xi^2}\left(\theta^1\right)\right]^2.$

Verifying the conditions in Proposition 6 therefore amounts to checking that the expected growth rate of the logarithm of wealth is higher for the agent who is at the brink of extinction. This is of course a natural condition for a competitive equilibrium, but the planner's problem was still required to determine these limiting growth rates.

3.4 Limiting behavior of asset prices

The results in the previous section establish that, as the Pareto share of one of the agents becomes negligible, agents make their portfolio and consumption-saving decisions as if they observed prices in a homogeneous economy populated only by the large agent. If there exists a nondegenerate stationary distribution of the Pareto share θ^1, then an agent will always have a nontrivial price impact in the future, even if his current Pareto share is negligible. The forward looking nature of the optimization problem then implies that this price impact should be taken into account when making current decisions. The results show that as the Pareto share of the agent vanishes, the time when his price impact becomes relevant is so distant that it is immaterial for current decisions.

The logic manifests itself in the behavior of the last term in the numerator of the portfolio share π^1 in equation (30). This term explicitly takes into account agent 1's knowledge about the impact of his portfolio decision on equilibrium prices. Since this term vanishes as $\theta^1 \searrow 0$, the agent understands that asymptotically the portfolio decisions made by agents of his type will not have any impact on local equilibrium price dynamics, and thus behaves as if he resided in an economy populated only by agent 2.

This implies that the survival question, whose answer only depends on the behavior at the boundaries, can be resolved by studying homogeneous economies with an infinitesimal price-taking agent. Even if the agent survives with probability one and has an impact on equilibrium prices in the long run, he does not influence *current* prices and returns.

The dynamics of the Pareto share (18) that has bounded drift and volatility coefficients is also informative about experiments that 'inject' infinitesimal heterogeneous agents into an initially homogeneous economy. By making the initial θ_0^1 arbitrarily close to zero, one can extend the time before the presence of the new agent becomes noticeable (measured, e.g., by sufficiently large deviations in prices or return distributions from their homogeneous economy counterparts) arbitrarily far into the future.

3.5 Survival regions

This section analyzes the regions of the parameter space in which agents with distorted beliefs survive or dominate the economy. It turns out that all four combinations outlined in Proposition 6 can occur, and Figure 1 visualizes the survival regions. Each panel fixes the belief distortions (u^1, u^2) and the volatility of aggregate endowment σ_y, and plots the regions in the risk aversion / inverse of IES $(1-\gamma, 1-\rho)$ plane. The results do not reveal what happens at the boundaries of the regions where conditions from Proposition 6 hold with equalities, but the survival characteristics inside the individual regions are well-defined.

3.5.1 Asymptotic results

It is useful to describe the asymptotic results as either risk aversion or intertemporal elasticity of substitution moves toward extreme values, holding the other parameters fixed.

Corollary 17 *Holding other parameters fixed, the survival restrictions imply the following asymptotic results.*

(a) *As risk aversion increases ($\gamma \searrow -\infty$), the agent who is relatively more optimistic about the growth rate of aggregate endowment always dominates in the long run.*

(b) *As agents become risk neutral ($\gamma \nearrow 1$), each agent dominates in the long run with a strictly positive probability.*

(c) *As intertemporal elasticity of substitution increases ($\rho \nearrow 1$), the relatively more optimistic agent always survives. The relatively more pessimistic agent survives (and thus a nondegenerate long-run equilibrium exists) when risk aversion is sufficiently small.*

(d) *As the intertemporal elasticity of substitution decreases to zero ($\rho \searrow -\infty$), a nondegenerate long-run equilibrium cannot exist.*

In order to shed more light on the results, it is useful to consider the discrete-time version of the recursive preferences, featured by Epstein and Zin (1989):

$$\tilde{V}_t = \left[\left(1 - e^{-\beta}\right)(C_t)^\rho + e^{-\beta}\left(E_t^Q\left[\left(\tilde{V}_{t+1}\right)^\gamma\right]\right)^{\frac{\rho}{\gamma}}\right]^{\frac{1}{\rho}}.$$

The risk aversion parameter γ drives the risk adjustment of the next-period continuation value $\tilde{V}_{t+1}$; and as risk aversion increases, the lower tail of the distribution of $\tilde{V}_{t+1}$ will contribute with an increasingly larger penalty to the expected value. When two agents differ in their beliefs, the more pessimistic agent assigns a higher probability to the tail events. Since the penalty increases as $\gamma \searrow -\infty$, he is willing to sacrifice an increasingly large amount of wealth in the remaining states in order to insure against the low-probability tail event that will ultimately drive him to extinction.

In the other extreme, when $\gamma \nearrow 1$ and agents become risk neutral with respect to intratemporal gambles, all that matters in the next period's contribution is the expected continuation value. Depending on the current distribution of wealth, there will be a belief-ratio threshold that divides the realizations of the shock tomorrow into two sets, where in

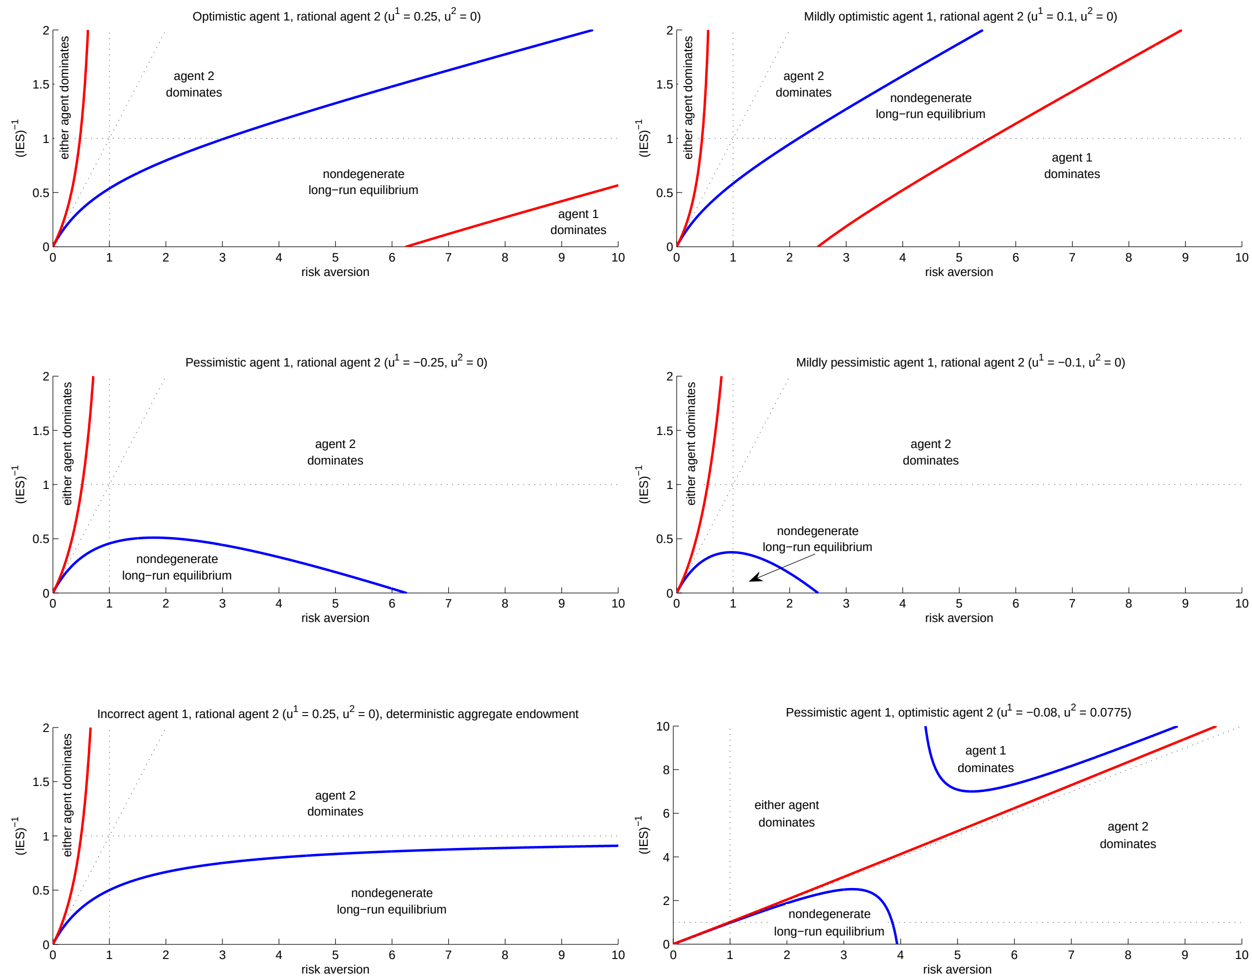

Figure 1: Survival regions for different parameterizations. All panels except the bottom left panel assume $\sigma_y = 0.02$. In the bottom left panel, aggregate endowment is deterministic, $\sigma_y = 0$. Belief distortion parameters u^n are shown in the titles of individual panels.

each set one of the agents appropriates all wealth while the other becomes immediately extinct.

A higher IES provides more incentives to substitute consumption across time and accept steeper consumption profiles. In an economy with $\rho > 0$ (IES > 1), a relatively more optimistic agent with negligible wealth who faces prices essentially determined by the more pessimistic agent is willing to postpone consumption into the future vis-à-vis higher expected returns on his portfolio, as indicated by the last term in the consumption-wealth ratio (31). As $\rho \nearrow 1$, this saving motive dominates, and the relatively more optimistic agent outsaves the other agent whenever the Pareto share of the relatively more optimistic agent becomes sufficiently small, thus guaranteeing his survival.

A similar mechanism operates when the Pareto share of the relatively more pessimistic agent becomes negligible. Observe that the last term in brackets in the consumption-wealth ratio (31), which dominates the saving decision when $\rho \nearrow 1$, is equal to

$$\frac{1}{2}\left(u^1 - u^2\right)\sigma_y\left(1 + \pi^1(0)\right).$$

If agent 1 is relatively more pessimistic, then $u^1 - u^2 < 0$, and thus $\pi^1(0) < -1$ is needed for the saving motive of agent 1 to increase as $\rho \nearrow 1$. In that case, despite being relatively more pessimistic about the growth rate of aggregate endowment, agent 1 becomes sufficiently *optimistic* about the return on his own portfolio, which contains a short position in the claim on aggregate endowment, so that he outsaves the relatively more optimistic agent 2 when the Pareto share of agent 1 declines to zero. This will happen when γ is not too large — a high risk aversion prevents the agents from taking sufficiently disparate portfolio positions that would imply a short stock position in the portfolio of the pessimistic agent.

Finally, when preferences of the agents become inelastic ($\rho \searrow -\infty$), formulas in Proposition 14 imply that the survival conditions cannot hold simultaneously. Inelastic preferences imply that the agents are unwilling to substantially change the slope of their consumption profiles; and the mechanism based on differences in saving rates, which operated for high IES, is largely absent. The consumption-wealth ratio when one agent has a negligible Pareto share is dominated by the second term in expression (31), which is common for both agents. An increased willingness to save when the Pareto share decreases is thus not strong enough to compensate for mistakes in portfolio allocation for at least one of the agents. When IES is sufficiently low, then a pessimistic agent 1 with relatively more distorted beliefs can dominate the economy in situations when the sum of the belief distortions of the two agents is not too large (the exact condition requires $u^1 + u^2 + 2\sigma_y > 0$).

3.5.2 Economic interpretation of the equilibrium price mechanism

Figure 1 documents the survival regions for different belief parameterizations. Previous literature on survival under separable preferences already confirmed that along the dotted diagonal, the agent with the smaller belief distortion dominates. This paper establishes that although this conclusion holds in a neighborhood of the diagonal, there are wide regions of the parameter space where nondegenerate long-run equilibria exist. Moreover, such equilibria generically arise for plausible parameterizations when risk aversion is larger than the inverse of IES that are typically used in the asset pricing literature.

The two boundaries in the top left panel which delimit the region with a nondegenerate stationary distribution of the Pareto share are asymptotically parallel as $\gamma \searrow -\infty$ with slope $2\sigma_y / (u^1 + u^2 + 2\sigma_y)$. The graphs confirm the asymptotic results from Section 3.5.1, where I explained why the relatively more optimistic agent dominates the economy if risk aversion is sufficiently high ($\gamma \searrow -\infty$). But why do we obtain an intermediate region in the parameter space where nondegenerate long-run equilibria exist? The existence of these equilibria critically depends on general equilibrium price effects.

Consider a parameterization that falls into the region where a nondegenerate long-run equilibrium exists in an economy populated by an optimistic and a rational agent, shown in the top two panels of Figure 1. Proposition 10 shows that as the Pareto share of one of the agents becomes negligible, prices are determined by the large agent. The optimistic agent holds a leveraged position in the stock, which is relatively overpriced due to his presence, but the extent of overpricing varies with the Pareto share. When the Pareto share of the optimistic agent is large, overpricing is large as well; investment in the overpriced asset slows down the growth rate of the optimistic agent's wealth, and allows the survival of the rational agent. On the other hand, when the Pareto share of the rational agent is large, overpricing disappears, and the leveraged high expected return strategy of the optimistic agent prevents his extinction.

The equilibrium price mechanism also explains why survival is not sufficient for dominance. Consider an economy populated by a pessimistic and a rational agent (the middle two panels of Figure 1) and the survival region when risk aversion is larger that the inverse of IES. Section 3.5.1 shows that a pessimistic agent is optimistic about the growth rate of his own wealth if his portfolio involves a short position in the stock. When the short position is sufficiently large and IES is larger than one, his consumption-wealth ratio can decline enough so that he is able to outsave the rational agent. However, when the Pareto share of the pessimistic agent is large, he determines equilibrium prices, and his wealth share invested in the stock approaches one. This makes the agent pessimistic again about the growth rate of his own portfolio, and the saving motive disappears.

The general equilibrium price effects thus play a central role in the construction of nondegenerate long-run equilibria. Partial equilibrium models with exogenous price dynamics that do not depend on wealth shares of individual agents cannot replicate this survival mechanism.

3.5.3 Comparative statics

Figure 1 also shows the sensitivity of the survival regions to changes in parameter values. Survival regions do not depend on the time preference parameter β and the growth rate of the economy μ_y, because these parameters influence the decision rules of the two agents in a symmetric way and offset each other in the difference in growth rates of individual wealth. The only remaining parameters to analyze are the belief distortions u^n of the two agents, and the volatility of aggregate endowment σ_y.

The belief distortions used in Figure 1 may be considered large — the incorrect agent misperceives the growth rate of the economy by $|u^1\sigma_y| = 0.005$, which is a quarter of a plausible value for the growth rate of $\mu_y = 0.02$. The top two panels illustrate how survival regions change when the magnitude of the belief distortion u^1 decreases in a situation when agent 2 is rational.

Decreasing the magnitude of the belief distortion in general improves the survival chances of the relatively more optimistic agent. The top right panel shows that decreasing the belief distortion of an optimistic agent shrinks the regions of the parameter space in which the rational agent dominates and where nondegenerate long-run equilibria exist. A decrease in $u^1 > 0$ leads to a lower leverage of agent 1 (see the asymptotic formula (32) for $\pi^1(\theta^1)$) and less overpricing of the risky asset in which the optimistic agent overinvests. Yet the stronger insurance motive of the relatively more pessimistic agent is still present. As u^1 declines to zero, the relative patience $\nu^2(\theta^1) - \nu^1(\theta^1)$ close to the boundaries in expressions (33) and (34) is dominated by the linear term $u^1\sigma_y$, but the threshold in the survival condition in Proposition 6 decreases quadratically. This makes it easier for agent 1 to survive, despite the fact that the speed of convergence to a steady state distribution of θ^1, influenced by the relative patience in the drift term of (18), may be slower.

Interestingly, the middle right panel of Figure 1 shows that decreasing the distortion of a pessimistic agent in general diminishes his survival chances when risk aversion is larger than the inverse of IES. In the previous sections, I argued that a pessimistic agent can coexist with a rational agent in the long run if he is able to outsave the rational agent when his Pareto share becomes small. This will occur in economies with IES > 1 in situations when the pessimistic agent is sufficiently optimistic about the growth rate of his own wealth, i.e., when he chooses a sufficiently large short position in the stock. But decreasing the magnitude of u^1 also decreases the agent's willingness to short the stock (see equation (32)), which subsequently diminishes the relative incentives of the pessimistic agent to save.

The region in which nondegenerate long-run equilibria exist can be expanded if we consider economies populated by an optimistic and a pessimistic agent (the bottom right panel of Figure 1, with a different scale on the vertical axis). In these situations, motives to misallocate assets arising from belief distortions work in opposite directions for the two agents.

For instance, a pessimistic agent can coexist with a rational agent ($u^1 < 0$, $u^2 = 0$) in the long run only when IES is larger than one (middle panels of Figure 1), although a nondegenerate long-run equilibrium can exist for IES smaller than one if we consider parameterizations where agent 2 is optimistic, albeit with a smaller belief distortion ($u^1 < -u^2 < 0$). With an at least somewhat optimistic agent 2, survival chances of the pessimistic agent no longer depend solely on his ability to outsave the other agent when preferences are elastic, but also on the willingness of the optimistic agent to overpay for the claim on the aggregate endowment.

The case of 'symmetric' optimism and pessimism, $0 < u^1 = -u^2$, which under CRRA preferences generates a rather delicate economy without a stationary distribution for the Pareto share yet with both agents surviving in the long run, is dissected in a straightforward way when $\rho \neq \gamma$. The parameter space is divided into four regions by the 45° line and a vertical boundary for γ that satisfies $(1-\gamma)\,\sigma_y = u^1$, and one of the four survival possibilities from Proposition 6 arises in each of the four regions. The online appendix analyzes this case in more detail.

Finally, the bottom left panel in Figure 1 illustrates that the analysis is still plausible when aggregate endowment is deterministic, $\sigma_y = 0$, as long as agents can write contracts on the realizations of W. Volatility of the aggregate consumption stream has an impact on the shape of the survival regions but is not central for the existence of parameterizations under which a nondegenerate long-run equilibrium exists. The essential component of the model is the existence of a betting mechanism with a probability distribution of outcomes about which the agents disagree. Not surprisingly, the bottom left panel in Figure 1 also corresponds to an economy where aggregate endowment is driven by a shock uncorrelated with W over which there is no disagreement.

3.6 Comparison to economies with only terminal consumption

In this paper, I analyze economies with intermediate consumption. Kogan, Ross, Wang, and Westerfield (2006) deal with a different framework with two agents endowed with CRRA preferences. In their economy, there is no intermediate consumption and the agents split and consume an aggregate dividend payoff at a terminal date T. The dividend evolves according to a geometric Brownian motion (1) as in this paper, and agents can continuously retrade claims on the terminal payoff during the lifetime of the economy. The notion of survival in that framework is captured by analyzing the limit of the consumption share distribution in a sequence of economies as $T \nearrow \infty$.

Without intermediate consumption, the agent's intertemporal decision is reduced to the maximization of the (risk-adjusted) expected growth rate of the portfolio. In this respect, the framework is similar to a model in this paper under unitary IES when agents' consumption-wealth ratio is constant and equal to β, and intermediate consumption has no impact on the difference of the wealth growth rates.

The difference that prevents a direct comparison of the results lies in the valuation of wealth. In the absence of intermediate consumption, Kogan, Ross, Wang, and Westerfield (2006) use the price of a bond maturing at time T as numeraire and define the initial wealth in the economy with horizon T as the time 0 price of the terminal payoff. (In this paper, this quantity corresponds to the price of a single cash flow from the aggregate endowment paid out at time T, scaled by the price of a bond with corresponding maturity.) Then they consider two approaches to survival analysis.

In the 'general equilibrium' approach, they study the limiting properties of the terminal consumption allocation obtained as a solution of a sequence of planner's problems as $T \nearrow \infty$. A critical assumption in this approach is the choice of the initial Pareto shares. These are chosen so that the initial wealth shares of the two agents are identical, which requires the initial Pareto share of the irrational agent to approach one as $T \nearrow \infty$. This mechanism reweighs the behavior of the tail and allows an optimistic agent to 'survive' in the sequence of planner's problems. In economies with intermediate consumption, consumption at distant dates contributes only little to the wealth levels, and thus the reweighting of initial Pareto shares in order to achieve equal initial wealth levels would have no effect on the survival results. Under the 'general equilibrium' notion of survival in Kogan, Ross, Wang, and Westerfield (2006), optimistic agents can survive when risk aversion is larger than one but the survival regions differ from the results in this paper under unitary IES ($\rho = 0$).

Kogan, Ross, Wang, and Westerfield (2006) contrast their 'general equilibrium' to a simplified approach that is analogous to the boundary analysis in this paper and that they call the 'partial equilibrium' method. This method constructs a homogeneous economy injected with an infinitesimal agent with different beliefs under the assumption that he does not affect local price dynamics. It turns out that this method delivers exactly the same survival regions as those derived in this paper under unitary IES.

Propositions 10 and 11 show that in the model with intermediate consumption considered in this paper, the return on aggregate wealth and prices of individual finite-horizon cash flows from the aggregate endowment converge to their homogeneous economy counterparts and thus the 'partial equilibrium' approach is actually the correct method for this paper under

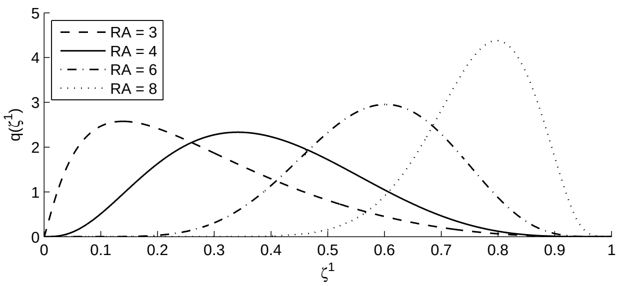

Figure 2: Stationary distributions for the consumption share of the agent with distorted beliefs $\zeta^1(\theta^1)$. All models are parameterized by $u^1 = 0.25$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.02$, and differ in levels of risk aversion.

general equilibrium. However, these results do not translate to the setup considered in Kogan, Ross, Wang, and Westerfield (2006). Although prices of individual cash flows from the aggregate endowment converge for every fixed $T \geq 0$, this convergence is not uniform on $T \in [0, \infty)$, which in general invalidates the result on converging returns and prices for the limit as $T \nearrow \infty$.

4 Dynamics of long-run equilibria

In Section 3, I derived parametric restrictions on the survival regions. However, even if a nondegenerate long-run equilibrium exists, the question remains whether this equilibrium delivers quantitatively interesting dynamics under which each of the agents can gain a significant wealth share. This section investigates numerically the equilibrium allocations and prices and their dynamics by solving the ODE (15) and decentralizing the allocations.

4.1 Consumption allocation

Figure 2 plots the densities $q(\zeta^1)$ for the stationary distribution of the consumption share in economies with an optimistic agent. The parameterizations[10] are chosen along a horizontal line in the top left panel of Figure 1. As risk aversion increases, the distribution of consumption shifts toward the optimistic agent, but the equilibria in general permit substantial variation over time in the consumption shares of the two agents.

The existence of nondegenerate long-run equilibria depends on the behavior of the relative patience $\nu^2(\theta^1) - \nu^1(\theta^1)$ in the neighborhood of the boundaries. Figure 3 displays three

[10] A full solution of the consumption dynamics requires setting additional parameters that do not influence the survival regions. I set $\beta = 0.05$ and $\mu_y = 0.02$. The high value for the time preference coefficient is chosen merely to assure that restrictions in Assumption 2 hold for all compared models.

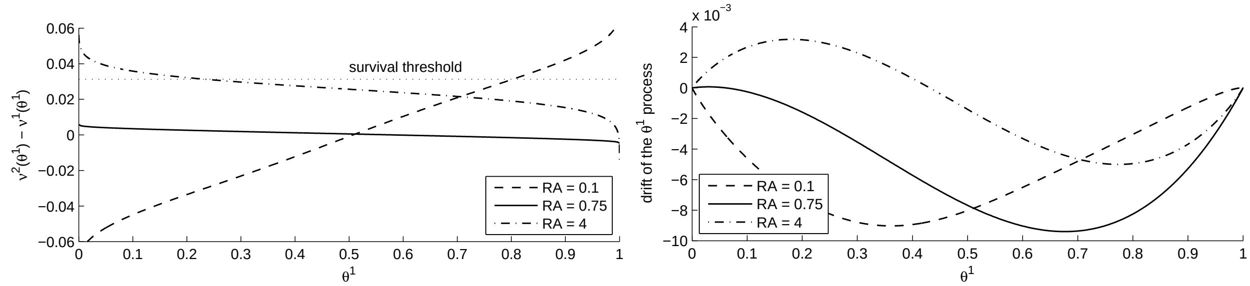

Figure 3: Relative patience $\left[\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)\right]/\left(u^1\right)^2$ (left panel) and the drift component of the Pareto share evolution $E\left[d\theta_t^1 \mid \mathcal{F}_t\right]/dt$ (right panel) as functions of the Pareto share θ^1. All models are parameterized by $u^1=0.25$, $u^2=0$, IES $=1.5$, $\beta=0.05$, $\mu_y=0.02$, $\sigma_y=0.02$, and differ in levels of risk aversion. The dotted horizontal line in the left panel represents the survival threshold $\frac{1}{2}\left(\left(u^1\right)^2-\left(u^2\right)^2\right)$ from Proposition 6.

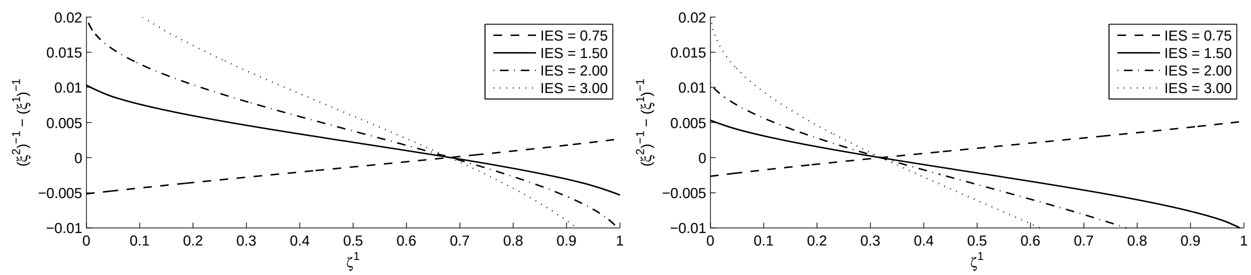

Figure 4: Difference in consumption-wealth ratios $\xi^n(\theta^1)$ as a function of the consumption share $\zeta^1(\theta^1)$. The left panel considers an optimistic agent 1 ($u^1=0.25$) while the right panel a pessimistic agent 1 ($u^1=-0.25$). The remaining parameters are $u^2=0$, RA $=2$, $\beta=0.05$, $\mu_y=0.02$, $\sigma_y=0.02$, and individual curves correspond to different levels of intertemporal elasticity of substitution.

different cases. The dashed line represents the low risk aversion case in which both attracting conditions from Proposition 6 hold and each of the agents dominates with a strictly positive probability. The solid line corresponds to a parameterization that is close to the CRRA case when only the survival condition for the rational agent 2 is satisfied (with CRRA preferences, the relative patience would be identically zero). Finally, a case for which both survival conditions hold is shown by the dot-dashed line.

Figure 3 also plots the impact of relative patience on the drift component of the Pareto share process. The drift vanishes at the boundaries and the boundaries are unattainable (a reflection of the Inada conditions), but sufficiently large positive (negative) slopes at the left (right) boundaries assure the existence of a nondegenerate stationary equilibrium of the Pareto share.

The essential components of the survival mechanism are the propensity to save and the portfolio allocation of the two agents. Figure 4 displays the differences in the consumption-

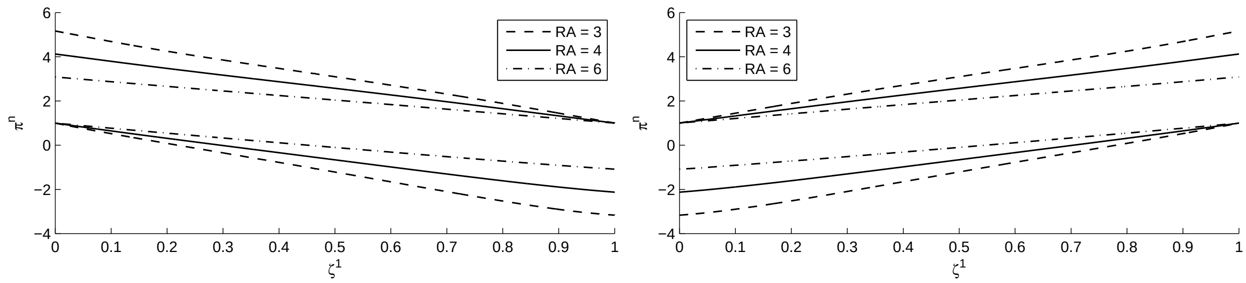

Figure 5: Wealth shares $\pi^n(\theta^1)$ of the two agents invested in the claim to aggregate endowment as functions of the consumption share $\zeta^1(\theta^1)$. The left panel considers an optimistic agent 1 ($u^1 = 0.25$) while the right panel a pessimistic agent 1 ($u^1 = -0.25$). The remaining parameters are $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.02$, and individual curves correspond to different levels of risk aversion. Wealth share curves originating at 1 for $\zeta^1\left(\theta^1\right) = 1$ ($\zeta^1\left(\theta^1\right) = 0$) belong to agent 1 (2).

wealth ratios $\left[\xi^n\left(\theta^1\right)\right]^{-1}$ of the two agents, which are primarily driven by the intertemporal elasticity of substitution. For the case of IES $= 1$, the difference is exactly zero since each agent consumes a fraction β of his wealth per unit of time. A higher IES improves the survival chances of the agent who is relatively more optimistic about the return on his own wealth, as he is willing to tilt his consumption profile more toward the future. Figure 4 captures this effect for both an optimistic agent 1 ($u^1 = 0.25$, left panel), as well as a pessimistic agent 1 ($u^1 = -0.25$, right panel).

The portfolio allocation mechanism is depicted in Figure 5 and is closely related to the behavior of the consumption-wealth ratios. The share of wealth invested in the risky asset is primarily driven by the risk aversion parameter γ. A higher risk aversion limits the amount of leverage. For the pessimistic agent, this implies that if risk aversion is high, he does not form a large enough short stock position that would make him sufficiently optimistic about the return on his own wealth and outsave the rational agent when IES > 1.

4.2 Evolution over time

In empirical applications, it may be advantageous if θ^1 converges to its stationary distribution from any initial condition fast enough, so that data observed over finite horizons are a representative sample of the stationary distribution. Proposition 6 gives sufficient conditions for the existence of a unique stationary distribution for θ^1 but it does not say anything about the rate of convergence.

It turns out that under the conditions in Proposition 6, convergence occurs at an exponential rate, so that the process θ^1 does not exhibit strong dependence properties, although

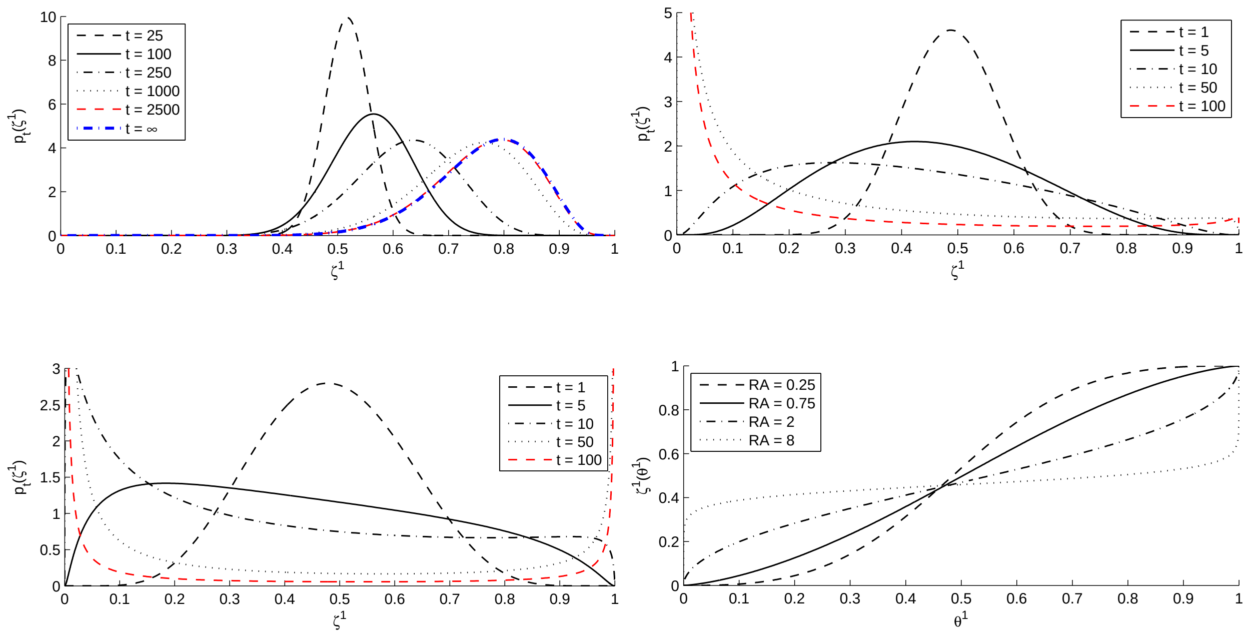

Figure 6: Distributions of $\zeta^1\left(\theta_t^1\right)$ conditional on $\zeta^1\left(\theta_0^1\right) = 0.5$, and the consumption share of agent 1, $\zeta^1\left(\theta^1\right)$, as a function of the Pareto share θ^1 (bottom right panel). In the top left panel, the economy has a nondegenerate long-run distribution. In the top right panel, agent 2 dominates, and in the bottom left panel, each agent dominates with a strictly positive probability. The parameters are $u^1 = 0.25$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.02$. Risk aversion is equal to 8 in the top left panel, 0.75 in the top right panel, and 0.25 in the bottom left panel.

the exponent may be small. Yan (2008) conducts numerical experiments under separable utility when one of the agents always vanishes, and shows that the rate of extinction can be very slow. The same quantitative result can hold under recursive preferences.

Having at hand numerical solutions for the evolution of θ^1 and the function $\zeta^1\left(\theta^1\right)$, one can investigate conditional distributions of $\zeta^1\left(\theta_t^1\right)$ conditional on θ_0^1 by solving the corresponding Kolmogorov forward equation

$$\frac{\partial q_t\left(\theta^1\right)}{\partial t} + \frac{\partial}{\partial \theta^1}\left[\theta^1 \mu_{\theta^1}\left(\theta^1\right) q_t\left(\theta^1\right)\right] - \frac{1}{2}\frac{\partial^2}{\partial\left(\theta^1\right)^2}\left[\left(\theta^1 \sigma_{\theta^1}\left(\theta^1\right)\right)^2 q_t\left(\theta\right)\right] = 0$$

for the conditional density $q_t\left(\theta^1\right)$ of θ_t^1 with the initial condition $q_0\left(\theta^1\right) = \delta_{\theta_0^1}\left(\theta^1\right)$, where δ is the Dirac delta function, and then transforming to obtain the conditional density for ζ^1

$$p_t\left(\zeta^1\left(\theta^1\right)\right) = q_t\left(\theta^1\right)\left[\frac{\partial \zeta^1}{\partial \theta^1}\left(\theta^1\right)\right]^{-1}.$$

Figure 6 considers the evolution of conditional densities for the consumption share in different economies. The speed of convergence depends on relative patience $\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)$

that governs the magnitude of the drift term of θ^1 and the shape of the function $\zeta^1(\theta^1)$ depicted in bottom right panel.

For high levels of risk aversion, convergence of the conditional distribution p_t is slow, due to the low slope of $\zeta^1(\theta^1)$. With a high level of risk aversion, agents are not willing to engage in large bets on the realizations of the Brownian motions W, and wealth and consumption shares evolve only slowly. In the example in Figure 6, it takes 2,500 periods until the density p_t is indistinguishable from the stationary density.

As risk aversion decreases, and agents are willing to bet larger portions of their wealth, the evolution of the conditional density p_t speeds up. In the middle panel of Figure 6, consumption is substantially skewed toward the dominating agent 2 already after 50 periods. In the bottom panel, when the risk aversion coefficient drops to 0.25 and each of the agents dominates with a strictly positive probability, the mass of the density quickly shifts toward both boundaries.

5 Extensions and concluding remarks

Before concluding, I consider two extensions of the analyzed model that involve Bayesian learning about the underlying model and representation of other preference structures as belief distortions. The online appendix outlines in more detail how to set up these problems within the framework of this paper.

5.1 The role of learning

The analysis in this paper focuses on the case of fixed belief distortions. Agents are firm believers in their probability models, and do not use new data to update their beliefs. A natural question is to ask what happens when agents are allowed to learn.

Blume and Easley (2006) provide a detailed analysis of the impact of Bayesian learning on survival under separable utility, and they are able to characterize the relationship between survival chances and the complexity of the learning problem. The central message arising from the analysis is that learning, which reduces belief distortions over time, in general aids survival of agents with incorrect beliefs.

It seems to be reasonable to expect that this insight should hold also under nonseparable preferences. Unfortunately, results presented in the previous analysis indicate that this logic is not generally correct. For instance, the middle panels of Figure 1 show that the survival region of a pessimistic agent can shrink if his belief distortion diminishes, and the pessimistic agent moves from a region of the parameter space where a nondegenerate long-run equilibrium exists to one where only the rational agent survives. Whether the pessimist can then learn quickly enough so that his beliefs converge to rational expectations at a rate that allows survival depends on the complexity of the learning problem, as shown by Blume and Easley (2006). As the beliefs converge, the evolution of the Pareto share process θ^1 settles. The limiting distribution of θ^1 as $t \nearrow \infty$ from which we can deduce the wealth and consumption distribution remains an open question.

5.2 Robust utility

The economic interpretation of the distortionary processes u^n is not limited to 'irrationality', and other preference specifications lead to representations which are observationally equivalent to belief distortions. Consider, for instance, an agent who believes that the model for the aggregate endowment dynamics is misspecified and views (1) only as a reference model that approximates the true dynamics, as in the robust utility models of Anderson, Hansen, and Sargent (2003) and Skiadas (2003). This class of models leads to a representation where agent n views as relevant the realization of the worst case scenario, characterized by the least favorable dynamics

$$\frac{dY_t}{Y_t} = \mu_y dt + \sigma_y \left(u_t^n dt + dW_t^n \right),$$

where W^n is a Brownian motion under Q_u^n associated with an endogenously determined distortionary process u^n. Epstein and Miao (2003) and Uppal and Wang (2003) construct models with ambiguity aversion where the optimal solution to the minimization problem involves a constant u^n, and thus exactly corresponds to the framework in this paper.

Under separable preferences, agents who fear misspecification more (and therefore assign a lower penalty θ to deviations from the reference model) choose a more distorted worst case scenario, which tends to worsen their survival chances. However, the results for constant belief distortions u^n indicate that survival chances of the more fearful agents may well look much better for appropriate nonseparable parameterizations of preferences. A detailed analysis of the dynamics of these models is left for future research.

5.3 Summary

Survival of agents with heterogeneous beliefs has been studied extensively under separable preferences. The main conclusion arising from the literature is a relatively robust argument in favor of the market selection hypothesis. Under complete markets and identical utility functions, a two-agent economy is dominated in the long run by the agent whose beliefs are closest to the true probability measure for a wide class of preferences and endowments. In particular, Kogan, Ross, Wang, and Westerfield (2009) show elegantly that this result holds, irrespective of the specification of the aggregate endowment process,[11] as long as relative risk aversion is bounded.

This paper shows that the robust survival result is specific to the class of separable preferences. Under nonseparable recursive preferences of the Duffie-Epstein-Zin type, nondegenerate long-run equilibria exist for a broad set of plausible parameterizations when risk aversion is larger than the inverse of the intertemporal elasticity of substitution. It is equally easy to construct economies dominated by agents with relatively more incorrect beliefs.

The analysis reveals the important role played by the interaction of risk aversion with respect to intratemporal gambles that determines risk taking, and intertemporal elasticity of substitution that drives the consumption-saving decision. Critical for obtaining the survival results, and in particular the nondegenerate long-run equilibria, are the general equilibrium price effects generated by the wealth dynamics.

The survival results are obtained by extending the planner's problem formulation of Dumas, Uppal, and Wang (2000) to a setting with heterogeneous beliefs. Long-run survival of the agents is determined by the dynamics of a stochastic process that models the Pareto share of one of the agents as the share becomes negligible. This dynamics can be characterized in closed form by studying the boundary behavior of a nonlinear ODE resulting from the planner's problem. This type of ODE arises in a wider class of recursive utility problems, so these results can be utilized in a broader variety of economic applications.

I provide in analytical form tight sufficient conditions that guarantee survival or extinction. These conditions can be interpreted as relative patience conditions analogous to those in Lucas and Stokey (1984). An agent survives in the long run if his relative patience becomes sufficiently large as his wealth share vanishes. However, in this framework, the dynamics of relative patience arises endogenously as an equilibrium outcome, and is not a direct property of agents' preferences. I also show that the survival conditions are equivalent to conditions on the limiting expected growth rates of the logarithm of individual wealth levels in a decentralized economy

These results are obtained for a two-agent economy with an aggregate endowment process that is specified as a geometric Brownian motion, but the theoretical framework can also be

[11]The survival results under separable utility thus also hold for 'exotic' endowment processes like the rare disaster framework in Chen, Joslin, and Tran (2010).

utilized to derive an analog HJB equation for multi-agent economies with more sophisticated Markov dynamics. In principle, the qualitative survival results should extend to a wider class of models with stable consumption growth dynamics, although the analysis of the existence of a stationary distribution for the Pareto share becomes more complicated in a multidimensional state space.

Importantly, the developed solution method is not limited to constant distortions and applies to a much wider class of preferences that are interpretable as deviations in beliefs. I outline how to use the method in a framework with model uncertainty and learning and in a model where agents are endowed with robust preferences. Solutions of these problems are left as open questions for future research. Similarly, formulas for survival regions can be extended by incorporating heterogeneity in preferences, as in Dumas, Uppal, and Wang (2000), in a straightforward way.

The bad news for the market selection hypothesis is in some sense good news for models with heterogeneous agents. Models with agents who differ in preferences or beliefs often have degenerate long-run limits in which only one class of agents survives. This paper shows that coupling belief heterogeneity (including preferences that can be interpreted as belief distortions) and recursive preferences with empirically plausible parameters leads to models in which the heterogeneity does not vanish over time.

A Proofs

Proof of Lemma 3. Schroder and Skiadas (1999) prove that $V^n(C^n)$ is concave. Consider the case $\alpha^1 \searrow 0$. Given optimal consumption streams $C^n(\alpha)$, we have

$$J_0(\alpha) = \alpha^1 V_0^1 \left(C^1(\alpha)\right) + \alpha^2 V_0^2 \left(C^2(\alpha)\right) \tag{35}$$

and since $V_0^1\left(C^1(\alpha)\right)$ is bounded from above as a function of α, it follows that

$$\alpha^1 V_0^1 \left(C^1(\alpha)\right) \overset{\alpha^1 \searrow 0}{\longrightarrow} v^1 \leq 0$$

and thus $J_0\left(0, \alpha^2\right) \leq \lim_{\alpha^1 \searrow 0} \alpha^2 V_0^2 \left(C^2(\alpha)\right) \leq \alpha^2 V_0^2(Y)$.

Assume suboptimal policies $\hat{C}^1\left(\alpha^1, \alpha^2\right) = \left(\alpha^1\right)^{\frac{1}{2|\gamma|}} Y$ and $\hat{C}^2\left(\alpha^1, \alpha^2\right) = \left(1 - \left(\alpha^1\right)^{\frac{1}{2|\gamma|}}\right) Y$. Then

$$\alpha^1 V_0^1 \left(\hat{C}^1\left(\alpha^1, \alpha^2\right)\right) = \left(\alpha^1\right)^{1 + \frac{1}{2}\frac{\gamma}{|\gamma|}} \gamma^{-1} Y_0^\gamma \tilde{V}^n \overset{\alpha^1 \searrow 0}{\longrightarrow} 0$$

and

$$\alpha^2 V_0^2 \left(\hat{C}^2\left(\alpha^1, \alpha^2\right)\right) = \alpha^2 \left(1 - \left(\alpha^1\right)^{\frac{1}{2|\gamma|}}\right)^\gamma \gamma^{-1} Y_0^\gamma \tilde{V}^n \overset{\alpha^1 \searrow 0}{\longrightarrow} \alpha^2 V_0^2(Y)$$

which implies $J_0\left(0, \alpha^2\right) \geq \alpha^2 V_0^2(Y)$. Therefore (13) holds, and the convergence of $C^2\left(\alpha^1, \alpha^2\right)$ is a direct consequence. ■

Proof of Proposition 4. The planner's problem has an appealing Markov structure. Denoting $\bar{\lambda} = \left(\bar{\lambda}^1, \bar{\lambda}^2\right)'$ and $u = \left(u^1, u^2\right)'$, the state vector is $Z = \left(\bar{\lambda}', Y\right)'$, and the planner's problem (11-12) leads to the Hamilton-Jacobi-Bellman equation for $J(Z)$,

$$0 \equiv \sup_{(C^1, C^2, \nu^1, \nu^2)} \sum_{n=1}^{2} \bar{\lambda}^n \left[F\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n} \nu^n\right] + J_y \mu_y Y + \frac{1}{2}\mathrm{tr}\left(J_{zz}\Sigma\right), \tag{36}$$

where

$$\Sigma = \begin{pmatrix} \left(\mathrm{diag}\left(\bar{\lambda}\right) u\right)\left(\mathrm{diag}\left(\bar{\lambda}\right) u\right)' & \left(\mathrm{diag}\left(\bar{\lambda}\right) u\right) \sigma_y Y \\ \sigma_y Y \left(\mathrm{diag}\left(\bar{\lambda}\right) u\right)' & \sigma_y^2 Y^2 \end{pmatrix}$$

and $\mathrm{diag}\left(\bar{\lambda}\right)$ is a 2×2 diagonal matrix with elements of $\bar{\lambda}$ on the main diagonal.

The maximization over $\left(\nu^1, \nu^2\right)$ in the HJB equation (36) can be solved separately. Under the optimal discount rate process ν^n for agent n,

$$f\left(C^n, J_{\bar{\lambda}^n}\right) \doteq \sup_{\nu^n} F\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n} \nu^n = \frac{\beta}{\rho}\left[\left(C^n\right)^\rho \left(\gamma J_{\bar{\lambda}^n}\right)^{1 - \frac{\rho}{\gamma}} - \gamma J_{\bar{\lambda}^n}\right]. \tag{37}$$

The function f is the aggregator in the stochastic differential utility representation of recursive preferences postulated by Duffie and Epstein (1992b). The online appendix gives more detail on this relationship. Optimal consumption shares ζ^n are given by the first-order conditions in the consumption allocation

$$\zeta^n \doteq \frac{C^n}{Y} = \frac{\left(\gamma J_{\bar{\lambda}^n}\right)^{\frac{1-\rho/\gamma}{1-\rho}} \left(\bar{\lambda}^n\right)^{\frac{1}{1-\rho}}}{\sum_{k=1}^{2} \left(\gamma J_{\bar{\lambda}^k}\right)^{\frac{1-\rho/\gamma}{1-\rho}} \left(\bar{\lambda}^k\right)^{\frac{1}{1-\rho}}},$$

where $J_{\bar{\lambda}^n}$ are agents' continuation values under the optimal consumption allocation.

The HJB equation (36) further implies that J is homogeneous degree one in $\bar{\lambda}$ and homogeneous degree γ in Y. The transformation of variables (14) leads to the guess

$$J(Z) = \gamma^{-1} Y^{\gamma} \theta^2 \tilde{J}\left(\theta^1\right) = \gamma^{-1} Y^{\gamma} \theta^2 \left[\theta^1 \tilde{J}^1\left(\theta^1\right) + \left(1-\theta^1\right) \tilde{J}^2\left(\theta^1\right)\right],$$

where $\tilde{J}^n\left(\theta^1\right)$ are continuation values of the two agents scaled by $\gamma^{-1}Y^{\gamma}$, defined in (16). The ODE for $\tilde{J}^n\left(\theta^1\right)$ then immediately follows. The continuity at the boundaries follows from Lemma 3.

In addition, the same logic and derivation of the HJB equation applies to multi-agent economies and more sophisticated Markov dynamics of the aggregate endowment process. In an N-agent economy, the state vector includes $N-1$ Pareto shares as state variables. The boundary conditions for $\theta^n = 0$, $n \in \{1, \ldots, N\}$ associated with the N-agent version of the ODE (15) are given by the solutions of $(N-1)$-agent economies that exclude agent n. In this way, solutions to multi-agent economies can be calculated by iteratively adding individual agents. ■

Proof of Proposition 6. Given an initial condition $\theta_0^1 \in (0,1)$, the process (18) lives on the open interval $(0,1)$ with unattainable boundaries (the preferences satisfy an Inada condition at zero). For any numbers $0 < a < b < 1$, the process θ^1 has bounded and continuous drift and volatility coefficients on (a,b), and the volatility coefficient is bounded away from zero. It is thus sufficient to establish the appropriate boundary behavior of θ^1 in order to make the process positive Harris recurrent (see Meyn and Tweedie (1993)). Since the process will also be φ-irreducible for the Lebesgue measure under these boundary conditions, there exists a unique stationary distribution.

Denote $\mu_\theta(\theta)$ and $\sigma_\theta(\theta)$ the drift and volatility coefficients in (18). The boundary behavior of the process θ^1 is captured by the scale measure $S : (0,1)^2 \to \mathbb{R}$ defined as

$$s(\theta) = \exp\left\{-\int_{\theta_0}^{\theta} \frac{2\mu_\theta(\tau)}{\sigma_\theta^2(\tau)} d\tau\right\} \qquad S[\theta_l, \theta_h] = \int_{\theta_h}^{\theta_l} s(\theta)\, d\theta$$

for an arbitrary choice of $\theta_0 \in (0,1)$, and the speed measure $M : (0,1)^2 \to \mathbb{R}$

$$m(\theta) = \frac{1}{\sigma_\theta^2(\theta)\, s(\theta)} \qquad M[\theta_l, \theta_h] = \int_{\theta_h}^{\theta_l} m(\theta)\, d\theta.$$

Karlin and Taylor (1981, Chapter 15) provide an extensive treatment of the boundaries.

The boundaries are nonattracting if and only if

$$\lim_{\theta_l \searrow 0} S[\theta_l, \theta_h] = \infty \quad \text{and} \quad \lim_{\theta_h \nearrow 1} S[\theta_l, \theta_h] = \infty \tag{38}$$

and this result is independent of the fixed argument that is not under the limit. With nonattracting boundaries, the stationary density will exist if the speed measure satisfies

$$\lim_{\theta_l \searrow 0} M[\theta_l, \theta_h] < \infty \quad \text{and} \quad \lim_{\theta_h \nearrow 1} M[\theta_l, \theta_h] < \infty, \tag{39}$$

again independently of the argument that is not under the limit.

In our case,

$$s(\theta) = \exp\left\{-\int_{\theta_0}^{\theta} \frac{2\left(\nu^2(\tau) - \nu^1(\tau)\right)}{\tau(1-\tau)\left(u^1-u^2\right)^2} d\tau\right\} s_{sep}(\theta),$$

where

$$s_{sep}(\theta) = \left(\frac{1-\theta}{1-\theta_0}\right)^{-\frac{2u^1}{u^1-u^2}} \left(\frac{\theta}{\theta_0}\right)^{\frac{2u^2}{u^1-u^2}} \tag{40}$$

is the integrand of the scale function in the separable case, when $\nu^2(\theta) - \nu^1(\theta) \equiv 0$.

For the left boundary, assume that in line with condition (i), there exist $\underline{\theta} \in (0,1)$ and $\underline{\nu} \in \mathbb{R}$ such that $\nu^2(\theta) - \nu^1(\theta) \geq \underline{\nu}$ for all $\theta \in (0, \underline{\theta})$. Taking $\theta_0 = \underline{\theta}$, the scale measure can be bounded as

$$\begin{aligned} S[\theta_l, \underline{\theta}] &\geq \int_{\theta_l}^{\underline{\theta}} \exp\left\{-\int_{\underline{\theta}}^{\theta} \frac{2\underline{\nu}}{\tau(1-\tau)\left(u^1-u^2\right)^2} d\tau\right\} \left(\frac{1-\theta}{1-\underline{\theta}}\right)^{-\frac{2u^1}{u^1-u^2}} \left(\frac{\theta}{\underline{\theta}}\right)^{\frac{2u^2}{u^1-u^2}} d\theta = \\ &= \int_{\theta_l}^{\underline{\theta}} \left(\frac{\theta}{\underline{\theta}}\right)^{\frac{2u^2}{u^1-u^2} - \frac{2\underline{\nu}}{(u^1-u^2)^2}} \left(\frac{1-\theta}{1-\underline{\theta}}\right)^{\frac{2\underline{\nu}}{(u^1-u^2)^2} - \frac{2u^1}{u^1-u^2}} d\theta \end{aligned}$$

The left limit in (38) thus diverges to infinity if

$$\frac{2u^2}{u^1-u^2} - \frac{2\underline{\nu}}{\left(u^1-u^2\right)^2} \leq -1,$$

which is satisfied when $\underline{\nu} \geq \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]$.

The argument for the right boundary is symmetric. Taking $\bar{\theta} \in (0,1)$ and $\bar{\nu} \in \mathbb{R}$ such that $\nu^2(\theta) - \nu^1(\theta) \leq \bar{\nu}$ for all $\theta \in (\bar{\theta}, 1)$, the calculation reveals that we require $\bar{\nu} \leq \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]$.

It turns out that the bounds implied by conditions (39) are marginally tighter. Following the same bounding argument as above, sufficient conditions for (39) to hold are

$$\underline{\nu} > \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right] \quad \text{and } \bar{\nu} < \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]. \tag{41}$$

The construction reveals that these bounds are also the least tight bounds of this type under which the proposition holds.

It is also useful to note that the unique stationary density $q(\theta)$ is proportional to the speed density $m(\theta)$. Finally, if the limits in Proposition 6 do not exist, they can be replaced with appropriate limits inferior and superior.

This discussion has sorted out case (a). Conditions (i') and (ii') are sufficient conditions for the boundaries to be attracting. Lemma 6.1 in Karlin and Taylor (1981) then shows that if the 'attracting' condition is satisfied for a boundary, then θ^1 converges to this boundary on a set of paths that has a strictly positive probability. This probability is equal to one if the other boundary is non-attracting. Combining these results, we obtain statements (b), (c), and (d). ■

Proof of Corollary 7. Assume without loss of generality that $|u^2| \leq |u^1|$. The sufficient part is an immediate consequence of Proposition 6. Under separable preferences, $\nu^2 - \nu^1 \equiv 0$, and thus if $|u^2| < |u^1|$ then conditions (i') and (ii) hold, and agent 2 dominates in the long run under P.

For the necessary part, when $u^2 = u^1$, then θ^1 is constant and both agents survive under P. When $-u^2 = u^1 = u$, then it follows from inspection of formula (40) in the proof of Proposition 6 that conditions (38) are satisfied and the boundaries are non-attracting. Lemma 6.1 in Karlin and Taylor (1981) then implies that both agents survive under P.

Note that even though both agents survive when $-u^2 = u^1$, the speed density $m(\theta) \propto \theta^{-1}(1-\theta)^{-1}$ is not integrable on $(0,1)$ and thus there does not exist a finite stationary measure.

The result on survival under measure Q^n follows from the fact that the evolution of Brownian motion W under the beliefs of agent n is $dW_t = u^n dt + dW_t^n$. Since the evolution of θ^1 completely describes the dynamics of the economy, substituting this expression into (18) and reorganizing yields the desired result. ■

Proof of Lemma 8. Lemma 3 implies that the planner's objective function can be continuously extended at $\theta^1 = 0$ by the continuation value for agent 2 living in a homogeneous economy. Expression (35) scaled by $(\alpha^1 + \alpha^2)\gamma^{-1}Y^\gamma$ leads to an equation in scaled continuation values

$$\tilde{J}(\theta^1) = \theta^1 \tilde{J}^1(\theta^1) + (1-\theta^1)\tilde{J}^2(\theta^1)$$

and the proof of Lemma 3 yields

$$\lim_{\theta^1 \searrow 0} \tilde{J}(\theta^1) = \lim_{\theta^1 \searrow 0} \tilde{J}^2(\theta^1) = \tilde{V}^2,$$

where $\tilde{V}^2$ is defined in (7). Since $\tilde{J}^2(\theta^1) = \tilde{J}(\theta^1) - \theta^1 \tilde{J}_{\theta^1}(\theta^1)$, then

$$\lim_{\theta^1 \searrow 0} \theta^1 \tilde{J}_{\theta^1}(\theta^1) = 0. \tag{42}$$

Further, consider the behavior of individual terms in ODE (15) as $\theta^1 \searrow 0$. Using expression (17), the first term is proportional to

$$\begin{aligned} \theta^1 \left(\zeta^1\left(\theta^1\right)\right)^\rho \left(\tilde{J}^1\left(\theta^1\right)\right)^{1-\frac{\rho}{\gamma}} &= \left(\theta^1\right)^{\frac{1}{1-\rho}} \left(\tilde{J}^1\left(\theta^1\right)\right)^{\frac{1-\rho/\gamma}{1-\rho}} \left[K\left(\theta^1\right)\right]^{-\rho} = \\ &= \zeta^1\left(\theta^1\right) \left[K\left(\theta^1\right)\right]^{1-\rho}, \end{aligned}$$

where $K\left(\theta^1\right)$ is the denominator in the formula for the consumption share (17), and $\lim_{\theta^1\searrow 0} K\left(\theta^1\right) = \left(\tilde{V}^2\right)^{\frac{1-\rho/\gamma}{1-\rho}}$, which is a finite value. Since $\lim_{\theta^1\searrow 0}\zeta^1\left(\theta^1\right) = 0$, the first term in (15) vanishes as $\theta^1 \searrow 0$. The sum of the second and third term converges to

$$\frac{\beta}{\rho}\left(\tilde{V}^2\right)^{1-\frac{\rho}{\gamma}} + \left(-\frac{\beta}{\rho} + \mu_y + u^2\sigma_y + \frac{1}{2}\left(\gamma-1\right)\sigma_y^2\right)\tilde{V}$$

and formula (7) implies that this expression is zero. Since the fourth term in (15) also converges to zero due to result (42), the last term in (15) must also converge to zero, or

$$\lim_{\theta^1\searrow 0}\left(\theta^1\right)^2 \tilde{J}_{\theta^1\theta^1}\left(\theta^1\right) = 0. \tag{43}$$

Finally, differentiate the PDE (15) by θ^1 and multiply the equation by θ^1. Using comparisons with results (42–43), the assumption that $\zeta^n\left(\theta^1\right)/\tilde{J}^n\left(\theta^1\right)^{1/\gamma}$ are bounded and bounded away from zero and $\lim_{\theta^1\searrow 0}\zeta^1\left(\theta^1\right) = 0$, it is possible to determine that all terms in the new equation containing derivatives of $\tilde{J}\left(\theta^1\right)$ up to second order vanish as $\theta^1 \searrow 0$. The single remaining term that contains a third derivative of $\tilde{J}\left(\theta^1\right)$ is multiplied by $\left(\theta^1\right)^3$ and must necessarily converge to zero as well, and thus

$$\lim_{\theta^1\searrow 0}\left(\theta^1\right)^3 \tilde{J}_{\theta^1\theta^1\theta^1}\left(\theta^1\right) = 0.$$

■

Proof of Lemma 9. Itô's lemma implies

$$\begin{aligned} d\tilde{J}^2\left(\theta_t^1\right) &= d\left[\tilde{J}\left(\theta_t^1\right) - \theta_t^1 \tilde{J}_{\theta^1}\left(\theta_t^1\right)\right] = \\ &= -\left(\theta_t^1\right)^2 \tilde{J}_{\theta^1\theta^1}\left(\theta_t^1\right)\frac{d\theta_t^1}{\theta_t^1} - \frac{1}{2}\left[\left(\theta_t^1\right)^2 \tilde{J}_{\theta^1\theta^1}\left(\theta_t^1\right) + \left(\theta_t^1\right)^3 \tilde{J}_{\theta^1\theta^1\theta^1}\left(\theta_t^1\right)\right]\left(\frac{d\theta_t^1}{\theta_t^1}\right)^2 \end{aligned}$$

and since the drift and volatility coefficients in the dynamics of θ^1 given by equation (18) are bounded by assumption, applying results from Lemma 8 proves the claim about the drift and volatility coefficients of $\tilde{J}^2\left(\theta^1\right)$ ($\tilde{J}^2$ itself converges to a nonzero limit so the scaling is innocuous).

Further notice that

$$d\tilde{J}^1\left(\theta_t^1\right) = d\left[\tilde{J}\left(\theta_t^1\right)+\left(1-\theta_t^1\right)\tilde{J}_{\theta^1}\left(\theta_t^1\right)\right] = -\left(\theta_t^1\right)^2\tilde{J}_{\theta^1\theta^1}\left(\theta_t^1\right)\frac{d\theta_t^1}{\theta_t^1}+ \quad (44)$$
$$+\frac{1}{2}\left[\left(\theta_t^1\right)^2\tilde{J}_{\theta^1\theta^1}\left(\theta_t^1\right)+\left(1-\theta_t^1\right)\left(\theta_t^1\right)^2\tilde{J}_{\theta^1\theta^1\theta^1}\left(\theta_t^1\right)\right]\left(\frac{d\theta_t^1}{\theta_t^1}\right)^2$$

and that

$$\frac{\zeta^1\left(\theta^1\right)}{\tilde{J}^1\left(\theta^1\right)^{\frac{1}{\gamma}}} = \left(\theta^1\right)^{\frac{1}{1-\rho}}\left(\tilde{J}^1\right)^{\frac{1-1/\gamma}{1-\rho}}K\left(\theta^1\right)^{-1} \quad (45)$$

is bounded and bounded away from zero by assumption. Denote the numerators of ζ^1 and ζ^2

$$Z^1\left(\theta^1\right) = \left(\theta^1\right)^{\frac{1}{1-\rho}}\left(\tilde{J}^1\left(\theta^1\right)\right)^{\frac{1-\rho/\gamma}{1-\rho}} \qquad Z^2\left(\theta^1\right) = \left(1-\theta^1\right)^{\frac{1}{1-\rho}}\left(\tilde{J}^2\left(\theta^1\right)\right)^{\frac{1-\rho/\gamma}{1-\rho}}.$$

Then $\zeta^2 = Z^2/\left(Z^1+Z^2\right)$ and, omitting arguments,

$$dZ^1 = \frac{1}{1-\rho}Z^1\frac{d\theta^1}{\theta^1}+\frac{1-\frac{\rho}{\gamma}}{1-\rho}Z^1\frac{d\tilde{J}^1}{\tilde{J}^1}+\frac{1}{2}\frac{\rho}{(1-\rho)^2}Z^1\left(\frac{d\theta^1}{\theta^1}\right)^2+$$
$$+\frac{1}{2}\frac{\left(\rho-\frac{\rho}{\gamma}\right)\left(1-\frac{\rho}{\gamma}\right)}{(1-\rho)^2}Z^1\left(\frac{d\tilde{J}^1}{\tilde{J}^1}\right)^2+\frac{1-\frac{\rho}{\gamma}}{(1-\rho)^2}Z^1\frac{d\theta^1}{\theta^1}\frac{d\tilde{J}^1}{\tilde{J}^1}$$

$$dZ^2 = -\frac{1}{1-\rho}Z^2\frac{\theta^1}{1-\theta^1}\frac{d\theta^1}{\theta^1}+\frac{1-\frac{\rho}{\gamma}}{1-\rho}Z^2\frac{d\tilde{J}^2}{\tilde{J}^2}+\frac{1}{2}\frac{\rho}{(1-\rho)^2}Z^2\left(\frac{\theta^1}{1-\theta^1}\right)^2\left(\frac{d\theta^1}{\theta^1}\right)^2+$$
$$+\frac{1}{2}\frac{\left(\rho-\frac{\rho}{\gamma}\right)\left(1-\frac{\rho}{\gamma}\right)}{(1-\rho)^2}Z^2\left(\frac{d\tilde{J}^2}{\tilde{J}^2}\right)^2-\frac{1-\frac{\rho}{\gamma}}{(1-\rho)^2}Z^2\frac{\theta^1}{1-\theta^1}\frac{d\theta^1}{\theta^1}\frac{d\tilde{J}^2}{\tilde{J}^2}.$$

Since the drift and volatility coefficients of $d\tilde{J}^2/\tilde{J}^2$ vanish as $\theta^1\searrow 0$, and $\lim_{\theta^1\searrow 0}Z^2\left(\theta^1\right) = \left(\tilde{V}^2\right)^{\frac{1-\rho/\gamma}{1-\rho}}$, the drift and volatility coefficients in the equation for dZ^2 vanish. In the equation for dZ^1, it remains to determine the behavior of terms containing $d\tilde{J}^1$ (the remaining contributions to drift and volatility terms converge to zero because $\lim_{\theta^1\searrow 0}Z^1\left(\theta^1\right)=0$):

$$\frac{Z^1}{\tilde{J}^1} = \theta^1\left[\left(\theta^1\right)^{\frac{1}{1-\rho}}\left(\tilde{J}^1\right)^{\frac{1-1/\gamma}{1-\rho}}\right]^{\rho},$$

where the term in brackets is bounded and bounded away from zero by utilizing (45). Using the first θ^1 to multiply the coefficients in $d\tilde{J}^1$ in formula (44), we conclude that the coefficients in

$Z^1 d\tilde{J}^1/\tilde{J}^1$ vanish as $\theta^1 \searrow 0$. Finally, the drift term arising from $\left(d\tilde{J}^1\right)^2$ vanishes, and

$$Z^1\left(\frac{d\tilde{J}^1}{\tilde{J}^1}\right)^2 = \frac{\left(\theta^1\right)^5\left(\tilde{J}_{\theta^1\theta^1}\right)^2}{\tilde{J}+\left(1-\theta^1\right)\tilde{J}_{\theta^1}}\left[\left(\theta^1\right)^{\frac{1}{1-\rho}}\left(\tilde{J}^1\right)^{\frac{1-1/\gamma}{1-\rho}}\right]^\rho\left(\frac{d\theta_t^1}{\theta_t^1}\right)^2.$$

Here, the last term has bounded drift, the second last term is bounded, and the first term converges to zero as $\theta^1 \searrow 0$, which can be shown by using the l'Hôpital's rule (the numerator converges to zero and the denominator to zero or $+\infty$, depending on the sign of γ):

$$\lim_{\theta^1\searrow 0}\frac{\left(\theta^1\right)^5\left(\tilde{J}_{\theta^1\theta^1}\right)^2}{\tilde{J}+\left(1-\theta^1\right)\tilde{J}_{\theta^1}} = \lim_{\theta^1\searrow 0}\frac{5\left(\theta^1\right)^4\tilde{J}_{\theta^1\theta^1}+2\left(\theta^1\right)^5\tilde{J}_{\theta^1\theta^1\theta^1}}{1-\theta^1}=0.$$

Thus all terms in the drift and volatility coefficients of dZ^1 vanish.

Applying Itô's lemma to ζ^2 yields

$$\begin{aligned} d\zeta^2 &= \frac{1}{Z^1+Z^2}dZ^2 - \frac{Z^2}{\left(Z^1+Z^2\right)^2}\left(dZ^1+dZ^2\right)+ \\ &\quad +\frac{Z^2}{\left(Z^1+Z^2\right)^3}\left(dZ^1+dZ^2\right)^2-\frac{1}{\left(Z^1+Z^2\right)^2}dZ^2\left(dZ^1+dZ^2\right) \end{aligned}$$

and the results on the behavior of dZ^1 and dZ^2 as $\theta^1 \searrow 0$ lead to the desired conclusion about the convergence of drift and volatility coefficients of $d\zeta^2$. ■

Proof of Proposition 10. Convergence of the risk-free interest rate follows from the direct calculation of

$$r(0) = \lim_{t\searrow 0} -\frac{1}{t}\log E\left[M_t^2 S_t^2(0) \mid \mathcal{F}_0\right]$$

where $S_t^2(0)$ is the limiting stochastic discount factor corresponding to the one prevailing in a homogeneous economy populated only by agent 2. Lemma 9 shows that the local behavior of S_t^2 converges to $S_t^2(0)$ as $\theta_0^1 \searrow 0$. Similarly, convergence of the wealth-consumption ratio follows from

$$\xi\left(\theta^1\right)=\xi^1\left(\theta^1\right)\zeta^1\left(\theta^1\right)+\xi^2\left(\theta^1\right)\zeta^2\left(\theta^1\right).$$

Since $\xi^n\left(\theta^1\right)$ are bounded and $\zeta^1\left(\theta^1\right)$ converges to zero, we have

$$\lim_{\theta^1\searrow 0}\xi\left(\theta^1\right)=\lim_{\theta^1\searrow 0}\xi^2\left(\theta^1\right)=\frac{1}{\beta}\left(\tilde{V}^2\right)^\rho,$$

where $\tilde{V}^2$ is given by (7).

In order to obtain the convergence of the infinitesimal return, observe that

$$\xi^1\left(\theta^1\right)\zeta^1\left(\theta^1\right)=\beta^{-1}\theta^1\tilde{J}^1\left(\theta^1\right)\left[Z^1\left(\theta^1\right)+Z^2\left(\theta^1\right)\right]^{\rho-1}$$

and

$$d\left[\theta^1\tilde{J}^1\left(\theta^1\right)\right]=\theta^1\tilde{J}^1\left(\theta^1\right)\frac{d\theta^1}{\theta^1}+\theta^1 d\tilde{J}^1\left(\theta^1\right)+\theta^1 d\tilde{J}^1\left(\theta^1\right)\frac{d\theta^1}{\theta^1}.$$

The drift and volatility coefficients of the first term on the right-hand side vanish as $\theta^1\searrow 0$ by the proof of Lemma 8, and the coefficients of the other two terms vanish by combining the results in that Lemma with equation (44). Further,

$$\begin{aligned} d\left\{\left[Z^1+Z^2\right]^{\rho-1}\right\} &= (\rho-1)\left[Z^1\left(\theta^1\right)+Z^2\left(\theta^1\right)\right]^{\rho-2}\left(dZ^1+dZ^2\right)+ \\ &+\frac{1}{2}(\rho-2)(\rho-1)\left[Z^1\left(\theta^1\right)+Z^2\left(\theta^1\right)\right]^{\rho-3}\left(dZ^1+dZ^2\right)^2 \end{aligned}$$

and since dZ^1 and dZ^2 have vanishing coefficients by the proof of Lemma 9 and the remaining terms are bounded, we obtain that $d\xi^1\left(\theta^1\right)\zeta^1\left(\theta^1\right)$ has vanishing drift and volatility coefficients as $\theta^1\searrow 0$. The same argument holds for $d\xi^2\left(\theta^1\right)\zeta^2\left(\theta^1\right)$, and thus $d\xi\left(\theta^1\right)$ has vanishing coefficients as well. Therefore all but the first term in

$$d\Xi_t=d\left[\xi\left(\theta^1_t\right)Y_t\right]=\Xi_t\frac{dY_t}{Y_t}+Y_t d\xi\left(\theta^1_t\right)+d\xi\left(\theta^1_t\right)dY_t$$

have coefficients that decline to zero as $\theta^1_t\searrow 0$, which proves the result. ∎

Proof of Proposition 11. The evolution of θ^1 given by equation (18) implies that for every fixed $t\geq 0$

$$\theta^1_0\searrow 0 \qquad\Longrightarrow\qquad \theta^1_t\to 0,\ P\text{-a.s.}$$

and thus also $\zeta^2\left(\theta^1_t\right)\to 1$ and $\tilde{J}^2\left(\theta^1_t\right)\to\tilde{V}^2$, P-a.s.[12] The last two terms in the expression for the stochastic discount factor, S^2_t, equation (24), converge to one, P-a.s., and since $\nu^2\left(\theta^1_s\right),\ 0\leq s\leq t$ also converges to $\nu^2(0)$ and is bounded, we have $S^2_t\xrightarrow{P}S^2_t(0)$. Consider a family of random variables $M^2_t S^2_t\left(\theta^1_0\right)$ indexed by the initial Pareto share θ^1_0. Since this family is uniformly integrable,

[12]This result becomes more transparent if we consider ζ^2 and $\tilde{J}^2$ as functions of $\log\theta^1$. The dynamics of $\log\theta^1$

$$\begin{aligned} d\log\theta^1_t &= \left(1-\theta^1_t\right)\left[\nu^2_t\left(\theta^1_t\right)-\nu^1\left(\theta^1_t\right)+\frac{1}{2}\left(\left(u^2\right)^2-\left(u^1\right)^2\right)-\frac{1}{2}\theta^1_t\left(u^1-u^2\right)^2\right]dt+ \\ &+\left(1-\theta^1_t\right)\left(u^1-u^2\right)dW_t \end{aligned}$$

has bounded drift and volatility coefficients and thus for $\forall\varepsilon>0$, $\forall k>0$, it is possible to achieve

$$P\left[\theta^1_t<k\right]=P\left[\log\theta^1_t<\log k\right]>1-\varepsilon$$

by setting $\log\theta^1_0$ sufficiently low.

then convergence in probability implies convergence in mean, and we obtain the convergence result for bond prices

$$E\left[M_t^2 S_t^2\left(\theta_0^1\right) \mid \mathcal{F}_0\right] \xrightarrow{\theta_0^1 \searrow 0} E\left[M_t^2 S_t^2(0) \mid \mathcal{F}_0\right].$$

The same argument holds for $M_t^2 S_t^2\left(\theta_0^1\right) Y_t$, which yields the result for the price of individual cash flows from the aggregate endowment. ∎

For the proof of the next lemma, the following result will be useful:

Lemma 18 *Let $f:\mathbb{R}\to\mathbb{R}$ be differentiable with a monotone first derivative in a neighborhood of $-\infty$ and have a finite limit $\lim_{x\to-\infty} f(x)$. Then $\lim_{x\to-\infty} f'(x)=0$.*

Proof of Lemma 12. Transformation (27) together with the previously used $\gamma V_t^1 = Y^\gamma \tilde{J}^1\left(\theta_t^1\right)$ imply that

$$\hat{V}^1\left(\theta^1\right) = \beta^\gamma \left(\frac{\tilde{J}^1\left(\theta^1\right)^{1/\gamma}}{\zeta^1\left(\theta^1\right)}\right)^{\gamma(1-\rho)}. \tag{46}$$

Think for a moment of $\hat{V}^1$ as a function of $\log\theta^1$, where we are interested in the limiting behavior as $\log\theta^1 \to -\infty$. We have

$$\theta^1 \hat{V}_{\theta^1}^1 = \hat{V}_{\log\theta^1}^1 \text{ and } \left(\theta^1\right)^2 \hat{V}_{\theta^1\theta^1}^1 = \hat{V}_{(\log\theta^1)^2}^1 - \hat{V}_{\log\theta^1}^1. \tag{47}$$

Differentiating repeatedly expression (46) and exploiting the local behavior of $\tilde{J}\left(\theta^1\right)$ as $\theta^1 \searrow 0$, we conclude that the assumptions of Lemma 18 hold, and thus both expressions in (47) converge to zero as $\theta^1 \searrow 0$. ∎

Proof of Lemma 13. Utilizing Lemma 12 to deduce which terms in ODE (29) vanish and Proposition 10 to determine the limiting values of the remaining coefficients, we obtain

$$\begin{aligned}\lim_{\theta^1\searrow 0} \beta^{\frac{1}{1-\rho}}\left(\hat{V}^1\left(\theta^1\right)\right)^{-\frac{\rho}{\gamma}\frac{1}{1-\rho}} &= \beta - \rho\left(\mu_y + u^2\sigma_y - \frac{1}{2}(1-\gamma)\left(\sigma_y\right)^2\right) - \\ &\quad -\frac{\rho}{1-\rho}\left[\left(u^1-u^2\right)\sigma_y + \frac{1}{2}\frac{\left(u^1-u^2\right)^2}{1-\gamma}\right],\end{aligned}$$

which is the limiting consumption-wealth ratio for agent 1. The formulas for the wealth share invested in the claim on aggregate consumption and the coefficients of the wealth process are obtained by plugging in the previous results into expressions (26) and (30). ∎

Proof of Proposition 14. Given convergence to the homogeneous economy counterpart, the expression for $\lim_{\theta^1\searrow 0}\nu^2\left(\theta^1\right)$ is given by equation (8). Utilizing the formula for the wealth-consumption ratio (20) and the result from Lemma 13 then yields

$$\lim_{\theta^1 \searrow 0} \nu^1\left(\theta^1\right) = \lim_{\theta^1 \searrow 0} \beta\frac{\gamma}{\rho} + (\rho-\gamma)\left[\xi^1\left(\theta^1\right)\right]^{-1} = \beta + (\gamma-\rho)\left(\mu_y + u^2\sigma_y - \frac{1}{2}(1-\gamma)\sigma_y^2\right) + \\ + \frac{\gamma-\rho}{1-\rho}\left[\left(u^1-u^2\right)\sigma_y + \frac{1}{2}\frac{\left(u^1-u^2\right)^2}{1-\gamma}\right].$$

The first two terms in the last expression are equal to the limit for $\nu^2\left(\theta^1\right)$, which yields the result for the difference of the discount rates. The expression for part (ii) is obtained by symmetry. ■

Proof of Corollary 15. The critical point is the limits for the consumption-wealth ratios as the Pareto share of one of the agents becomes small. Since the large agent's consumption-wealth ratio converges to that in a homogeneous economy, the relevant parameter restriction is the same as restriction (9) in Assumption 2. The consumption-wealth ratio of the small agent is given in expression (20), and restriction (10) in Assumption 2 assures that this quantity is strictly positive, and the wealth-consumption ratio finite. ■

Proof of Corollary 16. Utilize results in Proposition 13 and the fact that $\lim_{\theta^1 \searrow 0} \mu_{\Xi^2}\left(\theta^1\right) = \mu_y$ and $\lim_{\theta^1 \searrow 0} \sigma_{\Xi^2}\left(\theta^1\right) = \sigma_y$, then form the differences in the limiting expected logarithmic growth rates, and compare them to inequalities in Proposition 6. ■

Proof of Corollary 17. The results are obtained by taking limits of the expressions in Proposition 14. ■

References

Alchian, Armen A., 1950. "Uncertainty, Evolution, and Economic Theory." *Journal of Political Economy*. 58 (3): 211–221.

Anderson, Evan W., 2005. "The Dynamics of Risk-Sensitive Allocations." *Journal of Economic Theory*. 125 (2): 93–150.

Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent, 2003. "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection." *Journal of the European Economic Association*. 1 (1): 68–123.

Basak, Suleyman, and Domenico Cuoco, 1998. "An Equilibrium Model with Restricted Stock Market Participation." *Review of Financial Studies*. 11 (2): 309–341.

Bhamra, Harjoat S., and Raman Uppal, 2009. "Asset Prices with Heterogeneity in Preferences and Beliefs." Working paper.

Blume, Lawrence, and David Easley, 1992. "Evolution and Market Behavior." *Journal of Economic Theory*. 58 (1): 9–40.

Blume, Lawrence, and David Easley, 2006. "If You're so Smart, why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." *Econometrica*. 74 (4): 929–966.

Chen, Hui, Scott Joslin, and Ngoc-Khanh Tran, 2010. "Rare Disasters and Risk Sharing with Heterogeneous Beliefs." Working paper, MIT.

Colacito, Riccardo, and Mariano M. Croce, 2010. "Risk Sensitive Allocations with Multiple Goods: Existence, Survivorship, and the Curse of the Linear Approximation." Working paper.

Cuoco, Domenico, and Hua He, 2001. "Dynamic Aggregation and Computation of Equilibria in Finite-Dimensional Economies with Incomplete Financial Markets." *Annals of Economics and Finance*. 2: 265–296.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1991. "The Survival of Noise Traders in Financial Markets." *The Journal of Business*. 64 (1): 1–19.

Duffie, Darrell, and Larry G. Epstein, 1992a. "Asset Pricing with Stochastic Differential Utility." *The Review of Financial Studies*. 5 (3): 411–436.

Duffie, Darrell, and Larry G. Epstein, 1992b. "Stochastic Differential Utility." *Econometrica*. 60 (2): 353–394.

Duffie, Darrell, and Pierre-Louis Lions, 1992. "PDE Solutions of Stochastic Differential Utility." *Journal of Mathematical Economics.* 21 (6): 577–606.

Dumas, Bernard, Raman Uppal, and Tan Wang, 2000. "Efficient Intertemporal Allocations with Recursive Utility." *Journal of Economic Theory.* 93 (2): 240–259.

El Karoui, Nicole, Shige Peng, and Marie Claire Quenez, 1997. "Backward Stochastic Differential Equations in Finance." *Mathematical Finance.* 7 (1): 1–71.

Epstein, Larry G., 1987. "The Global Stability of Efficient Intertemporal Allocations." *Econometrica.* 55 (2): 329–355.

Epstein, Larry G., and Jianjun Miao, 2003. "A Two-Person Dynamic Equilibrium under Ambiguity." *Journal of Economic Dynamics and Control.* 27 (7): 1253–1288.

Epstein, Larry G., and Stanley E. Zin, 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica.* 57 (4): 937–969.

Epstein, Larry G., and Stanley E. Zin, 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *The Journal of Political Economy.* 99 (2): 263–286.

Evstigneev, Igor V., Thorsten Hens, and Klaus Reiner Schenk-Hoppé, 2006. "Evolutionary Stable Stock Markets." *Economic Theory.* 27 (2): 449–468.

Friedman, Milton, 1953. *Essays in Positive Economics*, University of Chicago Press, Chicago.

Gârleanu, Nicolae, and Stavros Panageas, 2010. "Young, Old, Conservative and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing." Working paper.

Geoffard, Pierre-Yves, 1996. "Discounting and Optimizing: Capital Accumulation Problems as Variational Minmax Problems." *Journal of Economic Theory.* 69 (1): 53–70.

Hansen, Lars Peter, 2004. "Comment on 'Backus, Routledge, and Zin (2004) Exotic Preferences for Macroeconomists'." *NBER Macroeconomics Annual.* 19: 391–405.

Harrison, J. Michael, and David M. Kreps, 1979. "Martingales and Arbitrage in Multiperiod Securities Markets." *Journal of Economic Theory.* 20 (3): 381–408.

Hommes, Cars H., 2006. "Heterogeneous Agent Models in Economics and Finance." in Leigh Tesfatsion, and Kenneth L. Judd (ed.), *Handbook of Computational Economics: Volume 2*chap. 23, pp. 1109–1186. Elsevier.

Jouini, Elyès, and Clotilde Napp, 2007. "Consensus Consumer and Intertemporal Asset Pricing with Heterogeneous Beliefs." *Review of Economic Studies*. 74 (4): 1149–1174.

Kan, Rui, 1995. "Structure of Pareto Optima When Agents Have Stochastic Recursive Preferences." *Journal of Economic Theory*. 66 (2): 626–631.

Karlin, Samuel, and Howard M. Taylor, 1981. *A Second Course in Stochastic Processes*, Academic Press.

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield, 2006. "The Price Impact and Survival of Irrational Traders." *The Journal of Finance*. 61 (1): 195–229.

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield, 2009. "Market Selection." Working paper.

Koopmans, Tjalling C., 1960. "Stationary Ordinal Utility and Impatience." *Econometrica*. 28 (2): 287–309.

Kreps, David M., and Evan L. Porteus, 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory." *Econometrica*. 46 (1): 185–200.

Lucas, Robert E., Jr., and Nancy L. Stokey, 1984. "Optimal Growth with Many Consumers." *Journal of Economic Theory*. 32 (1): 139–171.

Mazoy, José M., 2005. "Consumption Allocation Dynamics." PhD dissertation, University of Chicago.

Merton, Robert C., 1971. "Optimum Consumption and Portfolio Rules in a Continuous-time Model." *Journal of Economic Theory*. 3 (4): 373–413.

Meyn, Sean P., and Richard L. Tweedie, 1993. "Stability of Markovian Processes II: Continuous-Time Processes and Sampled Chains." *Advances in Applied Probability*. 25 (3): 487–517.

Morris, Stephen, 1995. "The Common Prior Assumption in Economic Theory." *Economics and Philosophy*. 11 (2): 227–253.

Negishi, Takashi, 1960. "Welfare Economics and Existence of an Equilibrium for a Competitive Economy." *Metroeconomica*. 12 (2–3): 92–97.

Sandroni, Alvaro, 2000. "Do Markets Favor Agents Able to Make Accurate Predictions?." *Econometrica*. 68 (6): 1303–1341.

Schroder, Mark, and Costis Skiadas, 1999. "Optimal Consumption and Portfolio Selection with Stochastic Differential Utility." *Journal of Economic Theory*. 89 (1): 68–126.

Skiadas, Costis, 1997. "Conditioning and Aggregation of Preferences." *Econometrica*. 65 (2): 347–367.

Skiadas, Costis, 2003. "Robust Control and Recursive Utility." *Finance and Stochastics*. 7 (4): 475–489.

Uppal, Raman, and Tan Wang, 2003. "Model Misspecification and Underdiversification." *The Journal of Finance*. 58 (6): 2465–2486.

Uzawa, Hirofumi, 1968. "Time Preference, the Consumption Function, and Optimum Asset Holdings." in J. N. Wolfe (ed.), *Value, Capital, and Growth: Papers in Honor of Sir John Hicks*Edinburgh University, Edinburgh.

Weil, Philippe, 1990. "Nonexpected Utility in Macroeconomics." *The Quarterly Journal of Economics*. 105 (1): 29–42.

Yan, Hongjun, 2008. "Natural Selection in Financial Markets: Does It Work?." *Management Science*. 54 (11): 1935–1950.