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

Jaroslav Borovička

New York University

jaroslav.borovicka@nyu.edu

July 26, 2013

Abstract

I study the long-run behavior of an economy with two types of agents who differ in their beliefs and are endowed with homothetic recursive preferences of the Duffie-Epstein-Zin type. Contrary to models with separable preferences in which the wealth of agents with incorrect beliefs vanishes in the long run, recursive preference specifications also lead to equilibria where both agents survive, or more incorrect agents dominate. 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 in which agents with differently accurate beliefs coexist in the long run, and show that these equilibria exist for broad ranges of 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.

*I am indebted to Lars Peter Hansen for his advice and continuous support. I appreciate helpful comments from Fernando Alvarez, David Backus, Gadi Barlevy, Katarína Borovičková, Hui Chen, Valentin Haddad, Narayana Kocherlakota, Virgiliu Midrigan, 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. 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.

Survival of agents with incorrect beliefs has been studied extensively in complete market models populated by agents endowed with separable preferences, and the literature has indeed found widespread support for the market selection hypothesis. If agents have identical preferences and relative risk aversion is bounded, then only agents with the most accurate beliefs survive in the long run in the sense that their wealth share does not converge to zero, regardless of the specification of the aggregate endowment process.

While these results look appealing, separable preferences are inconsistent even with many features of the asset pricing data, a fact reflected in different asset pricing 'puzzles'. Since the market selection hypothesis focuses on wealth dynamics of agents who trade in financial markets, it is sensible to study the survival problem for preferences which provide a better fit of the constructed models to empirically observed patterns in asset returns.

This paper shows that the above survival results are overturned when the assumption of separability of preferences is relaxed. I study a class of 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 disentangle the risk aversion with respect to intratemporal gambles from the intertemporal elasticity of substitution (IES), and include the separable, constant relative risk aversion (CRRA) utility as a special case. Thanks to the additional degree of flexibility, this class of preferences became the workhorse model used in the macroeconomics and asset pricing literature.

In addition, the decoupling of risk aversion and IES effectively separates two essential mechanisms for wealth accumulation that play an important role in the analysis of long-run wealth dynamics. Agents in an 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.

From the perspective of an agent who maximizes the present discounted value of future utility flow, rationality itself does not guarantee survival, nor do deviations from rational preferences imply extinction. Rationality may *facilitate* survival only insofar it prevents overconsumption and leads the agent to take appropriate bets. On the other hand, specific forms of deviations from rational beliefs may, at least in theory, provide even stronger incentives for survival in the long run, despite not being optimal in the rational sense.

I show that in the class of recursive preferences, there exist broad ranges of empirically plausible values for preference parameters under which agents with less accurate beliefs can survive or even dominate the economy. Perhaps most interestingly, agents with arbitrarily large belief distortions can coexist with rational agents in the long-run equilibrium under preference parameterizations adopted by a substantial share of the asset pricing literature. From the perspective of market selection, belief heterogeneity should thus be viewed as a natural long-run outcome.

While the literature has to a large extent focused on the characterization of optimal consumption allocations using a planner's problem, I also study the associated decentralization. This allows me to explain the crucial role of the equilibrium price mechanism in the long-run wealth dynamics. If an agent is to survive in the long run, then equilibrium prices in situations when her wealth share becomes small have to move in a direction that encourages her to accumulate wealth at a faster rate than the growth rate of aggregate wealth. Thanks to the analytical nature of the results, I am able to completely characterize the survival outcomes on the whole parameter space of the model.

I study an endowment economy populated by two classes of competitive agents (called, for simplicity, two agents) who differ in their beliefs about the growth rate of the stochastic aggregate endowment that follows a geometric Brownian motion. Agents are endowed with identical recursive preferences and trade in complete markets which, in the Brownian information setup, corresponds to dynamic trading in an infinitesimal risk-free bond and a risky claim to the aggregate endowment.

To shed more light on the mechanism that generates the long-run coexistence of agents with heterogeneous beliefs, it is illustrative to analyze the situation when the wealth of one of the agents becomes negligible. We want to study the incentives of this negligible agent for wealth accumulation that would prevent her wealth share from vanishing to zero. First notice that the wealth accumulation of the agent with the large wealth share is disciplined by market clearing — in the limit as the wealth share of the negligible agent converges to zero, the large agent has to hold the market portfolio and consume the aggregate endowment. Her wealth then grows at the same rate as the aggregate endowment and market prices can be inferred from the dynamics of her stochastic discount factor. The negligible agent then chooses an investment portfolio that overweighs positions in assets that are, according to her own beliefs, cheap and earn high expected *level* returns relative to their risk.

The first insight draws the distinction between expected level and logarithmic returns. While the optimal portfolio choice is determined by the tradeoff between the expected level return and the underlying volatility, survival chances depend on the expected *logarithmic* growth rate of wealth, and thus on the expected *logarithmic* return on the agent's portfolio. Volatile portfolios are therefore detrimental to survival, and the negligible agent will be more willing to choose a volatile portfolio if the risk aversion in the economy is low.

More precisely, holding other parameters fixed, there is always a level of risk aversion sufficiently low such that only one agent survives in the long run but it can be either of the two agents with a strictly positive probability, regardless of the relative accuracy of their beliefs. The reason is that as one of the agents observes a series of shock realizations that significantly decreases her wealth share, she starts choosing portfolios with very volatile returns which make the probability of further large wealth losses even more likely.

The second insight concerns the role of the IES parameter for the consumption-saving decision. The negligible agent chooses a portfolio that generates a high subjective expected level return, relative to its risk. While this portfolio can perform poorly under the true probability measure if the beliefs of the agent are distorted, it is the *subjective* expected level return that determines the agent's willingness to save.

When IES is higher than one, then the saving rate is an increasing function of the subjective expected level return on the agent's portfolio. As long as the negligible agent chooses a portfolio with a higher subjective expected level return than her large counterpart, then a high IES is conducive to the survival of the negligible agent who, by choosing a high saving rate, can compensate for the potentially inferior choice of her portfolio, and 'outsave' her extinction.

Observe that this consumption-saving mechanism under a high IES operates for the negligible agent, regardless of her identity. When the relative wealth shares of the agents switch, then it is again the negligible agent that can outsave her extinction by choosing a high saving rate vis-à-vis her high subjective expected level return. If this mechanism, which operates through endogenously determined equilibrium prices, is sufficiently strong, both agents can survive in the long run.

Finally, the third insight captures the role of risk premia and the associated advantage of optimistic agents who overweigh investment into risky assets with high expected returns. An increase in risk aversion reduces the amount of betting and the portfolios of the two agents become more alike, so that differences in beliefs that make the more incorrect agents choose portfolios with a suboptimal risk-return tradeoff diminish in importance. At the same time, the increase in risk aversion increases risk premia, so that the difference in expected returns on the portfolios of the two agents does not vanish. This benefits the relatively more optimistic agent who holds a larger share of her wealth invested in the risky asset. The consumption-saving mechanism is not sufficiently strong for high levels of risk aversion to make the less optimistic agent outsave her more optimistic counterpart.

As a consequence, holding other parameters fixed, there is always a level of risk aversion above which the relatively more optimistic agent dominates, even if her beliefs are relatively more incorrect.

These results indicate under which preference parameterizations can two agents with heterogeneous beliefs coexist in the long run. We are looking for parameterizations under which either of the agents chooses a high logarithmic rate of wealth accumulation whenever she becomes negligible and faces the risk of extinction. The results imply that risk aversion has to be sufficiently high to prevent excessive risk taking by the negligible agent but not excessive so as to motivate the relatively more pessimistic agent to choose a portfolio with a high subjective expected return when she becomes negligible. A sufficiently high IES then makes the negligible agent choose a high saving rate vis-à-vis the high subjective expected return on her portfolio, thus outsaving her extinction.

The encouraging observation is that this region of the parameter space is also considered by the macroeconomics and asset pricing literatures as empirically relevant. Moreover, the described mechanism does not hinge on belief distortions being small. On the contrary, when belief differences are large, the negligible agent can more easily construct portfolios with high perceived expected returns because asset prices determined by the wealth dynamics of the large agent are viewed as more incorrect by the negligible agent. Under a high IES, the high perceived expected return on her portfolio will make the negligible agent choose a high saving rate. I find that there is a broad region of the parameter space in which an arbitrarily incorrect agent coexists in the long run with an agent with perfectly accurate beliefs.

These results are also independent of the growth rate of the economy and are not driven by the particular nature of the stochastic process for the aggregate endowment. In fact, the results also hold for an economy with a deterministic aggregate endowment, as long as there is an observable shock (without any material impact on aggregate endowment) with a probability distribution about which the agents disagree that can be used by the two agents as a betting device.

The portfolio choice and consumption-saving decision mechanism eluded the attention of the survival literature because the long-run wealth dynamics in a complete-market economy when agents have separable utility can be solved for using a planner's problem without the need for a decentralization. Under CRRA preferences, risk aversion and IES are inversely related, and the two mechanisms described above offset each other. Increasing the risk aversion disciplines the portfolio choice of the negligible agent but the associated decrease in IES makes her decrease her saving rate, undermining her survival chances.

The description of the survival mechanism also highlights the critical role of the endogenously determined equilibrium price dynamics. In order for the two agents to coexist, equilibrium prices always have to be conducive to the survival of the negligible agent, and thus have to adjust when the roles of the two agents switch. I show that the behavior of the large agent is always disciplined by market clearing while the choices of the negligible agent can be inferred from a decision problem with prices determined by the wealth dynamics of the large agent.

The central message of this analysis is that the market selection hypothesis is not robust to departures from separable preferences. Recursive preferences provide an additional degree of freedom compared to the separable case that allows one to separate the portfolio allocation and the consumption-saving decision, two crucial aspects determining the rate of wealth accumulation. Contrary to the case of CRRA preferences, belief heterogeneity under plausibly parameterized recursive preferences is a pervasive long-run outcome.

1.1 Methodology and literature overview

The modern approach in the market survival literature 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.[1]

Subsequent research has shown that taking into account general equilibrium effects and intertemporal optimization of agents endowed with separable preferences eliminates much of the support for survival of agents with incorrect 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, the local properties of the utility function are

[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 (2013) provide a more general survey that also focuses on asset pricing implications of belief and preference heterogeneity. See also the discussion of price impact by Kogan, Ross, Wang, and Westerfield (2011) and portfolio impact by Cvitanić and Malamud (2011).

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.

immaterial for survival. 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 the survival of agents with relatively more incorrect beliefs. Kogan, Ross, Wang, and Westerfield (2011) show that a sufficient condition to prevent this outcome is the boundedness of the relative risk aversion function, i.e., a condition on the preferences being uniformly 'close' to the homothetic CRRA case. On the other hand, in this paper, preferences are homothetic, which assures that the survival results are not driven by exogenous differences in the local properties of the utility functions.[2]

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 analyze the survival mechanism 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 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 nondegen-

[2] The survival literature also focuses on other forms of heterogeneity. Yan (2008) and Muraviev (2013) construct 'survival indices' that combine the contribution of belief distortions and preference parameters and show that only agents with the lowest survival index can survive.

Market incompleteness or asymmetric information may be other ways how to counteract the extinction of agents with incorrect beliefs, as long they are judiciously chosen to prevent agents to place incorrect bets, see, e.g., Mailath and Sandroni (2003), Coury and Sciubba (2012), Cao (2010) or Cogley, Sargent, and Tsyrennikov (2013).

erate 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 a response to the equilibrium price dynamics 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. Guerdjikova and Sciubba (2010) study the interaction of expected-utility and smooth ambiguity averse agents. Branger, Dumitrescu, Ivanova, and Schlag (2011) analyze survival in long-run risk models with heterogeneous recursive preferences. 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. Sections 3 and 4 present 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 5. Section 6 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.[3]

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

[3] https://files.nyu.edu/jb4457/public/files/research/survival_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 (see, e.g., Huang and Pagès (1992)). Under the parameter restrictions below, constructed equilibria satisfy these assumptions.

The scalar aggregate endowment process Y follows a geometric Brownian motion

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

with given 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 that they use to assign probabilities to future events. I model the belief distortion of agent n using an adapted process u^n such that the process

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

is a martingale under P. The martingale M^n is called the Radon-Nikodým derivative or the belief ratio and defines the subjective probability measure Q^n that characterizes the beliefs of agent n. The Radon-Nikodým derivative measures the disparity between the subjective and true probability measures.

In order for the belief heterogeneity not to vanish in the long run, the measures P and Q^n cannot be mutually absolutely continuous.[4] However, given the construction of M^n, the restrictions of the measures P and Q^n, $n \in \{1, 2\}$ to $\mathcal{F}_t$ for every $t \geq 0$ are equivalent.[5] In other words, the agents agree with the data generating measure on zero-probability finite-horizon events. While a likelihood evaluation of 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 she cannot refute her 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 belief distortion process u^n has a clear economic interpretation. The Girsanov theorem implies that agent n, whose deviation from rational beliefs is described by u^n,

[4] Sandroni (2000) and Blume and Easley (2006) link absolute continuity of the subjective probability measures to merging of agents' beliefs.

[5] See, for example, Revuz and Yor (1999), Section VIII for details. The construction prevents arbitrage opportunities in finite-horizon strategies. The martingale representation theorem (e.g., Øksendal (2007), Theorem 4.3.4) implies that modeling belief distortions under Brownian information structures using martingales of the form (2) is essentially without loss of generality.

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 the growth rate of Y. When $\sigma_y = 0$, the distinction between optimism and pessimism 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. 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.[6]

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.[7] They show that recursive preferences can be represented as a solution to the maximization problem

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

[6] 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 rely on 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.

[7] 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.

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).

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

$$F(C, \nu) = \beta \frac{C^\gamma}{\gamma} \left(\frac{\gamma - \rho \frac{\nu}{\beta}}{\gamma - \rho} \right)^{1 - \frac{\gamma}{\rho}}, \tag{5}$$

with parameters satisfying $\gamma, \rho < 1$, and $\beta > 0$. Preferences specified by this felicity function[8] 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

$$\bar{\lambda}_t^n V_t^n = \sup_{\nu^n} E_t \left[\int_t^\infty \bar{\lambda}_s^n F(C_s^n, \nu_s^n) \, ds \right] \tag{6}$$

subject to

$$\frac{d\bar{\lambda}_t^n}{\bar{\lambda}_t^n} = -\nu_t^n dt + u_t^n dW_t, \quad t \geq 0; \quad \bar{\lambda}_0^n = 1. \tag{7}$$

The problem (6–7) 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,

[8]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.

and Wang (2000) — the only difference is that the modified discount factor process is not locally predictable.

The diffusion term $u^n_s dW_s$ 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 (7) indicates that under the true probability measure, positive realizations of dW_s increase the term $d\bar{\lambda}^n_s / \bar{\lambda}^n_s$, 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 $V^n_t(Y) = \gamma^{-1} Y^\gamma_t \tilde{V}^n$ where

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

with the associated discount rate

$$\nu^n = \frac{\beta}{\rho} \left(\gamma + (\rho - \gamma) \left(\tilde{V}^n \right)^{-\frac{\rho}{\gamma}} \right) = \beta + (\gamma - \rho) \left(\mu_y + u^n \sigma_y - \frac{1}{2} (1-\gamma) \sigma^2_y \right). \tag{9}$$

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

$$\beta > \max_n \rho \left(\mu_y + u^n \sigma_y - \frac{1}{2} (1-\gamma) \sigma^2_y \right), \tag{10}$$

$$\beta > \max_n \rho \left(\mu_c + u^{\sim n} \sigma_y - \frac{1}{2} (1-\gamma) \sigma^2_y \right) + \frac{\rho}{1-\rho} \left[\left(u^n - u^{\sim n} \right) \sigma_y + \frac{1}{2} \frac{\left(u^n - u^{\sim n} \right)^2}{1-\gamma} \right] \tag{11}$$

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 the 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

$$J_0(\alpha) = \sup_{(C^1, C^2, \nu^1, \nu^2)} \sum_{n=1}^{2} E_0 \left(\int_0^\infty \bar{\lambda}_t^n F(C_t^n, \nu_t^n) \, dt \right) \tag{12}$$

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

$$\frac{d\bar{\lambda}_t^n}{\bar{\lambda}_t^n} = -\nu_t^n dt + u_t^n dW_t, \;\; t \geq 0; \;\; \bar{\lambda}_0^n = \alpha^n \tag{13}$$

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 (12–13) is a straightforward extension when individual continuation values are well-defined.

The planner's problem (12–13) 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.[9,10]

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

[9] 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.

[10] 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 (12–13).

— the nonseparable preference structure and the belief distortion that drives the diffusion component in (13). Belief heterogeneity introduces an additional risk component $u_t^n dW_t$ arising through the stochastic reweighing of wealth shares which will have a direct impact on local risk prices.

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 6.

The planner's problem has an appealing Markov structure. Homogeneity of the planner's problem (12–13) 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 3 *The objective function for the planner's problem (12–13) is*

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

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

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

with boundary conditions $\tilde{J}(0) = \tilde{V}^2$ and $\tilde{J}(1) = \tilde{V}^1$, where $\tilde{V}^n$ are defined in (8). The functions $\tilde{J}^n(\theta^1)$ 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 in general have a 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. Indeed, the crucial part of the solution is the law of motion for the state variable θ^1 that dictates how the planner adjusts the weights of the two agents, and thus their current consumption and wealth, over time. In this respect, the only relevant force for survival is the willingness of the planner to increase the Pareto weight of the agent that becomes negligible and faces the risk of becoming extinct, and thus only the boundary behavior of $\tilde{J}(\theta^1)$ matters. 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 (12–13). 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 (2011) 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 when preferences are separable. Separable utility is obtained a special case of the variational utility (3–4) with an optimal discount rate choice $\nu^n = \beta$ where β is the time preference coefficient and the period utility function $F(C, \beta) \doteq U(C)$. The first-order condition for the planner's problem leads to the static equation

$$\alpha^1 M_t^1 U'\left(C_t^1\right) = \alpha^2 M_t^2 U'\left(C_t^2\right).$$

Survival analysis thus corresponds to analyzing a sequence of state- and time-indexed static problems that are interlinked only by the initial Pareto weights α^n, 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 with $\lim_{s\to\infty} M_{t+s}^1 = 0$ (P-a.s.) and $M_t^2 \equiv 1$, 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 $\lim_{s\to\infty} \zeta_{t+s}^1 / \zeta_{t+s}^2 = 0$ (P-a.s.).

When preferences are not separable, this straightforward link breaks down because marginal utilities of consumption also depend on continuation values and the first-order conditions involve 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 4 *Agent 1 becomes extinct along the path* $\omega \in \Omega$ *if* $\lim_{t\to\infty} \theta_t^1(\omega) = 0$. *Otherwise, agent 1 survives along the path* ω. *Agent 1 dominates in the long run along the path* ω *if* $\lim_{t\to\infty} \theta_t^1(\omega) = 1$.

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

Kogan, Ross, Wang, and Westerfield (2011) 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 (13). 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 $\theta_t^1 \equiv \alpha^1/(\alpha^1+\alpha^2)$. 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 (12–13) and are given by

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

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.[11] 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 5 *Define the following 'repelling' 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.*

(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 $q\left(\theta^1\right)$. The proof of Proposition 5 is based on the classification of boundary behavior of diffusion processes, discussed in Karlin and Taylor (1981). The four 'attracting' and 'repelling' conditions are only sufficient and their combinations stated in Proposition 5 are not exhaustive. However, the only unresolved cases are knife-edge cases involving equalities in the conditions of the Proposition, which are only of limited importance in the analysis below.

I call the difference in the discount rates $\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)$ *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 5 have an intuitive interpretation.

[11] The unattainability of the boundaries follows from the proof of Proposition 5.

Survival condition (i) states that agent 1 survives under the true probability measure even in cases when her beliefs are more distorted, $|u^1| > |u^2|$, as long as her relative patience becomes sufficiently high to overcome the distortion when her 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 the interaction of current discounting with the dynamics of the continuation values that reflects the behavior 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 (2011) show that in the economy presented in this paper under CRRA preferences, the agent whose beliefs are less distorted dominates in the long run under measure P. The conditions in Proposition 5 confirm these results as follows:

Corollary 6 *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 always 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 6 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 (2011), 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 6 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 5 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 $r_t = r\left(\theta_t^1\right)$ 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. This implies that an agent that becomes extinct in the long run also has no long-run *price impact* on the two assets that are traded.

These results are, however, even stronger because they also state that when the wealth of an agent becomes negligible, she has no impact on the current prices of the two assets even if she survives in the long run and her wealth recovers in the future. The ability to pin down asset returns when the wealth of one agent is negligible even though she may survive in the long run plays a crucial role in the analysis because it allows me to determine the wealth dynamics of the two agents in the proximity of the boundary by solving two straightforward portfolio choice problems. The solutions then yield in closed form the required limiting behavior of the discount rates ν^n from Proposition 5, and thus determine the survival outcomes. Further, the link between the decentralized solution and the planner's problem conditions from Proposition 5 reveals that the survival conditions 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

Without loss of generality, it is sufficient to focus on the case $\theta^1 \searrow 0$. Using the construction from Duffie and Epstein (1992a), the stochastic discount factor process for agent n under the subjective probability measure Q^n is given by

$$S_t^n = \exp\left(-\int_0^t \nu^n\left(\theta_s^1\right) ds\right)\left(\frac{Y_t}{Y_0}\right)^{\gamma-1}\left(\frac{\zeta^n\left(\theta_t^1\right)}{\zeta^n\left(\theta_0^1\right)}\right)^{\rho-1}\left(\frac{\tilde{J}^n\left(\theta_t^1\right)}{\tilde{J}^n\left(\theta_0^1\right)}\right)^{1-\frac{\rho}{\gamma}}. \tag{20}$$

For the large agent 2, Lemma 15 in the Appendix shows that $\lim_{\theta^1\searrow 0}\zeta^2\left(\theta^1\right) = 1$ and $\lim_{\theta^1\searrow 0}\tilde{J}^2\left(\theta^1\right) = \tilde{V}^2$ from equation (8), implying that $\lim_{\theta^1\searrow 0}\nu^2\left(\theta^1\right) = \nu^2$ which is given in (9). It remains to be shown that the local drift and volatility of the last two terms decline to zero as $\theta^1 \searrow 0$, which implies that the infinitesimal risk-free rate and the local price of risk converge to their counterparts from a homogeneous economy populated only by agent 2. 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 7 *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 process $d\Xi_t/\Xi_t = \mu_\Xi\left(\theta_t^1\right)dt + \sigma_\Xi\left(\theta_t^1\right)dt$ converge to their homogeneous economy counterparts:*

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

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

$$\left[\left[\xi\left(\theta_t^1\right)\right]^{-1} + \mu_\Xi\left(\theta_t^1\right)\right]dt + \sigma_\Xi\left(\theta_t^1\right)dW_t, \tag{21}$$

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 $\xi\left(\theta^1\right)$ converges as $\theta^1 \searrow 0$, yet the price dynamics are such that $\mu_\Xi\left(\theta^1\right)$ and $\sigma_\Xi\left(\theta^1\right)$ do not converge to μ_y and σ_y, respectively. The fact that this does not happen is closely linked to the dynamics of $\log\theta^1$ which has bounded drift and volatility coefficients. This ensures that the local variation in $\xi(\theta^1)$ becomes irrelevant as $\log\theta^1 \searrow -\infty$.

The results in Proposition 7 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 8 *For every fixed maturity t, the prices of a zero-coupon bond and a claim to a payout from the aggregate endowment stream (a consumption strip) converge to their homogeneous economy counterparts as $\theta^1 \searrow 0$.*

The agent with negligible wealth therefore has no price impact not only on the two assets that dynamically complete the market but also on every finite-maturity bond and consumption strip. Recall that the results from Proposition 7 and Corollary 8 do not assume that the agent with negligible wealth vanishes in the long run. The reason is that even if the agent survives, the logarithmic growth rates of her wealth are always bounded in this economy. This implies that the distribution of her wealth at a given future date t can be driven arbitrarily close to zero by driving to zero her current wealth level. Of course, ultimately, the surviving agent will recover from an arbitrarily small wealth level, so that the convergence in Corollary 8 is not uniform in the maturity horizon $t \in (0, \infty)$.

3.3.2 Decision problem of an agent with negligible wealth

Proposition 7 establishes that the actual *general equilibrium* price dynamics in the proximity of the boundary are locally the same as those in an economy populated only by agent 2. To conclude the argument, we need to infer the wealth dynamics for agent 1 that has negligible wealth. The marginal utility under recursive preferences is forward-looking and depends on agent's continuation value (see the stochastic discount factor specification (20)). If agent 1 ultimately survives, then she will always have a nontrivial price impact in the future, even if her current Pareto share is negligible. The forward looking nature of the optimization problem then implies that she should take this price impact into account when making her current portfolio and consumption-saving decisions. However, the result below shows that this impact of future price dynamics becomes immaterial as $\theta^1 \searrow 0$.

Proposition 9 *The consumption-wealth ratio of agent 1 converges to*

$$\lim_{\theta^1 \searrow 0} \left[\xi^n\left(\theta^1\right)\right]^{-1} = [\xi(0)]^{-1} - \frac{\rho}{1-\rho}\left[\left(u^1 - u^2\right)\sigma_y + \frac{1}{2}\frac{\left(u^1 - u^2\right)^2}{1-\gamma}\right] \tag{22}$$

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

$$\lim_{\theta^1 \searrow 0} \pi^1\left(\theta^1\right) = 1 + \frac{u^1 - u^2}{(1-\gamma)\,\sigma_y}. \tag{23}$$

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

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

Proposition 9 derives the consumption-saving decision (22) and portfolio allocation decision (23) relative to the same decisions of the large agent 2. Recall that agent's 2 choices agree in the limit as $\theta^1 \searrow 0$ with aggregate ones — her consumption-wealth ratio is equal to the aggregate ratio $[\xi(0)]^{-1}$ and she holds the market portfolio, $\lim_{\theta^1 \searrow 0} \pi^2\left(\theta^1\right) = 1$. As the formulas indicate, when $u^1 = u^2$ the agents are identical and their decisions and wealth dynamics coincide.

The logic of the proof relies on showing that the current continuation value dynamics of the agent with negligible wealth is not influenced by the fact that she may become non-negligible in the future and have impact on aggregate price dynamics. Since the drift and volatility coefficients of the logarithm of the Pareto share process $\log \theta^1$ are bounded, then by pushing the current θ^1 arbitrarily close to zero ($\log \theta^1$ arbitrarily far toward $-\infty$), one can extend the time before the presence of the agent 1 becomes noticeable from aggregate perspective (measured, e.g., by sufficiently large deviations in prices or return distributions from their homogeneous economy counterparts) arbitrarily far into the future.

The portfolio and consumption-saving decision of agent 1 as $\theta^1 \searrow 0$ thus coincides with a 'partial equilibrium' solution where agent 1 behaves as if she lived forever as an infinitesimal agent in a homogeneous economy populated only by the large 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 negligible agent survives with probability one and has an impact on equilibrium prices in the long run, these effects do not influence *current* prices, returns, and wealth dynamics.

3.3.3 Limiting relative patience and relationship to wealth growth

Importantly, the limiting discount rate $\nu^1\left(\theta^1\right)$ in Proposition 5 can be inferred from the portfolio problem (43–44), which leads to the statement of the main result of this section.

Proposition 10 *The expressions for the limiting behavior of the relative patience in Proposition 5 are*

$$\lim_{\theta^1 \searrow 0} \nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right) = \frac{\rho-\gamma}{1-\rho}\left[\left(u^1-u^2\right)\sigma_y + \frac{1}{2}\frac{\left(u^1-u^2\right)^2}{1-\gamma}\right], \tag{24}$$

$$\lim_{\theta^1 \nearrow 1} \nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right) = \frac{\rho-\gamma}{1-\rho}\left[\left(u^1-u^2\right)\sigma_y - \frac{1}{2}\frac{\left(u^1-u^2\right)^2}{1-\gamma}\right]. \tag{25}$$

Section 4 discusses which regions of the parameter space satisfy the individual survival and extinction conditions from Proposition 5. It remains for me to verify that the assumption about the boundedness of wealth consumption ratios indeed holds.

Corollary 11 *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 high.

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 12 *The survival conditions in part a) of Proposition 5 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 5 therefore amounts to checking that the expected growth rate of the logarithm of wealth (the drift coefficient for $d\log \Xi_t^n$) is higher for the agent who is at the brink of extinction.

3.3.4 Decomposition of wealth growth rates

The previous discussion identified the portfolio allocation and the consumption-saving decisions as the two mechanisms underlying wealth accumulation and long-run survival. It is therefore informative to decompose the growth rate of individual wealth from Corollary 12 into the contribution of the logarithmic return on the agent's portfolio, net of the rate of consumption represented by the consumption-wealth ratio,

$$d\log \Xi_t^n = d\log R_t^n - \left(\xi_t^n\right)^{-1} dt.$$

Denoting the drift coefficient on $d\log\Xi_t^n$ as $\tilde{\mu}_{\Xi^n}(\theta^1) = \mu_{\Xi^n}(\theta^1) - \frac{1}{2}\left[\sigma_{\Xi^n}(\theta^1)\right]^2$ and the drift coefficient on $d\log R_t^n$ (the expected logarithmic return) as $\tilde{\mu}_{R^n}(\theta^1)$, we can collect the results above to establish the following decomposition.

Proposition 13 *As $\theta^1 \searrow 0$, the difference in the logarithmic wealth growth rates between the agent with negligible wealth and the large agent can be written as*

$$\lim_{\theta^1\searrow 0}\left[\tilde{\mu}_{\Xi^1}\left(\theta^1\right) - \tilde{\mu}_{\Xi^2}\left(\theta^1\right)\right] = \lim_{\theta^1\searrow 0}\left[\tilde{\mu}_{R^1}\left(\theta^1\right) - \tilde{\mu}_{R^2}\left(\theta^1\right)\right] - \lim_{\theta^1\searrow 0}\left[\left(\xi^1\left(\theta^1\right)\right)^{-1} - \left(\xi^2\left(\theta^1\right)\right)^{-1}\right]$$

where the difference in the expected logarithmic portfolio returns is

$$\lim_{\theta^1\searrow 0}\left[\tilde{\mu}_{R^1}\left(\theta^1\right) - \tilde{\mu}_{R^2}\left(\theta^1\right)\right] = \underbrace{\frac{u^1 - u^2}{(1-\gamma)\,\sigma_y}}_{\substack{\text{difference in}\\ \text{portfolios}}}\underbrace{\left[(1-\gamma)\,\sigma_y^2 - u^2\sigma_y\right]}_{\text{risk premium}} - \underbrace{\frac{u^1 - u^2}{1-\gamma}\left(\sigma_y + \frac{1}{2}\frac{u^1 - u^2}{1-\gamma}\right)}_{\text{lognormal correction}}$$

and the difference in consumption rates is given by

$$\lim_{\theta^1\searrow 0}\left[\left(\xi^1\left(\theta^1\right)\right)^{-1} - \left(\xi^2\left(\theta^1\right)\right)^{-1}\right] = -\frac{1}{2}\frac{\rho}{1-\rho}\underbrace{\left[2\left(u^1 - u^2\right)\sigma_y + \frac{\left(u^1 - u^2\right)^2}{1-\gamma}\right]}_{\text{difference in subjective expected returns}}.$$

The difference in the expected logarithmic portfolio returns at the boundary only depends on the risk relative risk aversion $1-\gamma$, not on the parameter ρ that determines the IES. This difference consists of the familiar terms — the risk premium on the claim on aggregate consumption times the difference in the portfolio shares invested in the risky asset (recall that agent 2 holds the market portfolio while agent's 1 portfolio allocation is given by (23)). The risk premium consists of the standard rational expectations premium $(1-\gamma)\,\sigma_y^2$ and a 'mispricing' effect $-u^2\sigma_y$ (if the large agent 2 is optimistic, she overprices the risky asset which generates a lower expected return). However, it is the expected *logarithmic* return that drives survival, and thus a lognormal correction is necessary. This lognormal correction is the dominant force for survival when risk aversion declines to zero ($\gamma \nearrow 1$).

The difference in consumption rates also has two components. The term in brackets represents the difference between the expected portfolio return of agent 1 as perceived by agent 1, and the portfolio return of agent 2 as perceived by agent 2, informally

$$\frac{1}{dt}\left(E_t^{Q^1}\left[dR_t^1\right] - E_t^{Q^2}\left[dR_t^2\right]\right).$$

It is the subjective expected returns (computed under Q^n, not P) that enter the formula because the consumption-saving decision of the agent depends on the expected portfolio return as perceived by herself. When IES $= 1$ ($\rho = 0$), the consumption-wealth ratios of the two agents are identical and equal to β as in the case of myopic logarithmic utility, and the consumption-saving decision plays no role in the survival outcomes. When preferences are elastic (IES > 1, i.e., $1 > \rho > 0$), the saving rate is an increasing function of the subjective expected portfolio return and the difference in consumption rates is therefore negatively related to the difference in subjective expected returns — vis-à-vis a high expected return, the agent with elastic preferences decides to postpone consumption and tilt the consumption profile toward the future. This helps the agent with the higher expected subjective return outsave her extinction.

3.3.5 Dependence of survival results on individual parameters

It is worth noting that the survival results do not depend on the time-preference parameter β and the growth rate of the economy μ_y. Both these parameters influence individual preferences for saving versus consumption. Since the impact of these parameters on both agents' decisions is identical, they do not affect the difference in the rates of wealth accumulation. This would no longer be true if, for instance, the agents differed in the IES parameter.

Finally, using the results from Proposition 10 in the inequality conditions from Proposition 5 reveals that the survival regions depend on the ratios of parameters u^1/σ_y and u^2/σ_y, and not on the three parameters independently. This is an important insight that shows that what matters for survival in this economy is the importance of aggregate fundamental risk embedded in σ_y relative to the willingness of the agents to generate additional volatility in their individual consumption processes through betting, reflected in the magnitude of the belief distortions u^n. For instance, if agent 1 has correct beliefs, $u^1 = 0$, then the long-run survival outcome will be the same if we fix the belief distortion u^2 and make the aggregate endowment deterministic, $\sigma_y \searrow 0$, or if we fix σ_y and make the agent's beliefs infinitely incorrect. I will revisit this aspect of the survival results in the next section.

4 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 generated by the pair of inequalities in Proposition 5 do occur generically, 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 5 hold with equalities, but the survival characteristics inside the individual regions are well-defined.

4.1 Asymptotic results

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

Corollary 14 *Holding other parameters fixed, for any given pair of beliefs u^n, $n \in \{1,2\}$ and $\sigma_y > 0$, the survival restrictions imply the following asymptotic results.*

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

(b) *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.*

(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 provide the intuition underlying result **(a)**, consider the limiting case when agents are risk neutral ($\gamma = 1$). Then the felicity function $F(C,\nu)$ in (5) is linear in C, and the planner will choose a corner solution for the allocation of consumption in the next instant. There will be a cutoff in the distribution of the next-instant shock dW_t above which the planner allocates all consumption to the relatively more optimistic agent, and vice versa for states below the cutoff. Further, when the planner chooses zero consumption for the agent in the next instant, he will allocate zero consumption to that agent also at all subsequent dates and states. This result may seem puzzling but it is closely related to the exact role of the IES parameter, which captures the elasticity of substitution between current consumption and the expected risk-adjusted continuation value. When IES is finite ($\rho < 1$), then the only way how to optimally provide zero consumption in the next instant is to also provide zero continuation value in the same state, which also implies zero consumption at all subsequent dates and states (up to a set of paths of measure zero).[12] When risk aversion is strictly positive but sufficiently small, the same force will not operate instantaneously but over time in the long-horizon limit. The planner will optimally let the continuation values drift apart as he tolerates a wide distribution of future continuation values.

[12] A very similar mechanism underlies the results in Backus, Routledge, and Zin (2008). The discrete-time specification from Epstein and Zin (1989) adopted in their paper makes clear the role of IES parameter as the elasticity between current consumption and the expected risk-adjusted continuation value next period.

The interpretation of this result from the perspective of decentralized decision-making of the two agents has been already discussed in the introduction. Agents with a low risk aversion are willing to make large bets on the states next period. Once a sequence of unsuccessful bets reduces the wealth of one agent substantially, she is the only one who can continue to make large bets, as the agent with the dominant amount of wealth is disciplined by market clearing (prices have to adjust so as to make the large agent consume the aggregate endowment, and hold her portfolio entirely invested in the risky asset without any leveraging). Further betting of the small agent continues to be detrimental for her survival chances — despite the fact that she may earn a high expected level return on her portfolio, the expected logarithmic return is low due to the variance correction that comes from the high dispersion of future wealth that the agent tolerates. This leads to her extinction along a set of paths that has a strictly positive measure. Since events when either of the agents becomes sufficiently small recur with probability one, ultimately one of the agents becomes extinct with probability one, and each of the agents faces a strictly positive probability of extinction.

In the other extreme, when agents become very risk averse ($\gamma \searrow -\infty$, result **(b)**), the felicity function is highly curved and the provision of consumption in the low realizations of the shock becomes very costly for the planner. Since the relatively more pessimistic agent overweighs the probability of the low realizations, the planner will provide consumption to this agent in these low states but he will compensate the high cost of this consumption by providing a lower average promised utility across all states. In the decentralized equilibrium, this is manifested by the relatively more optimistic agent earning a strictly higher expected return on her portfolio that is overweighed in the high-return risky asset, despite the fact that betting motives that lead agents with inaccurate beliefs to allocate wealth incorrectly disappear as risk aversion increases.

Under separable preferences, the relatively more pessimistic agent would still survive (and dominate) in this situation if her beliefs are more accurate because an increase in risk aversion implies a decrease in IES. This agent would correctly understand that she pays for her insurance of the low states with a lower expected return, and the lower IES would also motivate her to decrease current consumption more. Under the recursive preference structure, the IES is fixed as risk aversion increases, and this saving channel will not outweigh the advantage of the relatively more optimistic agent.

Result **(c)** highlights the role of the consumption-saving mechanism. With a high IES, agents are willing to substantially decrease their consumption rate vis-à-vis an increase in the subjective expected return on their portfolio (see also the expression for the difference in consumption rates in Proposition 13 which scales the difference in subjective expected returns by the term $\rho / (1-\rho)$). Whenever an agent becomes small, she can choose a portfolio with a high subjective expected return while the market clearing mechanism forces her large counterpart to choose a portfolio that is close to the market portfolio. The high IES then gives a survival advantage to the small agent because it induces her to increase her saving rate in response to the high subjective expected return.

At the same time, risk aversion cannot be too high for this mechanism to be sufficiently strong. A high risk aversion discourages betting, and the incentives of the small agent to choose a sufficiently 'leveraged' portfolio with a high subjective expected return diminish.

Result **(d)** is a direct counterpart to **(c)**. When preferences of the agents become inelastic ($\rho \searrow -\infty$), formulas in Proposition 10 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. This result again shows the critical role of the consumption-saving decision and the endogenous equilibrium price dynamics in generating equilibria in which both agents survive in the long run. Partial equilibrium models with exogenous price dynamics that do not depend on wealth shares of individual agents cannot replicate this survival mechanism.

4.2 Comparative statics

The panels in Figure 1 depict the shape of the survival regions for different choices of the belief distortions. Agent 2 is assumed to have correct beliefs (except the last panel), while I vary the belief distortion of agent 1. The volatility of aggregate endowment is set to a plausible value of $\sigma_y = 0.02$. The magnitude of the belief distortion $u^1 = 0.1$ implies that the agent overestimates the annual growth rate of aggregate endowment by $u^1\sigma_y = 0.2\%$ (and, correspondingly, by 0.5% for $u^1 = 0.25$). The existing literature established that along the dotted diagonal, which represents the parameter combinations for separable CRRA preferences, the agent with more accurate beliefs (i.e., with a smaller $|u^n|$) dominates the economy in the long run.

4.2.1 Optimistic belief distortion

The first panel starts with a moderately optimistic agent 1. Consistent with the previous results, the correct agent 2 dominates in the long run in the neighborhood of the diagonal. Also, for very high levels of risk aversion, the relatively more optimistic agent 1 dominates the economy. However, there is a nontrivial intermediate region where both agents coexist in the long run. In this whole region, risk aversion is larger than the inverse of IES, which is a standard parametric choice in the asset pricing literature. The two boundaries in the top left panel which delimit this region are asymptotically parallel as $\gamma \searrow -\infty$ with slope $2\sigma_y / \left(u^1 + u^2 + 2\sigma_y\right)$. Finally, for very low levels of risk aversion, either agent can dominate

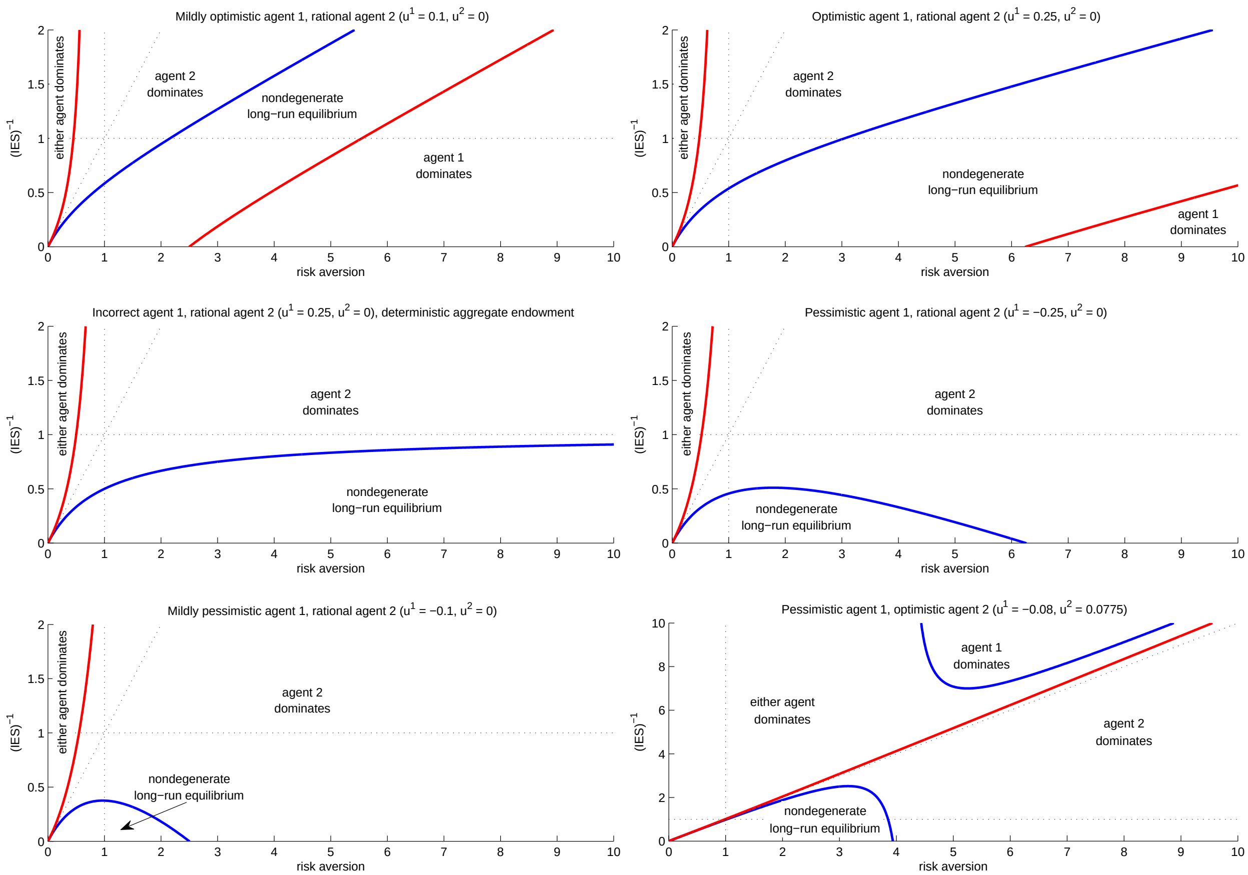

Figure 1: Survival regions for different parameterizations. All panels except the third panel assume $\sigma_y = 0.02$. In the third panel, aggregate endowment is deterministic, $\sigma_y = 0$, but the survival regions take the same shape for any parameterization with $u^2 = 0$ and $|u^1/\sigma_y| \to \infty$. Belief distortion parameters u^n are shown in the titles of individual panels.

in the long run with a strictly positive probability (result **(a)** from Corollary 14).

The second (top right) panel shows a comparison of the results when the optimism of the optimistic agent 1 increases to $u^1 = 0.25$. Due to the increased inaccuracy of beliefs of agent 1, the region where the correct agent 2 dominates expands. At the same time, the region in which both agents coexist in the long run expands as well.

The natural question is to ask whether the region in which both agents coexist vanishes as the inaccuracy of beliefs of agent 1 increases further. The answer is, maybe somewhat surprisingly, no. The third panel depicts the survival outcomes for the case $|u^1/\sigma_y| \to \infty$, which includes several rather different parameterizations. The first case is the case of extreme optimism, $u^1 \to \infty$. The region of the parameter space from the first two panels in which both agents survive shifts to the right and down, but converges to a nontrivial region that occupies most of the half-plane for IES > 1. Regardless how incorrectly optimistic agent 1 becomes, she still survives in a large set of plausible preference parameterizations.

4.2.2 Deterministic economy

However, the same region of coexistence of the two agents is obtained for the case of extreme pessimism ($u^1 \to -\infty$), or the case of an economy with a deterministic aggregate endowment ($\sigma_y = 0$). All these economies share the common feature that the magnitude of the risk premium associated with the claim on aggregate endowment, which is proportional to σ_y^2, becomes trivial compared to the amount of perceived mispricing that depends on the magnitude of the difference in belief distortions. Since the incorrect belief is manifested as a perceived drift term in the increment of the Brownian motion and this increment has a symmetric distribution, then in an economy where aggregate risk has a negligible pricing implication a positive and a negative drift distortion have the same (merely mirrored) distributional implications.

This is apparent in the case of a deterministic aggregate endowment. To interpret this economy, consider without loss of generality the case of $Y_t \equiv 1$. The agents observe the Brownian motion W and can contract upon its realizations. A possible decentralization of this economy involves an asset in unit supply that pays the safe aggregate endowment, and a risky asset in zero net supply that pays out according to the realizations of W. While the aggregate endowment is safe, W provides a lottery to the agents with a probability distribution about which they disagree, and they use their claims on the aggregate endowment to bet using this lottery. Because W is symmetric, it is irrelevant whether the agent's beliefs are distorted by u^n or $-u^n$.

As the wealth share of an agent converges to one, her consumption stream becomes deterministic which implies that risk is not priced in equilibrium. The zero net supply risky asset can be mispriced in equilibrium but the large agent holds only a negligible share of wealth invested in this asset due to market clearing, and thus her portfolio earns the risk-free return which is correctly perceived by the agent. The negligible agent can choose a risk-free portfolio as well but since she has different beliefs than the large agent, she perceives the risky asset as mispriced and chooses either a short or long position in the risky asset, depending on the direction of the mispricing. Such a portfolio generates a higher subjective expected return to the negligible agent than the risk-free return earned by the large agent. With IES > 1, this translates into a higher saving rate of the negligible agent, and if IES is sufficiently high, the high saving motive will always be strong enough to let the negligible agent outsave her extinction and survive in the long run.

4.2.3 Pessimistic belief distortion

The fourth and fifth graphs represent the survival results for a pessimistic agent 1. The sequence of the graphs can be interpreted as shrinking the magnitude of the agent's pessimism from an infinite belief distortion ($u^1 \to -\infty$) in the third graph toward smaller distortions in graphs four and five. Interestingly, the region in which the pessimistic agent survives shrinks as the belief distortion of the agent is reduced. This is in stark contrast to the conventional wisdom obtained from the study of survival under separable preferences. Not only can the incorrect agent survive in the long run, her survival chances are actually better with larger distortions.

Notice that the pessimistic agent invests a smaller share of her wealth into the risky asset, so she cannot benefit from the risky asset's higher expected return. At the same time, the pessimism seems to imply that her *subjective* expected return is even lower, so that she will not improve her survival chances by choosing a higher saving rate under IES > 1. But since the agent is pessimistic about the return on the asset, she actually is optimistic about the return on a short position in that asset. Observe that the last term in brackets in the consumption-wealth ratio (22), which dominates the saving decision of agent 1 when $\rho \nearrow 1$, is equal to

$$\frac{1}{2}\left(u^1 - u^2\right)\sigma_y\left(1 + \pi^1(0)\right). \tag{26}$$

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 dominate that of the large agent 2 as $\rho \nearrow 1$. While the short position in the risky asset earns a low objective expected return, a high IES will then generate a sufficiently strong offsetting saving motive that will allow the pessimistic agent to outsave her extinction.

The region in the fourth and fifth graph in which the two agents coexist does not include high levels of risk aversion and shrinks for smaller belief distortions. A high level risk aversion or a lower incentive to bet caused by a smaller belief distortion will prevent the small agent from choosing a sufficiently large short position in the risky asset which is, as shown in formula (26), necessary to generate the high subjective expected return needed for the saving mechanism to operate in favor of the pessimistic agent 1.

The above discussion also explains why the described mechanism cannot lead to the long-run dominance of the pessimistic agent. As the wealth share of the pessimistic agent approaches one, she can no longer hold a short position in the risky asset, and the effect of the saving mechanism generated through the high subjective expected return disappears.

Finally, there are also cases when the pessimistic agent dominates the economy even if her beliefs are more inaccurate than those of her counterpart, although the difference in belief accuracies of the two agents has to be only very small (the exact condition is $u^1 + u^2 + 2\sigma_y > 0$). The last graph of Figure 1 depicts such a case. Contrary to the previous results, the pessimistic agent can only survive in the empirically less plausible region when risk aversion is lower than the inverse of IES ($\gamma > \rho$). The online appendix discusses this case in more detail.

5 Dynamics of long-run equilibria

Propositions 5 and 10 in Section 3 derive parametric restrictions on the survival regions. However, even if a nondegenerate long-run equilibrium exists, the question remains whether this equilibrium delivers quantitatively interesting endogenous dynamics under which each of the agents can gain a significant wealth share. The derived analytical boundary conditions cannot answer this questions, as a full solution of the model is necessary to investigate the interior of the $(0,1)$ interval for the Pareto share. This section investigates numerically the equilibrium allocations and prices and their dynamics by solving the ODE (15) and the associated decentralization.

5.1 Stationary distributions and evolution over time

The top left graph of Figure 2 plots the densities $q\left(\zeta^1\right)$ for the stationary distribution of the consumption share ζ^1 of the optimistic agent 1 in economies with an optimistic and a correct agent. The parameterizations[13] are chosen along a horizontal line in the top right panel of Figure 1, corresponding to different levels of risk aversion. As risk aversion increases, the distribution of consumption shifts toward the optimistic agent (the results from the previous section show that for a sufficiently high risk aversion, the optimistic agent will ultimately dominate), but in all the parameterizations, both agents have quantitatively nontrivial consumption shares and the shape of the densities indicates that the equilibria permit substantial variation over time in these consumption shares.

In empirical applications, it is advantageous when the time-series of observable variables converge sufficiently quickly to their long-run distributions from any initial condition, so that data observed over finite horizons are a representative sample of the stationary distribution. For instance, 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. Proposition 5 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.

I show in the online appendix that under the conditions from Proposition 5, convergence for the state variable θ^1 occurs at an exponential rate, so that the process θ^1 does not exhibit strong dependence properties. At the same time, the exponent in the rate calculation can still be small, and I therefore conduct a numerical simulation. The remaining three graphs in Figure 2 plot the conditional distribution of the consumption share $\zeta^1\left(\theta^1_t\right)$ of the optimistic agent 1 conditional on $\zeta^1\left(\theta^1_0\right) = 0.5$ for different time horizons t. These are computed from

[13] 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.

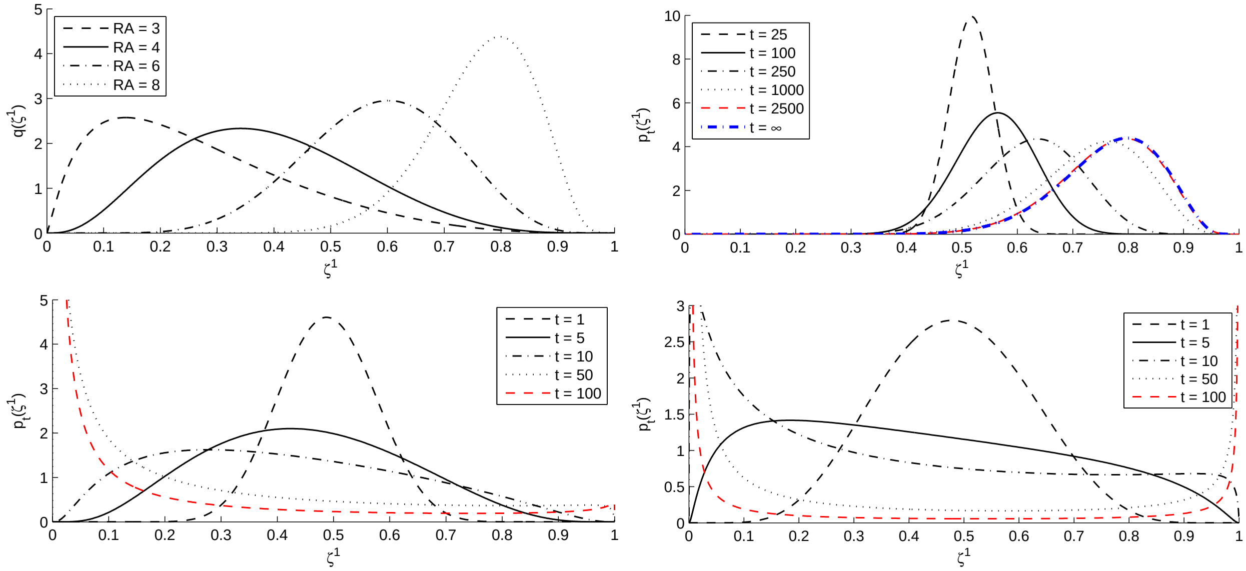

Figure 2: The top left panel depicts the stationary distributions for the consumption share $\zeta^1\left(\theta^1\right)$ of the agent with optimistically distorted beliefs. All models are parameterized by $u^1 = 0.25$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.02$, and differ in levels of risk aversion, shown in the legend. The remaining three panels plot the distributions of $\zeta^1\left(\theta_t^1\right)$ conditional on $\zeta^1\left(\theta_0^1\right) = 0.5$ for different time horizons t. In the top right panel (risk aversion $= 8$), the economy has a nondegenerate long-run distribution. In the bottom left panel (risk aversion $= 0.75$), the correct agent 2 dominates, and in the bottom right panel (risk aversion $= 0.25$), each agent dominates with a strictly positive probability.

the dynamics of the state variable θ^1 in equation (18) by solving the associated Kolmogorov forward equation

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

for the conditional density $p_t^\theta\left(\theta^1\right)$ of θ_t^1 with the initial condition $p_0^\theta\left(\theta^1\right) = \delta_{\theta_0^1}\left(\theta^1\right)$, where δ is the Dirac delta function, and then transforming to obtain the conditional density for ζ^1

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

The graphs show the evolution of the conditional distribution for three cases. In the top right graph, the conditional distribution converges to a nondegenerate stationary distribution and both agents survive. In the bottom left graph, the mass of the conditional distribution shifts to the left and agent 2 dominates. Finally, in the bottom right graph, the mass of the conditional distribution shifts out toward both boundaries, and either agent dominates with a strictly positive probability.

The speed of the evolution of the conditional distribution depends on the magnitude of the belief distortions and the level of risk aversion in the economy. When risk aversion is high, 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 2, 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.

While the evolution of the conditional density in Figure 2 may appear rather slow, the process can be accelerated substantially. One possible way is to increase the magnitude of the belief distortions but very large belief distortions may be rejected as empirically implausible.

A more fundamental argument relies on the appropriate interpretation of the modeled risk in this economy. In the model, the nature of risk is extremely simplistic, the agents can disagree only about the distribution of the aggregate shock. In reality, there are many other sources of aggregate or idiosyncratic risk about which the agents can disagree and write contracts on, and agents with heterogeneous beliefs would also find it optimal to introduce additional such betting devices, even if these are otherwise economically irrelevant. Fedyk, Heyerdahl-Larsen, and Walden (2013) show in an economy with CRRA preferences that if agents disagree about multiple sources of risk, the speed of extinction of the relatively more incorrect agent can be accelerated substantially. The same mechanism operates under recursive utility, increasing the magnitude of wealth fluctuations and the rate of convergence of $p_t(\zeta^1)$ to the stationary density $q(\zeta^1)$ when both agents survive in the long run.[14] The main message arising from these considerations is that the speed of extinction or rate of convergence to the stationary distribution in stylized models with very few sources of risk should not be viewed as a strong quantitative test of the model.

5.2 Survival forces

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 different cases. The dashed line represents the low risk aversion case in which both attracting conditions from Proposition 5 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

[14]The online appendix provides an example with two imperfectly correlated Brownian motions. One concern from the perspective of the survival results may be that belief distortions about multiple sources of risk can be reinterpreted as one large belief distortion. This view is, with some qualifications, correct but the survival results show that agents can coexist in the long run even under very large belief distortions.

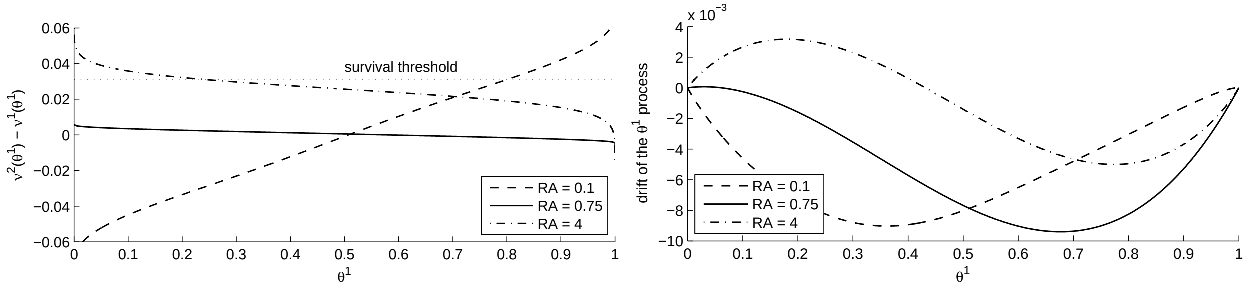

Figure 3: Relative patience $\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)$ (left panel) and the drift component of the Pareto share evolution $E\left[d\theta_t^1 \mid \mathcal{F}_t\right]/dt$ (right panel) as functions of the Pareto share θ^1. All models are parameterized by $u^1 = 0.25$, $u^2 = 0$, 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 5.

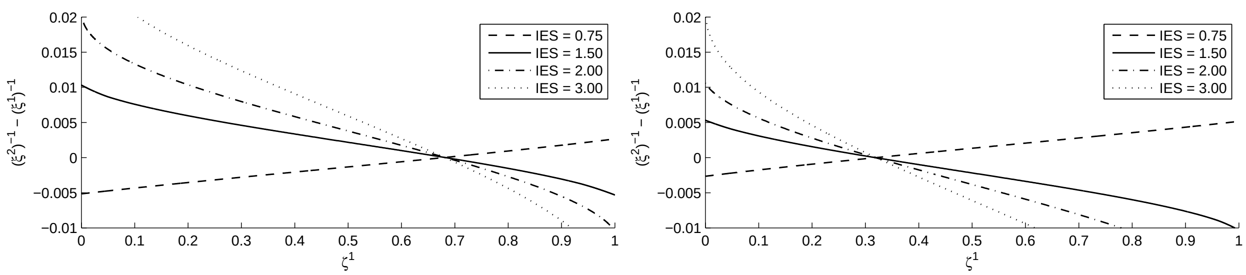

Figure 4: Difference in consumption-wealth ratios $\left(\xi^2\right)^{-1} - \left(\xi^1\right)^{-1}$ as a function of the consumption share ζ^1 of agent 1. The left panel considers an optimistic agent 1 ($u^1 = 0.25$) while the right panel a pessimistic agent 1 ($u^1 = -0.25$). The remaining parameters 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.

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 two 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-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 her wealth per unit of time. A higher IES improves the survival chances of the agent who is relatively more optimistic about the return on her own wealth, as she is willing to tilt her consumption profile more toward the future.

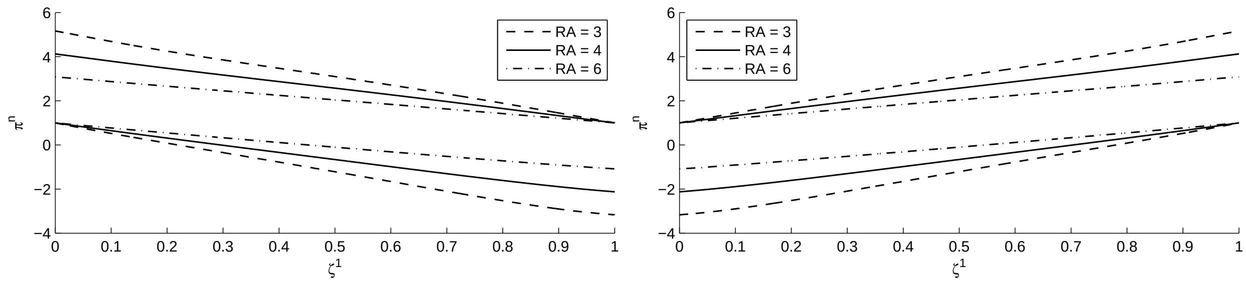

Figure 5: Wealth shares π^n of the two agents invested in the claim to aggregate endowment as functions of the consumption share ζ^1 of agent 1. The left panel considers an optimistic agent 1 ($u^1 = 0.25$) while the right panel a pessimistic agent 1 ($u^1 = -0.25$). The remaining parameters are $u^2 = 0$, 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 $\zeta^1 = 1$ ($\zeta^1 = 0$) belong to agent 1 (agent 2).

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. 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, she does not form a large enough short stock position that would make him sufficiently optimistic about the return on her own wealth and outsave the rational agent when IES > 1.

6 Extensions and concluding remarks

Before concluding, I briefly discuss 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.

6.1 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. This can be interpreted as the strongest form of incorrect beliefs, and a bias against survival of agents with whose beliefs are initially incorrect. A natural question is to ask what happens when agents are allowed to learn. Learning can be incorporated into the current framework by a introducing a law of motion that represents the Bayesian updating of the belief distortions u^n. These belief distortions become new state variables.

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, Figure 1 shows that the survival region of a pessimistic agent can shrink if her 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 her 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 agents learn and 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 and resolve the survival problem remains an open question.

6.2 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.

6.3 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 (2011) show elegantly that this result holds, irrespective of the specification of the aggregate endowment process,[15] 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 in which both agents coexist arise 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. These 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 similar to those in Lucas and Stokey (1984). An agent survives in the long run if her relative patience becomes sufficiently large as her wealth share vanishes. However, in this framework, the dynamics of relative patience arises endogenously as an equilibrium outcome, and is not a direct

[15] The survival results under separable utility thus also hold for 'exotic' endowment processes like the rare disaster framework in Chen, Joslin, and Tran (2012).

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 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.

Appendix

Before proving Proposition 3, I first discuss the behavior of the objective function at the boundaries. The planner's objective function is, for a given Y, 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 (12) thus exists. Since the continuation values are concave, first-order conditions are sufficient for the supremum problem. We have the following Lemma.

Lemma 15 *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 $\alpha^2 > 0$*

$$J_0\left(0, \alpha^2\right) \doteq \lim_{\alpha^1 \searrow 0} J_0\left(\alpha^1, \alpha^2\right) = \alpha^2 V_0^2(Y) \tag{27}$$

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

Proof of Lemma 15. 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

$$J_0(\alpha) = \alpha^1 V_0^1\left(C^1(\alpha)\right) + \alpha^2 V_0^2\left(C^2(\alpha)\right) \tag{28}$$

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 (27) holds, and the convergence of $C^2\left(\alpha^1, \alpha^2\right)$ is a direct consequence. ■

Proof of Proposition 3. 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 (12-13) leads to the Hamilton-Jacobi-Bellman equation for $J(Z)$,

$$0 \equiv \sup_{\left(C^1, C^2, \nu^1, \nu^2\right)} \sum_{n=1}^{2} \bar{\lambda}^n \left[F\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n} \nu^n\right] + J_y \mu_y Y + \frac{1}{2} \text{tr}\left(J_{zz} \Sigma\right), \tag{29}$$

where

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

and $\operatorname{diag}\left(\bar{\lambda}\right)$ 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 (29) can be solved separately. Under the optimal discount rate process ν^n for agent n,

$$f\left(C^n, J_{\bar{\lambda}^n}\right) \doteq \sup_{\nu^n} F\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n}\nu^n = \frac{\beta}{\rho}\left[\left(C^n\right)^\rho \left(\gamma J_{\bar{\lambda}^n}\right)^{1-\frac{\rho}{\gamma}} - \gamma J_{\bar{\lambda}^n}\right]. \tag{30}$$

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 (29) 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 15.

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 5. 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

$$\lim_{\theta_l\searrow 0}S[\theta_l,\theta_h]=\infty\quad\text{and}\quad\lim_{\theta_h\nearrow 1}S[\theta_l,\theta_h]=\infty\tag{31}$$

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

$$\lim_{\theta_l\searrow 0}M[\theta_l,\theta_h]<\infty\quad\text{and}\quad\lim_{\theta_h\nearrow 1}M[\theta_l,\theta_h]<\infty,\tag{32}$$

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

$$s_{sep}(\theta)=\left(\frac{1-\theta}{1-\theta_0}\right)^{-\frac{2u^1}{u^1-u^2}}\left(\frac{\theta}{\theta_0}\right)^{\frac{2u^2}{u^1-u^2}}\tag{33}$$

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 (31) 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 $\theta \in \left(\bar{\theta}, 1\right)$, 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 (32) are marginally tighter. Following the same bounding argument as above, sufficient conditions for (32) to hold are

$$\underline{\nu} > \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right] \quad \text{and } \bar{\nu} < \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]. \tag{34}$$

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 5 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 6. Assume without loss of generality that $\left|u^2\right| \leq \left|u^1\right|$. The sufficient part is an immediate consequence of Proposition 5. Under separable preferences, $\nu^2 - \nu^1 \equiv 0$, and thus if $\left|u^2\right| < \left|u^1\right|$ 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 (33) in the proof of Proposition 5 that conditions (31) 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 Proposition 7. The proof of the proposition relies on showing that they dynamics of the continuation values of the two agents in the proximity of the boundaries becomes degenerate in a specific sense. From this fact, I can infer the dynamics of the stochastic discount factor implied by the consumption process of the large agent and, consequently, the equilibrium price dynamics. I state the limiting properties of the continuation values separately in Lemmas 16 and 17.

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

$$\xi^n\left(\theta^1\right) = \frac{1}{\beta}\left(\frac{\tilde{J}^n\left(\theta^1\right)^{1/\gamma}}{\zeta^n\left(\theta^1\right)}\right)^\rho. \tag{35}$$

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 16 *The solution of the planner's problem implies that*

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

Proof. Lemma 15 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 (28) scaled by $\left(\alpha^1 + \alpha^2\right)\gamma^{-1}Y^\gamma$ leads to an equation in scaled continuation values

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

and the proof of Lemma 15 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 (8). Since $\tilde{J}^2\left(\theta^1\right) = \tilde{J}\left(\theta^1\right) - \theta^1 \tilde{J}_{\theta^1}\left(\theta^1\right)$, then

$$\lim_{\theta^1 \searrow 0} \theta^1 \tilde{J}_{\theta^1}\left(\theta^1\right) = 0. \tag{36}$$

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 (8) implies that this expression is zero. Since the fourth term in (15) also converges to zero due to result (36), the last term in (15) must also converge to zero, or

$$\lim_{\theta^1 \searrow 0} \left(\theta^1\right)^2 \tilde{J}_{\theta^1\theta^1}\left(\theta^1\right) = 0. \tag{37}$$

Finally, differentiate the PDE (15) by θ^1 and multiply the equation by θ^1. Using comparisons with results (36–37), 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

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

■

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

$$\begin{aligned} \frac{d\tilde{J}^n\left(\theta_t^1\right)}{\tilde{J}^n\left(\theta_t^1\right)} &= \mu_{\tilde{J}^n}\left(\theta_t^1\right) dt + \sigma_{\tilde{J}^n}\left(\theta_t^1\right) dW_t && (38) \\ \frac{d\zeta^n\left(\theta_t^1\right)}{\zeta^n\left(\theta_t^1\right)} &= \mu_{\zeta^n}\left(\theta_t^1\right) dt + \sigma_{\zeta^n}\left(\theta_t^1\right) dW_t, && (39) \end{aligned}$$

where the drift and volatility coefficients are functions of θ^1, and the results from Lemma 16 allow the characterization of their limiting behavior.

Lemma 17 *The coefficients in equations (38–39) for agent 2 satisfy*

$$\lim_{\theta^1 \searrow 0} \mu_{\tilde{J}^2}\left(\theta^1\right) = \lim_{\theta^1 \searrow 0} \sigma_{\tilde{J}^2}\left(\theta^1\right) = \lim_{\theta^1 \searrow 0} \mu_{\zeta^2}\left(\theta^1\right) = \lim_{\theta^1 \searrow 0} \sigma_{\zeta^2}\left(\theta^1\right) = 0. \tag{40}$$

Proof. 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 16.

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 16 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

$$d\tilde{J}^1\left(\theta_t^1\right) = d\left[\tilde{J}\left(\theta_t^1\right)+\left(1-\theta_t^1\right)\tilde{J}_{\theta^1}\left(\theta_t^1\right)\right] = -\left(\theta_t^1\right)^2\tilde{J}_{\theta^1\theta^1}\left(\theta_t^1\right)\frac{d\theta_t^1}{\theta_t^1}+ \\ +\frac{1}{2}\left[\left(\theta_t^1\right)^2\tilde{J}_{\theta^1\theta^1}\left(\theta_t^1\right)+\left(1-\theta_t^1\right)\left(\theta_t^1\right)^2\tilde{J}_{\theta^1\theta^1\theta^1}\left(\theta_t^1\right)\right]\left(\frac{d\theta_t^1}{\theta_t^1}\right)^2 \tag{41}$$

and that

$$\frac{\zeta^1\left(\theta^1\right)}{\tilde{J}^1\left(\theta^1\right)^{\frac{1}{\gamma}}} = \left(\theta^1\right)^{\frac{1}{1-\rho}}\left(\tilde{J}^1\right)^{\frac{1-1/\gamma}{1-\rho}}K\left(\theta^1\right)^{-1} \tag{42}$$

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 (42). Using the first θ^1 to multiply the coefficients in $d\tilde{J}^1$ in formula (41), 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

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

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 $d\zeta^2$. ■

We can now finally proceed with the proof of Proposition 7. Convergence of the risk-free interest rate follows from the direct calculation of

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

where $S^2_t(0)$ is the limiting stochastic discount factor corresponding to the one prevailing in a homogeneous economy populated only by agent 2. Lemma 17 shows that the local behavior of S^2_t converges to $S^2_t(0)$ as $\theta^1_0 \searrow 0$. 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 (8).

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 16, and the coefficients of the other two terms vanish by combining the results in that Lemma with equation (41). Further,

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

and since dZ^1 and dZ^2 have vanishing coefficients by the proof of Lemma 17 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

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

have coefficients that decline to zero as $\theta_t^1 \searrow 0$, which proves the result. ■

Proof of Proposition 8. The evolution of θ^1 given by equation (18) implies that for every fixed $t \geq 0$

$$\theta_0^1 \searrow 0 \quad \Longrightarrow \quad \theta_t^1 \to 0,\ P\text{-a.s.}$$

and thus also $\zeta^2\left(\theta_t^1\right)\to 1$ and $\tilde{J}^2\left(\theta_t^1\right)\to\tilde{V}^2$, P-a.s.[16] The last two terms in the expression for the stochastic discount factor, S_t^2, equation (20), converge to one, P-a.s., and since $\nu^2\left(\theta_s^1\right),\ 0\leq s\leq t$ also converges to $\nu^2(0)$ and is bounded, we have $S_t^2 \xrightarrow{P} S_t^2(0)$. Consider a family of random variables $M_t^2 S_t^2\left(\theta_0^1\right)$ indexed by the initial Pareto share θ_0^1. Since this family is uniformly integrable, then convergence in probability implies convergence in mean, and we obtain the convergence result for bond prices

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

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. ■

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

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

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

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

by setting $\log\theta_0^1$ sufficiently low.

Proof of Proposition 9. Agent 1, whose wealth Ξ^1 is close to zero, solves

$$\bar{\lambda}_t^1 V_t^1 = \max_{C^1,\pi^1,\nu^1} E_t\left[\int_t^\infty \bar{\lambda}_s^1 F\left(C_s^1,\nu_s^1\right) ds\right] \tag{43}$$

subject to (7) and the budget constraint,

$$\begin{aligned}\frac{d\Xi_t^1}{\Xi_t^1} &= \left[r\left(\theta_t^1\right) + \pi_t^1\left(\left[\xi\left(\theta_t^1\right)\right]^{-1} + \mu_\Xi\left(\theta_t^1\right) - r\left(\theta_t^1\right)\right) - \frac{C_t^1}{\Xi_t^1}\right] dt + \pi_t^1 \sigma_\Xi\left(\theta_t^1\right) dW_t = \\ &= \mu_{\Xi^1}\left(\theta_t^1\right) dt + \sigma_{\Xi^1}\left(\theta_t^1\right) dW_t\end{aligned} \tag{44}$$

where π^1 is the portfolio share invested in the risky asset. The local behavior of returns on the risk-free bond $r\left(\theta^1\right)$ and risky asset (21) as $\theta^1 \searrow 0$ is known from Proposition 7.

The homogeneity of the problem (43–44) motivates the guess

$$\gamma V_t^1 = \left(\Xi_t^1\right)^\gamma \hat{V}^1\left(\theta_t^1\right). \tag{45}$$

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

$$\Xi_t^1 = Y_t \zeta^1\left(\theta_t^1\right) \beta^{-\frac{1}{1-\rho}} \left[\hat{V}^1(\theta)\right]^{\frac{\rho}{\gamma}\frac{1}{1-\rho}}. \tag{46}$$

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)

$$\begin{aligned}0 &= \max_{C^1,\pi^1,\nu^1} \frac{1}{\rho}\beta^{\frac{1}{1-\rho}}\left(\hat{V}^1\right)^{1-\frac{\rho}{\gamma}\frac{1}{1-\rho}} + \hat{V}^1\left(-\frac{\beta}{\rho} + \mu_{\Xi^1} + u^1\sigma_{\Xi^1} - \frac{1}{2}(1-\gamma)\left(\sigma_{\Xi^1}\right)^2\right) + \\ &\quad + \hat{V}_{\theta^1}^1 \theta^1\left(\frac{1}{\gamma}\left(\mu_{\theta^1} + u^1\sigma_{\theta^1}\right) + \sigma_{\theta^1}\sigma_{\Xi^1}\right) + \hat{V}_{\theta^1\theta^1}^1\left(\theta^1\right)^2\frac{1}{2}\frac{1}{\gamma}\left(\sigma_{\theta^1}\right)^2,\end{aligned} \tag{47}$$

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

$$\begin{aligned}\frac{C_t^1}{\Xi_t^1} &= \beta^{\frac{1}{1-\rho}}\left(\hat{V}^1\left(\theta_t^1\right)\right)^{-\frac{\rho}{\gamma}\frac{1}{1-\rho}} \\ \pi_t^1 &= \frac{\left[\xi\left(\theta_t^1\right)\right]^{-1} + \mu_\Xi\left(\theta_t^1\right) + u_1\sigma_\Xi\left(\theta_t^1\right) - r\left(\theta_t^1\right) + \frac{\theta\hat{V}_{\theta^1}^1\left(\theta^1\right)}{\hat{V}^1\left(\theta^1\right)}\sigma_{\theta^1}\left(\theta_t^1\right)\sigma_{\Xi^1}\left(\theta_t^1\right)}{(1-\gamma)\left(\sigma_\Xi\left(\theta_t^1\right)\right)^2},\end{aligned} \tag{48}$$

where μ_{Ξ^1} and σ_{Ξ^1} are the drift and volatility coefficients on the right-hand side of (44), 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 (46).

The solution of this equation determines the consumption-wealth ratio of agent 1 and, consequently, the evolution of her wealth. While a closed-form solution of this equation is not available, it is again possible to characterize the asymptotic behavior as $\theta^1 \searrow 0$, established in Lemma 19. For the proof of that lemma, the following result, 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 $\lim_{x\to-\infty} f'(x) = 0$.*

Lemma 19 *The following results hold:*

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

Proof of Lemma 19. Transformation (45) together with the previously used $\gamma V_t^1 = Y^\gamma \tilde{J}^1\left(\theta_t^1\right)$ imply that

$$\hat{V}^1\left(\theta^1\right) = \beta^\gamma \left(\frac{\tilde{J}^1\left(\theta^1\right)^{1/\gamma}}{\zeta^1\left(\theta^1\right)}\right)^{\gamma(1-\rho)}. \tag{49}$$

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

$$\theta^1 \hat{V}^1_{\theta^1} = \hat{V}^1_{\log\theta^1} \text{ and } \left(\theta^1\right)^2 \hat{V}^1_{\theta^1\theta^1} = \hat{V}^1_{(\log\theta^1)^2} - \hat{V}^1_{\log\theta^1}. \tag{50}$$

Differentiating repeatedly expression (49) 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 (50) converge to zero as $\theta^1 \searrow 0$. ■

These results are similar to those in Lemma 16. They imply that the derivative terms in the ODE (47) 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.

We have thus pinned down the behavior of the last term in the numerator of the portfolio share π^1 in equation (48). This term explicitly takes into account agent 1's knowledge about the impact of her 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 her type will not have any impact on local equilibrium price dynamics, and thus behaves as if she resided in an economy populated only by agent 2.

We can now continue with the proof of Proposition 9. Utilizing Lemma 19 to deduce which terms in ODE (47) vanish and Proposition 7 to determine the limiting values of the remaining coefficients, we obtain

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

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 (44) and (48). ■

Proof of Proposition 10. 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 (9). Utilizing the formula for the wealth-consumption ratio (35) and the result from Lemma 9 then yields

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

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 11. 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 (10) in Assumption 2. The consumption-wealth ratio of the small agent is given in expression (35), and restriction (11) in Assumption 2 assures that this quantity is strictly positive, and the wealth-consumption ratio finite. ■

Proof of Corollary 12. Utilize results in Proposition 9 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 5. ■

Proof of Proposition 13. The difference in expected logarithmic returns is obtained by computing the limiting behavior of

$$\left(\pi^1 \left(\theta^1\right) - \pi^2 \left(\theta^1\right)\right) \left[\left[\xi \left(\theta^1\right)\right]^{-1} + \mu_\Xi \left(\theta^1\right) - r \left(\theta^1\right)\right] - \frac{1}{2} \left[\left(\pi^1 \left(\theta^1\right)\right)^2 - \left(\pi^2 \left(\theta^1\right)\right)^2\right] \left(\sigma_\Xi \left(\theta^1\right)\right)^2,$$

utilizing the results for $\theta^1 \searrow 0$ from Propositions 7 and 9. The first term above is the difference in the risk premium associated with the two portfolios, and the second term is the lognormal correction. The same propositions also contain the results for the consumption-wealth ratios of the two agents. ■

Proof of Corollary 14. The results are obtained by taking limits of the expressions in Proposition 10. ■

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.

Backus, David K., Bryan R. Routledge, and Stanley E. Zin, 2008. "Who Holds Risky Assets?." Mimeo.

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, 2013. "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.

Branger, Nicole, Ioana Dumitrescu, Vesela Ivanova, and Christian Schlag, 2011. "Preference Heterogeneity and Survival in Long-Run Risk Models." Working paper, Goethe University.

Cao, Dan, 2010. "Collateral Shortages, Asset Price and Investment Volatility with Heterogeneous Beliefs." MIT, job market paper.

Chen, Hui, Scott Joslin, and Ngoc-Khanh Tran, 2012. "Rare Disasters and Risk Sharing with Heterogeneous Beliefs." *Review of Financial Studies*. 25 (7): 2189–2224.

Cogley, Timothy, Thomas J. Sargent, and Viktor Tsyrennikov, 2013. "Wealth Dynamics in a Bond Economy with Heterogeneous Beliefs." Forthcoming in the Economic Journal.

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

Coury, Tarek, and Emanuela Sciubba, 2012. "Belief Heterogeneity and Survival in Incomplete Markets." *Economic Theory.* 49 (1): 37–58.

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.

Cvitanić, Jakša, and Semyon Malamud, 2011. "Price Impact and Portfolio Impact." *Journal of Financial Economics.* 100 (1): 201–225.

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.

Fedyk, Yurii, Heyerdahl-Larsen, and Johan Walden, 2013. "Market Selection and Welfare in a Multi-Asset Economy." *Review of Finance*. 17 (3): 1179–1237.

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.

Guerdjikova, Ani, and Emanuela Sciubba, 2010. "Survival with Ambiguity." .

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.

Huang, Chi-Fu, and Henri Pagès, 1992. "Optimal Consumption and Portfolio Policies with an Infinite Horizon: Existence and Convergence." *The Annals of Applied Probability*. 2 (1): 36–64.

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, 2011. "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.

Mailath, George J., and Alvaro Sandroni, 2003. "Market Selection and Asymmetric Information." *Review of Economic Studies.* 70 (2): 343–368.

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.

Muraviev, Roman, 2013. "Market Selection with Learning and Catching Up with the Joneses." *Finance and Stochastics.* 17 (2): 273–304.

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

Øksendal, Bernt, 2007. *Stochastic Differential Equations*, Springer Verlag, 6th edn.

Revuz, Daniel, and Marc Yor, 1999. *Continuous Martingales and Brownian Motion*, Springer-Verlag, 3rd edn.

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*pp. 65–84. 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.