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

Jaroslav Borovička

New York University

jaroslav.borovicka@nyu.edu

October 17, 2016

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 lead to long-run outcomes where both agents survive, or more incorrect agents dominate. 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. The results highlight a crucial interaction between risk sharing, speculative behavior and consumption-saving choice of agents with heterogeneous beliefs, and the role of equilibrium prices in shaping long-run outcomes.

*I am indebted to Lars Peter Hansen for his advice and continuous support. I appreciate helpful comments from Fernando Alvarez, David Backus, Gadi Barlevy, Anmol Bhandari, Katarína Borovičková, Hui Chen, Sebastian di Tella, Valentin Haddad, Narayana Kocherlakota, Virgiliu Midrigan, Alan Moreira, Stavros Panageas, Christopher Phelan, Thomas Sargent, José Scheinkman, 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

A growing body of empirical evidence documents systematic and persistent differences in portfolio returns and saving rates across agents. The evidence calls for theoretical models to analyze the sources of these differences and consequences for general equilibrium prices and evolution of the wealth distribution in the economy. This paper analyzes implications of heterogeneity in agents' beliefs as one plausible factor contributing to these phenomena. Taking this belief heterogeneity as given, I study the determinants of long-run wealth dynamics in a class of equilibrium economies populated by agents endowed with nonseparable recursive preferences.

These 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, 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 and the resulting ability to provide a better account of the patterns observed in asset return data, this class of preferences became the workhorse model used in the asset pricing literature.

I provide a complete analytical characterization of long-run outcomes in an endowment economy populated by two classes of competitive agents (called, for simplicity, two agents) who differ in their beliefs about the distribution of the stochastic aggregate endowment that follows a geometric Brownian motion. Agents are endowed with identical preferences and trade in complete markets.

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 prevail or even dominate the economy and hence affect equilibrium dynamics in the long run. Perhaps most interestingly, agents with arbitrarily large belief distortions can coexist with rational agents in the long-run equilibrium under preference parameterizations typically estimated in asset pricing models, with risk aversion sufficiently higher than the inverse of IES; see, e.g., the long-run risk literature initiated by Bansal and Yaron (2004). In contrast to a large literature on market selection initiated by Alchian (1950) and Friedman (1953), belief heterogeneity should thus be viewed as a natural long-run outcome.

The paper expounds the market forces generating these results. The long-run wealth distribution in the economy is determined by relative *logarithmic* wealth growth rates of the two agents. Agents in an economy accumulate wealth by holding portfolios with high expected *logarithmic* returns and by choosing a high saving rate.[1] The decoupling of risk aversion and IES effectively separates these two decisions. The portfolio choice is driven by interaction of risk aversion and belief distortions, while the consumption-saving decision is determined through the interaction between IES and perceived expected returns on the agent's portfolio. I establish how the portfolio and saving mechanisms emerge from equilibrium price dynamics and determine novel long-run outcomes

[1] To illuminate the difference between level and logarithmic returns, consider a simple case of a risk-neutral agent who starts with a given wealth of k dollars and engages in a sequence of \$1 coin flip bets that win with probability 0.5. While the net *level* return is zero and the agent views these bets as fair, she ultimately ends up with zero wealth with probability one. The expected *logarithmic* return on a sequence of these bets is negative and converges to minus infinity. This argument is closely related to the literature on growth-optimal portfolios, initiated by Kelly (1956) and Breiman (1960, 1961).

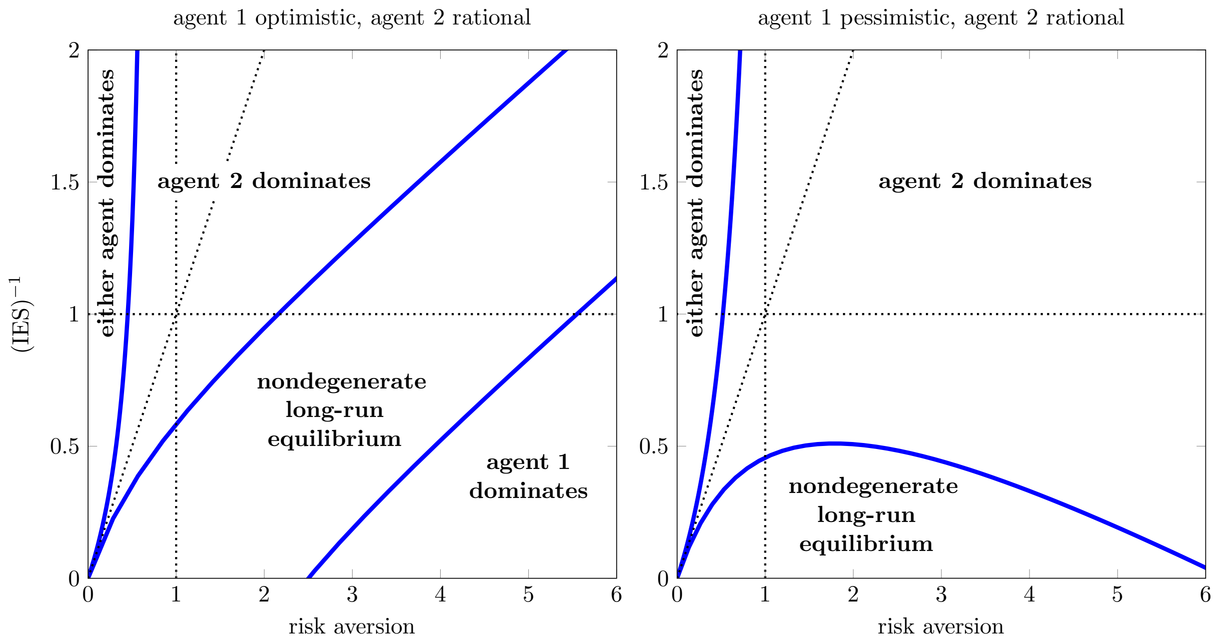

Figure 1: Survival regions for two pairs of agents' beliefs.

in the recursive preference setup.

The paper uncovers a crucial interaction between the use of risky assets for risk sharing and saving on the one hand and, on the other hand, as a speculative tool to trade on belief differences. Specifically, I identify three channels through which individual choices vis-à-vis equilibrium prices influence long-run wealth accumulation:

1. *Risk premium channel*: More optimistic agents hold larger positions in risky assets and thus benefit from high risk premia.

2. *Speculative volatility channel*: Speculative behavior arising from differences in beliefs makes agents choose portfolios with more volatile returns which lowers expected logarithmic returns due to Jensen's inequality.

3. *Saving channel*: Agents with a high perceived expected return on their portfolio choose a high (low) saving rate when IES is high (low) which aids (harms) their wealth growth.

Figure 1 provides an illustration of the results derived in the paper. Consider an economy where agent 2 has correct beliefs while agent 1 is optimistic (left panel) or pessimistic (right panel). Each panel shows long-run survival outcomes for economies with alternative combinations of preference parameters (which are always identical for both agents). Risk aversion is displayed on the horizontal axis, while the inverse of IES on the vertical axis. The dotted upward sloping line represents parameter combinations that correspond to CRRA preferences. The rational agent 2 always dominates in the neighborhood of the diagonal, which continuously extends existing survival results for separable preferences.

I now explore experiments that explain the role of the above three channels for the equilibrium wealth dynamics. These experiments rely on independently changing risk aversion and IES in the economy, and hence highlight the importance of recursive preferences in generating the novel long-run outcomes.

An increase in risk aversion (holding IES fixed, corresponding to a move to the right in the figure) discourages speculative behavior, as agents dislike volatile returns. At the same time, equilibrium risk premia in the economy increase. This increases relatively more the return on the portfolio of the more optimistic agent. When risk aversion is sufficiently high, the *risk premium channel* dominates. More precisely, holding other parameters fixed, there is always a level of risk aversion above which the more optimistic agent dominates. In the left panel of Figure 1 this is the optimistic agent 1, while in the right panel the relatively more optimistic agent is the rational agent 2.

Equilibrium asset price dynamics play a crucial role in how this channel helps preserve long-run heterogeneity. Consider a particular economy with an optimistic agent 1. When the optimistic agent holds only a negligible share of wealth, the risk premium is relatively high to make the rational agent 2 hold the market portfolio. As the optimistic agent accumulates wealth, the risk premium declines, weakening the risk premium channel. The general equilibrium price dynamics thus act as a balancing force, slowing down the rate of wealth accumulation of the optimistic agent. For a nontrivial set of moderately high values of the risk aversion parameter in Figure 1, this mechanism preserves a nondegenerate wealth distribution in the long run.

On the other hand, when risk aversion decreases (a move to the left in the figure), the *speculative volatility channel* determines long-run outcomes. Low risk aversion incentivizes risk taking, and agents choose 'speculative' portfolios with volatile returns that reflect the differences in their assessment of probabilities of future states. While the optimal Markowitz (1952)–Merton (1971) 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. Due to Jensen's inequality, volatile portfolios are detrimental to survival.

The equilibrium price dynamics generate a diverging force through the speculative volatility channel. Equilibrium asset prices have to induce the large agent to hold the market portfolio, while at those prices the negligible agent chooses a risky 'speculative' portfolio with volatile returns and hence a low *logarithmic* expected return. This tends to decrease the wealth growth rate of agents whose wealth share is already small, driving them to extinction, regardless of their belief distortion. More precisely, holding other parameters fixed, there is always a level of risk aversion sufficiently low such that, depending on the sequence of shock realizations, one of the agents will vanish in the long run but it can be either of the two agents with a strictly positive probability.

While the risk aversion parameter determines the portfolio choice decision, the IES parameter is crucial for the consumption-saving decision and impacts the *saving channel* of wealth accumulation. An increase in IES (a move down in the figure) increases relatively more the saving rate of the agent who perceives a higher subjective expected level return on her portfolio.

The saving channel under a high IES hence acts as a converging force that preserves long-run heterogeneity as long as the negligible agent (regardless whether it is agent 1 or 2) chooses a portfolio with a higher *subjective* expected level return than her large counterpart. Equilibrium asset prices have to adjust to induce the large agent to hold the market portfolio. When risk aversion is not excessively high, the negligible agent at these prices forms a speculative portfolio with large positions in assets which she believes are underpriced. Since this portfolio yields a high subjective return perceived by the agent, she consequently chooses a high saving rate and in this way 'outsaves' her extinction.

Figure 1 captures the survival outcomes in the plane of risk aversion/IES parameters for a particular choice of the magnitude of the belief distortions and aggregate uncertainty in the economy. I provide a complete analytical characterization of long-run outcomes for the whole parameter space and isolate the three channels described above. Several conclusions stand out.

First, the channels for the survival mechanism highlight the critical separate contributions of portfolio and consumption-saving decisions and their interaction with 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 wealth shares of the two agents switch. Recursive preferences play a crucial role in shaping these results.

Second, survival of agents with distorted beliefs is a robust outcome. Agents with incorrect beliefs can survive or dominate in bounded and unbounded economies populated by rational agents for a wide range of preference parameters. Moreover, these results do not hinge on belief distortions being small, or symmetric as in Scheinkman and Xiong (2003); in fact, as we will see, they hold for agents with arbitrarily large and arbitrarily asymmetric belief distortions.

Third, unlike in the separable utility case (Sandroni (2000), Blume and Easley (2006) or Yan (2008)), long-run prospects of optimistic and pessimistic agents differ and do not depend solely on the magnitude of the belief distortions. This reflects the asymmetric effects of the above channels on optimistic and pessimists.

Finally, equilibria in which agents with heterogeneous beliefs coexist in the long run occur for parameter combinations that are empirically relevant. In particular, risk aversion has to be sufficiently high to prevent the speculative volatility channel to dominate, and at the same time IES has to be sufficiently high to incentivize agents with a small wealth share to choose a high saving rate vis-à-vis the high subjective expected return on her portfolio, thus outsaving her extinction.

A large literature documents heterogeneity in beliefs of economic agents (Mankiw, Reis, and Wolfers (2003), Malmendier and Nagel (2016)), and its consequences for macroeconomic and asset price dynamics (Buraschi and Jiltsov (2006), Bachmann, Elstner, and Sims (2013), Giacoletti, Laursen, and Singleton (2015), or Barillas and Nimark (2016); see also Xiong (2013) for a survey in the asset pricing context). While I remain agnostic about the source of this belief heterogeneity and focus on its implications on economic outcomes, possible extensions that endogenize these subjective beliefs are discussed in concluding remarks.

Relative to the existing literature that studies long-run outcomes, I analyze in detail the role of the portfolio choice and consumption-saving decision mechanisms. The separable utility case that the literature primarily focused on can be solved by computing static optimal allocations in a planner's problem. The analysis of the interaction of equilibrium returns with individual decisions conducted in this paper provides insights for the determination of wealth dynamics that can be compared with the growing empirical evidence on individual decisions in financial markets and wealth accumulation (Calvet, Campbell, and Sodini (2009), Fagereng, Guiso, Malacrino, and Pistaferri (2016)).

The paper also contributes to the literature along the methodological and technical dimension. First, I provide a novel rigorous proof of the existence and properties of the continuous-time optimal allocation problem with heterogeneous agents endowed with recursive preferences, formulated as a dynamic problem with stochastic Pareto weights. Second, I prove that this class of problems can be studied by focusing on the boundary behavior of the economy when one of the agents becomes negligible, which is often significantly simpler—in my case, I obtain a complete analytical characterization of the survival results despite the fact that the economy does not have a closed-form solution. Finally, I explicitly link optimal allocations to decentralized portfolio decisions and equilibrium price dynamics. Since the methodological approach used in this paper differs significantly from the existing approaches in much of the survival literature, I provide a more detailed discussion and comparison to the literature later in the paper, in Section 7.

The rest of the paper is organized as follows. Section 2 outlines the economic environment and derives the planner's problem. The proof of the existence and differentiability of the solution is deferred to Appendix A. Sections 3 and 4 present the survival results in the form of tight analytical conditions for survival and extinction, followed by a discussion of asset price implications in Section 5. In order to illustrate the full dynamics of the model, Section 6 contains numerical analysis of consumption and price dynamics in the interior of the state space for specific example economies. Section 7 revisits the methodological contribution vis-à-vis the existing literature and Section 8 concludes. Appendix B contains further proofs omitted from the main text. Additional material that provides more detail and extends the analysis is available in the Online Appendix.[2]

2 Optimal allocations under heterogeneous beliefs

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

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

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

[2] https://www.borovicka.org/files/research/survival_heterogeneous_beliefs_online_appendix.pdf

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 univariate Brownian motion W.[3] The scalar aggregate endowment process Y follows a geometric Brownian motion

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

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. They agree on parameters μ_y and σ_y, observe realizations of W_t and Y_t but disagree about their distribution. 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.

Subjective beliefs are constructed so that the agents agree with the data generating measure on zero-probability finite-horizon events.[4] 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.

From now on, I assume that both agents have constant belief distortions u^n. These belief

[3] 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, individual endowments and permissible trading strategies, are adapted to $\{\mathcal{F}_t\}$ and satisfy usual 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.

[4] In order for the belief heterogeneity not to vanish in the long run, the measures P and Q^n cannot be mutually absolutely continuous. Sandroni (2000) and Blume and Easley (2006) link absolute continuity of the subjective probability measures to merging of agents' beliefs. 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 (e.g., Revuz and Yor (1999), Section VIII). The construction prevents arbitrage opportunities in finite-horizon strategies, and the Pareto optimal allocation can thus be decentralized using dynamic trading. 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. Online Appendix, Section OA.3, provides further details.

distortions have a clear economic interpretation. The Girsanov theorem implies that agent n, whose deviation from rational beliefs is described by u^n, views the evolution of the Brownian motion W as distorted by a drift component u^n, i.e., $dW_t = u^n dt + dW_t^n$, 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 distribution of future aggregate endowment. Conditional on time 0, agent n believes that the distribution of $\log Y_t$ is Normal with mean $\log Y_0 + (\mu_y + u^n \sigma_y)\, t$ and variance $\sigma_y^2 t$. When $\sigma_y = 0$, the distinction between optimism and pessimism loses its meaning but the survival problem is still nondegenerate when agents 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.[5]

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

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

subject to

$$d \log \lambda_t^n = -\nu_t^n dt, \quad \lambda_0^n > 0, \quad t \geq 0. \tag{4}$$

where ν^n is called the discount rate process, and λ^n the discount factor process. The felicity function

[5] 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 utilized in this paper. Similarly, the results from Duffie and Lions (1992) for Markov environments do not apply for all cases considered here. 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.

[6] Hansen (2004) offers a tractable summary of the link between recursive and variational utility. Interested readers may refer to the Online Appendix, Section OA.2, for a more detailed discussion.

$F(C,\nu)$ 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[7] 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 A.1 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 2.1 *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

$$\bar{\lambda}^n_t V^n_t(C^n) = \sup_{\nu^n} E_t \left[\int_t^\infty \bar{\lambda}^n_s F(C^n_s, \nu^n_s)\, ds \right] \tag{6}$$

subject to

$$d \log \bar{\lambda}^n_t = -\left(\nu^n_t + \frac{1}{2} (u^n)^2 \right) dt + u^n dW_t, \qquad \bar{\lambda}^n_0 > 0, \quad t \geq 0. \tag{7}$$

Problem (6)–(7) indicates that $F(C^n_t, \nu^n_t)$ can be viewed as a generalization of the period utility function with a potentially stochastic rate of time preference ν^n_t that depends on the properties of the consumption process and thus arises endogenously in a market equilibrium. Moreover, belief distortions are now fully incorporated in the framework of Dumas, Uppal, and Wang (2000)—the only difference is that the modified discount factor process is not locally predictable. The term $-\frac{1}{2}(u^n)^2$ in the drift of $\log \bar{\lambda}^n_t$ reflects the average bias arising from evaluating the utility flow under the subjective belief.

The diffusion term $u^n dW_t$ 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_t is shifted to the right—the agent effectively overweighs good realizations of dW_s. Formula (7) indicates that

[7] 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.

under the true probability measure, positive realizations of dW_t increase $\bar{\lambda}_t^n$, which implies that the optimistic agent discounts positive realizations of dW_t less than negative ones.

2.3 Planner's problem and optimal allocations

The problem of an individual agent (6)–(7) is homogeneous degree one in the modified discount factors and homogeneous degree γ in consumption. In the homogeneous economy, there exists a closed-form solution for the continuation value $V_t^n(Y) = Y_t^\gamma \bar{V}^n$, with $\bar{V}^n$ and the associated constant discount rate $\bar{\nu}^n$ given in the proof of Lemma A.3.

In the heterogeneous economy, I follow Dumas, Uppal, and Wang (2000) and introduce a fictitious planner who maximizes a weighted average of the continuation values of the two agents.[8] The pair of strictly positive initial Pareto weights $\bar{\lambda}_0 \doteq \left(\bar{\lambda}_0^1, \bar{\lambda}_0^2\right)$ determines the initial distribution of wealth. Define the consumption *shares* of the two agents as $\zeta^n \doteq C^n/Y$, $n \in \{1, 2\}$.

Definition 2.2 *The planner's* value function *is the solution to the problem*

$$J\left(\bar{\lambda}_0, Y_0\right) \doteq \sup_{(C^1,C^2)} \sum_{n=1}^{n} \bar{\lambda}_0^n V_0^n\left(C^n\right) = \sup_{(\zeta^1,\zeta^2,\nu^1,\nu^2)} \sum_{n=1}^{2} E_0\left(\int_0^\infty \bar{\lambda}_t^n Y_t^\gamma F\left(\zeta_t^n, \nu_t^n\right) dt\right) \tag{8}$$

subject to the law of motion for the modified discount factors (7) with initial conditions $\left(\bar{\lambda}_0^1, \bar{\lambda}_0^2\right)$, and the feasibility constraint $\zeta^1 + \zeta^2 \leq 1$.

The planner's problem is well-defined under a simple restriction on the parameters of the economy, imposed in Assumption A.1 in Appendix A. The restriction effectively states that the agents have to be sufficiently impatient (β is sufficiently high). Since the survival results do not depend on β, Assumption A.1 does not introduce substantial restrictions for the analysis of the problem.

The planner's problem (8) suggests that we can interpret the modified discount factor processes $\bar{\lambda}^n$ as stochastic Pareto weights. Indeed, if $\bar{\lambda}_0^n$ 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]

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.

[8] 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 (8) is a straightforward extension when individual continuation values are well-defined.

[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. 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.

2.4 Hamilton–Jacobi–Bellman equation

The planner's problem has an appealing Markov structure. Lemma A.3 in Appendix A shows that the value function (8) for the planner's problem at time t is homogeneous degree 1 in $\bar{\lambda} = (\bar{\lambda}^1, \bar{\lambda}^2)$, homogeneous degree γ in Y and can be written as

$$J\left(\bar{\lambda}_t, Y_t\right) = \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^{\gamma} \widetilde{J}\left(\theta_t\right)$$

where $\theta \doteq \bar{\lambda}^1 / \left(\bar{\lambda}^1 + \bar{\lambda}^2\right)$ represents the Pareto share of agent 1 and acts as the only relevant state variable in the problem.

As we will see in the next section, the dynamics of θ are central to the study of survival in this paper. Obviously, θ is bounded between zero and one. It will become clear that for strictly positive initial weights, the boundaries are unattainable, so that θ evolves on the open interval $(0, 1)$. Moreover, the planner's problem can be characterized as a solution to the Hamilton–Jacobi–Bellman equation for $\widetilde{J}(\theta)$. The proof of the following proposition together with further technical discussion is in Appendix A.

Proposition 2.3 *The Hamilton–Jacobi–Bellman equation*

$$\begin{aligned} 0 = \sup_{(\zeta^1,\zeta^2,\nu^1,\nu^2)} & \theta F\left(\zeta^1, \nu^1\right) + (1-\theta) F\left(\zeta^2, \nu^2\right) + \\ & + \left[-\theta\nu^1 - (1-\theta)\nu^2 + \left(\theta u^1 + (1-\theta) u^2\right)\gamma\sigma_y + \gamma\mu_y + \frac{1}{2}\gamma^2\sigma_y^2\right]\widetilde{J}(\theta) \\ & + \theta(1-\theta)\left[\nu^2 - \nu^1 + \left(u^1 - u^2\right)\gamma\sigma_y\right]\widetilde{J}_\theta(\theta) + \frac{1}{2}\theta^2(1-\theta)^2\left(u^1 - u^2\right)^2\widetilde{J}_{\theta\theta}(\theta) \end{aligned} \tag{9}$$

with boundary conditions $\widetilde{J}(0) = \bar{V}^2$ *and* $\widetilde{J}(1) = \bar{V}^1$ *has a unique bounded twice continuously differentiable solution such that* $J\left(\bar{\lambda}_t, Y_t\right) = \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^{\gamma} \widetilde{J}\left(\theta_t\right)$ *is the planner's value function.*

Equation (9) does not generally have a closed-form solution. However, the Pareto share θ of agent 1 remains the only state variable. Its dynamics dictate 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. Hence only the boundary behavior of the scalar Itô process θ matters. Despite the nonexistence of a closed-form solution for $\widetilde{J}(\theta)$, this boundary behavior can be characterized analytically by studying the limiting behavior of the objective function.[10]

[10]Equation (9) is not specific to the planner's problem (8). For instance, Gârleanu and Panageas (2015) 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 (9), 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 Long-run wealth distribution and survival

In this section, I formalize the exact relationship between survival and the boundary behavior of the Pareto share θ (Proposition 3.2), link it to the equilibrium dynamics of the wealth distribution (Proposition 3.3) and derive analytical formulas for the wealth dynamics at the boundaries (Proposition 3.4). These results provide a complete analytical characterization of survival outcomes in terms of fundamental parameters of the economy, and reveal the contribution of two crucial equilibrium forces—returns on agents' portfolios, and agents' saving rates.

This characterization of survival requires taking an approach that is different from the majority of the literature, which typically analyzes the asymptotic properties of relative entropy as a measure of disparity between subjective beliefs and the true probability distribution, and its convergence as $t \nearrow \infty$. I return to a more detailed comparison with this literature in Section 4.6.

Instead, I derive the local dynamics of the Pareto share θ and rely on its ergodic properties, which allow me to investigate the existence of a unique stationary distribution for θ 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 3.1 *Agent 1 becomes extinct along the path $\omega \in \Omega$ if $\lim_{t\to\infty} \theta_t(\omega) = 0$. Otherwise, agent 1 survives along the path ω. Agent 1 dominates in the long run along the path ω if $\lim_{t\to\infty} \theta_t(\omega) = 1$.*

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

Kogan, Ross, Wang, and Westerfield (2011) or Yan (2008) use the consumption shares ζ^n as a measure of survival. Since the consumption share is continuous and strictly increasing in θ and the limits are $\lim_{\theta\searrow 0} \zeta^1(\theta) = 0$ and $\lim_{\theta\nearrow 1} \zeta^1(\theta) = 1$ (see Remark OA.1 in the Online Appendix), the two measures are equivalent in this setting.

3.1 Long-run distribution of the Pareto share

Recall the dynamics of the modified discount factor processes $\bar{\lambda}^n$ in (7). To establish the survival results, it is convenient to consider the transformation $\vartheta \doteq \log(\theta/(1-\theta)) = \log\left(\bar{\lambda}^1/\bar{\lambda}^2\right)$. Then

$$d\vartheta_t = \left[\left(\nu_t^2 + \frac{1}{2}\left(u^2\right)^2\right) - \left(\nu_t^1 + \frac{1}{2}\left(u^1\right)^2\right)\right] dt + \left(u^1 - u^2\right) dW_t \doteq \mu_{\vartheta,t} dt + \sigma_{\vartheta,t} dW_t. \tag{10}$$

Under nonseparable preferences, the discount rates $\nu_t^n = \nu^n(\theta_t)$ are determined endogenously in the model as part of the solution to problem (8) and are given in the proof of Lemma A.8. Intuitively, one would expect a stationary distribution for θ 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 conditions on the drift coefficient $\mu_{\vartheta,t} = \mu_\vartheta(\theta_t)$:

Proposition 3.2 *Define the following 'repelling' conditions (i) and (ii), and their 'attracting' counterparts (i') and (ii'):*

$$
\begin{array}{llll}
(\mathbf{i}) & \lim_{\theta \searrow 0} \mu_\vartheta(\theta) > 0 & (\mathbf{i}') & \lim_{\theta \searrow 0} \mu_\vartheta(\theta) < 0 \\
(\mathbf{ii}) & \lim_{\theta \nearrow 1} \mu_\vartheta(\theta) < 0 & (\mathbf{ii}') & \lim_{\theta \nearrow 1} \mu_\vartheta(\theta) > 0
\end{array}
$$

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 (10), conditions (i) and (ii) are jointly sufficient for the existence of a unique stationary density $q(\theta)$. The proof of Proposition 3.2 is based on the Feller (1952, 1957) 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 3.2 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 $\nu^2(\theta) - \nu^1(\theta)$ *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 3.2 have an intuitive interpretation. Survival condition (i) states that agent 1 survives under the true probability measure even in cases when her beliefs are more distorted, $\left|u^1\right| > \left|u^2\right|$, as long as her relative patience becomes sufficiently high to overcome the distortion when her Pareto share vanishes. 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 equilibrium consumption streams.[11]

[11] 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.

3.2 Decentralization and equilibrium wealth dynamics

Proposition 3.2 states the survival conditions in terms of the endogenous discount rates ν^n. Now I link these conditions to the equilibrium wealth dynamics in the economy, and evaluate analytically the regions in the parameter space in which these conditions hold.

The proof strategy in this section relies on a decentralization argument and utilizes the asymptotic properties of the differential equation (9) for the planner's continuation value. The economy is driven by a single Brownian shock, and two suitably chosen assets that can be continuously traded are therefore sufficient to complete the markets in the sense of Harrison and Kreps (1979). Let the two traded assets be an infinitesimal risk-free bond in zero net supply that yields a risk-free rate $r_t = r(\theta_t)$ and a claim on the aggregate endowment with price $A_t = Y_t \xi(\theta_t)$, where $\xi(\theta)$ is the aggregate wealth-consumption ratio. Individual wealth levels are denoted $A_t^n = Y_t \zeta^n(\theta_t) \xi^n(\theta_t)$, where $\xi^n(\theta)$ are the individual wealth-consumption ratios. Individual wealth levels follow the law of motion $d \log A_t^n = \mu_{A^n}(\theta_t)\, dt + \sigma_{A^n}(\theta_t)\, dW_t$.

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 feature is closely related to the *price impact* that vanishing agents have on equilibrium asset prices, and I discuss the equilibrium asset price dynamics in detail in Section 5.

Proposition 3.3 *The boundary behavior of Pareto shares and agents' wealth satisfies*

$$\lim_{\theta \to \bar{\theta}} (1-\gamma)^{-1} \mu_{\vartheta}(\theta) = \lim_{\theta \to \bar{\theta}} \left[\mu_{A^1}(\theta) - \mu_{A^2}(\theta)\right], \qquad \bar{\theta} \in \{0, 1\}.$$

Survival conditions in Proposition 3.2 can thus be expressed in terms of relative wealth dynamics of the two agents.

Verifying the conditions in Proposition 3.2 therefore amounts to checking that the expected growth rate of the logarithm of wealth is higher for the agent who is at the brink of extinction. In Section 6.1, I revisit in more detail the complete dynamics of the Pareto share, discount rates, and consumption and portfolio choices in the interior of the state space.

The two central forces underlying wealth accumulation and long-run survival are agents' portfolio allocation and consumption-saving decisions. The rate of wealth accumulation can therefore be decomposed into the return on the agent's portfolio net of the consumption rate:

$$d \log A_t^n = d \log R_t^n - (\xi_t^n)^{-1}\, dt.$$

Both expressions on the right-hand side can be characterized analytically at the boundaries. Denoting $d \log R_t^n = \mu_{R^n}(\theta_t)\, dt + \sigma_{R^n}(\theta_t)\, dW_t$ where $\mu_{R^n}(\theta_t)$ is the expected logarithmic return on agent's n portfolio (and $\mu_{R^n}(\theta_t) + \frac{1}{2}\sigma_{R^n}^2(\theta_t)$ the expected level return), we can establish the following decomposition for the case when agent 1 becomes negligible ($\theta \searrow 0$). The case $\theta \nearrow 1$ is symmetric.

Proposition 3.4 *As $\theta \searrow 0$, the difference in the logarithmic wealth growth rates between the agent with negligible wealth and the large agent can be written as*

$$\lim_{\theta \searrow 0} \left[\mu_{A^1}(\theta) - \mu_{A^2}(\theta)\right] = \lim_{\theta \searrow 0} \left[\mu_{R^1}(\theta) - \mu_{R^2}(\theta)\right] + \lim_{\theta \searrow 0} \left[\left(\xi^2(\theta)\right)^{-1} - \left(\xi^1(\theta)\right)^{-1}\right]$$

where the difference in the expected logarithmic portfolio returns is

$$\lim_{\theta \searrow 0} \left[\mu_{R^1}(\theta) - \mu_{R^2}(\theta)\right] = \underbrace{\frac{u^1 - u^2}{(1-\gamma)\sigma_y}}_{\substack{\text{difference in} \\ \text{portfolios}}} \underbrace{\left[(1-\gamma)\sigma_y^2 - u^2\sigma_y\right]}_{\text{risk premium}} - \underbrace{\frac{u^1 - u^2}{1-\gamma}\left(\sigma_y + \frac{1}{2}\frac{u^1 - u^2}{1-\gamma}\right)}_{\text{volatility penalty}} \tag{11}$$

and the difference in consumption rates is given by

$$\lim_{\theta \searrow 0} \left[\left(\xi^2(\theta)\right)^{-1} - \left(\xi^1(\theta)\right)^{-1}\right] = \frac{1}{2}\frac{\rho}{1-\rho} \underbrace{\left[2\left(u^1 - u^2\right)\sigma_y + \frac{\left(u^1 - u^2\right)^2}{1-\gamma}\right]}_{\text{difference in subjective expected returns}}. \tag{12}$$

The proposition reveals a clear separation of the role of risk aversion and IES. The difference in the expected logarithmic portfolio returns at the boundary only depends on the relative risk aversion $1-\gamma$, not on the parameter ρ that determines the IES. The first term represents the *risk premium channel* — the risk premium on the claim on aggregate consumption times the difference in the portfolio shares invested in the risky asset, obtained in equation (16). The risk premium itself is composed 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 leads to a lower expected return). Since survival is driven by the expected *logarithmic* return, volatile portfolios are penalized by a lognormal correction, reflecting the *speculative volatility channel*. This volatility penalty is the dominant force for survival when risk aversion declines to zero ($\gamma \nearrow 1$).

The difference in consumption rates captures consists of two components. The term in brackets is 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,

$$\left[\mu_{R^1}(\theta_t) + \frac{1}{2}\sigma_{R^1}^2(\theta_t) + u^1\sigma_{R^1}(\theta_t)\right] - \left[\mu_{R^2}(\theta_t) + \frac{1}{2}\sigma_{R^2}^2(\theta_t) + u^2\sigma_{R^2}(\theta_t)\right].$$

Here, $\mu_{R^n}(\theta_t) + \frac{1}{2}\sigma_{R^n}^2(\theta_t)$ is the objective expected level return on agent's n portfolio, and $u^n\sigma_{R^n}(\theta_t)$ is the subjective bias. 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., $\rho > 0$), 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, reflecting the *saving channel* of survival.

3.3 Dependence of survival results on individual parameters

The results from Proposition 3.4 reveal 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 which 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 speculate, reflected in the magnitude of the belief distortions u^n. Large belief distortions encourage larger speculative portfolio positions with volatile returns and increase the role of the volatility penalty. Aggregate risk, on the other hand, discourages additional risk taking through speculation.

For instance, if agent 2 has correct beliefs, $u^2 = 0$, the long-run survival outcome is the same whether we fix the belief distortion u^1 and make aggregate endowment deterministic, $\sigma_y \searrow 0$, or if we fix σ_y and make the agent's beliefs infinitely incorrect. I revisit these aspects of the survival results in the next section.

The survival results also do not depend on the time-preference parameter β and the growth rate of the economy μ_y. Both these parameters affect individual consumption-saving decisions in the same way, and hence they have no impact on the difference in the rates of wealth accumulation. This would no longer be true if, for instance, the agents differed in the IES parameter.

In Section 4.4.3, I combine these insights and study an economy with constant aggregate endowment ($\mu_y = \sigma_y = 0$) to show that the survival results are not affected by the nonstationarity of aggregate endowment in a growing or decaying economy.

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 3.2 do occur in nontrivial parts of the parameters space.

Survival conditions in Proposition 3.4 depend only on parameters $\left(\gamma, \rho, u^1/\sigma_y, u^2/\sigma_y\right)$. Figure 2 provides a systematic treatment of the parameter space. Each panel plots the survival results in the 'risk aversion / inverse of IES' plane $(1-\gamma, 1-\rho)$ for different levels of belief distortions. To keep the discussion focused, I concentrate on the case when agent 2 has correct beliefs, $u^2 = 0$. The Online Appendix considers additional cases when beliefs of both agents are distorted but these are all special cases of Proposition 3.2. To get an idea about the magnitude of the belief distortions, recall that an agent with $u^1 = 0.1$ distorts the annual growth rate of aggregate endowment by $u^1\sigma_y$, e.g., believes it to be 2.2% instead of 2% when $\sigma_y = 0.02$.

The shaded area represents the parameter combinations for which a nondegenerate long-run equilibrium exists. The blue dashed lines in the graphs depict parameter combinations for which condition **(i)** in Proposition 3.2 holds with equality (as $\theta \searrow 0$), while the solid red lines capture the same situation for condition **(ii)** (as $\theta \nearrow 1$). The results do not reveal what happens at these boundaries but the long-run outcomes for the interiors of the individual regions are completely characterized by the conditions in Proposition 3.2. 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 $|u^n|$, in our case agent 2) dominates the economy in the long run.

It is useful to start by describing the asymptotic results as either risk aversion or intertemporal elasticity of substitution moves toward extreme values, holding other parameters fixed. These limiting cases isolate the role of the individual survival channels outlined in the introduction. Section 4.4 then analyzes the interaction of these forces in the whole parameter space.

Corollary 4.1 *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 IES 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 IES decreases to zero ($\rho \searrow -\infty$), a nondegenerate long-run equilibrium cannot exist.*

4.1 Low risk aversion and the speculative volatility channel

In order to provide the intuition underlying result **(a)**, consider first the limiting case when agents are risk neutral ($\gamma = 1$). Then the felicity function $F(C, \nu)$ in (5) is linear in C, and agents choose to play a one-shot lottery with all their wealth. The more optimistic agent wins in states with a high realization of the next-period shock, while the other agent wins in states with a low realization. The cutoff is determined so that both agents are willing to participate (the agent with more wealth faces a higher probability of winning). After this one-shot lottery, the losing agent immediately becomes extinct, consuming zero at all subsequent dates.[12]

When the agents are close to risk neutral ($\gamma \nearrow 1$), neither of the agents becomes extinct in finite time, but the same force, reflecting the *speculative volatility channel*, dominates the long-run dynamics. In this case, the wealth dynamics in Proposition 3.4 are dominated by the term

[12] 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). 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 the IES parameter as the elasticity between current consumption and the expected risk-adjusted continuation value next period.

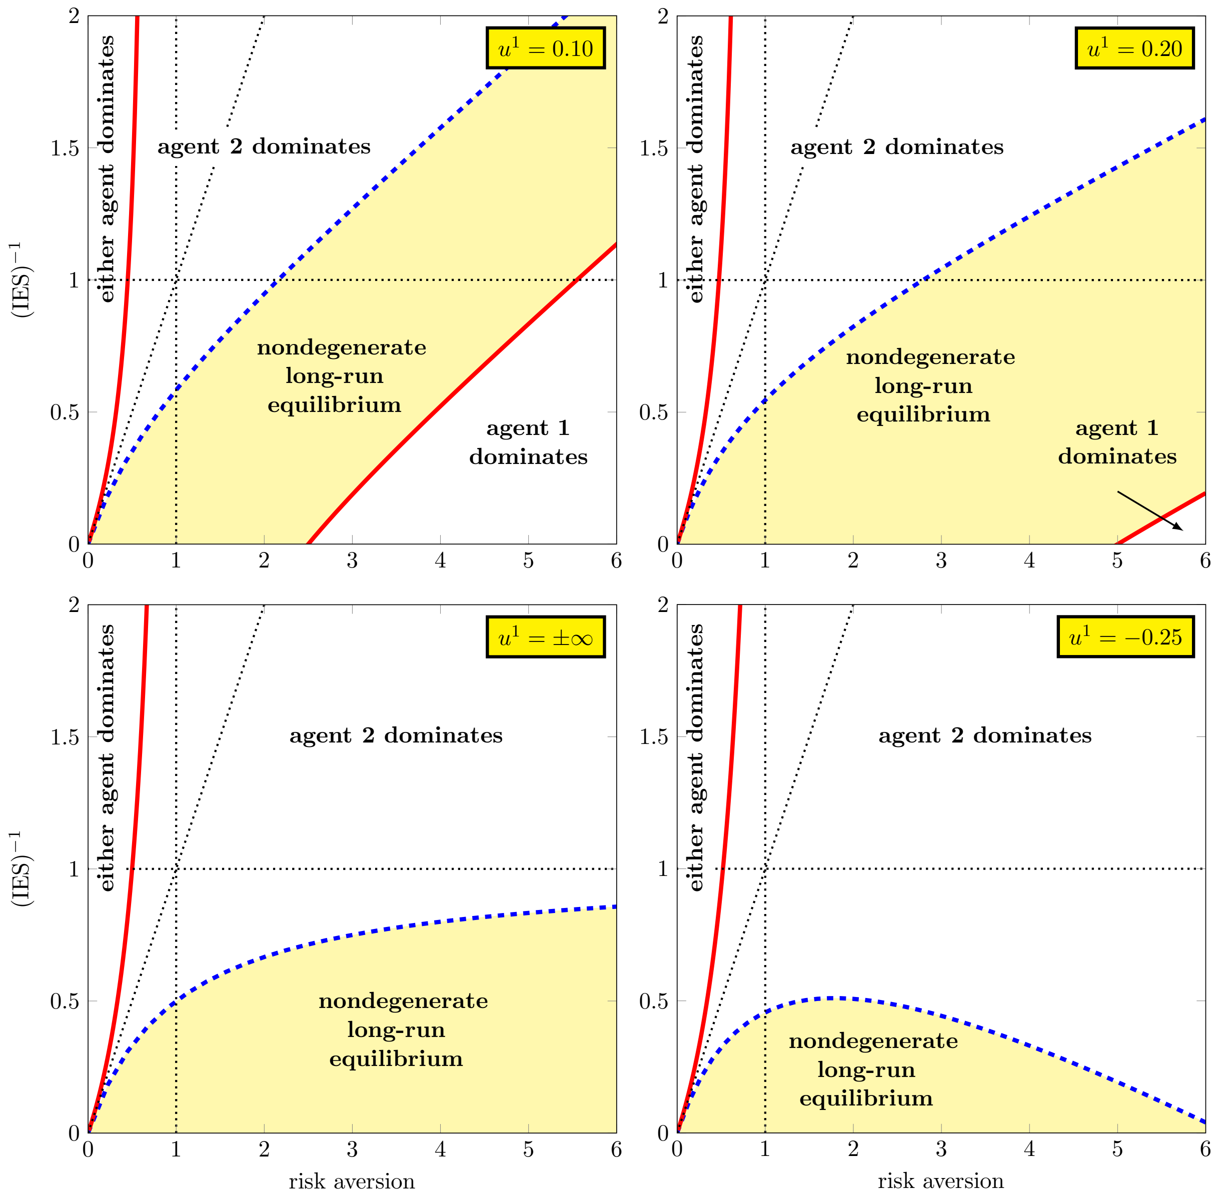

Figure 2: Survival regions for different belief distortions of agent 1 (see legend of each plot). Agent 2 always has correct beliefs, $u^2 = 0$, and the volatility of aggregate endowment is $\sigma_y = 0.02$.

$-\frac{1}{2}\left[\left(u^1 - u^2\right) / (1-\gamma)\right]^2$ in the volatility penalty. This term is always negative and therefore generates an attracting force that drives the distribution toward the boundary, irrespective of the identity of the agent.

Economically, once a sequence of unsuccessful bets reduces the wealth of one agent substantially, equilibrium prices have to adjust to make the large agent hold approximately the market portfolio, which prevents her from taking risky asset positions with volatile returns. At these prices, the negligible agent 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. When risk aversion is low, this 'speculative' position in the negligible agent's portfolio is large, the portfolio return very volatile, and the expected *logarithmic* return on such a portfolio very low. This leads to her extinction on 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.

4.2 High risk aversion and the risk premium channel

In the other limit, when agents become very risk averse ($\gamma \searrow -\infty$, result **(b)**), they put a high value on insuring states with low aggregate shock realizations. Since the relatively more pessimistic agent places a higher probability on these states, she is insured in equilibrium by the relatively more optimistic agent who holds a larger share of her wealth in the risky asset. The price of this insurance is the foregone risk premium in the risky asset.

Expression (11) shows that as risk aversion $1-\gamma$ increases, the risk premium grows linearly, and the pessimistic agent responds by insuring less the adverse states. This is reflected in the vanishing difference in the portfolio positions. However, the total cost of this insurance, given by the product of the risk premium and the difference in the portfolios, converges to a nonzero constant. Since the portfolio positions of the two agents converge to each other as risk aversion increases, their volatilities converge as well, and the volatility penalty vanishes. For high risk aversion levels, the *risk premium channel* dominates, and the more optimistic agent earns a strictly higher expected logarithmic return on her portfolio.

Now pick an economy where a more optimistic agent 1 survives in the long run (in the top two panels in Figure 2, those are economies to the right of the dashed blue line, which also satisfy condition **(i)** from Proposition 3.2). As the more optimistic agent accumulates a larger wealth share, the risk premium declines and the risk premium channel weakens. The general equilibrium price dynamics act as a balancing force, slowing down the rate of wealth accumulation of the optimistic agent. For a nontrivial set of moderately high values of the risk aversion parameter, this mechanism preserves a nondegenerate wealth distribution in the long run.

4.3 IES and the saving channel

Result **(c)** highlights the role of the *saving channel*. 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 3.4 which scales the difference in subjective expected returns by the term $\rho/(1-\rho)$). Whenever an agent becomes negligible, she can choose a 'speculative' portfolio with a high subjective expected return while the market clearing mechanism forces her large counterpart to hold 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. Equilibrium asset prices again generate a balancing force that contributes to long-run survival of both agents.

At the same time, risk aversion cannot be too high for this mechanism to be sufficiently strong. A high risk aversion discourages speculation, 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 3.4 imply that the survival conditions **(i)** and **(ii)** from Proposition 3.2 cannot hold simultaneously. Agents with inelastic preferences *decrease* their saving rate in response to a higher subjective expected return on their portfolio, and the saving channel operates in the opposite direction, as an extinction force for the small agent.

4.4 Asymmetry between optimistic and pessimistic distortions

Survival chances of agents endowed with separable preferences depend solely on the accuracy of their beliefs (Sandroni (2000), Blume and Easley (2006), Yan (2008)). Under recursive preferences, this is no longer true, and long-run wealth accumulation of optimists and pessimists differs. This is the consequence of asymmetric effects of optimistic and pessimistic beliefs on portfolio choice and saving decisions, which recursive preferences allow to separate. In the section, I study in detail the contribution of the individual survival channels to these outcomes.

4.4.1 Isolating the survival channels

Alternative parameter combinations allow us to isolate the three survival channels. When IES $= 1$ ($\rho = 0$), agents' saving rates are constant and equal to each other (each agent's consumption rate is equal to β), and the saving channel is inoperational. The top two panels in Figure 2 show that the risk premium channel alone can lead to long-run survival of an incorrectly optimistic agent 1 when risk aversion is sufficiently high. However, in line with previous discussion, it cannot generate survival of an incorrectly pessimistic agent in the presence of an agent with correct beliefs (bottom right panel).

When $\sigma_y = 0$, the rational risk premium is zero and the risk premium channel is shut down. We will explain in Section 4.4.3 that survival outcomes are in this case described by the bottom left panel of Figure 2. The isolated saving channel then generates long-run survival of the incorrect agent 1 only when IES > 1.

When IES $= 1$ in addition to $\sigma_y = 0$, only the speculative volatility channel is present and agent with incorrect beliefs can never survive with probability one in the presence of a correct agent. When the correct agent is large, risk premia are zero, so the choice of a volatile speculative portfolio by the negligible agent can only lead to a loss in terms of the expected logarithmic return due to the volatility penalty. Since the saving rates are constant as well, the speculative volatility channel is the sole force contributing to the extinction of the incorrect agent.

4.4.2 Optimistic belief distortion

We can now more systematically explore Figure 2. The first panel starts with a moderately optimistic agent 1. The correct agent 2 dominates in the long run in the neighborhood of the dotted diagonal, extending the results for the CRRA case continuously in the parameter space. The graph also confirms all four asymptotic results from Corollary 4.1.

At the same time, there is a nontrivial intermediate region (depicted as shaded in the graph) 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 parameteric 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 $2\sigma_y / \left(u^1 + u^2 + 2\sigma_y\right)$.

As we increase the optimism of agent 1 (second panel in Figure 2), the lines delimiting the shaded region rotate clockwise. The area in which agent 2 dominates expands, reflecting an increase in inaccuracy of agent 1's beliefs, but the region in which both agents coexist in the long run never vanishes.

In fact, as $u^1 \nearrow \infty$ and agent 1 becomes infinitely optimistically biased, we obtain the third panel in Figure 2. The optimistic agent 1 never dominates the economy but there is a large set of parameter combinations for which both agents coexist in the long run. The dashed line delineating this set converges to IES $= 1$ as risk aversion increases. The shaded region includes plausible parameterizations used in asset pricing models; for instance, much of the long-run risk literature initiated by Bansal and Yaron (2004) advocates IES significantly above one and risk aversion above five.

4.4.3 Economy with constant aggregate endowment

In economies in the bottom left panel of Figure 2, incentives to speculate, driven by the magnitude of the difference in belief distortions, are arbitrarily large relative to aggregate risk in the economy. Given $u^2 = 0$, we have seen in Section 3.3 that survival results do not depend on u^1 and σ_y independently but only on the ratio u^1/σ_y. Survival results for the case $u^1/\sigma_y \to \infty$ thus equivalently describe economies with $u^1 \to \infty$ or $\sigma_y \to 0$.

To illuminate this mechanism, consider the limiting case of an economy without aggregate risk, $\sigma_y = 0$, with u^1 being an arbitrary nonzero belief distortion and $u^2 = 0$. Since survival results do not depend on μ_y, we can take $\mu_y = 0$ and hence consider an economy with constant aggregate endowment, $Y_t = \overline{Y}$. Agents in this economy trade for purely speculative motives, see, e.g, Brunnermeier, Simsek, and Xiong (2014). Importantly, this experiment also shows that the survival results in this paper do not hinge upon the economy being unbounded (see Section 7 for a more detailed discussion) and are driven solely by the characteristics of Epstein–Zin preferences.[13]

The long-run survival results for this economy are perfectly equivalent to the bottom left panel in Figure 2. As agent 1 becomes negligible, agent 2 has to hold the market portfolio. Since this portfolio corresponds to a claim on the deterministic consumption stream, her consumption becomes deterministic and risk premia converge to zero (see also the pricing results in Section 5). From the perspective of agent 1, the claim on W now offers a high perceived return and, with IES > 1, this translates into a higher saving rate of the negligible agent. When IES is sufficiently high, the high

[13] The case without aggregate risk requires a clarification regarding the contract space in the decentralized economy. A natural decentralization in the economy with $\sigma_y > 0$ involves a claim to aggregate endowment with unit supply and an infinitesimal risk-free claim in zero net supply. In order to allow agents to trade on their heterogeneous beliefs regarding the probability distribution of the Brownian motion W when $\sigma_y = 0$, a suitable decentralization involves a claim on W in zero net supply and a risk-free claim with supply $\overline{Y}$. This decentralization is explained in more detail in Section OA.3.3 of the Online Appendix.

saving motive is always strong enough to let the negligible agent outsave her extinction and survive in the long run. Section 4.5 also analyzes this economy under symmetric belief distortions.

4.4.4 Pessimistic belief distortion

The third panel in Figure 2 also represents the case when $u^1/\sigma_y \to -\infty$, i.e., the case of an infinitely pessimistic agent 1. Recall that the limit $\left|u^1/\sigma_y\right| \to \infty$ corresponds to a situation where the role of aggregate risk vanishes relative to the speculative motives generated by belief heterogeneity. In this limit, the agents are speculating on the realizations of the Brownian shock W without distinguishing 'good' and 'bad' aggregate states. Because this shock is symmetric, it does not matter whether agent 1 is 'optimistic' and speculates on right-tail realizations of the Brownian shock or 'pessimistic' and speculates on left-tail realizations. This logic is most clearly visible in the case with deterministic aggregate endowment, $\sigma_y = 0$, where the survival results are the same for an arbitrary value of $u^1 \neq 0$.

What happens when the magnitude of pessimism decreases and u^1 starts moving from $-\infty$ closer to zero? The change in the survival regions is represented by a move from the third to the fourth panel of Figure 2. The region in which the pessimistic agent 1 survives actually *shrinks*.

Since the pessimistic agent invests a smaller share of her wealth in the risky asset, she cannot benefit from the risky asset's higher expected return through the risk premium channel. At the same time, a long position in the risky asset would also imply that her *subjective* expected return is even lower and she will not improve her survival chances by choosing a higher saving rate under IES > 1. However, a sufficiently strong speculative motive induces the pessimistic agent to short the risky asset, and makes her in fact *optimistic* about the return on such a portfolio. As we will see in Section 5, the term in brackets in the consumption rate difference (12), which dominates the saving decisions when $\rho \nearrow 1$, is equal to

$$\left(u^1 - u^2\right) \sigma_y \left(1 + \pi^1(0)\right) \tag{13}$$

where $\pi^1(0)$ is equal to agent 1's risky portfolio share. 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 can generate a sufficiently strong offsetting saving motive that will allow the pessimistic agent to outsave her extinction.

The region in the fourth graph in which the two agents coexist does not include high levels of risk aversion and shrinks for smaller belief distortions. A high level of risk aversion or a lower incentive to speculate caused by a smaller belief distortion prevent the small agent from choosing a sufficiently large short position in the risky asset which is, as shown in formula (13), 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

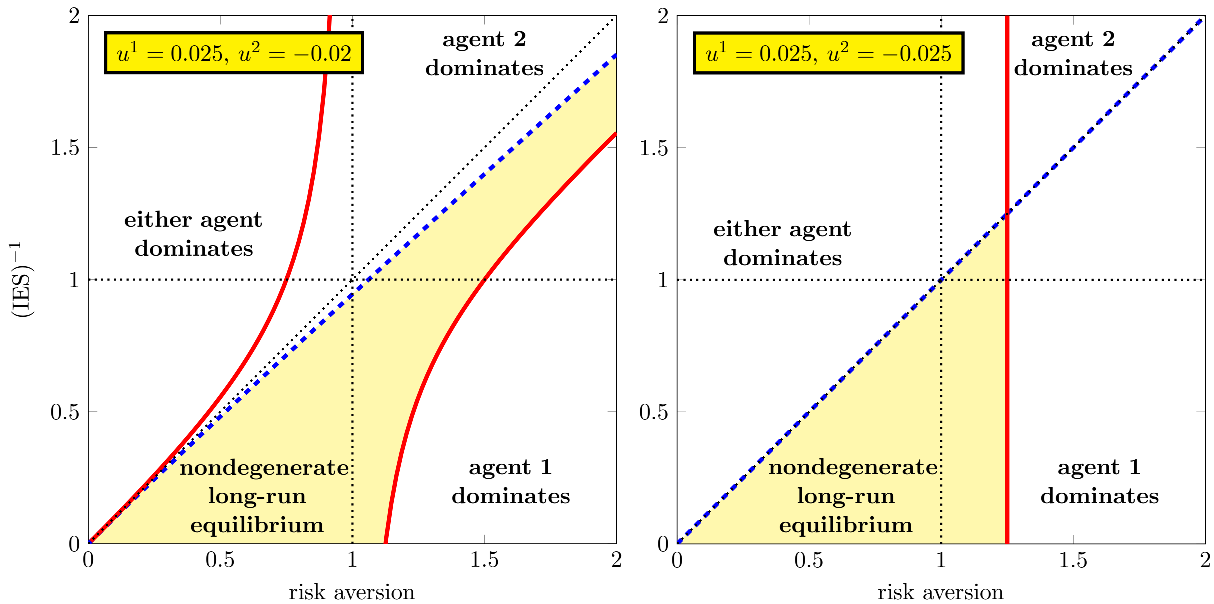

Figure 3: Survival regions for an optimistic agent 1 and a pessimistic agent 2 (see legend of each plot). Volatility of aggregate endowment is $\sigma_y = 0.02$.

generated through the high subjective expected return disappears.

4.5 Symmetric belief distortions

Figure 2 analyzes economies where agent 2 has correct beliefs. Another interesting case arises when the two agents have equal magnitudes of belief biases with opposite signs. Figure 3 captures the case of an optimistic agent 1 ($u^1 = 0.025$) and a pessimistic agent 2. In the left panel, beliefs of agent 2 are somewhat less biased ($u^2 = -0.02$), while in the right panel, the magnitude of belief biases is equal for both agents.

As in the previous analysis, all four combinations of survival outcomes occur in relevant parts of the state space. In the right panel, the parameter space is exactly separated along the CRRA preference parameterizations (dotted blue line), and along the solid red vertical line that lies at the level of risk aversion equal to u^1/σ_y. Therefore, as the magnitude of the belief distortion $u^1 = -u^2$ increases, or as aggregate uncertainty vanishes and we converge to the setup from Section 4.4.3, the vertical line shifts to the right and the region with a nondegenerate long-run equilibrium expands.[14] Section OA.4 in the Online Appendix provides more detail.

4.6 Separable 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)

[14] Equilibria along the dotted blue line in the right panel of Figure 3 are analyzed in the proof to Corollary 4.2. In this specific situation with CRRA preferences and symmetric distortions, none of the agents becomes extinct but nonetheless a nondegenerate long-run distribution for θ does not exist.

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 3.2 confirm these results as follows:

Corollary 4.2 *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 θ in (10) do not depend on the characteristics of the endowment process. 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

$$\bar{\lambda}_0^1 M_t^1 U'\left(C_t^1\right) = \bar{\lambda}_0^2 M_t^2 U'\left(C_t^2\right).$$

Survival analysis in the separable case thus corresponds to analyzing a sequence of state- and time-indexed static problems that are interlinked only by the initial Pareto weights $\bar{\lambda}_0^n$. When 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 $\lim_{t\to\infty} M_t^1 = 0$ (P-a.s.) and $M_t^2 \equiv 1$, and thus $\lim_{t\to\infty} U'\left(C_t^1\right)/U'\left(C_t^2\right) = +\infty$ (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 $\lim_{t\to\infty} \zeta_t/(1-\zeta_t) = 0$ (P-a.s.). Kogan, Ross, Wang, and Westerfield (2011) analyze this case for a general class of period utility functions.

5 Asset prices and price impact

Kogan, Ross, Wang, and Westerfield (2006, 2011) distinguish between survival of agents with incorrect beliefs and their impact on equilibrium prices. This leads to the following two distinct questions. First, do agents with incorrect beliefs have an impact on asset prices in the long run? Second, if an agent currently holds a negligible wealth share, does she have any impact on *current* asset prices and returns, even if she potentially survives in the long run?

The results in this section reveal that as the Pareto share of one of the agents becomes negligible, current asset prices and infinitesimal asset returns converge to those which prevail in a homogeneous economy populated by the agent with the large Pareto share, regardless whether the negligible agent ultimately survives or vanishes. This directly implies that an agent who becomes extinct in the long run also has no long-run *price impact*. On the other hand, an agent who survives will have an impact on asset prices in the future when her wealth share recovers, even if her current wealth share and price impact may be negligible.

The ability to pin down asset returns when the wealth of one agent is negligible plays a crucial role in establishing the analytical results in Proposition 3.4 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 following Proposition summarizes the limiting pricing implications as the wealth share of one of the agents becomes arbitrarily small. Without loss of generality, it is sufficient to focus on the case $\theta \searrow 0$.

Proposition 5.1 *As $\theta \searrow 0$, the infinitesimal risk-free rate $r(\theta)$, the aggregate wealth-consumption ratio $\xi(\theta)$, and the coefficients of the aggregate wealth process $d\log A_t = \mu_A(\theta_t)\, dt + \sigma_A(\theta_t)\, dt$ converge to their homogeneous economy counterparts:*

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

Further, the infinitesimal logarithmic return on the claim on aggregate wealth,

$$
d\log R_t \doteq \left[[\xi(\theta_t)]^{-1} + \mu_A(\theta_t)\right] dt + \sigma_A(\theta_t)\, dW_t, \tag{14}
$$

has coefficients that converge as well.

The proof is provided in Appendix B and is based on the characterization of the dynamics of the equilibrium stochastic discount factor as $\theta \searrow 0$. Marginal utility under recursive preferences is forward-looking and depends on agent's continuation value (see the stochastic discount factor specification (33)), so that future equilibrium consumption dynamics affect the current local evolution of the stochastic discount factor. A crucial step involves showing that the dynamics of the relative Pareto share $\vartheta(\theta)$ in (10) have bounded drift and volatility coefficients. This implies that the rate of wealth accumulation of the negligible agent is sufficiently slow so that even in the case she survives, her potential future impact on the economy is sufficiently distant to be inconsequential for the current evolution of the stochastic discount factor and asset prices. In addition, the agent with negligible wealth has no current price impact not only on the two assets that dynamically complete the market but also on every finite-maturity bond and consumption strip.

Corollary 5.2 *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 $\theta \searrow 0$.*

5.1 Decision problem of an agent with negligible wealth

Proposition 5.1 establishes that the actual *general equilibrium* price dynamics and choices of the large agent in the proximity of the boundary are locally the same as those in an economy populated only by the large agent 2. To conclude the argument, we need to infer the wealth dynamics for agent 1 that has negligible wealth. As in the case of the large agent, the impact of future equilibrium consumption dynamics on the current decisions of the negligible agent becomes immaterial as $\theta \searrow 0$, despite the nonseparability of preferences.

Proposition 5.3 *The consumption-wealth ratio of agent 1 converges to*

$$\lim_{\theta \searrow 0} \left[\xi^1(\theta)\right]^{-1} = \left[\xi(0)\right]^{-1} - \frac{1}{2}\frac{\rho}{1-\rho}\left[2\left(u^1 - u^2\right)\sigma_y + \frac{\left(u^1 - u^2\right)^2}{1-\gamma}\right] \tag{15}$$

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

$$\lim_{\theta \searrow 0} \pi^1(\theta) = 1 + \frac{u^1 - u^2}{(1-\gamma)\sigma_y}. \tag{16}$$

Proposition 5.3 derives the consumption-saving decision (15) and portfolio allocation decision (16) relative to the same decisions of the large agent 2. Recall that agent 2's choices agree in the limit as $\theta \searrow 0$ with aggregate ones—equilibrium prices have to adjust so that her consumption-wealth ratio is equal to the aggregate ratio $[\xi(0)]^{-1}$ and she holds the market portfolio, $\pi^2(0) = 1$.

At these equilibrium prices, agent 1's consumption and portfolio choices deviate from the aggregate dynamics according to formulas (15) and (16). When these deviations lead to a high saving rate or a high logarithmic return on agent 1's portfolio, they can prevent her extinction. As the formulas indicate, when $u^1 = u^2$ the agents are identical and their decisions and wealth dynamics coincide. A relatively more optimistic agent ($u^1 > u^2$) chooses a larger share of her wealth to be invested in the risky asset, and chooses to save more (less) when IES > 1 (< 1).

These portfolio and consumption-saving decision of agent 1 as $\theta \searrow 0$ coincide with a 'partial equilibrium' solution where agent 1 locally behaves as if she lived forever as an infinitesimal agent in a homogeneous economy populated only by the large agent 2. The logic of the proof relies on showing that by pushing the current θ arbitrarily close to zero, 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.

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.

6 Equilibrium dynamics and evolution of wealth

Propositions 3.2 and 3.4 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. We start with the following observation.

Proposition 6.1 *When agent n survives in the long run, she attains an arbitrarily large wealth share $A^n/A \in (0,1)$ and consumption share $\zeta^n \in (0,1)$ with probability one at some future date t.*

This result is a consequence of sufficient mixing in the Pareto share process θ. However, assessing whether the agent also *typically* holds a large wealth share requires a full numerical solution of

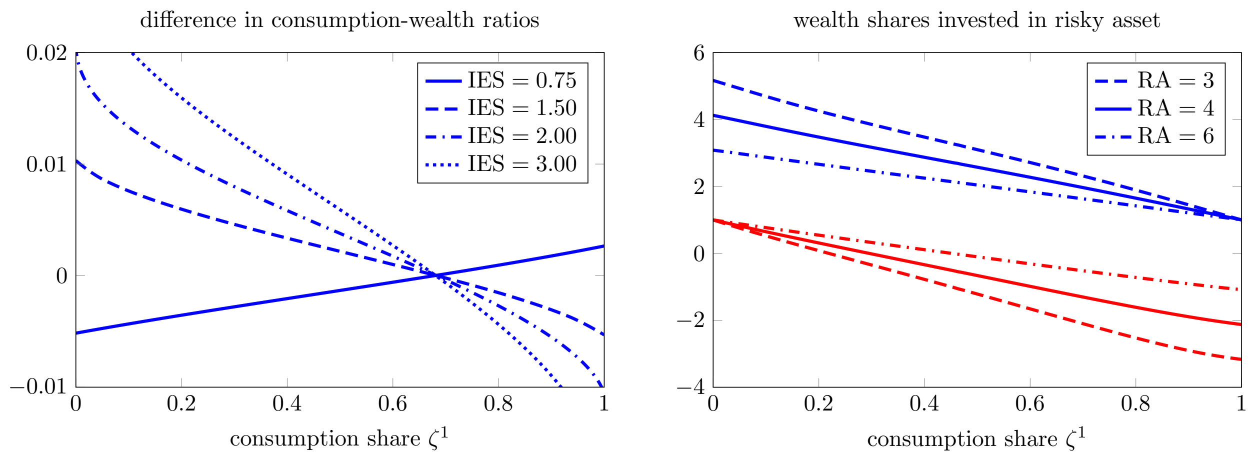

Figure 4: *Left panel*: Difference in consumption-wealth ratios $\left(\xi^2\right)^{-1} - \left(\xi^1\right)^{-1}$ as a function of the consumption share ζ^1 of agent 1, plotted for different levels of intertemporal elasticity of substitution. The remaining parameters are $u^1 = 0.25$, $u^2 = 0$, RA $= 2$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.02$. *Right panel*: 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, plotted for different levels of risk aversion. The remaining parameters are $u^1 = 0.25$, $u^2 = 0$, 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).

the model in the interior of the state space. This section investigates the equilibrium allocations and agents' decisions by numerically solving the planner's problem (9) and the associated decentralization. I show using a series of examples that agents with incorrect beliefs can indeed have a quantitatively substantial impact on the wealth dynamics, and discuss the dependence of these dynamics on the preference and belief distortion parameters.

6.1 Survival forces in the interior of the state space

The two essential components of the survival mechanism identified in Section 3.2 are the propensity to save and the portfolio allocation of the two agents. The left panel in Figure 4 displays the effect of propensity to save in the form of difference in the consumption-wealth ratios $[\xi^n(\theta)]^{-1}$, which are primarily driven by the IES. For the case of IES $= 1$, the difference is zero since each agent consumes a fraction β of her wealth per unit of time, and the saving channel is inactive. 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. In the graph, high levels of IES are conducive to survival of both agents — the difference in consumption rates $\left(\xi^2\right)^{-1} - \left(\xi^1\right)^{-1}$ is positive when agent 1 is negligible, and negative when agent 2 is negligible.

The portfolio allocation mechanism is depicted in the right panel of Figure 4. The share of wealth invested in the risky asset is primarily driven by the risk aversion parameter. The graph shows the optimistic agent 1's wealth share π^1 invested in the risky asset in blue (upper three lines), and π^2 in red (lower three lines). As the consumption share of agent n converges to 1, her portfolio allocation $\pi^n \to 1$, reflecting the fact that the large agent's portfolio position must converge to the market portfolio.

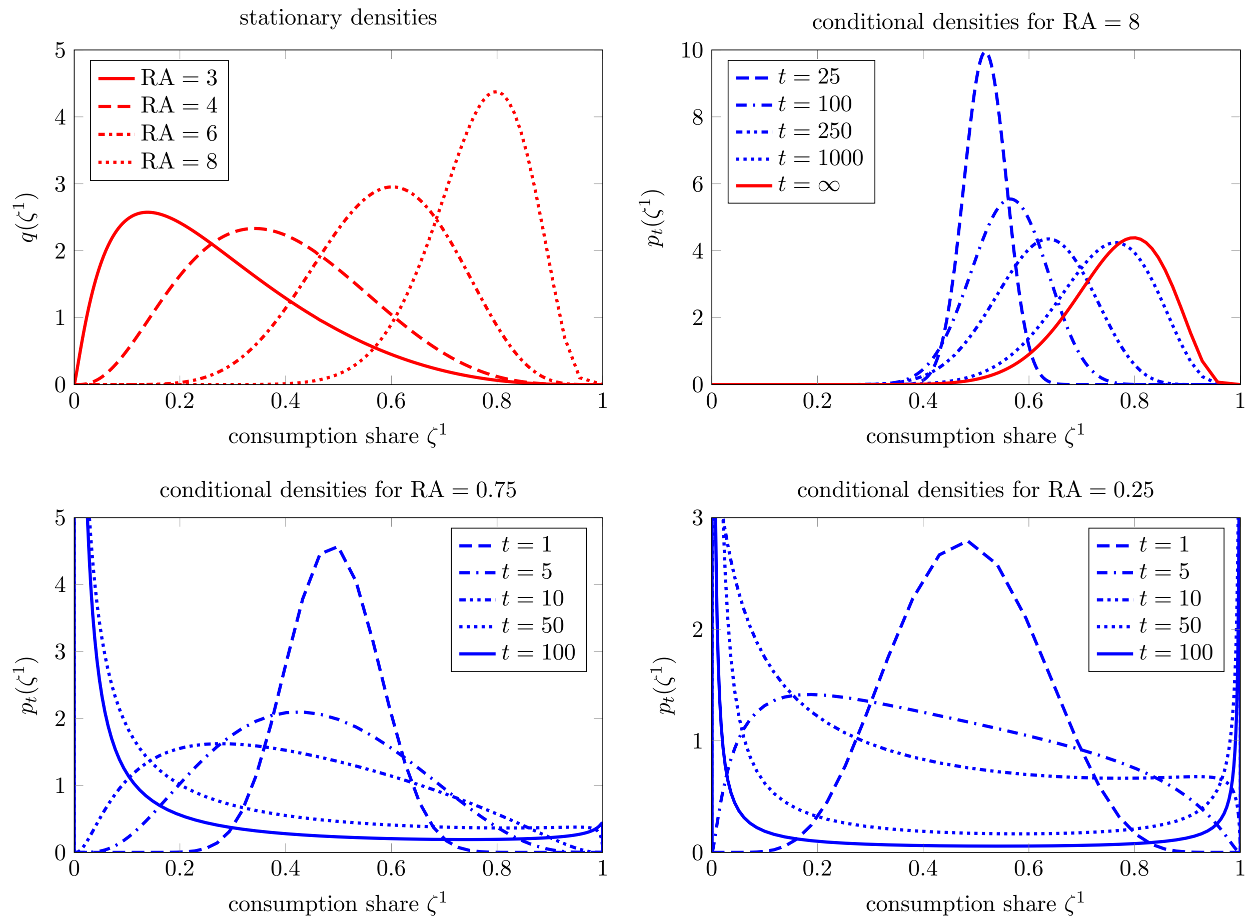

Figure 5: The top left panel depicts the stationary distributions for the consumption share $\zeta^1(\theta)$ of the agent with optimistically distorted beliefs. All models are parameterized by $u^1 = 0.25$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.02$, and differ in levels of risk aversion, shown in the legend. The remaining three panels plot the distributions of $\zeta^1(\theta_t)$ conditional on $\zeta^1(\theta_0) = 0.5$ 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.

A higher risk aversion (dash-dotted lines) limits the amount of leverage, and the portfolio positions are closer to one. This in turn limits the impact of speculative motives on market outcomes, and the role of the risk premium channel increases. Notice that some degree of speculative behavior is necessary for the survival of a pessimistic agent—when risk aversion is high, she does not choose a sufficiently large short stock position that would make her sufficiently optimistic about the return on her own wealth and outsave the rational agent when IES > 1.

6.2 Stationary distributions and evolution over time

The full solution of the model allows us to study the evolution of the wealth distribution over time. I start in the top left graph of Figure 5 with the densities $q\left(\zeta^1\right)$ for the stationary distribution of the consumption share ζ^1 in example economies where both agents survive, for the case of an optimistic agent 1 and correct agent 2 and alternative levels of risk aversion. Proposition 6.1 already revealed that these densities have a full support on $(0, 1)$.

The graphs in Figure 2 showed that increasing risk aversion improves the survival chances of the optimistic agent 1. The top left graph of Figure 5 provides a complementary perspective. As we increase risk aversion, the distribution of the consumption share shifts in favor of the optimistic agent, due to the stronger risk premium channel.

Several observations emerge. First, when both agents survive in the long run, the more incorrect agent can plausibly own and consume a substantial share of aggregate endowment. Second, the shape of the stationary densities for the consumption share indicates that long-run equilibria permit substantial variation over time in these consumption shares. Finally, the same survival channels that generate different types of long-run survival outcomes also act in favor of individual agents within the interior of the state space.

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 3.2 gives sufficient conditions for the existence of a unique stationary distribution for θ but it does not say anything about the rate of convergence.

I show in the Online Appendix that under the conditions from Proposition 3.2, convergence for the state variable θ occurs at an exponential rate, so that the process θ 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 5 plot the conditional distribution of the consumption share $\zeta^1(\theta_t)$ of the optimistic agent 1 conditional on $\zeta^1(\theta_0) = 0.5$ for different time horizons t. These are computed from the dynamics of the state variable θ in equation (10) by solving the associated Kolmogorov forward equation

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

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

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

The graphs show the evolution of the conditional distribution for three cases. In the top right graph, the conditional distribution converges to a nondegenerate long-run 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 motion W, and wealth and consumption shares evolve only slowly. In the example in Figure 5, it takes roughly 2,500 periods until the density p_t is visually indistinguishable from the stationary density. As risk aversion decreases, and agents are willing to speculate more, the evolution of the conditional density p_t speeds up.

While the evolution of the conditional density in Figure 5 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 and agents 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 $p_t\left(\zeta^1\right)$ to the stationary density $q\left(\zeta^1\right)$ when both agents survive in the long run.[15] 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.

7 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.[16]

[15] 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.

[16] 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

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, local properties of the utility function are 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 survive.

With unbounded aggregate endowment, local properties of the utility function become an additional survival factor. Even if preferences are identical across agents, 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 elegantly 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.[17] 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. Section 4.4.3 also documents that unbounded aggregate endowment is not essential for the survival mechanisms analyzed in this economy.

Survival analysis under separable preferences corresponds to studying a sequence of time- and state-indexed static problems that are only interlinked through the initial marginal utility of wealth. Nonseparability of preferences breaks this straightforward link, and I therefore develop a different method that is more suitable for this environment. I utilize the planner's problem derived in Dumas, Uppal, and Wang (2000) and extend it to include heterogeneity in beliefs. While 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 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. 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

portfolio rule selection.

[17] The survival results under separable utility thus also hold for 'exotic' endowment processes like the rare disaster framework in Chen, Joslin, and Tran (2012). 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), Beker and Chattopadhyay (2010), Coury and Sciubba (2012), Cao (2013a) or Cogley, Sargent, and Tsyrennikov (2013).

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. Beker and Espino (2011) consider separable preferences under a richer set of beliefs. 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. Branger, Dumitrescu, Ivanova, and Schlag (2011) analyze survival in long-run risk models with heterogeneous recursive preferences. Bhandari (2015) studies a problem where heterogeneous beliefs arise endogenously in a setup with robust preferences of Hansen and Sargent (2001). Massari (2016) links survival conditions to a comparison of agents' beliefs and equilibrium prices.

In a related paper, Guerdjikova and Sciubba (2015) study a heterogeneous-preference economy populated by expected-utility and smooth ambiguity averse agents. They show that smooth ambiguity averse agents can dominate in the economy if their preferences satisfy a decreasing absolute ambiguity aversion condition. However, they do not identify cases where both types of agents coexist in the long run, nor do they analyze economic decisions in the decentralized economy. Finally, Dindo (2015) formulates a discrete-time problem analogous to the one studied in this paper, and is able to derive analytical results for particular configurations of preference parameters using different techniques, confirming that the survival results are not specific to the continuous-time Brownian information framework.

8 Concluding remarks

This paper analyzes portfolio and consumption saving decisions and their implications for long-run wealth dynamics in an economy populated by agents with heterogeneous beliefs. It shows that the robust survival results favoring agents with most accurate beliefs found in the existing literature are specific to the class of separable preferences. Under nonseparable recursive preferences of the Duffie–Epstein–Zin type, long-run outcomes in which heterogeneity prevails and agents with incorrect beliefs affect equilibrium dynamics arise for a broad set of plausible parameterizations when risk aversion is larger than the inverse of the intertemporal elasticity of substitution.

The analysis reveals the important distinct roles played by risk aversion with respect to intratemporal gambles that determines risk taking, and intertemporal elasticity of substitution that drives the consumption-saving decision. In particular, the paper uncovers the complex interaction between risk sharing and speculative motives of agents with heterogeneous beliefs, and their joint impact on the saving behavior. Speculative behavior by an agent with incorrect beliefs distorts her rationally optimal risk-return tradeoff. This can aid wealth accumulation if it leads the agent to choose a portfolio with a higher expected logarithmic return. Similarly, a higher *subjective* expected return implies a higher saving rate when IES is sufficiently high.

Critical for obtaining the survival results, and in particular the nondegenerate long-run equilibria, are the general equilibrium price dynamics generated by shifts in the wealth distribution. When the wealth share of one of the agents becomes negligible, prices have to adjust so that the large agent holds the market portfolio and consumes the aggregate endowment. When these prices lead the small agent to choose a portfolio and a saving rate that lead to a high wealth growth rate, she can prevent her extinction. I show that preference parameterizations that are conducive to long-run coexistence of both agents involve risk aversion higher than the inverse of IES, in line with parameter combinations used in the asset pricing literature.

The focus on individual decisions also allows one to map theoretical results to the growing empirical evidence on differences in portfolio choice and saving rates (Calvet, Campbell, and Sodini (2009), Fagereng, Guiso, Malacrino, and Pistaferri (2016)), and the implications for wealth accumulation (Kopczuk (2015), Saez and Zucman (2016)). Kendall and Oprea (2016) complement this evidence with results from laboratory experiments that link belief biases, portfolio choice and consumption-saving decisions.

In this way, the paper contributes to the theoretical literature analyzing various factors underlying long-run wealth heterogeneity, see, e.g., Benhabib, Bisin, and Zhu (2011), Gabaix, Lasry, Lions, and Moll (2015), or the survey in Benhabib and Bisin (2016). In particular, it complements alternative models of incorrect beliefs and limited investor sophistication (Kacperczyk, Nosal, and Stevens (2015), Lusardi, Mitchell, and Michaud (2016)). While this paper analytically exposes the main economic forces through which belief heterogeneity impacts long-run wealth accumulation, further quantitative analysis is left for future work.

On the technical side, I provide a novel existence proof of a classical solution to the Hamilton–Jacobi–Bellman equation associated with the planner's problem and its equivalence with the planner's value function. The survival results are tightly linked to the analytical characterization of the boundary behavior of this equation and agents' portfolio and consumption choices in the associated decentralized economy. Since this type of ODE arises in a wider class of recursive utility problems, these techniques can be utilized in a broader variety of economic applications.

These analytical results are obtained for a two-agent economy with an aggregate endowment process that is specified as a geometric Brownian motion. However, the economic forces underlying the long-run wealth dynamics are not specific to this environment and extend to more general environments with richer consumption dynamics and endogenously determined beliefs arising from learning or concerns for model misspecification. I discuss some of these extensions in the Online Appendix.

This paper is agnostic about the source of belief heterogeneity and focuses on its implications for long-run wealth accumulation. Further research on the origins of this belief heterogeneity, arising perhaps from informational frictions, ambiguity aversion or versions of rational inattention will also help us better understand the wealth dynamics in the economy.

Appendix

The Appendix provides proofs of the propositions from the main text. Lengthier detailed calculations are deferred to Sections OA.7 and OA.8 of the Online Appendix.

A Proof of Proposition 2.3

I prove the proposition through a sequence of lemmas. The proof builds on results from Fleming and Soner (2006), Pham (2009) and Strulovici and Szydlowski (2014). The framework differs, however, along important dimensions, in particular the endogenously determined discount rate and vanishing volatility at the boundaries, so that it requires a separate treatment.

Section A.1 (Lemmas A.3 and A.6) establishes elementary properties of the value function. In Section A.2, I formulate the corresponding Hamilton–Jacobi–Bellman equation. Proving the existence and properties of the solution of this HJB equation is complicated by the fact that the volatility of the Pareto share process θ vanishes at the boundaries of the interval $\theta \in [0,1]$. In Section A.3, I therefore formulate an auxiliary problem on the interval $[\varepsilon, 1-\varepsilon]$. Lemma A.8 proves the existence, uniqueness and differentiability of the solution to this problem. In Section A.4 (Corollary A.9), I extend the solution to the interval $[0,1]$ through a limiting argument. Finally, in Section A.5 (Lemma A.11), I prove the usual verification theorem.

The planner's problem is well-defined when the following restriction on the parameters holds.

Assumption A.1 *The parameters in the model satisfy the restrictions*

$$\beta > \max_n \rho \left(\mu_y + u^n \sigma_y + \frac{1}{2}\gamma\sigma_y^2 \right), \tag{17}$$

$$\beta > \max_n \rho \left(\mu_y + u^{\sim n} \sigma_y + \frac{1}{2}\gamma\sigma_y^2 \right) + \frac{\rho}{1-\rho} \left[\left(u^n - u^{\sim n}\right)\sigma_y + \frac{1}{2}\frac{\left(u^n - u^{\sim n}\right)^2}{1-\gamma} \right] \tag{18}$$

for $n \in \{1,2\}$ 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 is a sufficient condition assuring that the wealth-consumption ratio is asymptotically well-behaved in the survival proofs when the agent becomes infinitesimally small. Both conditions are restrictions on the time-preference parameter of the agents and can always be jointly satisfied by making the agents sufficiently impatient. For instance, when IES $= 1$, these conditions amount to $\beta > 0$. However, since the survival results do not depend on β, Assumption A.1 does not introduce substantial restrictions for the analysis of the problem. Section OA.4.1 in the Online Appendix provides more detail.

A.1 Properties of the value function

We start with definitions and some elementary properties of the value function.

Definition A.2 *The planner's control $a = \left(C^1, C^2, \nu^1, \nu^2\right)$ is admissible if $C^1 + C^2 = Y$ and, for $n \in \{1,2\}$, the Pareto weight processes $\bar{\lambda}^n$ given by*

$$d\log \bar{\lambda}_t^n = -\left(\nu_t^n + \frac{1}{2}\left(u^n\right)^2\right)dt + u^n dW_t$$

have a unique strong solution and

$$V_t^n\left(C^n,\nu^n\right)\doteq E_t\left[\int_t^{\infty}\frac{\bar{\lambda}_s^n}{\bar{\lambda}_t^n}\left|F\left(C_s^n,\nu_s^n\right)\right|ds\right]<+\infty.$$

The set of admissible controls is denoted $\mathcal{A}$.

The value function satisfies the following homogeneity property.

Lemma A.3 *The value function (8) satisfies $J\left(\bar{\lambda}_t,Y_t\right)=\left(\bar{\lambda}_t^1+\bar{\lambda}_t^2\right)Y_t^{\gamma}\widetilde{J}\left(\theta_t\right)$ where $\widetilde{J}\left(\theta_t\right)$ is a bounded function of the Pareto share $\theta_t=\bar{\lambda}_t^1/\left(\bar{\lambda}_t^1+\bar{\lambda}_t^2\right)$.*

Proof. See Section OA.7 in Online Appendix. ■

Given $\nu\doteq\left(\nu^1,\nu^2\right)$, the optimal choice of ζ^n in (8) only involves static optimization and yields

$$\begin{aligned}h^0\left(\nu,\theta\right)&\doteq\max_{\zeta^1+\zeta^2=1}\theta F\left(\zeta^1,\nu^1\right)+\left(1-\theta\right)F\left(\zeta^2,\nu^2\right)=\\&=\frac{\beta}{\gamma}\left[\left(\theta\left(\frac{\gamma-\rho\frac{\nu^1}{\beta}}{\gamma-\rho}\right)^{1-\frac{\gamma}{\rho}}\right)^{\frac{1}{1-\gamma}}+\left(\left(1-\theta\right)\left(\frac{\gamma-\rho\frac{\nu^2}{\beta}}{\gamma-\rho}\right)^{1-\frac{\gamma}{\rho}}\right)^{\frac{1}{1-\gamma}}\right]^{1-\gamma}.\end{aligned}\tag{19}$$

Hence we can focus on admissible controls ν with $\zeta\left(\nu\right)$ implied by (19), and on planner's indirect utility flow $h^0\left(\nu,\theta\right)$. The structure of the problem implies that the optimal Markov control of the planner is of the form $\nu_t^n=\nu^n\left(\theta_t\right)$, $n\in\{1,2\}$. Throughout the proof, I impose the following restriction on the underlying discount rate processes ν^n.

Assumption A.4 *The discount rates $\nu_t^n=\nu^n\left(\theta_t\right)$, $n\in\{1,2\}$ are bounded functions of θ on $[0,1]$ that are Lipschitz continuous, and there $\exists\varepsilon>0$ such that*

$$\frac{\gamma-\rho\nu_t^n/\beta}{\gamma-\rho}>\varepsilon.$$

Later I verify that this assumption holds for the optimal discount rate process on every interval $[\varepsilon,1-\varepsilon]$. I postpone the verification of this restriction at the boundaries as $\varepsilon\searrow0$ to Appendix B where I characterize the boundary behavior of the economy in more detail and explicitly calculate the limit. The results in Appendix B show that the bounds imposed in Assumption A.4 correspond to the assumption that agents' wealth-consumption ratios are bounded and bounded away from zero. The subsequent characterization of the optimal control implies that the optimal policy necessarily satisfies these assumptions.

Lemma A.5 *If a is admissible and satisfies Assumption A.4, then*

$$E_t\left[\int_t^{\infty}\frac{\bar{\lambda}_s^n}{\bar{\lambda}_t^n}\left(\frac{Y_s}{Y_t}\right)^{\gamma}ds\right]<+\infty,\qquad n\in\{1,2\}\tag{20}$$

and

$$\lim_{\tau\to\infty}E_t\left[\frac{\bar{\lambda}_\tau^n}{\bar{\lambda}_t^n}\left(\frac{Y_\tau}{Y_t}\right)^{\gamma}\right]=0,\qquad n\in\{1,2\}\,.\tag{21}$$

Proof. See Section OA.7 in Online Appendix. ■

The following lemma characterizes the limits of $\widetilde{J}\left(\theta_t\right)$ at the boundaries.

Lemma A.6 *The planner's value function $J\left(\bar{\lambda}_t, Y_t\right)$ is continuously extended at the boundaries as $\bar{\lambda}_t^1 \searrow 0$ or $\bar{\lambda}_t^2 \searrow 0$ by the continuation values from the homogeneous agent economies. E.g., for $\bar{\lambda}_t^2 > 0$,*

$$J\left(0, \bar{\lambda}_t^2, Y_t\right) \doteq \lim_{\bar{\lambda}_t^1 \searrow 0} J\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right) = \bar{\lambda}_t^2 V_t^2\left(Y\right). \tag{22}$$

Further, the optimal choice of consumption $C_u^1\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right)$ for agent 1 and time $u \geq t$ satisfies

$$\lim_{\bar{\lambda}_t^1 \searrow 0} C_u^1\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right) = 0 \qquad P\text{-a.s.} \tag{23}$$

The case $\bar{\lambda}_t^2 \searrow 0$ is symmetric.

Proof. See Section OA.7 in Online Appendix. ■

A vanishing Pareto weight on agent 1 thus leads to a vanishing consumption level (23) for every given time $u \geq t$. However, convergence in (23) is not uniform in u. Importantly, this argument therefore does not prevent the possibility that for a given arbitrarily small Pareto weight, agent's consumption recovers in the future. A direct consequence of result (22) is

$$\lim_{\theta \searrow 0} \widetilde{J}\left(\theta\right) = \bar{V}^1 \qquad \lim_{\theta \nearrow 1} \widetilde{J}\left(\theta\right) = \bar{V}^2. \tag{24}$$

A.2 The Hamilton–Jacobi–Bellman equation

Denoting $\bar{\lambda} = \left(\bar{\lambda}^1, \bar{\lambda}^2\right)'$ and $u = \left(u^1, u^2\right)'$, the state vector is $Z = \left(\bar{\lambda}', Y\right)'$. This suggests that the planner's problem (8) leads to the Hamilton–Jacobi–Bellman equation for $J\left(\bar{\lambda}, Y\right)$,

$$0 = \sup_{a \in \mathcal{A}} \sum_{n=1}^{2} \bar{\lambda}^n \left[F\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n} \nu^n\right] + J_y \mu_y Y + \frac{1}{2} \mathrm{tr}\left(J_{zz} \Sigma\right), \tag{25}$$

where

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

and $\mathrm{diag}\left(\bar{\lambda}\right)$ is a 2×2 diagonal matrix with elements of $\bar{\lambda}$ on the main diagonal. Using the conjecture $J\left(\bar{\lambda}, Y\right) = \left(\bar{\lambda}^1 + \bar{\lambda}^2\right) Y^\gamma \widetilde{J}\left(\theta\right)$ reduces the problem to the ordinary differential equation for $\widetilde{J}\left(\theta\right)$ given by (9) with boundary conditions $\widetilde{J}\left(0\right) = \bar{V}^2$ and $\widetilde{J}\left(1\right) = \bar{V}^1$, as determined by Lemma A.6. Further define

$$\begin{aligned} h^1\left(\nu, \theta\right) &\doteq -\theta \nu^1 - \left(1-\theta\right) \nu^2 + \left(\theta u^1 + \left(1-\theta\right) u^2\right) \gamma \sigma_y + \gamma \mu_y + \frac{1}{2} \gamma^2 \sigma_y^2 \\ h^2\left(\nu, \theta\right) &\doteq \theta\left(1-\theta\right)\left[\nu^2 - \nu^1 + \left(u^1 - u^2\right) \gamma \sigma_y\right] \\ h^3\left(\theta\right) &\doteq \frac{1}{2} \theta^2 \left(1-\theta\right)^2 \left(u^1 - u^2\right)^2 \end{aligned}$$

Together with $h^0\left(\nu, \theta\right)$ from (19), the HJB equation (9) can be written as

$$0 = \sup_{\nu} h^0\left(\nu, \theta\right) + h^1\left(\nu, \theta\right) \widetilde{J}\left(\theta\right) + h^2\left(\nu, \theta\right) \widetilde{J}_\theta\left(\theta\right) + h^3\left(\theta\right) \widetilde{J}_{\theta\theta}\left(\theta\right)$$

with boundary conditions $\widetilde{J}\left(0\right) = \bar{V}^2$ and $\widetilde{J}\left(1\right) = \bar{V}^1$. Under Assumption A.4, all functions h^j are bounded and Lipschitz.

The goal is to show that there exists a unique twice continuously differentiable solution of this equation that corresponds to the value function. Once the solution of the HJB equation is characterized, we prove that it corresponds to the value function. In order to do that, the stochastic process for θ_t needs to be well-defined. An application of Itô's lemma to $\theta_t = \bar{\lambda}_t^1 / \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right)$ yields

$$\begin{aligned} d\theta_t = {} & \theta_t \left(1-\theta_t\right) \left[\nu_t^2 - \nu_t^1 + \left(\theta_t u^1 + \left(1-\theta_t\right) u^2\right) \left(u^2 - u^1\right)\right] dt + \\ & + \theta_t \left(1-\theta_t\right) \left(u^1 - u^2\right) dW_t. \end{aligned} \tag{26}$$

Lemma A.7 *Under Assumption A.4, the stochastic differential equation (26) has a unique strong solution.*

Proof. Under Assumption A.4, the drift and volatility coefficients in (26) are Lipschitz and bounded, so that a unique strong solution exists (see, e.g., Pham (2009), Theorem 1.3.15). ■

A.3 An auxiliary problem

Consider the following auxiliary planner's problem with suboptimal control. Fix $\varepsilon \in \left(0, \frac{1}{2}\right)$. When $\theta_t = \bar{\lambda}_t^1 / \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) \in (\varepsilon, 1-\varepsilon)$, the planner exercises local control optimally, given her value function. When θ_t hits the boundary ε (a symmetric argument holds for $1-\varepsilon$), the planner is restricted to fix consumption shares of the two agents to $\bar{\zeta}^n(\varepsilon)$ that are the optimal static consumption shares from the proof of Lemma A.3 and keep them fixed forever. Formally, the auxiliary problem for a fixed ε and $\theta_t \in [\varepsilon, 1-\varepsilon]$ is given by

$$J^{\varepsilon}\left(\bar{\lambda}_t, Y_t\right) = \sup_{a \in \mathcal{A}} E_t \left[\sum_{n=1}^{2} \int_t^{\tau^{\varepsilon}} \bar{\lambda}_s^n F\left(C_s^n, \nu_s^n\right) ds + J^{\varepsilon}\left(\bar{\lambda}_{\tau^{\varepsilon}}, Y_{\tau^{\varepsilon}}\right)\right] \tag{27}$$

where τ^{ε} is a stopping time given by $\tau^{\varepsilon} = \inf\left\{s \geq t : \theta_s \notin (\varepsilon, 1-\varepsilon)\right\}$ and the continuation value at the stopping time is established in Lemma A.3 as $J^{\varepsilon}\left(\bar{\lambda}_{\tau^{\varepsilon}}, Y_{\tau^{\varepsilon}}\right) = \left(\bar{\lambda}_{\tau^{\varepsilon}}^1 + \bar{\lambda}_{\tau^{\varepsilon}}^2\right) Y_{\tau^{\varepsilon}}^{\gamma} \bar{J}\left(\theta_{\tau^{\varepsilon}}\right)$ with $\theta_{\tau^{\varepsilon}} \in \{\varepsilon, 1-\varepsilon\}$. The value function of the auxiliary problem also satisfies the homotheticity property and

$$J^{\varepsilon}\left(\bar{\lambda}_t, Y_t\right) = \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^{\gamma} \widetilde{J}^{\varepsilon}\left(\theta_t\right), \qquad \theta_t \in (\varepsilon, 1-\varepsilon)$$

where $\widetilde{J}^{\varepsilon}\left(\theta_t\right)$ is analogous to $\widetilde{J}\left(\theta_t\right)$ from Lemma A.3. The solution can be continuously extended for $\theta \in [0,1] \setminus [\varepsilon, 1-\varepsilon]$ using $\widetilde{J}^{\varepsilon}(\theta) = \bar{J}(\theta)$.

Lemma A.8 *The Hamilton–Jacobi–Bellman equation for the auxiliary problem*

$$0 = \sup_{\nu} h^0(\nu, \theta) + h^1(\nu, \theta) \widetilde{J}^{\varepsilon}(\theta) + h^2(\nu, \theta) \widetilde{J}_{\theta}^{\varepsilon}(\theta) + h^3(\theta) \widetilde{J}_{\theta\theta}^{\varepsilon}(\theta) \tag{28}$$

with boundary conditions $\widetilde{J}^{\varepsilon}(\varepsilon) = \bar{J}(\varepsilon)$ and $\widetilde{J}^{\varepsilon}(1-\varepsilon) = \bar{J}(1-\varepsilon)$ has a unique twice continuously differentiable solution on $[\varepsilon, 1-\varepsilon]$.

Proof. See Section OA.7 in Online Appendix. ■

A.4 Solution to the Hamilton–Jacobi–Bellman equation

We want co characterize the limiting solution of a sequence of the auxiliary problems for $\left\{\varepsilon^k\right\}_{k=1}^{\infty}$ as $\varepsilon^k \searrow 0$. First notice that the boundary condition $\bar{J}\left(\varepsilon^k\right)$ converges to $\bar{V}^2$ (and $\bar{J}\left(1-\varepsilon^k\right)$ to $\bar{V}^1$) as $\varepsilon^k \searrow 0$, i.e., to the limiting points of the value function given by (24). We want to establish convergence of the sequence $\widetilde{J}^{\varepsilon^k}(\theta)$ to $\widetilde{J}(\theta)$ on every interval $[\varepsilon, 1-\varepsilon]$, $\varepsilon > 0$.

The difficulty is the vanishing coefficient $h^3(\theta)$ as θ approaches 0 or 1. While the coefficients $k^j(\nu,\theta)$ are bounded for every fixed $\left[\varepsilon^k, 1-\varepsilon^k\right]$, they are not uniformly bounded across all such intervals as $\varepsilon^k \searrow 0$. However, there is a suitable transformation of variables. Define $\vartheta(\theta) = \log(\theta/(1-\theta)) \in (-\infty, +\infty)$ as in (10) and $\widehat{J}(\vartheta(\theta)) \doteq \widetilde{J}(\theta)$. The HJB equation (9) can be written as a differential equation for $\widehat{J}(\vartheta)$ given by

$$0 = \sup_{\nu} \hat{h}^0(\nu,\vartheta) + \hat{h}^1(\nu,\vartheta)\widehat{J}(\vartheta) + \hat{h}^2(\nu,\vartheta)\widehat{J}_{\vartheta}(\vartheta) + \hat{h}^3(\vartheta)\widehat{J}_{\vartheta\vartheta}(\vartheta) \tag{29}$$

with $\theta(\vartheta) = \exp(\vartheta)/(1+\exp(\vartheta))$ and

$$\begin{aligned}
\hat{h}^0(\nu,\vartheta) &\doteq h^0(\nu,\theta(\vartheta)) \\
\hat{h}^1(\nu,\vartheta) &\doteq -\theta(\vartheta)\nu^1 - (1-\theta(\vartheta))\nu^2 + \left(\theta(\vartheta)u^1 + (1-\theta(\vartheta))u^2\right)\gamma\sigma_y + \gamma\mu_y + \frac{1}{2}\gamma^2\sigma_y^2 \\
\hat{h}^2(\nu,\vartheta) &\doteq \nu^2 - \nu^1 + \left(u^1-u^2\right)\gamma\sigma_y + \frac{1}{2}\left(u^1-u^2\right)^2(2\theta(\vartheta)-1) \\
\hat{h}^3(\vartheta) &\doteq \frac{1}{2}\left(u^1-u^2\right)^2.
\end{aligned}$$

The boundary conditions for the problem are given by $\lim_{\vartheta\to-\infty}\widehat{J}(\vartheta) = \bar{V}^2$ and $\lim_{\vartheta\to+\infty}\widehat{J}(\vartheta) = \bar{V}^1$. Under Assumption A.4, the coefficients $\hat{k}^j(\nu,\vartheta) = k^j(\nu,\vartheta)/k^3(\vartheta)$ for $j = 0,1,2$ are bounded for $\vartheta \in (-\infty,\infty)$.

The HJB equation (29) thus satisfies conditions of the proof in Strulovici and Szydlowski (2014), Appendix B.3, that extends the solution of the HJB equation on a sequence of bounded domains to an unbounded limit. Rather than repeating the proof here, I note that the structure of (29), in particular the bounded coefficients $\hat{k}^j$, implies that Lemma 8 in Strulovici and Szydlowski (2014) is satisfied, for instance, with the function $\phi(z) = Kz$ for K sufficiently large, and that the functions $\widehat{J}$, $\widehat{J}_{\vartheta}$, $\widehat{J}_{\vartheta\vartheta}$ are bounded. An application of the Arzelà–Ascoli theorem implies that there is a uniformly convergent subsequence of solutions $\widehat{J}^{\varepsilon^\kappa}$ on the interval $[\varepsilon, 1-\varepsilon]$ for every $\varepsilon > 0$.

Strulovici and Szydlowski (2014) then use the following diagonalization argument. Start with interval $\left[\varepsilon^1, 1-\varepsilon^1\right]$. Find the uniformly convergent subsequence of $\widehat{J}^{\varepsilon^k}(\vartheta)$ on $\left[\varepsilon^1, 1-\varepsilon^1\right]$ and denote its limit $w^1(\vartheta)$. Now take interval $\left[\varepsilon^2, 1-\varepsilon^2\right]$ and find a subsequence of the first subsequence that converges on $\left[\varepsilon^2, 1-\varepsilon^2\right]$. Denote the solution $w^2(\vartheta)$ and notice that $w^1(\vartheta) = w^2(\vartheta)$ for $\vartheta \in \left[\varepsilon^1, 1-\varepsilon^1\right]$. Continue iteratively and define the limiting solution as follows: for $\vartheta \in \left[\varepsilon^k, 1-\varepsilon^k\right] \setminus \left[\varepsilon^{k-1}, 1-\varepsilon^{k-1}\right]$, set $\widehat{J}(\vartheta) \doteq w^k(\vartheta)$.

Corollary A.9 *The limiting solution $\widetilde{J}(\theta) = \widehat{J}(\vartheta(\theta))$ constructed in this way exists, is twice continuously differentiable and uniquely solves the Hamilton–Jacobi–Bellman equation (9).*

The following lemma is useful as a clarification for the intuition for why the solution of the HJB equation (9) can be defined through the limit of solutions on closed subintervals.

Lemma A.10 *Let $\left\{\varepsilon^k\right\}_{k=1}^{\infty}$ satisfy $\varepsilon^k \searrow 0$. Under Assumption A.4, the sequence of stopping times $\left\{\tau^{\varepsilon^k}\right\}_{k=1}^{\infty}$ in the auxiliary problem (27) is almost surely diverging, $P\left(\tau^{\varepsilon^k} \overset{k\to\infty}{\longrightarrow} +\infty\right) = 1$.*

Proof. We use the transformation $\vartheta(\theta) = \log\frac{\theta}{1-\theta}$. In the state space represented by ϑ, the sequence of stopping times $\left\{\tau^k\right\}_{k=1}^{\infty}$ corresponds to a sequence of first crossing times of thresholds $\pm\bar{\vartheta}^k$ as $\bar{\vartheta}^k \nearrow +\infty$. Since $\vartheta(\theta)$ follows (10), it is an Itô process with bounded coefficients, for which the claim of the lemma is a standard result. ■

As we move the boundary ε^k in the auxiliary problem closer to zero, the crossing time of this boundary diverges to infinity. With discounting (under Assumption A.1) and under a uniform bound on the boundary values, their contribution to the value function for a given initial value θ_0 vanishes as $\varepsilon^k \searrow 0$.

While boundedness and the Lipschitz property stated in Assumption A.4 hold for $\nu^n(\theta)$ for every given auxiliary problem (for every fixed ε^k), they may not hold in the limit as $\varepsilon^k \searrow 0$. I prove that Assumption A.4 holds also in the limit in Appendix B, by obtaining closed form solutions for these limits.

A.5 Verification theorem

The last step is to verify that the solution of the Hamilton–Jacobi–Bellman equation yields the value function. This is a standard verification argument.

Lemma A.11 *The function $J\left(\bar{\lambda}_t, Y_t\right) = \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^{\gamma} \widetilde{J}(\theta_t)$, where $\widetilde{J}(\theta)$ is the solution of the Hamilton–Jacobi–Bellman equation (9), coincides with the value function (8).*

Proof. See Section OA.7 in Online Appendix. ■

B Characterization of the boundary behavior

This section contains proofs of propositions that characterize the boundary behavior of the economy as $\theta \to \{0,1\}$. I start with Proposition 3.2, then move ahead to prove propositions from Section 5 and finally return to prove the remaining statements.

Proof of Proposition 3.2. Given an initial condition $\theta_0 \in (0,1)$, the process θ given by (26) 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 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 in order to make the process positive Harris recurrent (see Meyn and Tweedie (1993)). Since the process is also 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 (10). The boundary behavior of the process θ is captured by the scale function $S : (0,1)^2 \to \mathbb{R}$ defined as

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

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

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

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

The boundaries are nonattracting if and only if

$$\lim_{\theta_l \searrow 0} S[\theta_l, \theta_h] = \infty \quad \text{and} \quad \lim_{\theta_h \nearrow 1} S[\theta_l, \theta_h] = \infty, \tag{30}$$

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

$$\lim_{\theta_l \searrow 0} M\left[\theta_l, \theta_h\right] < \infty \quad \text{and} \quad \lim_{\theta_h \nearrow 1} M\left[\theta_l, \theta_h\right] < \infty, \tag{31}$$

again independently of the argument that is not under the limit. In our case,

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

where

$$s_{sep}\left(\theta\right) = \left(\frac{1-\theta}{1-\theta_0}\right)^{-\frac{2u^1}{u^1-u^2}} \left(\frac{\theta}{\theta_0}\right)^{\frac{2u^2}{u^1-u^2}} \tag{32}$$

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

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 $\nu^2\left(\theta\right) - \nu^1\left(\theta\right) \geq \underline{\nu}$ for all $\theta \in (0, \underline{\theta})$. Taking $\theta_0 = \underline{\theta}$, the scale function can be bounded as

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

The left limit in (30) 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 $\nu^2\left(\theta\right) - \nu^1\left(\theta\right) \leq \bar{\nu}$ 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 (31) are marginally tighter. Following the same bounding argument as above, sufficient conditions for (31) to hold are

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

The construction also reveals that these bounds are the least tight bounds of this type under which the proposition holds. It is useful to note that the unique stationary density $q\left(\theta\right)$ is proportional to the speed density $m\left(\theta\right)$. Finally, if the limits in Proposition 3.2 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 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 Proposition 5.1. The proof relies on showing that the dynamics of the continuation values

V_t^n in the proximity of the boundaries become degenerate in a specific sense. From this fact, I can infer the behavior of the stochastic discount factor implied by the consumption process of the large agent and, consequently, the equilibrium price dynamics.

Without loss of generality, consider the boundary as $\theta \searrow 0$. Using the construction from Duffie and Epstein (1992a), the stochastic discount factor process for the 'large' agent 2 under the subjective probability measure Q^2 is given by

$$S_t^2 = \exp\left(-\int_0^t \nu^2(\theta_s)\,ds\right)\left(\frac{Y_t}{Y_0}\right)^{\gamma-1}\left(\frac{\zeta^2(\theta_t)}{\zeta^2(\theta_0)}\right)^{\rho-1}\left(\frac{\widetilde{J}^2(\theta_t)}{\widetilde{J}^2(\theta_0)}\right)^{1-\frac{\rho}{\gamma}}. \tag{33}$$

The proof of Lemma A.6 shows that $\lim_{\theta\searrow 0}\zeta^2(\theta) = 1$ and $\lim_{\theta\searrow 0}\widetilde{J}^2(\theta) = \bar{V}^2$, implying that $\lim_{\theta\searrow 0}\nu^2(\theta) = \bar{\nu}^2$, with $\bar{V}^2$ and $\bar{\nu}^2$ given in the proof of Lemma A.3. Homotheticity of preferences implies that individual wealth-consumption ratios are given by

$$\xi^n(\theta) = \frac{1}{\beta}\left(\frac{\gamma\widetilde{J}^n(\theta)}{\zeta^n(\theta)^\gamma}\right)^{\frac{\rho}{\gamma}} \qquad n\in\{1,2\} \tag{34}$$

and since $\lim_{\theta\searrow 0}\xi^2(\theta) = \lim_{\theta\searrow 0}\xi(\theta)$, we obtain the limiting behavior of the aggregate wealth-consumption ratio using the above limits for $\zeta^2(\theta)$ and $\widetilde{J}^2(\theta)$.

It remains to be shown that the local drift and volatility in the local evolution of the last two terms of (33) decline to zero as $\theta \searrow 0$, so that

$$d\log S_t^2 \doteq \mu_{S^2}(\theta_t)\,dt + \sigma_{S^2}(\theta_t)\,dW_t = \left[\mu_{S^2}(\theta_t) + u^2\sigma_{S^2}(\theta_t)\right]dt + \sigma_{S^2}(\theta_t)\,dW_t^n$$

with

$$\lim_{\theta\searrow 0}\mu_{S^2}(\theta) = -\bar{\nu}^2 + (\gamma-1)\,\mu_y, \qquad \lim_{\theta\searrow 0}\sigma_{S^2}(\theta) = (\gamma-1)\,\sigma_y. \tag{35}$$

The results for the limiting behavior of the risk-free rate and return on the aggregate wealth then follow directly from the local behavior of the stochastic discount factor in (35). Details of the calculations proving (35) are provided in Section OA.8 of the Online Appendix. ■

Proof of Proposition 5.2. The evolution of θ given by equation (26) implies that for every fixed $t \geq 0$

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

and thus also $\zeta^2(\theta_t) \to 1$ and $\widetilde{J}^2(\theta_t) \to \bar{V}^2$, P-a.s.[18] The last two terms in the expression for the stochastic discount factor S_t^2, equation (33), converge to one, P-a.s., and since $\nu^2(\theta_s)\,,\ 0 \leq s \leq t$ 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 $M_t^2 S_t^2(\theta_0)$ indexed by the initial Pareto share θ_0. Since this family is uniformly integrable, then convergence in probability implies convergence in mean, and we obtain the convergence result for bond prices

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

[18]This result becomes more transparent if we consider ζ^2 and $\widetilde{J}^2$ as functions of $\vartheta_t = \vartheta(\theta_t)$. The dynamics of ϑ_t in (10) have bounded drift and volatility coefficients and thus for $\forall\varepsilon > 0$, $\forall k > 0$, it is possible to achieve

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

by setting $\log\theta_0$ sufficiently low.

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

Proof of Proposition 5.3. Agent 1, whose wealth A^1 is close to zero, solves

$$\bar{\lambda}_t^1 V_t^1 = \max_{(C^1,\pi^1,\nu^1)} E_t \left[\int_t^\infty \bar{\lambda}_s^1 F\left(C_s^1, \nu_s^1\right) ds \right] \tag{36}$$

subject to (7) and the budget constraint,

$$\begin{aligned} \frac{dA_t^1}{A_t^1} &= \left[r(\theta_t) + \pi_t^1 \left([\xi(\theta_t)]^{-1} + \mu_A(\theta_t) + \frac{1}{2} [\sigma_A(\theta_t)]^2 - r(\theta_t) \right) - \frac{C_t^1}{A_t^1} \right] dt + \pi_t^1 \sigma_A(\theta_t) dW_t = \\ &= \mu_{A^1}(\theta_t) dt + \sigma_{A^1}(\theta_t) dW_t \end{aligned} \tag{37}$$

where π^1 is the portfolio share invested in the risky asset. The solution of this equation determines the consumption-wealth ratio of agent 1 and, consequently, the evolution of her wealth. The local behavior of returns on the risk-free bond $r(\theta)$ and risky asset (14) as $\theta \searrow 0$ is known from Proposition 5.1.

The proof of this proposition is based on showing that the limit of the optimal consumption, portfolio and discount rate choice $\left(C^1, \pi^1, \nu^1\right)$ as $\theta \searrow 0$ is equivalent to taking the limit for equilibrium returns $(r, \xi, \mu_A, \sigma_A)$ as $\theta \searrow 0$ first, and then computing the optimal individual choice of agent 1 under the limiting returns. Details of this proof are provided in Section OA.8 of the Online Appendix. ■

It remains for me to verify that the boundedness assumption for wealth-consumption ratios indeed holds.

Corollary B.1 *Under parameter restrictions in Assumption A.1, the wealth-consumption ratios are bounded and bounded away from zero.*

Proof of Corollary B.1. 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 (17) in Assumption A.1. The consumption-wealth ratio of the small agent is derived in the proof of Proposition 5.3. Restriction (18) in Assumption A.1 assures that this quantity is strictly positive, and the wealth-consumption ratio finite. This confirms that the discount rate restrictions in Assumption A.4 hold as well. ■

Proof of Corollary 3.3. Utilize results in Proposition 5.3 and the fact that $\lim_{\theta \searrow 0} \mu_{A^2}(\theta) = \mu_y$ and $\lim_{\theta \searrow 0} \sigma_{A^2}(\theta) = \sigma_y$, then form the differences in the limiting expected logarithmic growth rates, and compare them to inequalities in Proposition 3.2. ■

Proof of Proposition 3.4. The difference in expected logarithmic returns is obtained by computing the limiting behavior of

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

utilizing the results for $\theta \searrow 0$ from Propositions 5.1 and 5.3. The first term above is the difference in the risk premium associated with the two portfolios, and the second term is the volatility penalty. The same propositions also contain the results for the consumption-wealth ratios of the two agents. ■

Proof of Corollary 4.1. The results are obtained by taking limits of expressions in Proposition 3.4. ■

Proof of Corollary 4.2. Assume without loss of generality that $\left|u^2\right| \leq \left|u^1\right|$. The sufficient part is an immediate consequence of Proposition 3.2. 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 is constant and both agents survive under P. When $-u^2 = u^1 = u$, then it follows from inspection of formula (32) in the proof of Proposition 3.2 that conditions (30) 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 completely describes the dynamics of the economy, substituting this expression into (10) and reorganizing yields the desired result. ■

Proof of Proposition 6.1. Since θ is an Itô diffusion with bounded coefficients and volatility that is bounded away from zero on $(\varepsilon, 1-\varepsilon)$, $\forall \varepsilon > 0$, the stationary density, if it exists, is unique and has a full support. Hence every point in $(0,1)$ is visited with probability one, given any initial condition $\theta_0 \in (0,1)$. ■

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.

Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims, 2013. "Uncertainty and Economic Activity: Evidence from Business Survey Data." *American Economic Journal: Macroeconomics*. 5 (2): 217–249.

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

Bansal, Ravi, and Amir Yaron, 2004. "Risk for the long run: A potential resolution of asset pricing puzzles." *The Journal of Finance*. 59 (4): 1481–1509.

Barillas, Francisco, and Kristoffer P. Nimark, 2016. "Speculation and the Term Structure of Interest Rates." Forthcoming in Review of Financial Studies.

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

Beker, Pablo F., and Chattopadhyay, 2010. "Consumption Dynamics in General Equilibrium: A Characterization When Markets Are Incomplete." *Journal of Economic Theory*. 145 (6): 2133–2185.

Beker, Pablo F., and Emilio Espino, 2011. "The Dynamics of Efficient Asset Trading with Heterogeneous Beliefs." *Journal of Economic Theory*. 146 (1): 189–229.

Benhabib, Jess, and Alberto Bisin, 2016. "Skewed Wealth Distributions: Theory and Empirics." Working paper, New York University.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu, 2011. "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Many Agents." *Econometrica*. 79 (1): 123–157.

Bhamra, Harjoat S., and Raman Uppal, 2013. "Asset Prices with Heterogeneity in Preferences and Beliefs." Working paper.

Bhandari, Anmol, 2015. "Doubts, Asymmetries, and Insurance." Job market 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.

Breiman, Leo, 1960. "Investment Policies for Expanding Businesses Optimal in a Long-Run Sense." *Naval Research Logistics Quarterly.* 7 (4): 647–651.

Breiman, Leo, 1961. "Optimal Gambling Systems for Favorable Games." in *Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability*pp. 63–68. University of California Press.

Brunnermeier, Markus K., Alp Simsek, and Wei Xiong, 2014. "A Welfare Criterion for Models with Distorted Beliefs." *Quarterly Journal of Economics.* 129 (4): 1753–1797.

Buraschi, Andrea, and Alexei Jiltsov, 2006. "Model Uncertainty and Option Markets with Heterogeneous Beliefs." *The Journal of Finance.* 61 (6): 2841–2897.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2009. "Fight or Flight? Portfolio Rebalancing by Individual Investors." *Quarterly Journal of Economics.* 124 (1): 301–348.

Cao, Dan, 2013a. "Speculation and Financial Wealth Distribution under Belief Heterogeneity." .

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.

Dindo, Pietro, 2015. "Survival in Speculative Markets." Working paper, Scuola Superiore Sant'Anna.

Duffie, Darrell, and Larry G. Epstein, 1992a. "Asset Pricing with Stochastic Differential Utility." *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 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.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, 2016. "Heterogeneity and Persistence in Returns to Wealth." .

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

Feller, William, 1952. "The Parabolic Differential Equations and the Associated Semi-Groups of Transformations." *Annals of Mathematics, Second Series*. 55 (3): 468–519.

Feller, William, 1957. "On Boundaries and Lateral Conditions for the Kolmogorov Differential Equations." *Annals of Mathematics, Second Series*. 65 (3): 527–570.

Fleming, Wendell H., and Halil Mete Soner, 2006. *Controlled Markov Processes and Viscosity Solutions*, Springer, 2nd edn.

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

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, 2015. "The Dynamic of Inequality." NBER Working paper W21363.

Gârleanu, Nicolae, and Stavros Panageas, 2015. "Young, Old, Conservative and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing." *Journal of Political Economy*. 123 (3): 670–685.

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

Giacoletti, Marco, Kristoffer T. Laursen, and Kenneth J. Singleton, 2015. "Learning, Dispersion of Beliefs, and Risk Premiums in an Arbitrage-free Term Structure Model." .

Guerdjikova, Ani, and Emanuela Sciubba, 2015. "Survival with Ambiguity." *Journal of Economic Theory*. 155: 50–94.

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

Hansen, Lars Peter, and Thomas J. Sargent, 2001. "Acknowledging Misspecification in Macroeconomic Theory." *Review of Economic Dynamics*. 4 (3): 519–535.

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.

Kacperczyk, Marcin, Jaromir Nosal, and Luminita Stevens, 2015. "Investor Sophistication and Capital Income Inequality." .

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.

Kelly, John L., 1956. "A New Interpretation of Information Rate." *IRE Transactions on Information Theory*. 2 (3): 185–189.

Kendall, Chad, and Ryan Oprea, 2016. "Are Biased Beliefs Fit to Survive? An Experimental Test of the Market Selection Hypothesis." .

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

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield, 2011. "Market Selection." Working paper.

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

Kopczuk, Wojciech, 2015. "What Do We Know about the Evolution of Top Wealth Shares in the United States?." *Journal of Economic Perspectives*. 29 (1): 47–66.

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.

Lusardi, Annamaria, Olivia S. Mitchell, and Pierre-Carl Michaud, 2016. "Optimal Financial Knowledge and Wealth Inequality." Forthcoming in Journal of Political Economy.

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

Malmendier, Ulrike, and Stefan Nagel, 2016. "Learning from Inflation Experiences." *Quarterly Journal of Economics*. 131 (1): 53–87.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers, 2003. "Disagreement about Inflation Expectations." *NBER Macroeconomics Annual*. 18: 209–248.

Markowitz, Harry, 1952. "Portfolio Selection." *Journal of Finance*. 7 (1): 77–91.

Massari, Filippo, 2016. "Trading in the Market: A Necessary and Sufficient Condition for a Trader to Vanish." Working paper, University of New South Wales.

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.

Pham, Huyên, 2009. *Continuous-time Stochastic Control and Optimization with Financial Applications*, Springer.

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

Saez, Emmanuel, and Gabriel Zucman, 2016. “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data.” *Quarterly Journal of Economics*. 131 (2): 519–578.

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

Scheinkman, José A., and Wei Xiong, 2003. “Overconfidence and Speculative Bubbles.” *Journal of Political Economy*. 111 (6): 1183–1219.

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.

Strulovici, Bruno, and Martin Szydlowski, 2014. “On the Smoothness of Value Functions and the Existence of Optimal Strategies.” .

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.

Xiong, Wei, 2013. “Bubble, Crises, and Heterogeneous Beliefs.” in *Handbook on Systemic Risk*chap. 24, pp. 663–713. Cambridge University Press.

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