Survival and stationary equilibria with heterogeneous beliefs under recursive preferences*

Jaroslav Borovička

University of Chicago

borovicka@uchicago.edu

July 15, 2010

Abstract

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

*I am indebted to Lars Peter Hansen for his advice and continuous support. I appreciate helpful comments from Fernando Alvarez, Katarína Borovičková, Alan Moreira, Harald Uhlig, Mark Westerfield, and the participants of the Economic dynamics working group at the University of Chicago. All typos and errors are mine.

1 Introduction

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

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

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

Sandroni (2000) and Blume and Easley (2006) show that agents with incorrect beliefs also vanish in economies where aggregate endowment is bounded from above and away from zero, with preferences satisfying much milder conditions. This result is again based on the same idea. The boundedness of aggregate endowment limits the differences in local properties of the utility functions that can arise between rich and poor agents, and belief distortions thus cannot be overcome by the local characteristics of preferences.

This paper analyzes the robustness of these results with respect to departures from time-separability of preferences. In order to focus solely on the impact of belief heterogeneity, I endow agents with identical homothetic recursive preferences of the Duffie-Epstein-Zin type. These preferences can be viewed as a generalization of the CRRA utility that allows to disentangle risk aversion and IES. Homotheticity assures that survival results are not driven

by exogenous differences in the local properties of the utility functions.

The decoupling of risk aversion and IES proves to have a crucial impact. Consider the case of an agent who is optimistic about the growth rate of aggregate endowment stream, and an asset which pays the aggregate endowment stream as dividend. An increase in risk aversion limits the willingness of the agent to take excessively large levered positions in the asset vis-à-vis its perceived underpricing. Leverage increases optimism about the growth rate of the portfolio, which in turns increases consumption today in line with the permanent income model. A higher risk aversion and less leverage therefore imply more frugal behavior of the optimistic agent.

Similarly, an increase in IES increases the willingness of the agent who is optimistic about the growth rate of his portfolio to tilt the consumption profile toward the future, thus increasing his saving rate.

When preferences are time-separable, an increase in the risk aversion is coupled with a decrease in IES. Promoting more frugal behavior of an optimistic agent by limiting the leveraged positions through an increase in risk aversion is always accompanied by a decrease in IES, implying less incentives for tilting the consumption profile toward the future and saving his way out of extinction. These two effects always offset each other as we change the coefficient of relative risk aversion, which leads to the extinction results described above. Similar results also hold for a pessimistic agent who becomes optimistic about the growth rate of his portfolio once he takes a short position in the risky asset.

When the two effects can be disentagled, I find that increasing the risk aversion independently of IES (or, equivalently, increasing the IES independently of the risk aversion coefficient) improves the conditions for the long-run survival of agents with incorrect beliefs.

More precisely, I study a two-agent continuous-time endowment economy with complete markets. Aggregate endowment is modeled as a geometric Brownian motion, and agents have common homothetic recursive Duffie-Epstein-Zin preferences and differ only in their beliefs. I find that both agents survive and a nondegenerate stationary distribution exists for wide regions of the parameter space when risk aversion is larger than the inverse of the IES. Moreover, the rate of convergence to the stationary distribution is exponential. In the opposite case, only one agent survives but each agent may survive with a strictly positive probability. The results hold both for optimistic as well as pessimistic deviations from rationality, although quantitative outcomes in terms of average long-run wealth or consumption shares captured by the agent with beliefs distorted toward optimism or pessimism differ substantially.

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

1.1 Methodology and literature overview

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

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

With unbounded aggregate endowment, intertemporal elasticity of substitution becomes an additional survival factor, as shown in Yan (2008). In a growing economy, a low IES is a disadvantage, since it diminishes the willingness of an agent to tilt the consumption profile toward the future. In Yan's model, agents have homothetic CRRA utility, so once preferences are identical across agents, only belief disparities determine survival.

Kogan, Ross, Wang, and Westerfield (2009) show that imposing identical preferences across agents is not sufficient for belief disparities to be the sole determinant of survival. Even if agents are given identical utility functions, these functions may have very different local properties conditional on the level of consumption or wealth. When aggregate endowment is unbounded, these differences may outweigh the divergence in beliefs, and lead to survival of agents with relatively more incorrect beliefs. A sufficient condition to prevent this outcome

is the boundedness of the relative risk aversion coefficient. This condition can be interpreted as a bound on deviations of the utility function from homotheticity.

This paper shows that support in favor of the market survival hypothesis weakens considerably once the assumption of time-separability in preferences is relaxed. I focus on the recursive preferences axiomatized by Kreps and Porteus (1978), and developed by Epstein and Zin (1989) and Weil (1990) in discrete time, and by Duffie and Epstein (1992b) in continuous time. These preferences are homothetic, allow to disentagle the risk aversion with respect to intratemporal gambles from the intertemporal elasticity of substitution, and include the CRRA utility as a special case. Thanks to the additional degree of flexibility, this class of preferences is widely used in the asset pricing literature to provide a better fit of the constructed models to empirically observed patterns in asset returns.

The agents assess future distributions of underlying economic variables according to their own subjective beliefs which may differ from the objective probabilities. Agents 'agree to disagree' about their beliefs, and do not view the belief differences as a result of information asymmetries. Consequently, there is no room for strategic trading behavior and prices do not convey any additional information on the evolution of the economy. Competitive equilibria with complete markets can be therefore characterized by solving a planner's problem, and subsequently decentralizing the allocations.

While the market survival literature generally investigates the asymptotic behavior of the ratio of subjective beliefs to the true probability distribution, I develop a method that is more suitable for the nonseparable nature of the preferences. I utilize the planner's problem derived in Dumas, Uppal, and Wang (2000) and extend it to include heterogeneity in beliefs. The solution of the planner's problem involves endogenously determined processes that can be interpreted as stochastic Pareto weights.

The analysis of market survival then corresponds to investigating the long-run behavior of scaled Pareto weights. I present tight sufficient conditions for the existence of nondegenerate stationary distributions, and for dominance and extinction, and investigate numerically in which regions of the parameter space are these conditions satisfied.

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

Heterogeneous agent models populated by agents endowed with recursive preferences have been analyzed in the literature along several dimensions. Lucas and Stokey (1984)

study such models under certainty, and link the existence of a stable interior steady state in a simple two-agent economy to a sufficient condition called the increasing marginal impatience. This condition postulates that agents discount future less as they become poorer. I show that my analysis crucially depends on a similar quantity that I call *relative impatience*. 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, preferences used in this paper are homothetic and the coefficient of pure time preference is constant. However, differences between the agents in the preference for consuming early or delaying consumption and accumulating wealth are determined endogenously as an equilibrium outcome driven by belief differences. These changes in relative impatience are specific to nonseparable preferences and vanish as we move to the time-separable case.

Anderson (2005) studies Pareto optimal allocations under heterogeneous recursive preferences in a discrete time setup using similar methods but he does not consider survival under belief heterogeneity. Mazoy (2005) discusses long-run consumption dynamics when agents differ in their IES. Isaenko (2008) analyzes the behavior of the term structure of interest rates when agents have differently parameterized Kreps-Porteus utilities. However, none of these papers treats systematically the case of belief heterogeneity. This work aims at filling this gap.

1.2 Organization of the paper

The paper is organized at follows. Section 2 outlines the economic environment, provides a theoretical exposition to recursive preferences, and derives the planner's problem that is central to the analysis. Section 3 discusses the survival results. Sufficient conditions for survival and extinction are provided, followed by a discussion of the numerical analysis. Section 4 decentralizes the allocations obtained a solution to the planner's problem and provides economic intuition underlying the survival results. Section 5 outlines extensions of the developed framework to the case of model uncertainty with learning, and to robust preferences. Section 6 concludes. The Appendix contains derivations omitted from the main text.

2 Optimal allocations under heterogeneous beliefs

I analyze the dynamics of equilibrium allocations in a continuous-time endowment economy populated by two infinitely-lived agents. Each agent ranks the available consumption streams using a version of time-nonseparable preferences axiomatized in discrete-time finite-period environment by Kreps and Porteus (1978), and extended to continuous time in the work of Duffie and Epstein (1992b). In addition, agents differ in their beliefs about the distribution of future quantities. I refer to the heterogeneous economy as to one where both agents have strictly positive wealth shares. A homogeneous economy is populated by a single agent only.

Without introducing any specific market structure, I assume that markets are dynamically complete in the sense of Harrison and Kreps (1979). Moreover, agents are firm believers in their probability models, so that there is no strategic trading behavior. This allows to sidestep the problem of directly calculating the equilibrium by considering a planner's problem. The discussion of market survival then amounts to the analysis of the dynamics of Pareto weights associated with this planner's problem. Optimal allocations and continuation values generate a valid stochastic discount factor and a replicating trading strategy for the decentralized equilibrium.

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

2.1 Information structure and beliefs

The stochastic structure of the economy is given by a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ with an augmented filtration defined by a family of σ-algebras $\{\mathcal{F}_t\}$, $t \geq 0$ generated by a univariate Brownian motion W. Given the continuous-time nature of the problem, equalities are meant in the appropriate almost sure sense. I also assume that all processes, in particular belief distortions and permissible trading strategies, satisfy regularity conditions like square integrability over finite horizons, so that stochastic integrals are well defined and pathological cases are avoided.

The scalar aggregate endowment process Y satisfies

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

with constant parameters μ_y and σ_y.

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

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

where u^n is an adapted process. The martingale M^n measures the disparity between the subjective and true probability measure and is commonly called the belief ratio. The main results of the paper are developed using a constant u^n but the computational strategy allows to incorporate more general distortion processes. Section 5 discusses endogenously determined distortions in models with learning and robust utility.

The Girsanov theorem implies that agent n whose beliefs are distorted by M^n views the evolution of the Brownian motion W as distorted by a drift component u^n

$$dW_t = u_t^n dt + dW_t^n \tag{3}$$

where W^n is a Brownian motion under Q^n. Consequently, the aggregate endowment is perceived to contain an additional drift component $u^n\sigma_y$, and u^n can be interpreted as a degree of optimism or pessimism about Y.

2.2 Recursive utility

Agents endowed with time-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 time-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 time-separability assumption by axiomatizing discrete-time preferences where temporal resolution of uncertainty matters. While intratemporal lotteries in the Kreps-Porteus axiomatization still satisfy the von Neumann-Morgenstern expected utility axioms, intertemporal lotteries cannot be in general reduced to atemporal ones. Kreps and Porteus motivated preference for early resolution of uncertainty as a re-

duced form for an underlying auxiliary decision model in which resolving the uncertainty early allows the agent to take utility-improving actions that lie outside of the main model.

The representation result in Kreps and Porteus (1978) shows how to characterize the preference relation using a recursion in which the continuation value at a given point in time is calculated by aggregating the contribution of consumption today and of the expected continuation value tomorrow using a nonlinear function, called the aggregator.

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.

The recursive utility process $V^n(C^n)$ of agent n associated with a given consumption process C^n is defined as

$$V_t^n(C^n) \equiv E_t^{Q^n}\left[\int_t^\infty f(C_s^n, V_s^n)\, ds\right] \tag{4}$$

where f is the aggregator that links together consumption C_s^n at time $s \in [t, \infty)$ with the continuation value V_s^n. Agents prefer early resolution of uncertainty when the aggregator is convex in the expected continuation value. Time-separability in preferences is achieved as a special case when the aggregator is linear in the expected continuation value and additive in the contribution of the two components.

Existence and concavity of the recursive utility process are often difficult to prove. Duffie and Epstein (1992b) focus on the finite-horizon case and prove concavity only for a concave aggregator f. Appendix C in their paper discusses the infinite-horizon case but the sufficient conditions are too strict for this paper. However, the Markov structure of the problem allows me to utilize the infinite-horizon extensions demonstrated in Duffie and Lions (1992). Schroder and Skiadas (1999) prove that $V^n(C^n)$ is concave even when f is convex in its second argument, a case that is central in this work.

I am interested in a particular functional form of the aggregator that is the continuous-time counterpart of the constant elasticity of substitution preferences used in Epstein and Zin (1989, 1991)

$$f(C, V) = \frac{\beta}{\rho}\left[(C)^\rho (\gamma V)^{1-\frac{\rho}{\gamma}} - (\gamma V)\right] \tag{5}$$

with parameters satisfying $\gamma, \rho < 1$, and $\beta > 0$. Preferences specified by this form of the aggregator[1] are homothetic, exhibit a constant relative risk aversion with respect to

[1] The cases of $\rho \to 0$ and $\gamma \to 0$ can be obtained as limits with an ordinally equivalent aggregator in which

intratemporal wealth gambles $\alpha = 1-\gamma$, and (under intratemporal certainty) a constant intertemporal elasticity of substitution $\eta = \frac{1}{1-\rho}$. Parameter β is the time-preference coefficient. Assumption 2 below restricts parameters to assure sufficient discounting for the continuation values to be finite.

In the case when $\gamma = \rho$, the utility reduces to the time-separable CRRA utility with the coefficient of relative risk aversion α. Preferences given by (4–5) thus can be viewed as a natural extension of the CRRA utility case that allows to disentangle the bond between intratemporal risk aversion and intertemporal elasticity of substitution. Appendix A shows how to derive the aggregator f as the limiting continuous-time case of the discrete-time preferences formulated by Epstein and Zin (1989).

Duffie, Geoffard, and Skiadas (1994) were the first to study optimal and equilibrium allocations with stochastic recursive utility specified in (4). Dumas, Uppal, and Wang (2000) offer a different way of defining the recursive utility that is more convenient for the purposes of this paper.[2] They show that the recursive utility process V^n can be equivalently represented as a solution to the minimization problem

$$\lambda_t^n V_t^n = \inf_{\nu^n} E_t^{Q^n}\left[\int_t^\infty \lambda_s^n F\left(C_s^n, \nu_s^n\right) ds\right] \tag{6}$$

subject to

$$\begin{aligned} \frac{d\lambda_s^n}{\lambda_s^n} &= -\nu_s^n ds, \qquad s \geq t \\ \lambda_0^n &= 1 \end{aligned} \tag{7}$$

where ν^n is called the discount rate process, and λ^n the discount factor process. The felicity function F and the aggregator f are linked through the Legendre transformation

$$F\left(C,\nu\right) \equiv \inf_{V\in R}\left[f\left(C,V\right)+\nu V\right] \tag{8}$$

$$f\left(C,V\right) \equiv \sup_{\nu\in R}\left[F\left(C,\nu\right)-\nu V\right] \tag{9}$$

The transformation (8-9) assumes that f is convex in its second argument. When f is concave, it suffices to swap the sup and inf operators in the above definitions.

For the case of the Duffie-Epstein-Zin preferences (5), taking the first order condition in

γu is replaced with $1+\gamma u$. To save notation, the analysis in the text is formally limited to $\rho \neq 0 \neq \gamma$ but the results hold for the limiting cases as well.

[2]This specification is based on the more general variational utility approach studied by Geoffard (1996) in the deterministic case, and El Karoui, Peng, and Quenez (1997) in stochastic environment.

(8) leads to

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

Formula (6) together with an application of the Girsanov theorem suggest that it is advantageous to combine the contribution of the discount factor process λ^n and the martingale M^n that specifies the belief distortion in (2):

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

$$\bar{\lambda}^n \doteq \lambda^n M^n. \tag{10}$$

Applying Itô's lemma to (10) leads to a minimization problem under the true probability measure

$$\bar{\lambda}_t^n V_t^n = \inf_{\nu} E_t \left[\int_t^{\infty} \bar{\lambda}_s^n F\left(C_s^n, \nu_s^n\right) ds \right] \tag{11}$$

subject to

$$\begin{aligned} \frac{d\bar{\lambda}_s^n}{\bar{\lambda}_s^n} &= -\nu_s^n dt + u_s^n dW_s, \qquad s \geq t \\ \bar{\lambda}_0^n &= 1. \end{aligned} \tag{12}$$

The problem (11-12) indicates that $F(C,\nu)$ can be viewed as a generalization of the period utility function with a potentially stochastic rate of time preference ν that depends on the properties of the consumption process, and thus arises endogenously in a market equilibrium. Moreover, belief distortions are now fully incorporated in the framework of Dumas, Uppal, and Wang (2000) — the only difference is that the modified discount factor process in not locally predictable.

The diffusion term $u_s^n dW_s$ has an intuitive interpretation. Consider an optimistic agent with $u^n > 0$. This agent's beliefs are distorted in that the mass of the distribution of dW_s is shifted to the right — the agent effectively overweighs good realizations of dW_s. Formula (12) indicates that under the true probability measure, positive realizations of dW_s increase the term $d\bar{\lambda}_s/\bar{\lambda}_s$ which implies that the optimistic agent discounts positive realizations of dW_s less than negative ones.

From the perspective of the utility-maximizing agent, assigning a higher probability to an event and lower discounting of the utility contribution of this event have the same effect.

In fact, equation (6) suggests that we can understand the belief distortion as a preference shock, and view $\bar{\lambda}F(C,\nu)$ as a state-dependent felicity function. However, interpreting the martingale M^n as a belief distortion is more appealing since it bears a clearer economic meaning, separating the structure of beliefs and preferences.

2.3 Planner's problem and optimal allocations

Although observed equilibrium prices indicate that beliefs differ across agents, the agent does not ascribe the belief differences to superior information available to other market participants. While a likelihood evaluation of the past observed data reveals that the view of an agent with distorted beliefs becomes less and less likely to be correct as time passes, absolute continuity of the measure Q^n with respect to P implies that he can never refute his view of the world as impossible. Since agents are firm believers in their subjective models, they agree to disagree about the interpretation of commonly observed endowment dynamics, and prices do not convey additional information.

In this environment, I can follow Dumas, Uppal, and Wang (2000) and utilize market completeness to identify competitive equilibria in the heterogeneous agent economy with optimal solutions to a planner's problem. A decentralization of the solution then provides the market prices and replicating trading strategies.

Notice first that the problem of an individual agent (6–7) is homogeneous degree one in the discount factors, and homogeneous degree γ in consumption. In the homogeneous economy, there exists a closed-form solution for the continuation values $V_t^n(Y) = \bar{\lambda}_t^n Y_t^\gamma \tilde{V}^n$ where

$$\tilde{V}^n = \frac{1}{\gamma}\left(\frac{\beta}{\beta - \rho\left(\mu_c + u^n\sigma_c + \frac{1}{2}(\gamma-1)\sigma_c^2\right)}\right)^{\frac{\gamma}{\rho}} \tag{13}$$

The following assumption is needed for the continuation values in the homogeneous economies to be well-defined:

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

$$\beta > \max_n \rho\left(\mu_c + u^n\sigma_c + \frac{1}{2}(\gamma-1)\sigma_c^2\right).$$

In the heterogeneous economy, the fictitious planner maximizes a weighted average of the continuation values of the two agents. Given a pair of strictly positive initial Pareto weights

$\alpha = (\alpha^1, \alpha^2)$, the planner's time-0 objective function $J_0(\alpha)$ is the solution to the problem

$$J_0(\alpha) = \sup_{(C^1,C^2)} \inf_{(\nu^1,\nu^2)} \sum_{n=1}^{2} E_0 \left(\int_0^\infty \bar{\lambda}_s^n F(C_s^n, \nu_s^n)\, ds \right) \tag{14}$$

subject to the law of motion for the modified discount factors

$$\begin{aligned} \frac{d\bar{\lambda}_s^n}{\bar{\lambda}_s^n} &= -\nu_s^n dt + u_s^n dW_s, \qquad s \geq 0 \\ \bar{\lambda}_0^n &= \alpha^n \end{aligned} \tag{15}$$

for $n \in \{1, 2\}$, and the feasibility constraint

$$C^1 + C^2 \leq Y. \tag{16}$$

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 (14-16) is a straightforward extension when individual continuation values are well defined. The planner's objective function is bounded from above by the weighted average of continuation values from the homogeneous economies

$$J_0(\alpha) \leq \alpha^1 V_0^n(Y) + \alpha^2 V_0^n(Y)$$

and the supremum in (14) thus exists. Since the continuation values are concave, first-order conditions are sufficient for the supremum problem. The following Lemma describes the behavior of the objective function at the boundaries.

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

$$J_0(0, \alpha^2) \doteq \lim_{\alpha^1 \searrow 0} J_0(\alpha^1, \alpha^2) = \alpha^2 V_0^2(Y) \tag{17}$$

and

$$\lim_{\alpha^1 \searrow 0} C^2(\alpha^1, \alpha^2) = Y \tag{18}$$

The case $\alpha^2 \searrow 0$ is symmetric.

Proof. Consider the case $\alpha^1 \searrow 0$. Given optimal consumption streams $C^n(\alpha)$, we have

$$J_0(\alpha) = \alpha^1 V_0^1\left(C^1(\alpha)\right) + \alpha^2 V_0^2\left(C^2(\alpha)\right)$$

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

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

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

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

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

and

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

which implies $J_0(0, \alpha^2) \geq \alpha^2 V_0^2(Y)$. Therefore (17) holds, and (18) is a direct consequence. ■

The planner's problem (14-16) suggests that we can interpret the discount factor processes $\bar{\lambda}^n$ as stochastic Pareto weights. Indeed, if $\bar{\lambda}_0^n = \alpha^n$ are the initial weights, then $\bar{\lambda}_t^n$ are the consistent state-dependent weights for the continuation problem of the planner at time t.

The evolution of the weights involves the drift component ν^n and thus can be only determined in equilibrium unless agent n's preferences are time separable in which case $\nu^n = \beta$. The variation in Pareto weights arises from the interaction of two components in the model — the time-nonseparable preference structure as in Lucas and Stokey (1984) and Dumas, Uppal, and Wang (2000), and the belief distortion that drives the diffusion component in (15).[3]

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 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

[3]The belief heterogeneity introduces an additional risk component arising through the stochastic reweighting of wealth shares. The diffusive component in the weight dynamics will have a direct impact on local risk prices. Notice that other sources of heterogeneity, including differences in risk aversion and IES parameters, do not lead to a diffusion component in the dynamics of the Pareto weights (15), and reweighting therefore has no impact on local risk prices in these cases.

problem, as long as the distorting processes u^n are functions of the existing state variables.[4]

The only cost of introducing belief distortions into the problem comes from the fact that the modified discount processes are not locally smooth but contain the diffusion component $u_t^n dW_t$ (unless the agent has the correct beliefs). In the Markov setup below, this will be reflected by the occurrence of additional second partial derivative terms.

More generally, the social planner can choose to maximize welfare as the weighted average of utilities evaluated as distorted relative to any subjective measure, as long as the absolute continuity assumption is satisfied, and the distorting martingales M^n are properly constructed relative to the chosen measure. Then the discount rate process ν^n of the agent whose belief distortion coincides with the distortion of the social planner will be locally predictable.

2.4 Hamilton-Jacobi-Bellman equation

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

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

$$0 \equiv \sup_{(C^1,C^2)} \inf_{(\nu^1,\nu^2)} \sum_{n=1}^{2} \bar{\lambda}^n \left[F\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n}\nu^n\right] + J_y \mu_y Y + \frac{1}{2}\mathrm{tr}\left(J_{zz}\Sigma\right) \tag{19}$$

where

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

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

The minimization problem in the HJB equation (19) can be solved separately by recalling the Legendre transformation (9). Under the optimal discount rate process ν^n for agent n,

$$f\left(C^n, J_{\bar{\lambda}^n}\right) \equiv F\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n}\nu^n. \tag{20}$$

Optimal consumption shares ζ^n are given by the first-order conditions in the maximiza-

[4]This is not the case when learning is introduced but even then a tractable recursive formulation is at hand, as shown in Section 5.

tion problem

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

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

The HJB equation (19) further implies that J is homogeneous degree one in $\bar{\lambda}$, and homogeneous degree γ in Y. This suggests a transformation of variables

$$\theta^1 = \frac{\bar{\lambda}^1}{\bar{\lambda}^1 + \bar{\lambda}^2} \qquad \theta^2 = \bar{\lambda}^1 + \bar{\lambda}^2 \tag{21}$$

and the following guess

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

where $\tilde{J}^n\left(\theta^1\right)$ are scaled continuation values of the two agents

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

The single state variable θ^1 represents the Pareto share of agent 1. The dynamics of θ^1 is central to the study of survival in this paper. Obviously, θ^1 is bounded between zero and one. It will become clear that for strictly positive initial weights the boundaries are unattainable, so that θ^1 evolves on the open interval $(0,1)$.

These considerations, together with Lemma 3, imply the following Proposition:

Proposition 4 *The objective function for the planner's problem (14-16) is*

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

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

$$\begin{aligned} 0 \;=\; & \theta^1 \frac{\beta}{\rho} \left(\zeta^1\right)^\rho \left(\tilde{J}^1\right)^{1-\frac{\rho}{\gamma}} + \left(1-\theta^1\right) \frac{\beta}{\rho} \left(1-\zeta^1\right)^\rho \left(\tilde{J}^2\right)^{1-\frac{\rho}{\gamma}} + \\ & + \left(-\frac{\beta}{\rho} + \mu_c + \frac{1}{2}(\gamma-1)\sigma_c^2 + \left(\theta^1 u^1 + \left(1-\theta^1\right) u^2\right)\sigma_c\right) \tilde{J} + \\ & + \theta^1\left(1-\theta^1\right)\left(u^1 - u^2\right)\sigma_c \tilde{J}_{\theta^1} + \frac{1}{2}\frac{1}{\gamma}\left(1-\theta^1\right)^2 \left(\theta^1\right)^2 \left(u^1-u^2\right)^2 \tilde{J}_{\theta^1\theta^1} \end{aligned} \tag{22}$$

with boundary conditions $\tilde{J}(0) = \tilde{V}^2$ *and* $\tilde{J}(1) = \tilde{V}^1$ *where* $\tilde{V}^n$ *are defined in (13).*

Proof. The proof follows from the preceding derivations. ■

Unfortunately, equation (22) does not have a general closed-form solution. However, the Pareto share θ^1 of agent 1 remains the only state variable which considerably simplifies numerical solutions.

The same solution method and derivation of the HJB equation applies to multi-agent economies and more sophisticated Markov dynamics of the aggregate endowment process. In an N-agent economy, the state vector includes $N-1$ Pareto shares as state variables. The boundary conditions for $\theta^n = 0$, $n \in \{1, \ldots, N\}$ associated with the N-agent version of the PDE (22) are given by the solution of an $(N-1)$-agent economy that excludes agent n. In this way, solutions to multi-agent economies are calculated by iteratively adding individual agents.

3 Survival

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

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

Since prices are observed in equilibrium, agents have to agree on them. When markets are complete, this implies that the object

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

has to be equalized across agents, P-a.s. Deviations in beliefs thus have to be offset by reciprocal deviations in marginal utilities. If agent 1 has a constant belief distortion $u^1 \neq 0$ and agent 2 is rational, then $\lim_{s\to\infty} M_{t+s}^1 = 0$ (P-a.s.), and thus $\lim_{s\to\infty} U'\left(C_{t+s}^1\right)/U'\left(C_{t+s}^2\right) = +\infty$ (P-a.s.). For a class of utility functions that includes the CRRA utility case $\gamma = \rho$ in this paper, this implies $\lim_{s\to\infty} \zeta_{t+s}^1/\zeta_{t+s}^2 = 0$ (P-a.s.).

When preferences are not time-separable, this straightforward link breaks down because

marginal utilities also depend on continuation values. Since these continuation values are not available in closed form, they have to be solved for numerically. Studying asymptotic behavior with numerically calculated objects can be treacherous due to the propagation of the approximation errors, and thus a different approach is called for. I derive the local dynamics of the Pareto share θ^1, and then investigate the existence of a unique stationary distribution.

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

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

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

The definition of extinction and survival differs from those in Kogan, Ross, Wang, and Westerfield (2009) or Yan (2008) who use the consumption share ζ^1. However, these definitions are equivalent in this setting:

Lemma 6 *Definition 5 remains unchanged when the Pareto share θ^1 is replaced with the consumption share ζ^1.*

Proof. The consumption share

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

is continuous and monotonically increasing in θ^1 and

$$\lim_{\theta^1 \searrow 0} \zeta^1 = 0 \quad \text{and} \quad \lim_{\theta^1 \nearrow 1} \zeta^1 = 1.$$

■

3.1 Dynamics of the Pareto share and stationary distributions

Recall the dynamics of the modified discount factor processes $\bar{\lambda}^n$ in (15). An application of Itô's lemma to $\theta^1 = \bar{\lambda}^1 / \left(\bar{\lambda}^1 + \bar{\lambda}^2\right)$ yields

$$\begin{aligned} d\theta_t^1 &= \theta_t^1 \left(1 - \theta_t^1\right) \left[\nu_t^2 - \nu_t^1 + \left(\theta_t^1 u^1 + \left(1 - \theta_t^1\right) u^2\right) \left(u^2 - u^1\right)\right] dt + \\ &\quad + \theta_t^1 \left(1 - \theta_t^1\right) \left(u^1 - u^2\right) dW_t \end{aligned} \tag{23}$$

When $u^1 = u^2$, the dynamics of the Pareto share is degenerate, and the Pareto share is constant, $\theta_t^1 \equiv \alpha^1 / \left(\alpha^1 + \alpha^2\right)$. In this case, the individual consumption dynamics is identical to that of an economy populated by two rational agents. In what follows, I abstract from this situation, and assume $u^1 \neq u^2$.

The drift term of the Pareto share contains the difference of the discount rates $\nu^2 - \nu^1$ which I call the *relative patience* of agent 1. When preferences are non-separable, the discount rates are determined endogenously in the model as a solution to problem (14–16), and are given by the first-order condition in problem (8):

$$\nu^n = \frac{\beta}{\rho} \left(\gamma + (\rho - \gamma) \left(\zeta^n\right)^\rho \left(\tilde{J}^n\right)^{-\frac{\rho}{\gamma}}\right) \tag{24}$$

The discount rates ν^n are twice continuously differentiable functions of the state variable θ^1, and thus θ^1 is an Itô process on the open interval $(0, 1)$ with continuous drift and volatility coefficients.[5] Intuitively, one would expect a stationary distribution to exist if the process exhibits sufficient pull toward the center of the interval when close to the boundaries. This is formalized in the following Proposition:

Proposition 7 *Given the dynamics of the Pareto share (23), it is sufficient for the existence of a unique stationary density $q(\theta)$ that the following two conditions hold*

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

(ii) $\limsup_{\theta^1 \nearrow 1} \left[\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)\right] < \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]$.

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

Proof. Given an initial condition $\theta_0^1 \in (0, 1)$, the process (23) lives on the open interval $(0, 1)$ with unattainable boundaries (the aggregator satisfies an Inada condition at zero).

[5] The unattainability of the boundaries follows from the proof of Proposition 7.

For any numbers $0 < a < b < 1$, the process θ^1 has bounded and continuous drift and volatility coefficients on (a, b), and the volatility coefficient is bounded away from zero. It is thus sufficient to establish appropriate boundary behavior of θ^1 in order to make the process positive Harris recurrent (see Meyn and Tweedie (1993)). Since the process will also be φ-irreducible for the Lebesgue measure under these boundary conditions, there exists a unique stationary distribution.

Denote $\mu_\theta(\theta)$ and $\sigma_\theta(\theta)$ the drift and volatility coefficients in (23). The boundary behavior of the process θ^1 is captured by the scale measure $S : (0,1)^2 \to \mathbb{R}$ defined as

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

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

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

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

The boundaries are nonattracting if and only if

$$\lim_{\theta_l \searrow 0} S[\theta_l, \theta_h] = \infty \quad \text{and} \quad \lim_{\theta_h \nearrow 1} S[\theta_l, \theta_h] = \infty \tag{25}$$

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

$$\lim_{\theta_l \searrow 0} M[\theta_l, \theta_h] < \infty \quad \text{and} \quad \lim_{\theta_h \nearrow 1} M[\theta_l, \theta_h] < \infty \tag{26}$$

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

In our case,

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

where

$$s_{TS}(\theta) = \left(\frac{1-\theta}{1-\theta_0}\right)^{-\frac{2u^1}{u^1-u^2}} \left(\frac{\theta}{\theta_0}\right)^{\frac{2u^2}{u^1-u^2}} \tag{27}$$

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

For the left boundary, assume that in line with condition (i) there exist $\underline{\theta} \in (0,1)$ and

$\underline{\nu} \in \mathbb{R}$ such that

$$\nu^2(\theta) - \nu^1(\theta) \geq \underline{\nu}$$

for all $\theta \in (0, \underline{\theta})$. Taking $\theta_0 = \underline{\theta}$, the scale measure can be bounded as

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

The left limit in (25) thus diverges to infinity if

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

which is satisfied when

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

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

$$\bar{\nu} \leq \frac{1}{2}\left[(u^1)^2 - (u^2)^2\right].$$

It turns out that the bounds implied by conditions (26) are marginally tighter. Following the same bounding argument as above, sufficient conditions for (26) to hold are

$$\begin{aligned} \underline{\nu} &> \frac{1}{2}\left[(u^1)^2 - (u^2)^2\right] \\ \bar{\nu} &< \frac{1}{2}\left[(u^1)^2 - (u^2)^2\right] \end{aligned} \tag{28}$$

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

It is also useful to note that the unique stationary density $q(\theta)$ is proportional to the speed density $m(\theta)$. ■

In empirical applications, it may be preferable for the economy to converge to the stationary distribution fast enough from any initial condition, so that data observed over finite horizons are a representative sample of the stationary distribution. Proposition 7 gives sufficient conditions for the existence of a unique stationary distribution but it does not say

anything about the rate of convergence. It turns out that under the conditions in Proposition 7, convergence occurs at an exponential rate.

Proposition 8 *Under the sufficient conditions in Proposition 7, the process θ^1 is ρ-mixing with an exponential decay rate, i.e. there exist constants $B > 0$ and $\delta \in (0,1)$ such that for any square-integrable function $\phi \in L^2$ where*

$$L^2 = \left\{ \phi : (0,1) \to \mathbb{R} : \|\phi\| \doteq \left(\int_0^1 |\phi(\theta)|^2 q(\theta)\, d\theta \right)^{\frac{1}{2}} < \infty \right\}$$

we have

$$\sup_{\|\phi\|=1} \left\| E\left[\phi\left(\theta_t^1\right) \mid \theta_0^1 = \theta_0 \right] - \int_0^1 \phi(\theta)\, q(\theta)\, d\theta \right\| \doteq \rho_t \leq B e^{-\delta t}.$$

Proof. Chen, Hansen, and Carrasco (2010) show that under the conditions that stated by Proposition 7, the sufficient conditions for exponential convergence in L^2 norm are

$$\liminf_{\theta \searrow 0} \left(\frac{\mu_\theta(\theta)}{|\sigma_\theta(\theta)|} - \frac{|\sigma'(\theta)|'}{2} \right) > 0 \quad \text{and} \quad \liminf_{\theta \nearrow 1} \left(\frac{\mu_\theta(\theta)}{|\sigma_\theta(\theta)|} - \frac{|\sigma'(\theta)|'}{2} \right) < 0$$

These conditions are satisfied by imposing the same bounds as in (28). An analogous calculation to that in the proof of Proposition 7 reveals that the bounds are again the least tight bounds of this type. ■

Conditions in Proposition 7 can be checked numerically from the solution of the planner's problem. These conditions have an intuitive interpretation. For instance, condition (i) states that agent 1 can survive under the true probability measure even in cases when his beliefs are more distorted, $|u^1| > |u^2|$, as long as his relative patience $\nu^2 - \nu^1$ becomes sufficiently high as his consumption share vanishes. In this way, despite the fact that his investment decisions are skewed toward making incorrectly optimistic or pessimistic bets, his savings are sufficiently high to overcome the losses from these incorrect bets, at the cost of lowering his current consumption. A high propensity to save implies a rightward pull on θ^1, preventing convergence of θ^1 to zero.

When the assumptions of Proposition 7 are not satisfied, then either of the agents can dominate in the long run along some paths. The following Proposition summarizes the result from Proposition 7, and provides additional characterization of situations that do not lead to a nondegenerate stationary distribution.

Corollary 9 *Define the 'attracting' conditions*

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

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

as alternatives to 'repealing' conditions (i) and (ii) in Proposition 7. 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*

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

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 $\omega \in S^2$.

Proof. Statement (a) is a direct consequence of Proposition 7. Conditions (i') and (ii') are sufficient conditions for the boundaries to be attracting. Lemma 6.1 in Karlin and Taylor (1981) then shows that if the 'attracting' condition is satisfied for a boundary then θ^1 converges to this boundary on a set of paths that has a strictly positive probability. This probability is equal to one if the other boundary is non-attracting. Combining these results, we obtain statements (b), (c), and (d). ■

The four 'attracting' and 'repealing' conditions are only sufficient and their combinations stated in Proposition 9 are not exhaustive. There are delicate cases involving nonstrict inequalities, which are however not relevant for the numerical analysis below.

3.2 Change of measure

The developed survival criteria are stated from the perspective of a rational agent. However, agents whose beliefs differ from the true probability measure evaluate their suvival chances differently. Although both agents understand that the optimal (and equilibrium) allocations are given as a solution to the problem outlined in Proposition 4, they differ in their view about the future consumption dynamics.

Lemma 10 *Agent n views the dynamics of the economy as if the belief distortions were given by*

$$(u^n)_{(n)} = 0 \qquad (u^{\sim n})_{(n)} = u^{\sim n} - u^n$$

where $\sim n$ indexes the other agent in the economy and $\left(u^k\right)_n$ are the beliefs of agent k from the standpoint of agent n.

Proof. The evolution of Brownian motion W under the beliefs of agent n is

$$dW_t = u^n dt + dW_t^n$$

Since the evolution of θ^1 completely describes the dynamics of the economy, substituting this expression into (23) and reorganizing yields the desired result. ■

The Lemma implies in particular that Proposition 7 applies for the survival considerations under a subjective probability measure Q^n, as long as u^k are replaced with $\left(u^k\right)_n$ for $k = 1, 2$. Further, Proposition 9 can be modified in the same way to obtain statements about dominance under Q^n.

3.3 CRRA preferences

The framework introduced in this paper includes as a special case the time-separable constant relative risk aversion preferences when $\gamma = \rho$. Yan (2008) or Kogan, Ross, Wang, and Westerfield (2009) show that in the economy presented in this paper, the agent whose beliefs are less distorted dominates in the long run under measure P. We can confirm this result using the statements in Proposition 9.

Corollary 11 *Under time-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.*

Proof. Assume without loss of generality that $|u^2| \leq |u^1|$. The sufficient part is an immediate consequence of Proposition 9. Under time-separable preferences, $\nu^2 - \nu^1 \equiv 0$, and thus if $|u^2| < |u^1|$ then conditions (i') and (ii) hold, and agent 2 dominates in the long run under P.

For the necessary part, when $u^2 = u^1$, then θ^1 is constant and both agents survive under P. When $-u^2 = u^1 = u$, then it follows from inspection of formula (27) in the proof of Proposition 7 that conditions (25) are satisfied and the boundaries are non-attracting.

Lemma 6.1 in Karlin and Taylor (1981) then implies that both agents survive under P. This completes the proof.

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. ■

Corollary 12 *Under time-separable CRRA preferences, agent n survives under measure Q^n, and dominates in the long run under Q^n if and only if $u^n \neq u^{\sim n}$.*

Proof. This is an immediate consequence of Lemma 10 and Corollary 11. ■

Under time-separable CRRA preferences, the dynamics of the Pareto share (23) does not depend on the characteristics of the endowment process. The survival result in Corollary 11 thus extends to an arbitrary adapted aggregate endowment process Y that satisfies elementary integrability conditions, including a constant one, as long as the two agents can write contracts on the realizations of the Brownian motion W. It is a special case of the analysis in Kogan, Ross, Wang, and Westerfield (2009) who show that this survival result holds, under mild conditions, for any time-separable preferences with bounded relative risk aversion. In this sense, the time-separable environment generates a robust result about the extinction of agents whose beliefs are relatively imprecise.

3.4 The time-nonseparable case

When preferences are not separable in time, the evolution of the Pareto share θ^1 depends on the dynamics of the endowment process through the endogeneously determined discount rates ν^n. Proposition 7 states sufficient conditions for the existence of a stationary distribution but it does not indicate whether there actually exist parameter configurations under which these conditions are satisfied, and whether these potential configurations are empirically sensible. In this section, I investigate this question numerically. I solve for the optimal allocations for a wide range of risk aversion and IES levels, $\left(1-\gamma, (1-\rho)^{-1}\right)$, and then check whether the dynamics of the Pareto share satisfies conditions for the existence of a stationary distribution that I developed in Section 3.1.

In all the numerical calculations, agent 2 is assumed to be rational, $u^2 = 0$. The time-preference coefficient is set to $\beta = 0.05$ in order to assure sufficient discounting across all parameter combinations but the survival results are not sensitive to a particular choice of β.

The aggregate endowment process is parameterized by $\mu_y = 0.02$ and $\sigma_y = 0.015$ in the baseline scenario.

The main findings are summarized in Figure 1. The graphs depict the countour lines for the mean consumption share of agent 1 under the stationary distribution, $E\left[\zeta_t^1\right]$, if this stationary distribution exists. In the top panel, agent 1 is optimistic with $u^1 = 0.05$. The bottom panel captures the results for a pessimistic agent, $u^1 = -0.05$. The dotted diagonal line shows parameter combinations when $\gamma = \rho$, corresponding to economies where agents are endowed with time-separable CRRA utility.

The magnitude of the distortions is relatively small — it corresponds to a 3.75 percent relative error in the perception of the growth rate of the aggregate endowment. This magnitude is chosen only as an illustration, without having a particular empirical counterpart in mind. The sensitivity of the results with respect to the chosen parameters is discussed below. In particular, I show that the misperception of the growth rate, $u^1\sigma_y$, is not a key predictor of the mean consumption share of the incorrect agent.

The results indicate that a nondegenerate stationary distribution exists when the coefficient of relative risk aversion is sufficiently larger than the inverse of the intertemporal elasticity of substitution ($\gamma < \rho$), at least for the analyzed range of parameters. An analytical proof of this numerical result remains an open question, though. In particular, it is not possible to 'prove' numerically whether it is sufficient for the existence of a nondegenerate stationary distribution that $\gamma < \rho$, or that a stricter condition is required, $\gamma < \rho - \varepsilon(\rho)$ for a strictly positive function $\varepsilon(\rho)$ (given all other parameters in the economy). I was not able to establish the existence of such a strictly positive function. In the neighborhood of the set of time-separable parameterizations $\gamma = \rho$, the consumption share of agent 1 becomes negligible, and numerically indistinguishable from zero. In Section 4 I provide an argument that explains why increasing the risk aversion or the IES relative to the time-separable case increases the survival chances of the agent with distorted beliefs but this argument is not able to dissect the question about $\varepsilon(\rho)$.

Figure 4 plots the densities $q(\zeta^1)$ for the stationary distributions in several economies with an optimistic agent for which a stationary distribution exists. As risk aversion increases, the mass of the density shifts to the right, and the optimistic agent is indeed able to acquire a substantial fraction of the aggregate consumption.

When risk aversion is smaller than the inverse of the IES ($\gamma > \rho$), conditions (i') and (ii') from Corollary 9 are satisfied numerically, and one of the agents dominates in the long run, each on a set of paths of strictly positive probability. A similar cautious remark regarding

the general validity of the numerical results when γ is close to ρ applies as above.

Time-separable utility ($\gamma = \rho$) can therefore be viewed as a threshold scenario characterized by the almost sure long-run dominance of the rational agent. This result is not robust to sufficiently large deviations from time-separability, and deviations in opposite directions from $\gamma = \rho$ yield qualitatively different results. In particular, nondegenerate stationary distributions exist in the empirically plausible cases when $\gamma < \rho$.

The qualitative survival results hold regardless of whether agent 1 is optimistic or pessimistic but the mean consumption share of agent 1 differ substantially. While the optimistic agent is able to acquire a substantial share of aggregate endowment in a wider region of the parameter space, the pessimistic agent has a significantly smaller consumption share under the stationary distribution. The reason is the correlation of the aggregate endowment growth with the shock that the two agents bet on, and the next section elaborates on this observation. When aggregate endowment risk vanishes but the agents can still write contracts on the outcome of the shock, optimism and pessimism become completely symmetric. Figure 3 plots the countour lines for the mean consumption share of agent 1 under the stationary distribution when $\sigma_y = 0$ — the graph is identical for $u^1 = \pm 0.05$.

4 Saving and portfolio decisions in a decentralized economy

In order to shed light on the economic mechanism that generates the survival results described in the preceding section, I decentralize the solution of the planner's problem, and investigate the saving and portfolio decisions of the two agents. The separation of the solution into a consumption-saving problem and a portfolio allocation problem is a recursive-utility extension of Karatzas, Lehoczky, and Shreve (1987) and Cox and Huang (1989) analyzed in Schroder and Skiadas (1999).

Using the construction from Duffie and Epstein (1992a), the stochastic discount factor process is given by

$$S_t^n = \exp\left(-\int_0^t \nu_s ds\right)\left(\frac{Y_t}{Y_0}\right)^{\gamma-1}\left(\frac{\zeta_t^n}{\zeta_0^n}\right)^{\rho-1}\left(\frac{\tilde{J}_t^n}{\tilde{J}_0^n}\right)^{1-\frac{\rho}{\gamma}}$$

and its evolution can be written as

$$\frac{dS_t^n}{S_t^n} = \mu_{s^n}\left(\theta_t^1\right) dt + \sigma_{s^n}\left(\theta_t^1\right) dt.$$

Since each of the agents prices asset payoffs under his subjective probability measure, the quantities

$$\frac{M_{t+s}^n}{M_t^n}\frac{S_{t+s}^n}{S_t^n}$$

have to be equalized across $n \in \{1, 2\}$, P-a.s. under complete markets. Agents have to agree on the prices of traded assets but the subjective and objective evaluations of risk prices will differ. In particular, while each agent subjectively perceives the local risk price to be σ_{s^n}, the local risk price under the true probability measure is $\sigma_{s^n} + u^n$, and this quantity is identical for $n \in \{1, 2\}$.

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). Given the scaling of the economy, it is straightforward to conjecture that asset returns will be functions of the Markov state θ^1. Let the two traded assets be an infinitesimal risk-free bond in zero net supply that yields a risk-free rate $r\left(\theta^1\right)$, and a stock representing the aggregate wealth that pays the aggregate endowment flow as dividend. The stock price evolves as

$$\frac{d\Xi_t}{\Xi_t} = \mu_\Xi\left(\theta_t^1\right) dt + \sigma_\Xi\left(\theta_t^1\right) dW_t$$

The parameters of the return processes are determined numerically from the calculated asset prices.

It is convenient to define the aggregate wealth-endowment ratio $w = \Xi/Y$ and the ratios of individual wealth to aggregate endowment $w^n = \Xi^n/Y$. These ratios can be calculated numerically using the optimal consumption shares and continuation values obtained in the planner's problem (22), and an application of Ito's lemma to $w^n\left(\theta^1\right)$ yields the corresponding drift and volatility coefficients $\mu_{w^n}\left(\theta^1\right)$ and $\sigma_{w^n}\left(\theta^1\right)$.

On the other hand, a standard portfolio argument implies that (omitting the dependence on θ^1)

$$\begin{aligned}\frac{dw_t^n}{w_t^n} &= \left(r_t + \pi_t^n\left[\mu_{\Xi,t} - r_t - \sigma_{\Xi,t}\sigma_y\right] - \mu_y + \sigma_y^2 + \frac{\pi_t^n}{w_t} - \frac{\zeta_t^n}{w_t^n}\right) dt + \\ &\quad + \left(\pi_t^n\sigma_{\Xi,t} - \sigma_y\right) dW_t\end{aligned}$$

where π^n is the wealth share of agent n invested in the risky asset, π^n/w the dividend flow (relative to individual wealth) accrued from holding the stock, and ζ^n/w^n the consumption flow relative to individual wealth. A comparison of the drift and volatility coefficients with those obtained numerically above yields the wealth shares π^n invested in the risky asset.

4.1 The risk-sharing mechanism

The key insight to understanding the survival mechanism is that while the volatility coefficient of the subjective stochastic discount factor σ_{s^n} increases with the degree of risk aversion $1-\gamma$, this coefficient under the true probability measure, $\sigma_{s^n}+u^n$, contains a constant distortion term. The distortions thrust a larger relative wedge between the two subjective discount factors when risk aversion is low. An optimistic (pessimistic) agent will thus accept relatively more of the procyclical (countercyclical) variation in the consumption process in an economy where the risk aversion is relatively low (despite the fact that both agents are equally risk averse). Figure 7 documents this by plotting the wealth share of agent 1 invested in the risky asset. The deviation from one is the larger the lower is the level of risk aversion.

An optimistic agent 1 with a levered stock position and a pessimistic agent 1 with a short stock position are both optimistic regarding the growth rate of their portfolio. Although they agree with the rational agent 2 on the equilibrium price of their wealth, their optimistic belief leads them to consume more from their wealth. The behavior of the consumption-wealth ratio is depicted in Figure 8. A lower risk aversion leads to a higher leverage for an optimistic agent (or a larger short position in case of a pessimistic agent), which translates to a higher degree of optimism and a higher consumption-wealth ratio for small levels of the consumption share.

The behavior of the consumption-wealth ratio then decides the survival chances of the two agents. If the consumption-wealth ratio is sufficiently low for an agent as his consumption share declines, he will save his way out of extinction. Considering that the agent is optimistic about his own wealth growth, he needs, under his subjective beliefs, to tilt the consumption profile toward the future in order to survive, and he is more willing to do so under a higher intertemporal elasticity of substition.

4.2 Sensitivity to the choice of parameters

Under time-separable preferences, analytical results presented in Section 3 show that agent 1 cannot escape extinction. An increase in risk aversion of the CRRA utility does not overturn

this outcome — the lower relative discrepancy between the volatility terms in the stochastic discount factor σ_{s^n} is offset by a decrease in the IES which diminishes the willingness to tilt the consumption profile toward the future.

The analysis in this paper shows that the extinction result from the time-separable case is a specific scenario. Starting from the time-separable case, agent 1 improves his survival chances when risk aversion $1-\gamma$ increases, alleviating the relative discrepancy in σ_{s^n}, and when the IES $(1-\rho)^{-1}$ increases, allowing to tilt the consumption profile toward the future and outsave extinction.

This argument does not imply that the mean consumption share of the incorrect agent has to increase monotonically with risk aversion and IES. For instance, a pessimistic agent becomes optimistic about the growth rate of his portfolio only when he shorts the stock. However, in an economy with high risk aversion the pessimistic agent may not have a chance to take such a short position unless his wealth share is sufficiently small, because the rational agent is not willing to accept the counterbalancing leveraged position. Such an effect is pronounced in the bottom panel of Figure 1 — for high levels of IES, the mean consumption share of the pessimistic agent is strongly non-monotonic in the degree of risk aversion.

When the belief distortion increases in magnitude, one would expect that the incorrect agent loses his consumption share, as he makes more incorrect bets. This is to a great extent confirmed by Figure 2 which replicates Figure 1 for larger magnitudes of the distortions. However, there is also a counteracting effect — an agent with a larger distortion may also be more optimistic about the growth rate of his portfolio, and avoid declines in his consumption by increasing his propensity to save already at higher wealth shares. As expected, I find the latter effect to dominate for some parameterizations in the region of the parameter space with a high IES. The bottom panel in Figure 2 displays the consumption shares for a more pessimistic agent than the one in the bottom panel of Figure 1. The region in the bottom part of the graph, corresponding to high IES, in which the incorrect agent can gain a relatively large average consumption share expands to the right.

An increase in the volatility of the aggregate endowment process σ_y has a similar effect as a decrease in the magnitude of the belief distortion. These two adjustments both decrease the relative importance of the belief distortion in the volatility of the stochastic discount factor under the true probability measure, $\sigma_{s^n} + u^n$. For the investigated parameterizations, proportional marginal increases in both σ_y and u^1 roughly offset each other and have very little overall impact on the mean consumption share of agent 1.

In this respect, the misperception of the growth rate of aggregate endowment, $u^1\sigma_y$, is

not a good predictor for the mean consumption share of agent 1. A marginal increase in $u^1\sigma_y$ caused by an increase in u^1 has an analogous effect as a marginal decrease caused by a decrease in σ_y. Economies with different misperceptions $u^1\sigma_y$ produce plots very similar to those in Figure 1, as long as the ratio u^1/σ_y is held constant.[6]

The discussion in this section can be completed by noting that the described results are not sensitive to changes in the time-preference coefficient β and the growth rate μ_y. These parameters influence the willingness of the agents to tilt the consumption profiles but since the changes in these parameters influence both agents in a symmetric way, they have to a great extent an offsetting character.

4.3 Comparison to economies with only terminal consumption

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

Without intermediate consumption, the agent's intertemporal decision is reduced to the maximization of the (risk-adjusted) expected growth rate of the portfolio, i.e. the agent behaves as an agent with unitary elasticity of substitution. Although the differences between the two frameworks are substantial, the results presented in this paper suggest that the survival chances of the incorrect agent will increase with the degree of risk aversion.

Indeed, Kogan, Ross, Wang, and Westerfield (2006) find that the incorrect agent survives (and actually dominates the economy) for moderate degrees of optimism when the risk aversion is larger than one, and the region of belief distortions that guarantees survival widens as risk aversion increases. However, in their model, economies in which both agents survive correspond only to knife-edge parameterizations. This is due to the absent consumption-saving decision that enables the agents to save their way out of extinction in situations when their wealth share becomes small.

[6]This statement is correct for the range of examined parameterizations, and is not a result that is necessarily true in general.

5 Extensions

In this section, I consider two extensions of the analyzed model. The first extension introduces learning and leads to an endogenously varying belief distortions u^n. The second extension incorporates robust utility models.

While I do not solve these variants, I describe the solution method and suggest interesting open questions. Answering these questions is left for future research.

5.1 Model uncertainty and learning

The survival analysis in the previous text assumed a constant belief distortion u^n. However, the framework developed in Section 2 includes more general processes that can be used to model the distortions. This allows to incorporate agents who learn about the true mean growth rate μ_y of aggregate endowment as they receive new information about the evolution of the economy.

There are various ways of introducing learning into this model. One is to specify for agent n a continuous prior $F_0^n(\mu)$ on $\mathbb{M} \subseteq \mathbb{R}$ such that $\mu_y \in \text{supp } F_0^n$, and update the prior as new information arrives. The disadvantage of this approach for a numerical implementation are unclear boundary conditions at the boundaries of $\mathbb{M}$.

Instead, I assume that the agent has in mind a set of K models that differ in the mean growth rate. The set of models is represented by a vector of distorting components $u^n = (u_k^n)_{k=1}^K$, with the true model being ordered first, i.e. $u_1^n = 0$. At time t, the agent assigns a probability distribution $p_t^n = (p_{kt}^n)_{k=1}^K$ to this vector. The vector p_0^n denotes the prior distribution independent of the realizations of the Brownian motion W. In order to avoid pathologies, I assume $p_{k0}^n > 0$ for all $k \in \{1, \ldots, K\}$.

As in the previous text, agents agree to disagree about the subjective probability measures Q^n. Each agent understands from observed prices that their beliefs differ, he however does not ascribe this differences to superior information of the other agent. Agents therefore only learn from the observed realizations of the aggregate endowment process.

In the setup with time-separable utility, the aggregator $f(C, V)$ in (4) is additive and linear in V, and the law of iterated expectations can be utilized to solve the problem of a Bayesian learner in two steps. First, calculate the continuation values in the recursive formula (4) conditional on a particular model, and then integrate out across models. This two-step solution works because posterior distributions of a Bayesian learner are martingales under the subjective probability measure of the learner.

This method cannot be used when $f(C,V)$ is not separable. Instead, I will show how to approach the problem in a similar manner as one with a constant (or, more generally, exogenously specified) distortion. I construct the appropriate distorting martingale that accounts for model uncertainty. The marginal distorted measure, integrated out across models, is again absolutely continuous with respect to the true probability measure P. As a result, a modified discount factor can be defined as in (12), and the solution method for the problem (14–16) applies.

Recall that under model k, agent n perceives the trend component of the aggregate endowment process (1) to be $\mu_{y,k}^n = \mu_y + u_k^n \sigma_y$. It is well-known from the literature on Bayesian updating (see Wonham (1964)) that the evolution of the probability distribution across models for a Bayesian learner follows

$$dp_t^n = \Delta\left(p_t^n\right)\left(\frac{dY_t}{Y_t} - \left(\mu_y^n\right)' p_t^n dt\right) \tag{29}$$

where

$$\Delta\left(p_t^n\right) = \left|\sigma_y\right|^{-2}\left(\operatorname{diag}\left(p_t^n\right) - p_t^n\left(p_t^n\right)'\right)\mu_y^n$$

is the regresssion coefficient in the regression of the true state on the evolution of the observed variable under the agent's information set, and $\operatorname{diag}(p)$ is a diagonal matrix with elements of vector p on the main diagonal.

The agent perceives the local trend component of the evolution of Y_t to be $Y_t\left(\mu_y^n\right)' p_t^n$, and thus

$$Y_t - \int_0^t Y_s\left(\mu_y^n\right)' p_s^n ds$$

is a martingale under Q^n. This leads to the construction of a Brownian motion W^n under Q^n defined by

$$dW_t^n \equiv \frac{\frac{dY_t}{Y_t} - \left(\mu_y^n\right)' p_t^n dt}{\sigma_y} = -\left(u^n\right)' p_t^n dt + dW_t$$

The Brownian motion W that is a martingale under the true measure is distorted by the trend component $\left(u^n\right)' p_t^n$ under the subjective measure. The martingale

$$M_t^n = \exp\left(\int_0^t -\frac{1}{2}\left[\left(u^n\right)' p_s^n\right]^2 ds + \int_0^t \left(u^n\right)' p_s^n dW_s\right)$$

is therefore the distorting martingale that replaces (2) in the case of a learning agent. The agent acts as if there was a time-varying average distortion process $\bar{u}_t^n = \left(u^n\right)' p_t^n$. The optimization problem (14–16) is extended by the filtering equation (29) and the evolution of

the modified discount factor becomes

$$\frac{d\bar{\lambda}_t^n}{\bar{\lambda}_t^n} = -\nu_t^n dt + (u^n)' \, p_t^n dW_t$$

Conjecturing a new Markov state

$$Z = \left(\bar{\lambda}', Y, (p^1)', (p^2)'\right)'$$

it is possible to derive a new version of the HJB equation (19) on a multidimensional but compact set with well defined boundary conditions that can be built up sequentially from solutions of lower-dimensional problems. The algorithm is the same as in the case of a multi-agent economy described in Section 2.4.

What impact does the introduction of a learning agent 1 have on the survival problem analyzed in Section 3.4? It is reasonable to expect that both agents survive in the long run for parameterizations satisfying $\gamma < \rho$, at least as long as they survive conditionally on each distorted model u_k^1 that agent 1 has in mind. Under learning, beliefs of agent 1 converge to the true probability distribution, which improves the survival chances of agent 1. As the beliefs converge, the evolution of the process θ^1 settles. The limiting distribution of θ^1 as $t \to \infty$ from which we can deduce the wealth and consumption distribution remains an open question.

When $\gamma > \rho$, the survival problem is more complicated. Under a constant belief distortion of agent 1, each of the agents may dominate the economy with a strictly positive probability. Survival under learning will depend on the relative rates of learning and extinction. Again, this problem is an open question for future research.

5.2 Robust utility

Consider an agent who believes that the model for the aggregate endowment dynamics is misspecified and views (1) only as a reference model that approximates the true dynamics. Anderson, Hansen, and Sargent (2003) and Skiadas (2003), among others, suggest to model the misspecification by modifying the continuation value problem (6) as

$$\lambda_t^n V_t^n = \inf_{u^n, \nu^n} E_t^{Q_u^n} \left[\int_t^\infty \lambda_s^n \left[F\left(C_s^n, \nu_s^n\right) + \frac{\theta_s}{2} \left|u_s^n\right|^2 \right] ds \right]$$

subject to

$$\begin{aligned} \frac{d\lambda_s^n}{\lambda_s^n} &= -\nu_s^n ds, \qquad s \geq t \\ \lambda_0^n &= 1 \end{aligned}$$

The measure Q_u^n is specified by the Radon-Nikodým derivative (2) and the explicit subindex expresses the fact that the minimization problem also includes the choice of the appropriate subjective measure. The set of permissible processes u^n needs to satisfy some regularity conditions like square integrability.

The minimization over u^n expresses agent's fear about the realization of the worst case scenario, characterized by the least favorable dynamics

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

where W^n is a Brownian motion under Q_u^n. At the same time, the agent understands that specifications which are statistically easy to discriminate from the approximate dynamics are unlikely to be correct, and thus large distortions are penalized by the penalty process $\frac{\theta}{2}\left|u^n\right|^2$. Anderson, Hansen, and Sargent (2003) consider a constant θ, while Maenhout (2004) makes θ proportional to the continuation value V to retain homogeneity of the optimization problem.

Except the penalty process in the objective function and the endogenous choice of the distortion process u^n, the calculation of the continuation value is analogous to that introduced in Section 2. Optimal allocations in an economy with two agents endowed with robust preferences are then found by solving an suitably modified planner's problem (14–16).

Under time-separable preferences, agents who fear misspecification more (and therefore assign a lower penalty θ to deviations from the reference model) choose a more distorted worst case scenario, which worsens their survival chances.[7] However, the results for constant belief distortions u indicate that survival chances of the more fearful agents are likely to improve for appropriate non-separable parameterizations of preferences.

This characterization of robust decision-making suggests that it is possible to understand robust preferences merely as a specific belief distortion. Such a simplification however conceals the economic content of the decision-theoretic foundations underlying the robust preferences. Reverting the argument, the framework introduced in this paper can be used to

[7] An exact statement about survival naturally depends on the model and the choice of the process θ for each of the agents.

analyze stationary equilibria in heterogeneous agent economies endowed with a much wider class of preferences than the constant belief distortions that I focused at in Section 3.

6 Conclusion

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

This paper shows that the robust survival result is specific to the class of time-separable preferences. I analyze a two-agent continuous-time economy with iid aggregate endowment growth populated by agents endowed with recursive preferences of the Duffie-Epstein-Zin type. I show that nondegenerate stationary distributions exist for a broad set of parameterizations when risk aversion is larger than the inverse of the intertemporal elasticity of substitution. In this region, the high risk aversion relative to the time-separable case prevents an optimistic (pessimistic) agent from taking excessively large long (short) positions in the asset that pays the aggregate endowment as dividend. In this way, he does not become overly optimistic about the growth of his portfolio which would otherwise lead to excessively high current consumption rate and loss of wealth in the long run. Similarly, a high IES relative to the time-separable case induces the agent to tilt the consumption profile toward the future, and outsave extinction.

An opposite result occurs when risk aversion is lower than the inverse of the IES. When the wealth share of an agent declines, he takes excessively large positions in the risky asset, becomes overly optimistic about the growth rate of his portfolio, increases his consumption beyond a sustainable level and becomes extinct. Since this can happen to either of the agents, each agent can dominate the economy in the long run with a strictly positive probability.

In this sense, the robust survival result from the time-separable case is a threshold scenario in the class of recursive preferences which may only extend to a neighborhood of the time-separable case when the mean consumption share of the distorted agent becomes numerically indistinguishable from zero. In the class of recursive preferences, the market selection hypothesis fails for sufficiently large deviations of the risk aversion coefficient from

the inverse of the IES. Also, the results highlight the different roles of risk aversion and IES in the portfolio allocation and consumption-saving decisions.

The survival results are obtained by extending the solution method of Dumas, Uppal, and Wang (2000) to a setting with heterogeneous beliefs. Long-run survival of the agents is determined by the behavior of a stochastic process that models the Pareto share of one of the agents. I provide tight sufficient conditions that guarantee survival or extinction. Unfortunately, these conditions need to be checked numerically as part of the solution of the model. These conditions can be interpreted as relative patience conditions analogous to those in Lucas and Stokey (1984). An agent survives in the long run if his relative patience becomes sufficiently large as his wealth share vanishes. However, in this framework the relative patience arises endogenously as an equilibrium outcome, and is not a direct property of agents' preferences.

These results are obtained for a two-agent economy with an aggregate endowment process that is specified as a geometric Brownian motion but the theoretical framework can also be utilized to derive an analog HJB equation for multi-agent economies with more sophisticated Markov dynamics. In principle, the qualitative survival results should extend to a wider class of models with stable consumption growth dynamics, although the analysis of the existence of a stationary distribution for the Pareto share becomes more complicated in a multidimensional state space.

Importantly, the developed solution method is not limited to constant distortions and applies to a much wider class of preferences that are interpretable as deviations in beliefs. I outline how to use the method in a framework with model uncertainty and learning, and in a model where agents are endowed with robust preferences. Solutions of these problems are left as open questions for future research.

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

A Epstein-Zin preferences in continuous time

In this section, I show informally how to derive the continuous-time recursive preferences of Kreps and Porteus (1978) as a limiting case of the Epstein and Zin (1989) preferences.

The continuation value process $\tilde{V}$ for an agent with Epstein-Zin preferences is given by

$$\begin{aligned} \tilde{V}_t &= \left[\left(1-e^{-\beta}\right)(C_t)^{\rho} + e^{-\beta}\tilde{\mathcal{R}}_t\left(\tilde{V}_{t+1}\right)^{\rho}\right]^{\frac{1}{\rho}} \\ \tilde{\mathcal{R}}_t\left(\tilde{V}_{t+1}\right) &= \left(E_t\left[\left(\tilde{V}_{t+1}\right)^{\gamma}\right]\right)^{\frac{1}{\gamma}} \end{aligned} \tag{30}$$

where $\rho < 1$ and $\gamma < 1$ are parameters with the same interpretation as in (5). For the sake of simplicity, I omit the situations when $\gamma = 0$ or $\rho = 0$ but these can be treated as appropriate limiting cases. Observe that when $\gamma = \rho$, the preferences reduce to the time-separable CRRA utility.

Specification (30) does not directly yield an aggregator of the form (5) as its limit since the certainty equivalence $\tilde{\mathcal{R}}_t(V_{t+1}) = h^{-1}(E_t[h(V_{t+1})])$ is not linear in V. As shown by Duffie and Epstein (1992b), this nonlinearity leads in continuous-time case to a compensation using a variance multiplier that introduces an additional term to the recursive utility definition (4).

To avoid this issue, it is advantageous to consider an ordinally equivalent transformation of the utility process

$$V_t = \frac{1}{\gamma}\left(\tilde{V}_t\right)^{\gamma} \tag{31}$$

which leads to

$$V_t = \frac{1}{\gamma}\left[\left(1-e^{-\beta}\right)(C_t)^{\rho} + e^{-\beta}\left(\gamma E_t V_{t+1}\right)^{\frac{\rho}{\gamma}}\right]^{\frac{\gamma}{\rho}}. \tag{32}$$

This transformation reduces the certainty equivalence to an expectation[8].

Instead of using a discrete time interval of length one, take a time step of length ε and analyze the limit as $\varepsilon \to 0$. Express $E_t[V_{t+\varepsilon}]$ from (32) to obtain

$$E_t[V_{t+\varepsilon}] = \left[e^{\beta\varepsilon}(V_t)^{\frac{\rho}{\gamma}} - \left(e^{\beta\varepsilon}-1\right)\gamma^{-\frac{\rho}{\gamma}}(C_t)^{\rho}\right]^{\frac{\gamma}{\rho}}.$$

[8]Notice that γ and V will always have the same sign, so that $(\gamma E_t V_{t+1})^{\frac{\rho}{\gamma}}$ is well-defined.

Applying the L'Hospital rule leads to

$$\begin{aligned}
\lim_{\varepsilon \searrow 0} \frac{E_t\left[V_{t+\varepsilon}\right] - V_t}{\varepsilon} &= \lim_{\varepsilon \searrow 0} \frac{\left[e^{\beta\varepsilon}\left(V_t\right)^{\frac{\rho}{\gamma}} - \left(e^{\beta\varepsilon} - 1\right)\gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right]^{\frac{\gamma}{\rho}} - V_t}{\varepsilon} = \\
&= \lim_{\varepsilon \searrow 0} \frac{\gamma}{\rho}\left[e^{\beta\varepsilon}\left(V_t\right)^{\frac{\rho}{\gamma}} - \left(e^{\beta\varepsilon} - 1\right)\gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right]^{\frac{\gamma}{\rho}-1} \cdot \\
&\qquad \cdot \left(\beta e^{\beta\varepsilon}\left(V_t\right)^{\frac{\rho}{\gamma}} - \beta e^{\beta\varepsilon}\gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right) \\
&= \beta\frac{\gamma}{\rho}\left(V_t\right)^{1-\frac{\rho}{\gamma}}\left(\left(V_t\right)^{\frac{\rho}{\gamma}} - \gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right) = \\
&= -\frac{\beta}{\rho}\frac{\left(C_t\right)^{\rho} - \left(\gamma V_t\right)^{\frac{\rho}{\gamma}}}{\left(\gamma V_t\right)^{\frac{\rho}{\gamma}-1}} \equiv \\
&\equiv -f\left(C_t, V_t\right) \qquad (33)
\end{aligned}$$

Integrating this expression over time and taking expectations yields

$$E_t\left[\int_t^{\infty} -f\left(C_s, V_s\right) ds\right] = \lim_{T\to\infty} E_t\left[V_T\right] - V_t$$

which, assuming the transversality condition $\lim_{T\to\infty} E_t\left[V_T\right] = 0$, implies formula (4).

References

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

Anderson, Evan W., 2005. "The dynamics of risk-sensitive allocations." *Journal of Economic Theory*. 125 (2): 93–150.

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

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

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

Chen, Xiaohong, Lars Peter Hansen, and Marine Carrasco, 2010. "Nonlinearity and Temporal Dependence." *Journal of Econometrics*. 155 (2): 155–169.

Cox, John C., and Chi-fu Huang, 1989. "Optimal consumption and portfolio policies when asset prices follow a diffusion process." *Journal of Economic Theory*. 49 (1): 33–83.

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

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

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

Duffie, Darrell, Pierre-Yves Geoffard, and Costis Skiadas, 1994. "Efficient and equilibrium allocations with stochastic differential utility." *Journal of Mathematical Economics*. 23 (2): 133–146.

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., 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.

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

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

Harrison, J. Michael, and David M. Kreps, 1979. "Martingales and arbitrage in multiperiod securities markets." *Journal of Economic Theory*. 20 (3): 381–408.

Isaenko, Sergei, 2008. "The Term Structure of Interest Rate in a Pure Exchange Economy where Investors have Heterogeneous Recursive Preferences." *The Quarterly Review of Economics and Finance*. 48 (3): 457–481.

Karatzas, Ioannis, John P. Lehoczky, and Steven E. Shreve, 1987. "Optimal Portfolio and Consumption Decisions for a "Small Investor" on a Finite Horizon." *SIAM Journal on Control and Optimization*. 25 (6): 1557–1586.

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

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

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

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.

Maenhout, Pascal J., 2004. "Robust Portfolio Rules and Asset Pricing." *The Review of Financial Studies*. 17 (4): 951–983.

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

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.

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

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

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

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

Wonham, W. Murray, 1964. "Some Applications of Stochastic Differential Equations to Optimal Nonlinear Filtering." *SIAM Journal on Control and Optimization*. 2 (3): 347–369.

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

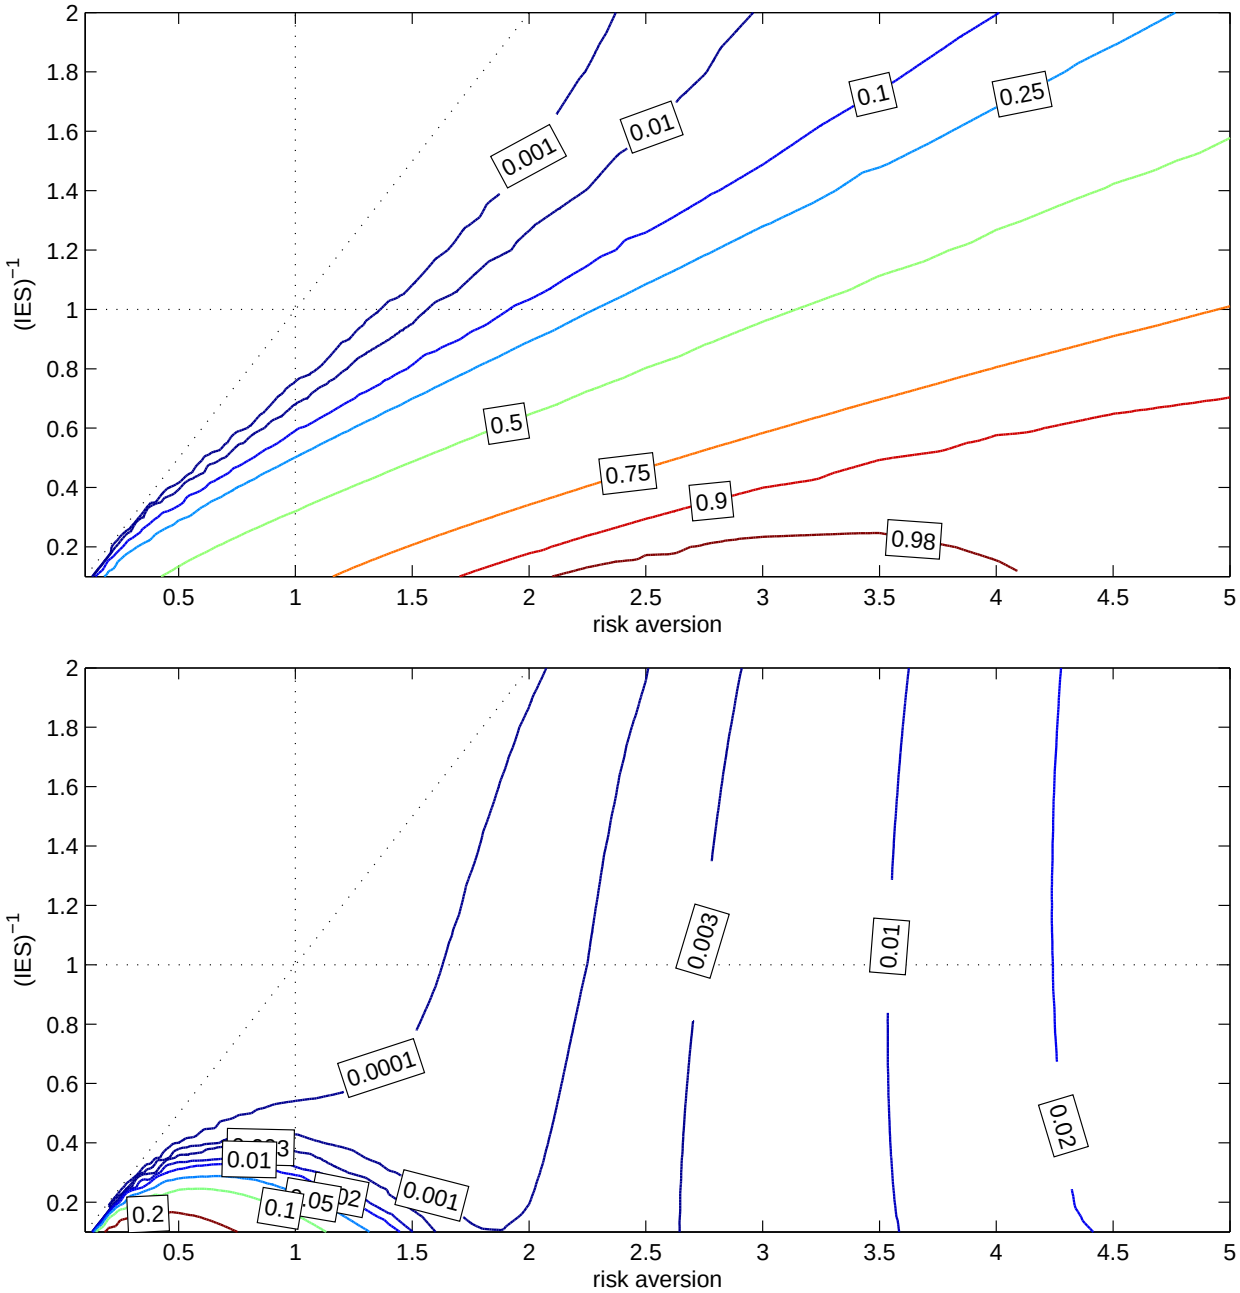

Figure 1: Average consumption share of the agent with distorted beliefs for an optimistic (top panel) and a pessimistic (bottom panel) distortion. The graphs depict the contour lines for $E\left[\zeta_t^1\right]$ under the stationary distribution in the $(1-\gamma, (1-\rho)^{-1})$ space. Belief distortions are $u^1 = 0.05$, $u^2 = 0$ for the top panel, and $u^1 = -0.05$, $u^2 = 0$ for the bottom panel. Remaining parameters are $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.015$.

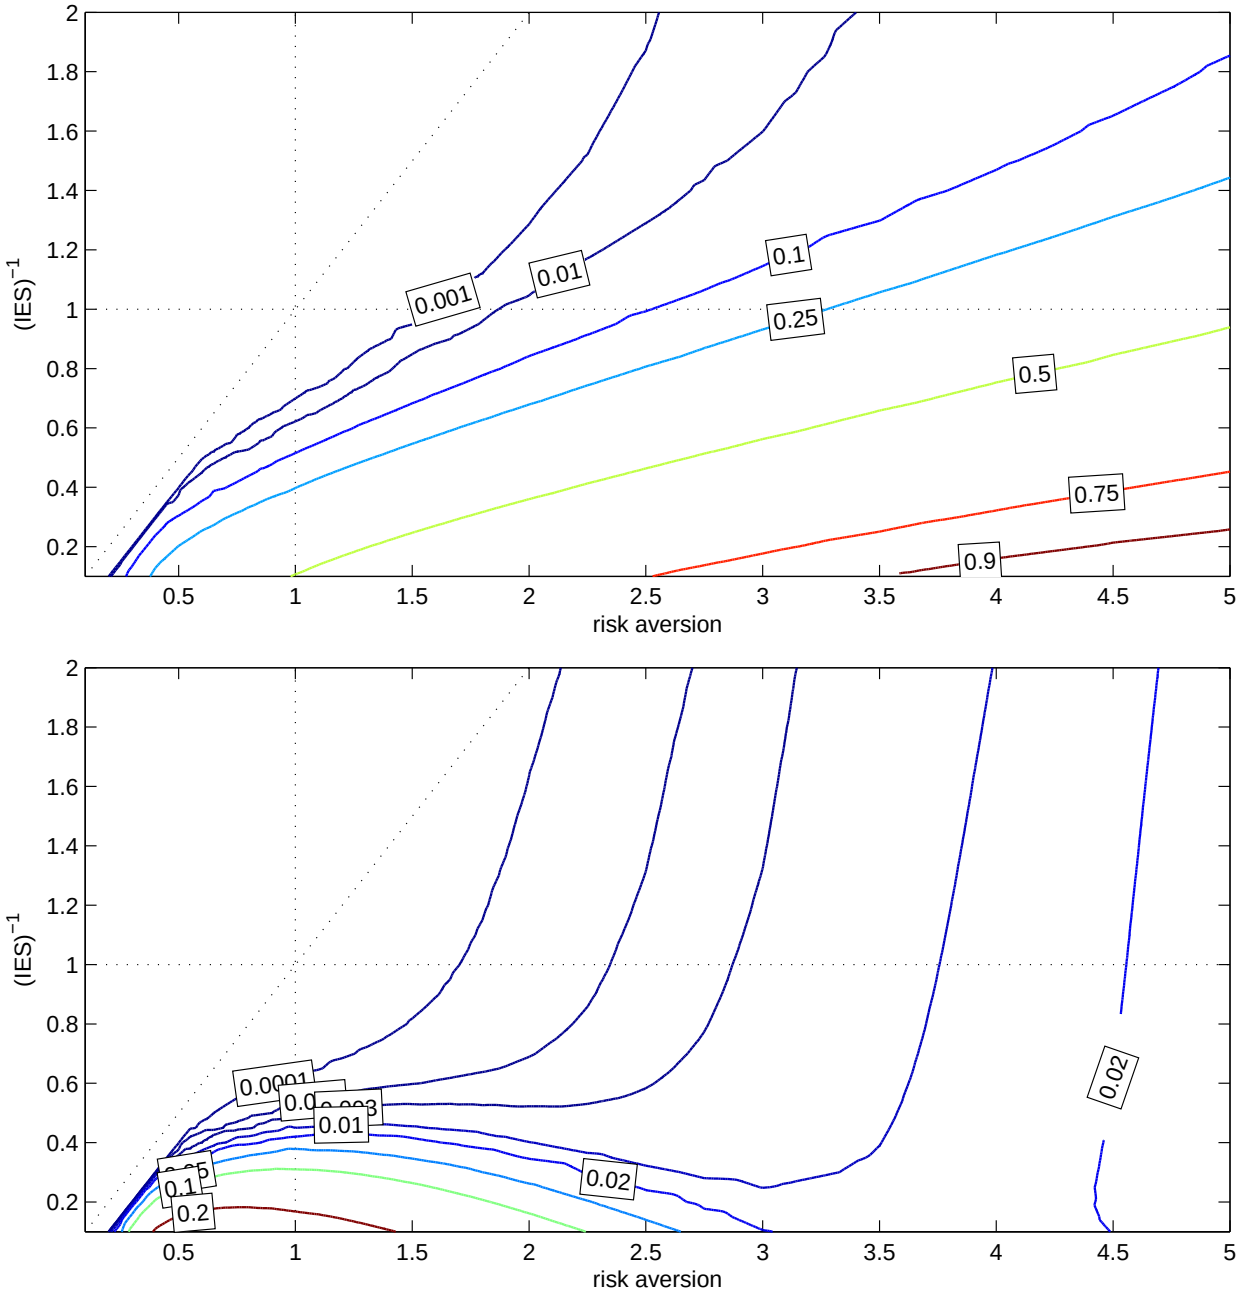

Figure 2: Average consumption share of the agent with distorted beliefs for an optimistic (top panel) and a pessimistic (bottom panel) agent with a larger distortion than in Figure 1. The graphs depict the contour lines for $E\left[\zeta_t^1\right]$ under the stationary distribution in the $(1-\gamma, (1-\rho)^{-1})$ space. Belief distortions are $u^1 = 0.1$, $u^2 = 0$ for the top panel, and $u^1 = -0.1$, $u^2 = 0$ for the bottom panel. Remaining parameters are $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.015$.

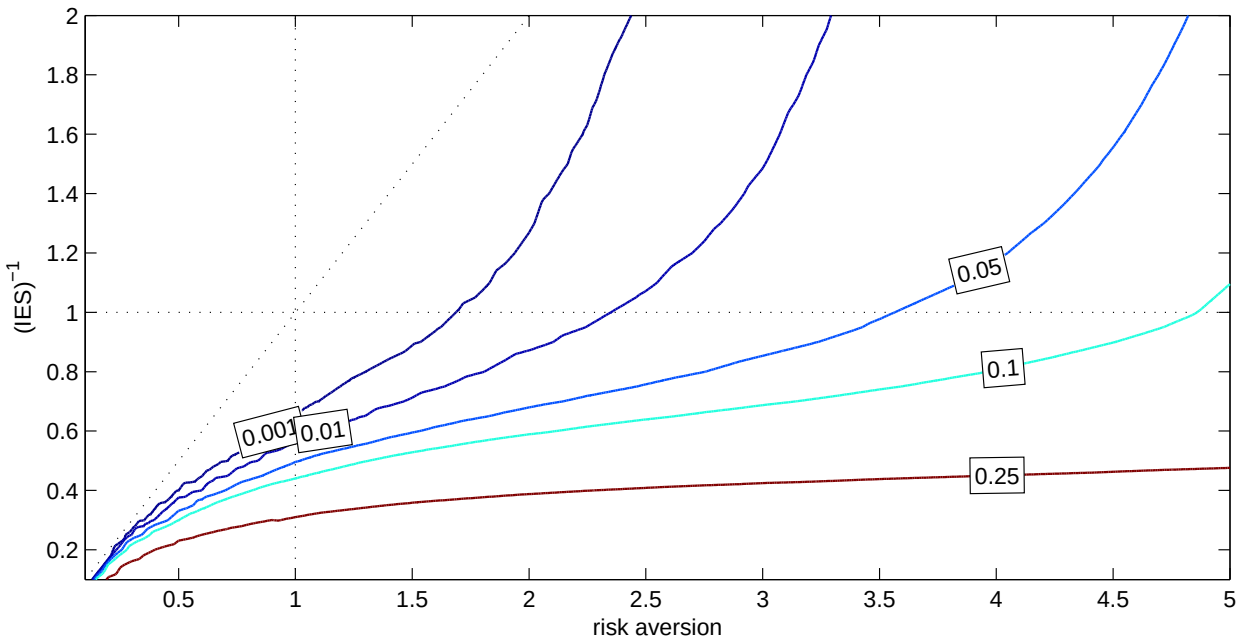

Figure 3: Average consumption share of the agent with distorted beliefs in a model with deterministic growth. The graph depicts the contour lines for $E\left[\zeta_t^1\right]$ under the stationary distribution in the $(1-\gamma, (1-\rho)^{-1})$ space, and is identical for symmetric optimistic and pessimistic distortions $|u^1| = 0.05$, $u^2 = 0$. Remaining parameters are $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0$.

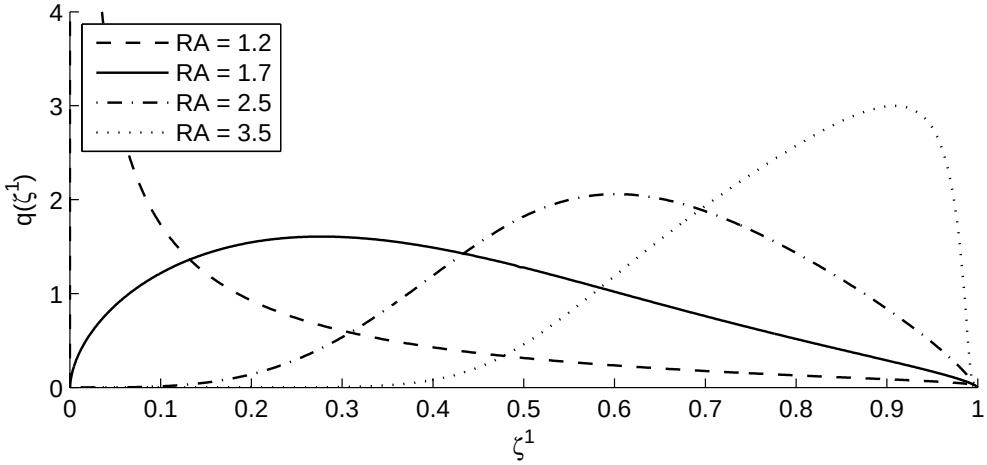

Figure 4: Stationary distributions for the consumption share of the agent with distorted beliefs ζ^1. All models are parameterized by $u^1 = 0.05$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.015$, and differ in levels of risk aversion.

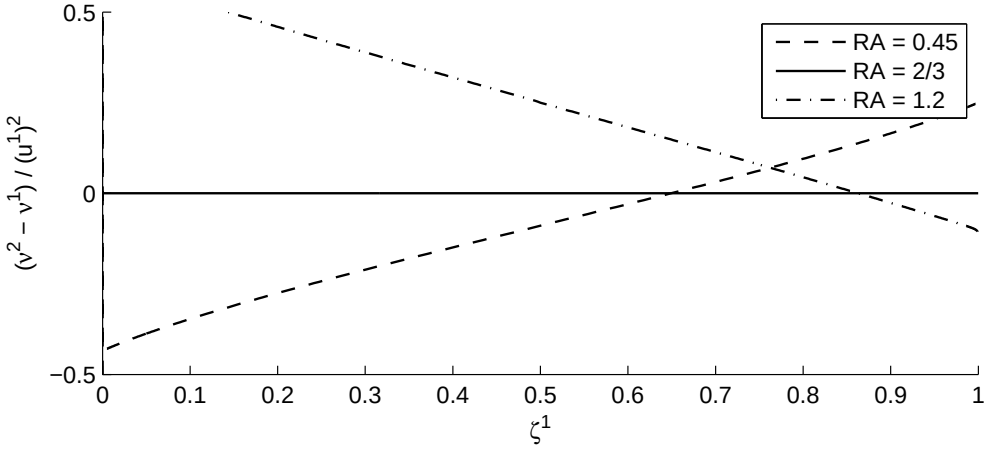

Figure 5: Scaled relative patience $[\nu^2(\theta^1) - \nu^1(\theta^1)] / (u^1)^2$ as a function of the consumption share ζ^1. All models are parameterized by $u^1 = 0.05$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.015$, and differ in levels of risk aversion. The solid line corresponds to the time-separable case.

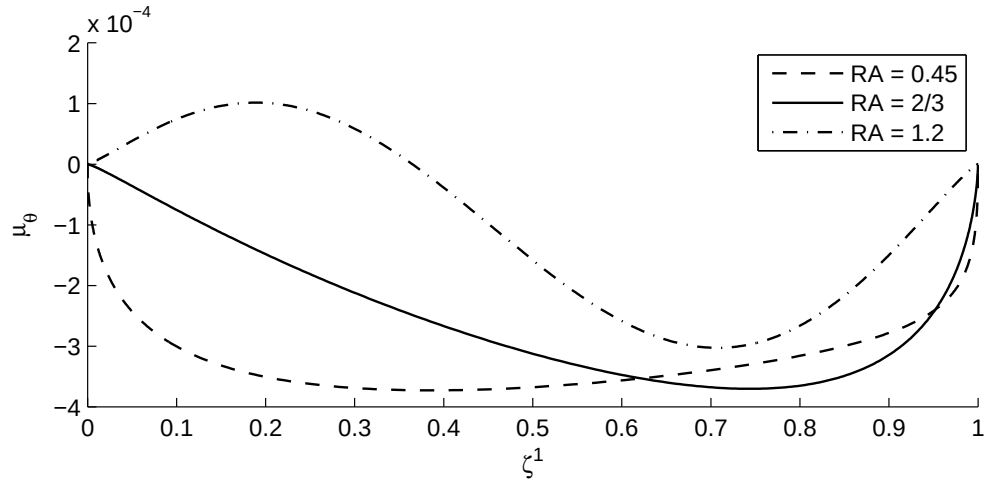

Figure 6: Drift component of the Pareto share evolution $E_t[d\theta_t^1]/dt$ as a function of the consumption share ζ^1. All models are parameterized by $u^1 = 0.05$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.015$, and differ in levels of risk aversion. The solid line corresponds to the time-separable case.

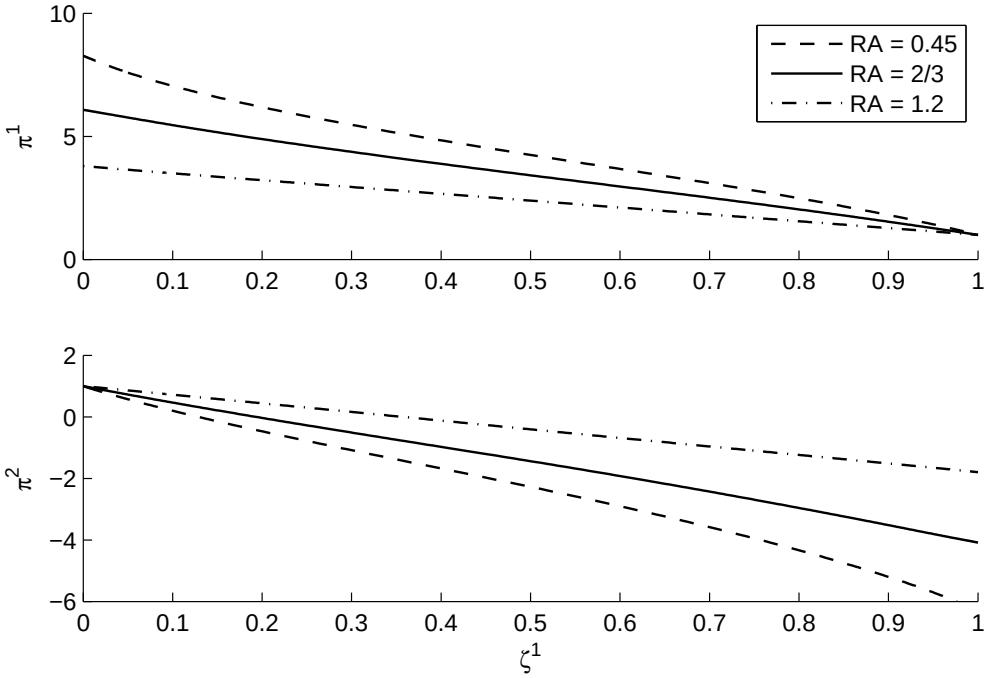

Figure 7: Wealth shares π^n invested in the risky asset as a function of the consumption share ζ^1 of agent 1. All models are parameterized by $u^1 = 0.05$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.015$, and differ in levels of risk aversion. The solid line corresponds to the time-separable case.

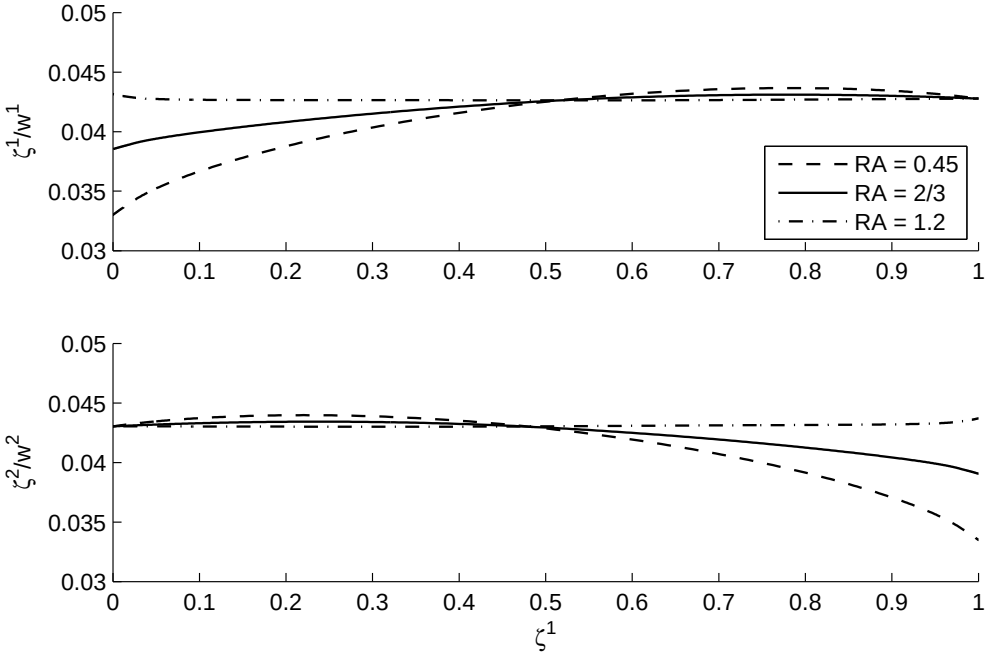

Figure 8: Propensity to consume from wealth ζ^n/w^n a function of the consumption share ζ^1 of agent 1. All models are parameterized by $u^1 = 0.05$, $u^2 = 0$, IES = 1.5, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.015$, and differ in levels of risk aversion. The solid line corresponds to the time-separable case.