Heterogeneous beliefs under recursive preferences*

Jaroslav Borovička

University of Chicago
borovicka@uchicago.edu

April 10, 2009

Abstract

We analyze the impact of heterogeneous beliefs in economies where agents have recursive preferences of the Kreps-Porteus type and differ in their subjective beliefs about the distributions of future quantities. We show that the set of Pareto optimal allocations is given as a solution to a planner's problem with Pareto weights evolving according to an Itô process. In a Markov environment, belief heterogeneity can be incorporated without the need for additional state variables. Using a parsimonious economy, we show that agents with distorted beliefs are not driven out of the market for an empirically relevant range of parameters when preferences are nonseparable across time. Return decomposition reveals nonlinearities in contributions of belief heterogeneity across payoff horizons.

*Work in progress. I am indebted to Lars Peter Hansen for his advice and continuous support. I appreciate helpful comments from John Cochrane, George Constantinides, John Heaton, Katarína Kotovová, Alan Moreira, Harald Uhlig, and the participants of the Economic dynamics working group at the Department of Economics, University of Chicago. All typos and errors are mine.

1 Introduction

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

Survival of irrational agents has been studied in a variety of setups. For instance, De Long, Shleifer, Summers, and Waldmann (1990, 1991) study the survival of traders in a partial equilibrium setup and find that irrational noise traders can outgrow their rational counterparts and dominate the market. On the other hand, Blume and Easley (2006) find that in a general equilibrium model with intermediate consumption, the irrational traders always vanish when markets are complete. In addition, the article also shows that speed of learning matters, and that agents who learn faster have better survival prospects. In a similar manner, Sandroni (2000) also shows in a dynamic bounded economy that agents who make inaccurate predictions are driven out of the market. Contrary to this, Kogan, Ross, Wang, and Westerfield (2006) show that in a model with a growing economy (albeit with a very specific structure), irrational traders can survive even in a dynamically complete market.

Despite a substantial amount of contributions to this field, the message about the survival of irrational traders is mixed. Thus there is no clear answer under what general conditions the market selection hypothesis holds.

Moreover, even when markets are able to weed out irrational traders, it is not straightforward to see how quickly (and if at all) the market prices converge to their rational expectations counterparts when the wealth of the irrational agents converges to zero. Fama and French (2007) provide a qualitative result in their analysis of the CAPM model without complete agreement:

> "When informed investors hold the lion's share of invested wealth, large portfolio distortions by the misinformed can leave the portfolio of the informed close to the market portfolio, and the price effects of misinformed beliefs are minor."

However, quantitative implications or the generality of the statement are not obvious. For instance, Kogan, Ross, Wang, and Westerfield (2006) produce results that show rather striking nonlinearities in the relationship between the wealth of the irrational traders and their impact on prices as their wealth share converges to zero:

> "The price impact of irrational traders does not rely on their long-run survival, and they can have a significant impact on asset prices even when their wealth becomes negligible."

Based on their results, Kogan, Ross, Wang, and Westerfield (2006) argue that survival and price impact are two concepts that are not identical. However, their analysis is cast in a rather specific setup. They consider a finite-time economy with only intermediate consumption and only terminal payoff which is distributed among two agents who different in their subjective beliefs about future distributions. The long-run characteristics are then calculated as limits of sequences of differing finite-time economies. Their model thus constitutes a very specific and a rather unusual case.

In our work, we deviate from the assumption of time-separable preferences, and analyze a class of continuous-time models with agents possesssing nonseparable recurive utility of the Kreps and Porteus (1978) type. In particular, we focus on the continuous-time version of the CES specification used by Epstein and Zin (1989) and Duffie and Epstein (1992b). These preferences can be viewed as a generalization of the constant relative risk aversion utility that allows to decouple the link between the risk aversion with respect to intratemporal gambles and the intertemporal elasticity of substitution.

The agents assess the future distributions of underlying economic variables according to their own subjective beliefs which may differ from the objective probabilities. Since agents in our framework are firm believers in their subjective models and these beliefs are public knowledge, there is no space for learning from prices or observed allocations. Therefore, we can characterize the equilibria by solving a planner's problem in the spirit of Dumas, Uppal, and Wang (2000), and then decentralizing the economy. Dumas, Uppal, and Wang (2000) argue that optimal allocations for agents with heterogeneous recursive preferences can be described in terms of Pareto weights that evolve over time but are locally predictable. We extend their analysis by introducing belief heterogeneity, and show that in this case, the Pareto weights are given by Itô processes with nonzero volatilities.

In a Markov environment, we derive the Hamilton-Jacobi-Bellman equation for the planner's value function which also yields the optimal consumption allocation and the rules for the evolution of the Pareto weights.

Part of the solution of the planner's problem are state-dependent discount rates which the agents use to discount future continuation values. We show that survival in our environment is determined by the behavior of the difference in these discount rates as the share of one of the agents becomes negligible. In the time-separable case, these discount rates are constant and correspond to the rate of time preference, which in turn implies that in an environment where agents differ solely in their beliefs, only agents with least distorted beliefs survive in the long run. When agents face time-nonseparabilities in preferences are present, this does not have to be the case. In particular, we show that when agents are sufficiently risk averse relative to the power utility case (holding the intertemporal elasticity of substitution as given), they may survive even if their beliefs are relatively more distorted. This is the case in an empirically relevant region of parameter values.

We use the stochastic discount factor generated by the allocations obtained from the planner's problem to analyze the term structure of risk prices in the spirit of Hansen (2008). We show that agents with distorted beliefs retain a strong impact on the valuation of cash flows at distant horizons, even if their share in the economy becomes very small. However, horizons at which these effects become quantitatively large, are beyond empirical relevance in our parsimonious model.

1.1 Literature review

The literature that analyzes the survival of nonrational agents in dynamic economies is directly influenced by the market selection hypothesis of Alchian (1950) and Friedman (1953). In a similar manner, Fama (1965) argues that investors with consistently incorrect predictions will be ultimately driven out of the market. Despite the long history that this hypothesis has, the issue of survival in dynamic economies has not been completely resolved, and the known results depend on the type of problem considered. De Long, Shleifer, Summers, and Waldmann (1990, 1991) show that in a partial equilibrium setup with fixed saving rules, nonrational agents may survive, and may even drive out the rational agents. This result was overturned by Blume and Easley (1992) who show that once we allow for endogenous saving rules, the somewhat paradoxical argument made by De Long, Shleifer, Summers, and Waldmann is no longer valid.

The most exhaustive treatment so far is one by Sandroni (2000) who shows that in a bounded economy, the agents whose beliefs are 'closest to the truth' in a certain sense, survive, while other agents are driven out of the market. The quantity that determines survival is the entropy of agents' beliefs. Sandroni (2000) proves that incorrect beliefs play a similar role as agents' impatience, which is also a factor that influences survival. Blume and Easley (2006) show that the survival of agents with correct beliefs is closely linked to

Pareto optimality of the equilibria. For instance, agents with incorrect beliefs can survive in incomplete markets.

Although Sandroni (2000) provides quite a general treatment of agents' survival, his analysis excludes the case of growing or decaying economies, and only focuses on time-separable utility. He focuses directly on equilibrium allocations, and bases his results on intertemporal comparisons of marginal utilities for a given agent. His results require that the period utility functions satisfy Inada conditions while equilibrium ratios marginal utilities remain bounded, which in turn implies the need for the aggregate endowment to be bounded from above and away from zero.

Our analysis is different in spirit. We sidestep a direct characterization of the equilibrium by considering a planner's problem. The primary object of our analysis is thus the intratemporal comparison of marginal utilities across agents, which does not require the boundedness of the aggregate endowment as a precondition. This allows us to naturally incorporate growing or decaying economies.

Moreover, a formulation of the planner's problem in the spirit of Dumas, Uppal, and Wang (2000) allows us to incorporate time-nonseparable preferences. Our primary focus is on the 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. Thanks to this additional degree of flexibility, the recursive preferences of the Kreps-Porteus type were applied widely in the asset pricing literature to provide a better fit of the asset pricing models to empirically observed patterns of asset returns.

Anderson (1998) studies Pareto optimal allocations under heterogeneous recursive preferences in a discrete time setup, while Mazoy (2005) considers the continuous-time environment. Isaenko (2004) analyzes the behavior of the term structure when agents have differently parametrized Kreps-Porteus utilities. However, none of the papers we are aware of 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 provides a theoretical exposition to recursive preferences, derives the planner's problem for an economy with agents equipped with heterogeneous beliefs, and characterizes the optimal allocations. Section 3 provides a discussion of survival and provides a sufficient condition for the existence of a nondegenerate stationary distribution. Section 4 analyzes the asset pricing implications of belief heterogeneity. Section 5 illustrates a way of how to incorporate learning into our framework. 6 concludes.

Detailed proofs of the claims and some auxiliary results are deferred to the Appendix.

2 Optimal allocations under heterogeneous beliefs

We analyze the dynamics of equilibrium allocations and asset prices in a continuous-time endowment economy. Following the literature on market survival[1], the economy is populated by N 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' beliefs about the distribution of future quantities are distorted by a term that can be interpreted as a deviation from rational expectations.

Without introducing any specific market structure, we assume that markets are dynamically complete in the sense of Harrison and Kreps (1979). Moreover, the model structure including the belief distortions is public knowledge, so that there is no space for strategic trading behavior exploiting the differences in beliefs. This, together with standard conditions under which the welfare theorems hold, allows us 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 the planner's problem. Equilibrium allocations and continuation values enable us to construct a valid stochastic discount factor using results from Duffie and Epstein (1992a).

In this section, we characterize agents' preferences and belief distortions, and lay out the planner's problem. We 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. In a Markov environment, we derive the corresponding Hamilton-Jacobi-Bellman equation and provide intuition on how to find optimal allocations in an N-agent economy by inductively solving economies with $1, 2, \ldots, N$ agents. We then set up an example 2-agent economy that will serve as a laboratory for our survival and asset pricing analysis.

2.1 Characterization of recursive utility

Economic agents endowed with time-separable preferences reduce intertemporal compound lotteries (different payoff streams allocated over time) to atemporal simple lotteries that

[1]LITERATURE HERE

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 reduced 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. 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, as is familiar from dynamic programming.

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 which is also the object of our analysis.

We start with a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ with filtration defined by a family of σ-algebras $\{\mathcal{F}_t\}$, $t \geq 0$ generated by a multivariate Brownian motion W. Throughout the text, we will assume that the described processes are adapted to the filtration, and satisfy the usual regularity conditions like square integrability. A detailed characterization of the regularity conditions applicable to our environment can be found in Duffie and Epstein (1992b), Duffie, Geoffard, and Skiadas (1994), Schroder and Skiadas (1999) or Dumas, Uppal, and Wang (2000).

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

$$V_t(C) \equiv E_t\left[\int_t^T f(C_s, V_s)\, ds\right] \qquad V_T = 0 \tag{1}$$

where $T \in [t, +\infty]$ (the limiting case $T \to \infty$ requires some additional regularity conditions, see for instance Appendix C in Duffie and Epstein (1992b)). The function f is the aggregator that links together the consumption C_s at time $s \in [t, T]$ with the continuation value V_s at the same time.

The recursive preference setup includes the case of time-separable utility when $f(c, u) = U(c) - \beta u$ for some function $U(c)$. We are interested in another particular functional form of the aggregator that is the continuous-time counterpart of the constant elasticity of subsitution preferences used in Epstein and Zin (1989, 1991):

$$\begin{aligned} f(c, u) &= \frac{\beta}{\rho} \frac{(c)^{\rho} - (\gamma u)^{\frac{\rho}{\gamma}}}{(\gamma u)^{\frac{\rho}{\gamma}-1}} = \frac{\beta}{\rho} \left[(c)^{\rho} (\gamma u)^{1-\frac{\rho}{\gamma}} - (\gamma u) \right] \\ \gamma, \rho &< 1 \end{aligned} \tag{2}$$

Preferences specified by this form of the aggregator[2] are homothetic, exhibit a constant relative risk aversion with respect to intratemporal wealth gambles $\alpha = 1 - \gamma$, and (under intratemporal certainty) a constant intertemporal elasticity of substitution (IES) $\eta = \frac{1}{1-\rho}$.

In the case when $\gamma = \rho$, the utility reduces to the time-separable power utility case with the coefficient of relative risk aversion α (and IES $\frac{1}{\alpha}$). Preferences specified by (2) thus can be viewed as a natural extension of the power utility case that allows to disentangle the bond between intratemporal risk aversion and intertemporal elasticity of substitution.

This flexibility became one of the reasons

Appendix B shows how to derive the aggregator $f(c, u)$ as the limiting continuous-time case of the discrete-time preferences formulated by Epstein and Zin (1989).

The disadvantage of representation (1) is that the utility process is defined using a recursive formula. Dumas, Uppal, and Wang (2000) show that we can characterize the preferences by defining a conceptually different, yet equivalent, process[3] V with a given felicity function[4] F if there exists a *discount rate process* ν such that (V, ν) solve

$$\inf_{\nu} \{F(C_t, \nu) - \nu V_t\} \tag{3}$$

[2]The cases of $\rho \to 0$ and $\gamma \to 0$ can be obtained as limiting cases with an ordinally equivalent aggregator in which we replace γu with $1 + \gamma u$.

[3]Dumas, Uppal, and Wang (2000) call this process V the stochastic variational utility process. Since we will operate under assumptions under which the stochastic variational utility and recursive utility processes are equivalent, we will not make a formal terminological distinction between the two.

[4]This concept was first applied by Geoffard (1996) in the deterministic environment, and subsequently by El Karoui, Peng, and Quenez (1997) in the stochastic case.

and

$$V_t = E_t \left\{ \int_t^T e^{\int_t^s -\nu_\tau d\tau} F(C_s, \nu_s)\, ds \right\} \tag{4}$$

Processes $V(C)$ and V coincide if the felicity function F and the aggregator f are linked through the Legendre transformation

$$F(c,\nu) \equiv \inf_{u \in R} \left[f(c,u) + \nu u \right] \tag{5}$$

$$f(c,u) \equiv \sup_{\nu \in R} \left[F(c,\nu) - u\nu \right] \tag{6}$$

The transformation (5) and (6) holds when f is convex in its second argument[5]. 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 (2), taking the first order condition in (5) 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 (4) suggests that we can view F as a generalization of the period utility function with a potentially stochastic rate of time discount ν that needs to be determined endogenously.

Lemma 1.1 in Dumas, Uppal, and Wang (2000) shows that for a given consumption process C, the pair (V,ν) solves the problem (3, 4) if and only if it solves the problem

$$\lambda_t V_t = \inf_{\nu} E_t \left[\int_t^T \lambda_s F(C_s, \nu_s)\, ds \right] \tag{7}$$

subject to

$$\begin{aligned} \frac{d\lambda_s}{\lambda_s} &= -\nu_s ds \\ \lambda_0 &= 1 \end{aligned} \tag{8}$$

where λ is called the *discount factor process.* In what follows, we will show how to interpret the discount factor processes λ for individual agents as processes characterizing the Pareto weights of the agents in a planner's problem, and how to extend them to incorporate our notion of belief distortions.

[5]This case, corresponding to preference for early resolution of uncertainty, is the empirically more relevant case in asset pricing applications. It roughly corresponds to increasing the risk aversion while holding the intertemporal elasticity of substitution fixed.

2.2 Heterogeneous beliefs

The rational expectations hypothesis relies on the fact that economic agents know the probability model that underlies the dynamics of the economy, and thus the expectation in (7) that represents agents' preferences over streams of current and future consumption is calculated using the true probability measure. The purpose of this work is to assess the impact of deviations from this hypothesis in an environment when time-nonseparabilities in preferences matter.

Consider again the recursive utility process V, this time for an agent who evaluates the probability distributions of future events using a different (subjective) probability measure Q

$$\lambda_t V_t = \inf_{\nu} E_t^Q \left[\int_t^T \lambda_s F\left(C_s, \nu_s\right) ds \right] \tag{9}$$

We assume that the measure Q representing the subjective beliefs of the agent is absolutely continuous with respect to the true measure P at every given time s. Then we can represent the expectation in (9) as

$$\lambda_t V_t \equiv \inf_{\nu} E_t \left[\int_t^T \frac{M_s}{M_t} \lambda_s F\left(C_s, \nu_s\right) ds \right] \tag{10}$$

where M is a martingale that represents the Radon-Nikodym derivative of the measure Q with respect to the measure P. This time, the expectation is evaluated using the true probability measure.

Having in mind the application of the Girsanov theorem, and in line with Kogan, Ross, Wang, and Westerfield (2006), we model the belief distortion by taking M to be of the form

$$M_t = \exp\left(-\frac{1}{2} \int_0^t |u_s|^2 \, ds + \int_0^t u_s dW_s \right) \tag{11}$$

where u is an adapted process.

It is important to emphasize that the structure of the beliefs is assumed to be common knowledge, so that there is no space for potential strategic behavior that would exploit information asymmetries. In fact, formula (10) suggests that we can view the belief distortion as a modification of preferences or a preference shock. One way to proceed thus would be to redefine the felicity function F by incorporating the martingale process M, and consider an economy with 'heterogeneous preferences'. The interpretation of the martingale process M as a belief distortion is more appealing, though. The Girsanov theorem implies that an agent whose beliefs are distorted by M views the evolution of the Brownian motion W as having a trend component u

$$dW_t = u_t dt + d\tilde{W}_t$$

where $\tilde{W}$ is a Brownian motion under Q. In this way, u captures the degree of optimism or pessimism about W. While an interpretation of M as a preference shock is correct, understanding M as a belief distortion is more illuminative a provides an interpretable structure in the separation of beliefs and preferences.

We interpret a constant u as the case when the agent is convinced that her view of the world is the correct one. While a likelihood evaluation of the past data will reveal that her view of the world becomes less and less likely to be correct during time, absolute continuity of the measure Q with respect to P implies that she can never refute her view of the world as impossible.

In Section 5, we show how to use this framework to incorporate learning. In particular, we will consider a situation where the agent considers two models to which he assigns nonzero probabilities of being correct — the true model, and a distorted model with $u \neq 0$. The agent learns by observing the realizations of the stochastic process underlying the economy, and updates the probability distribution across the two models. Effectively, she acts as if there was a time-varying 'average distortion' process $\bar{u}$ the evolution of which we will be able to characterize. We thus stay in the same analytical framework even if we incorporate learning and endogenize the distortion.

In light of equation (9), we extend the specification of the discount factor process λ.

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

$$\bar{\lambda} = \lambda M. \tag{12}$$

Applying Itô's lemma to (12) yields

$$\frac{d\bar{\lambda}_t}{\bar{\lambda}_t} = -\nu_t dt + u_t dW_t. \tag{13}$$

Notice that the original discount factor process (8) was locally smooth. The belief distortion introduces a diffusion term to the modified discount factor process, but otherwise leaves the structure of the problem unchanged.

The diffusion term $u_t dW_t$ has an intuitive interpretation. Consider an optimistic agent with $u_t > 0$. This agent's beliefs are distorted in that the mass of the distribution of dW_t is shifted to the right — the agent effectively overweighs good realizations of dW_t. Formula (13) indicates that under the true probability measure, positive realizations of dW_t increase the term $d\bar{\lambda}_t/\bar{\lambda}_t$ which implies that the optimistic agent discounts positive realizations of dW_t less than the negative ones. From the position 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.

2.3 Planner's problem and optimal allocations

We now consider an economy populated by N agents who rank their consumption streams using (possibly different) recursive preferences represented by aggregators f^n, or, equivalently, by felicity functions F^n. Given our assumptions on market and information structure, welfare theorems hold and we can identify competitive equilibria with solutions to a planner's problem, as in Dumas, Uppal, and Wang (2000). A decentralization of the planner's problem then provides us with market prices.

The objective of the fictitious planner is to choose individual consumption streams C^n in order to maximize 'aggregate welfare'

$$\sum_{n=1}^{N} \alpha^n V_0^n \left(C^n\right) \tag{14}$$

subject to the feasibility constraint

$$\sum_{n=1}^{N} C^n \leq e$$

where e is the total endowment process, V^n are the continuation values of individual agents evaluated under the subjective beliefs (9), and $\alpha = \left(\alpha^1, \ldots, \alpha^n\right)$ are the weights assigned to each of the agents that sum up to one.

The use of the fictitious planner is motivated purely by the link between competitive equilibria and solutions to the planner's problem. We do not intend to claim that the solutions to the planner's problem are truly optimal in the sense of welfare maximization. The discussion of welfare questions in our model is necessarily tainted by issues concerning the proper way of evaluating the utilities of agents with distorted beliefs by the benevolent planner who is aware of these distortions. These questions are far outside the scope of this article. We merely conclude that the planner's problem involving the subjective continuation values (9) is the right one to obtain competitive equilibria in our economy.

The approach in formulating the planner's problem that we take here closely follows

Dumas, Uppal, and Wang (2000). We extend their analysis by considering agents with explicitly modeled belief distortions. The following Proposition replicates their Theorem 2.1 for the case of heterogeneous beliefs.

Proposition 2 *Let f^n be the generators of the recursive utilities for agents $n = 1, \ldots, N$, and F^n the corresponding felicity functions. Let Q^n be the subjective probability measures of the agents, and*

$$M_t^n = \left(\frac{dQ^n}{dP}\right)_t$$

the martingale characterizing the Radon-Nikodym derivative of measure P^n with respect to the true probability measure P. Then the solution of the planner's problem with Pareto weights $\alpha = \left(\alpha^1, \ldots, \alpha^n\right)$ is given by

$$J_t = \sup_C \hat{J}_t\left(C\right) = \sup_C \inf_\nu \sum_{n=1}^N E_t\left(\int_t^T \bar{\lambda}_s^n F^n\left(C_s^n, \nu_s^n\right) ds\right) \tag{15}$$

at time 0 *subject to*

$$\begin{aligned} \frac{d\bar{\lambda}_t^n}{\bar{\lambda}_t^n} &= -\nu_t^n dt + u_t^n dW_t \\ \bar{\lambda}_0^n &= \alpha^n \qquad\qquad\qquad n = 1, \ldots, N \end{aligned} \tag{16}$$

and the supremum is taken over all consumption allocation processes $C = \left(C^1, \ldots, C^N\right)$ that satisfy the feasibility constraint

$$\sum_{n=1}^N C^n \leq e$$

where e is the aggregate endowment process.

The proof essentially follows from Lemma 1.1 and Theorem 2.1 in Dumas, Uppal, and Wang (2000), and is omitted here. The difference lies in the use of the modified discount factor which is not locally smooth.

The methodological advantage of this approach to solving the model with belief heterogeneity is twofold. First, even with the introduction of belief distortions, we stay in the class of recursive utilities. Second, the redefinition of the discount factor processes keeps 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. We will see in short that we can do the same with the modified discount factor processes $\bar{\lambda}^n$. Belief distortions thus do not introduce any additional state variables into the problem, as long as the distorting processes u^n are functions of the existing state variables[6].

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. 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 Deriving the Bellman equation in a Markov economy

We now make the previous theoretical derivation more specific. In particular, we will assume a Markov economy with an underlying state X_t whose evolution is described by an K-dimensional Itô process

$$dX_t = \mu_x\left(X_t, t\right) + \sigma_x\left(X_t, t\right) dW_t$$

The process X_t can determine the behavior of the level of aggregate endowment process $e_t = e\left(X_t\right)$ or its local evolution

$$\frac{de_t}{e_t} = \mu_e\left(X_t\right) dt + \sigma_e\left(X_t\right) dW_t$$

in which case e_t becomes another state variable (which we, with a slight abuse of notation, implicitly merge with the vector X_t).

An introspection of the problem (15) indicates that the welfare J will be a function of the vector $\left(X, \bar{\lambda}, t\right)$. Compared with the homogeneous beliefs case, the belief heterogeneity does not introduce any new state variables. The problem now reads

$$J\left(X, \bar{\lambda}, t\right) \equiv \sup_C \inf_\nu \sum_{n=1}^N E_t\left(\int_t^T \bar{\lambda}_s^n F^n\left(C_s^n, \nu_s^n\right) ds\right) \tag{17}$$

[6]This is not the case when we introduce learning, but we will still offer a tractable recursive formulation even in that case.

subject to

$$\begin{aligned}
\sum_{n=1}^{N} C_s^n &\leq e\left(X_s\right) \\
\frac{d\bar{\lambda}_t^n}{\bar{\lambda}_t^n} &= -\nu_t^n dt + u_t^n dW_t \\
\bar{\lambda}_t^n &= \bar{\lambda}^n \qquad\qquad n = 1, \ldots, N \\
J\left(x', \bar{\lambda}', T\right) &= 0 \qquad\qquad \forall x', \bar{\lambda}'
\end{aligned}$$

where the last condition can, under regularity conditions, be replaced by a transversality condition if $T = +\infty$.

An application of the Itô's lemma on equation (17) leads to the Hamilton-Jacobi-Bellman equation (omitting the arguments and time indices)

$$0 \equiv \sup_{C} \inf_{\nu} \sum_{n=1}^{N} \bar{\lambda}^n \left[F^n\left(C^n, \nu^n\right) - J_{\bar{\lambda}^n} \nu^n\right] + J_t + J_x \cdot \mu_x + \frac{1}{2} \operatorname{tr}\left(J_{zz} \Sigma\right) \tag{18}$$

where

$$\begin{aligned}
Z &= \begin{pmatrix} \bar{\lambda} \\ X \end{pmatrix} \\
\Sigma &= \begin{pmatrix} \left(\operatorname{diag}\left(\bar{\lambda}\right) u\right)\left(\operatorname{diag}\left(\bar{\lambda}\right) u\right)' & \left(\operatorname{diag}\left(\bar{\lambda}\right) u\right) \sigma_x' \\ \sigma_x \left(\operatorname{diag}\left(\bar{\lambda}\right) u\right)' & \sigma_x \sigma_x' \end{pmatrix}
\end{aligned}$$

where $\operatorname{diag}\left(\bar{\lambda}\right)$ is an $N \times N$ diagonal matrix with elements of $\bar{\lambda}$ on the diagonal, and $\left(\operatorname{diag}\left(\bar{\lambda}\right) u\right)$ thus represents the local volatility matrix of the process $\bar{\lambda}$.

As we have already indicated before, compared to the homogeneous beliefs case, we yield extra terms in the second derivative matrix J_{zz}, in particular the derivative submatrices $J_{\bar{\lambda}\bar{\lambda}}$ and $J_{\bar{\lambda}x}$, which come from the fact that $\bar{\lambda}$ is not locally predictable.

The inf problem in the Bellman equation (18) can be solved for separately in the same manner as in Dumas, Uppal, and Wang (2000) by recalling the Legendre transformation (6). Thus under the optimal discount rate process ν^n, we have

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

and then (18) can be written as

$$0 \equiv \sup_{C} \sum_{n=1}^{N} \bar{\lambda}^n f^n \left(C^n, J_{\bar{\lambda}^n}\right) + J_t + J_x \cdot \mu_x + \frac{1}{2} \mathrm{tr} \left(J_{zz} \Sigma\right) \tag{19}$$

which is a nonlinear second order partial differential equation that can be generally solved for numerically.

If we consider the infinite lifetime case with $T \to \infty$ and further assume that the utilities, the dynamics of the economy, and the belief distortions are time invariant, then the Bellman equation (19) is time invariant, which implies $J_t = 0$. This is the case we will deal with below.

2.5 Solution of the model

In this section, we sketch the approach for solving the model. We derive the optimal intratemporal allocations, show a transformation of variables that leads to a PDE defined on the interval $[0, 1]$ in the dimensions of the state space corresponding to the modified discount factors $\bar{\lambda}$, and formulate some boundary conditions.

2.5.1 Optimal intratemporal consumption allocation

Recalling the functional form for $f^n \left(C^n, J_{\bar{\lambda}^n}\right)$ from (2), we can solve the sup problem in (19) subject to the feasibility constraint

$$\sum_{n=1}^{N} C^n \leq e\left(X\right).$$

Claim 3 *If all agents have the same intertemporal elasticity of substitution, i.e.* $\rho^n = \rho$ *for* $n = 1, \ldots, N$*, then the optimal intratemporal consumption allocation is given by*

$$C^{n*} = e\left(X\right) \frac{\left(\gamma^n J_{\bar{\lambda}^n}\right)^{\frac{\rho/\gamma^n - 1}{\rho - 1}} \left(\bar{\lambda}^n \beta^n\right)^{\frac{1}{1-\rho}}}{\sum_{n=1}^{N} \left(\gamma^n J_{\bar{\lambda}^n}\right)^{\frac{\rho/\gamma^n - 1}{\rho - 1}} \left(\bar{\lambda}^n \beta^n\right)^{\frac{1}{1-\rho}}}. \tag{20}$$

The proof is straightforward, and the algebraic details are contained in Appendix A. Observe that this is in fact not a closed-form solution since it depends on the yet unknown derivatives of the value function $J_{\bar{\lambda}^n}$. In what follows, we omit the star superscript in cases where it is obvious that we mean the optimal level of consumption.

Further notice that the Bellman equation (19) with the optimal consumption allocation (20) permits a solution that is homogeneous of degree 1 (HD1) in the vector $\bar{\lambda}$. Under this guess, $J_{\bar{\lambda}^n}$ is HD0, so that the sum in (19) is HD1, and so are the partial derivatives J_t and J_x. Also notice that the product $J_{zz}\Sigma$ will be HD1, and this property is carried over by the trace operator.

By Euler's theorem, the homogeneity property allows us to write the solution as

$$J\left(X, \bar{\lambda}, t\right) = \sum_{n=1}^{N} \bar{\lambda}^n J_{\bar{\lambda}^n}\left(X, \bar{\lambda}, t\right) \tag{21}$$

Equation (21) allows us to identify the partial derivatives $J_{\bar{\lambda}^n}$ with the scaled continuation values V^n of the individual agents.

2.5.2 Transformation of variables

Since closed form solutions to the PDE (19) are in general not available, a numerical solution is required. In the subsequent numerical application, we use the finite difference method — we overlay the state space in variables in $\left(X, \bar{\lambda}\right)$ with a discrete grid, and approximate the partial derivatives in (19) by differences calculated in the grid nodes, thus yielding a system of algebraic equations. To increase the efficiency of the algorithm, we transform the Bellman equation into one in Pareto weights, rather than in discount factors $\bar{\lambda}$.

One of the disadvantages of the given functional form of the value function for the numerical solution is the fact that the discount factors $\bar{\lambda}$ are not bounded from above. Since we are not interested in the particular scale of the discount factors $\bar{\lambda}$, we can define the following transformation of variables $\theta = g\left(\bar{\lambda}\right)$

$$\begin{aligned} \theta^n &= g^n\left(\bar{\lambda}\right) = \frac{\bar{\lambda}^n}{\sum_{j=1}^{N} \bar{\lambda}^j} \qquad n = 1, \ldots, N-1 \\ \theta^N &= g^N\left(\bar{\lambda}\right) = \sum_{j=1}^{N} \bar{\lambda}^j \end{aligned} \tag{22}$$

and subsequently define

$$J\left(X, \bar{\lambda}, t\right) \equiv \tilde{J}\left(X, \theta, t\right).$$

Observe that since J is HD1 in $\bar{\lambda}$, then $\tilde{J}$ is HD1 in θ^N. Thus the variable θ^N serves as a scaling factor that can be used to rescale the Pareto weights, and its level is inconsequential for consumption allocation rules or asset pricing. In our numerical solution, we will fix

$\theta^N = 1$.

Further, the transformation allows us to interpret the scaled discount factors $\theta^1, \ldots, \theta^{N-1}$ as Pareto weights for agents $1, \ldots, N-1$ (with agent N having Pareto weight $1 - \sum_{n=1}^{N-1} \theta^n$). Thus, instead an N-dimensional vector of variables $\bar{\lambda} \in (\mathbb{R}^+)^N$, we have $N-1$ variables $\left(\theta^1, \ldots, \theta^{N-1}\right) \in [0,1]^{N-1}$. Using the above transformation, we reduced the dimensionality of the problem by one dimension, and compactified a subspace of the state space.

Claim 4 *Using the transformation of variables (22), the partial differential equation (19) becomes*

$$\begin{aligned} 0 \;=\; & \theta^N \sum_{n=1}^{N} \Theta_{nn} f^n \left(C^{n*}, \frac{\partial g'}{\partial \bar{\lambda}^n} \tilde{J}_\theta \right) + \tilde{J}_t + \tilde{J}_x \cdot \mu_x + \\ & + \frac{1}{2} \mathrm{tr} \left(\left(\frac{\partial g'}{\partial \bar{\lambda}} \tilde{J}_{\theta\theta} \frac{\partial g}{\partial \bar{\lambda}'} + \mathcal{D}\left(g, \tilde{J}_\theta\right) \right) \left(\theta^N \Theta u\right) \left(\theta^N \Theta u\right)' \right) + \\ & + \mathrm{tr} \left(\left(\frac{\partial g'}{\partial \bar{\lambda}} \tilde{J}_{\theta x} \right) \sigma_x \left(\theta^N \Theta u\right)' \right) + \frac{1}{2} \mathrm{tr} \left(\tilde{J}_{xx} \sigma_x \sigma_x' \right) \end{aligned} \tag{23}$$

where

$$\Theta \equiv \begin{bmatrix} \theta^1 & 0 & \ldots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & \theta^{N-1} & 0 \\ 0 & \ldots & 0 & 1 - \sum_{n=1}^{N-1} \theta^n \end{bmatrix}.$$

The algebraic derivation is contained in Appendix A. Notice that the homogeneity properties of partial derivatives imply that we can indeed use the scaling $\theta^N = 1$ without loss of generality.

2.5.3 Boundary conditions

Let us focus now specifically on the time-invariant problem. The boundary conditions in X will depend on the particular form of the dynamic evolution that we prescribe. Often, these will take the form of derivative conditions. In general, the region for X can be unbounded, so we instead limit ourselves to a 'large enough' bounded region (often a rectangle)

$$\mathcal{X} \subset \mathbb{R}^K$$

and impose boundary conditions at the boundary $\partial \mathcal{X}$ as if they were imposed at infinity. Our hope is that choosing the approximating boundary conditions that are 'far enough' will be numerically inconsequential for the region for which we are interested in the solution.

The boundary conditions in the θ subspace can be determined iteratively. This conclusion is based on the following claim.

Claim 5 *The objective function J of the social planner is continuous at $\theta^i = 0$ for all $i = 1, \ldots, N-1$.*

As we have argued, we impose $\theta^N = 1$. Let us illustrate the approach on a 2-agent economy. The boundary conditions at $\theta^1 = 0$ and $\theta^1 = 1$ are determined as solutions to representative agent economies with either agent 1 or agent 2. These correspond to solving the (time-invariant) partial differential equations

$$0 \equiv \bar{\lambda}^n f^n \left(C, J_{\bar{\lambda}^n}\right) + J_x \cdot \mu_x + \frac{1}{2}\mathrm{tr}\left(J_{zz}\Sigma\right) \qquad n = 1, 2 \tag{24}$$

where

$$\begin{aligned} \frac{1}{2}\mathrm{tr}\left(J_{zz}\Sigma\right) &= \frac{1}{2}\mathrm{tr}\left(J_{xx}\sigma_x\sigma_x'\right) + \mathrm{tr}\left(J_{\bar{\lambda}^n x}\sigma_x\bar{\lambda}^n\left(u^n\right)'\right) + \\ &\quad + \frac{1}{2}\mathrm{tr}\left(J_{\bar{\lambda}^n\bar{\lambda}^n}\bar{\lambda}^n\left(u^n\right)\bar{\lambda}^n\left(u^n\right)'\right) \end{aligned} \tag{25}$$

However, notice that J_x is HD1 and thus, at the boundaries

$$J_x = \sum_{j=1}^{2} \bar{\lambda}^j J_{\bar{\lambda}^j x} = \bar{\lambda}^n J_{\bar{\lambda}^n x} = J_{\bar{\lambda}^n x}$$

since $\bar{\lambda}^n = 1$ and $\bar{\lambda}^{-n} = 0$, and also

$$\begin{aligned} J &= \sum_{j=1}^{2} \bar{\lambda}^j J_{\bar{\lambda}^j} \\ J_{\bar{\lambda}^n} &= \frac{1}{\bar{\lambda}^n}\left(J - \bar{\lambda}^{-n} J_{\bar{\lambda}^{-n}}\right) \\ J_{\bar{\lambda}^n\bar{\lambda}^n} &= -\frac{1}{\left(\bar{\lambda}^n\right)^2}\left(J - \bar{\lambda}^{-n} J_{\bar{\lambda}^{-n}}\right) + \frac{1}{\bar{\lambda}^n}\left(J_{\bar{\lambda}^n} - \bar{\lambda}^{-n} J_{\bar{\lambda}^{-n}\bar{\lambda}^n}\right) = \\ &= -J + J_{\bar{\lambda}^n} = 0 \end{aligned}$$

This simplifies the terms in (25), and we are able to rewrite the partial differential equations in (24) as

$$0 \equiv \bar{\lambda}^n f^n\left(C, J\right) + J_x \cdot \mu_x + \frac{1}{2}\mathrm{tr}\left(J_{xx}\sigma_x\sigma_x'\right) + J_x\sigma_x\left(u^n\right)' \qquad n = 1, 2$$

Solutions to these one-agent economies give us boundary conditions for the two-agent economy at $\theta^1 = 0$ and $\theta^1 = 1$. Iteratively, solutions to all economies with the number of agents up to $(N-1)$ constructed by leaving out all possible subsets of agents from the N-agent economy determine the boundary conditions in the θ subspace of the N-agent economy.

2.6 A two-agent economy with iid consumption growth

In our applications, we primarily focus on a two-agent economy with iid endowment growth

$$\frac{de_t}{e_t} = \mu_e dt + \sigma_e dW_t \tag{26}$$

and constant belief distortions $u_t^n \equiv u^n$, $n = 1, 2$. This evolution of the aggregate endowment corresponds to the dynamics studied in Kogan, Ross, Wang, and Westerfield (2006), although their economy has a substantially different structure which we discuss in more detail below. In order to focus our analysis on the differences in beliefs, we assume that the aggregators f^n (and, consequently, the felicity functions F^n) are the same for all agents.

We will argue that survival only depends on the belief distortions through the difference $u^2 - u^1$, while the asset pricing results can be calculated relative to beliefs with an arbitrary distortion. We therefore make agent 2 rational, $u^2 = 0$, without loss of generality.

While in our numerical calculations we assume that μ_e and σ_e are constant, they can in general depend on the state variable X as well. We expect that at least as long as the dynamics of X is stable, making μ_e and σ_e state-dependent will not have an impact on our conclusions about survival. Strictly speaking, our numerical results treat only the constant coefficient case (26), though.

With the dynamics of the aggregate endowment specifed by (26), we can guess the form of the solution to the PDE (23) as

$$J = \theta^2 e^{\gamma} \tilde{J}\left(\theta^1\right)$$

which allows us to divide the PDE by $\theta^2 e^{\gamma}$, and reduce the problem to a nonlinear ordinary differential equation for the unknown normalized value function $\tilde{J}\left(\theta^1\right)$. Given the independence of the scale of aggregate consumption, we will characterize consumption allocations in terms of the consumption shares

$$\zeta^i\left(\theta^1\right) = \frac{C^i\left(\theta^1\right)}{e}.$$

The boundary conditions at $\theta^1 = 0, 1$ are given by the normalized value functions $\tilde{J}^{1,i}$ in the single-agent economies. The value functions for the single-agent economies can be in this case solved for analytically by guessing the solution $J^{1,n} = \bar{\lambda} e^{\gamma} \tilde{J}^{1,n}$ to the single-agent version of the PDE (24), and then calculating

$$\tilde{J}^{1,n} = \frac{1}{\gamma} \left(\frac{\beta}{\beta - \rho \left(\mu_c + u^n \sigma_c + \frac{1}{2} \left(\gamma - 1 \right) \sigma_c^2 \right)} \right)^{\frac{\gamma}{\rho}}.$$

In our simulation, we fix the parameters of the endowment dynamics to be $\mu_e = 0.02$ and $\sigma_e = 0.015$, We set the time-discount factor of the agents to $\beta = 0.01$, and study results for different configurations of the measures of risk-aversion $1 - \gamma$, intertemporal elasticity of substitution $\frac{1}{1-\rho}$, and belief distortion of the agent 1, u^1. These values constitute a rough calibration approximately in line with commonly used parametrizations of the aggregate consumption dynamics and agents' preferences, but we do not aspire to provide any particular fit to empirical data. Our interest lies in the relative differences in results for different parametrizations.

To obtain a quantitative notion for the size of the belief distortion, observe that an application of the Girsanov theorem, together with equation (26), implies that the agent with distorted beliefs views the trend component of the aggregate endowment as $\mu_e + \sigma_e u^1$. We use $u^1 = \pm 0.1$; the optimistic and pessimistic agents thus expect the average consumption growth to be 0.0215 and 0.0185, respectively.

3 Survival

In this section, we study the long-run survival chances of agents with distorted beliefs. We base our analysis on the 2-agent economy specified in Section 2.6, with the second agent being rational, i.e. $u^2 = 0$.

We have argued in the previous section that we can interpret the processes $\theta^1 = \bar{\lambda}^1 / \left(\bar{\lambda}^1 + \bar{\lambda}^2 \right)$ and $1 - \theta^1$ as time-varying Pareto shares of the two agents in the planner's problem. These quantities also motivate our definition of survival.

Definition 6 *Agent 1 becomes extinct if*

$$\lim_{t \to \infty} \theta_t^1 = 0 \qquad P\text{-a.s.}$$

Agent 1 survives in the long-run if she does not become extinct.

This definition is borrowed from Kogan, Ross, Wang, and Westerfield (2009) with the exception that we use the Pareto share θ^1 rather than the consumption share ζ^1. However, given that $\theta^1 \to 0$ also implies $\zeta^1 \to 0$ in our case (see Figure 1), the two concepts of survival actually coincide here.

Recall the dynamics of the modified discount factor processes $\bar{\lambda}^n$ in (16). An application of Itô's lemma to $\theta^1 = \bar{\lambda}^1 / \left(\bar{\lambda}^1 + \bar{\lambda}^2\right)$ yields

$$d\theta_t^1 = \theta_t^1 \left(1 - \theta_t^1\right) \left(\nu_t^2 - \nu_t^1 - \theta_t^1 \left(u^1\right)^2\right) dt + \theta_t^1 \left(1 - \theta_t^1\right) u^1 dW_t. \tag{27}$$

Observe that the drift term of the Pareto share contains the difference of the discount rates $\nu_t^2 - \nu_t^1$ which is in the case of non-separable preferences determined endogenously in the model as a solution to problem (15)-(16). Since the general case does not lead to closed-form solutions, we have to rely on numerical methods.

The discount rates ν_t^n are deterministic functions of the state variable θ_t^1, and thus θ^1 is an Itô process on the interval $[0, 1]$ with decaying variance at the boundaries. If there exists a nondegenerate stationary density $q\left(\theta^1\right)$ for θ^1, then it has to satisfy (see Karlin and Taylor (1981, p.220)) the differential equation

$$\frac{q'\left(\theta^1\right)}{q\left(\theta^1\right)} = \frac{2\mu_\theta\left(\theta^1\right) - \left(\sigma_\theta^2\left(\theta^1\right)\right)'}{\sigma_\theta^2\left(\theta^1\right)} \tag{28}$$

where μ_θ and σ_θ^2 are the trend and volatility components of the process

The existence of a stationary density implies that agent 1 survives in the long run. If a nondegenerate stationary density does not exist, then we can base our decision about survival on the pull of the process θ^1.

3.1 Power utility

Recall that in the time-separable case, we have $\nu_t^1 = \nu_t^2 = \beta$, and the irrational agent is losing her Pareto share in expectation. This leads us to the following proposition.

Proposition 7 *Under time-separable preferences (when $\gamma = \rho$), the agent with distorted beliefs becomes extinct, i.e. her Pareto share converges to* 0 *almost surely.*

Proof. Plugging in the components of the θ^1 process into (28) together with $\nu_t^1 = \nu_t^2 = \beta$ implies

$$\frac{q'\left(\theta^1\right)}{q\left(\theta^1\right)} = -\frac{2}{\theta^1} \tag{29}$$

Integrating the equation leads to $q\left(\theta^1\right) = k\left(\theta^1\right)^{-2}$, which is not a proper density on $[0,1]$.

The form of $q\left(\theta^1\right)$ suggests that the conditional distribution of θ^1 shifts toward zero. The fact that the expected share of the irrational agent converges to zero follows from taking conditional expectations of equation (27), thus obtaining the local trend, and integrating it over time. Since zero is the essential infimum of the process θ^1, convergence of the mean to zero together with the boundedness of the process implies convergence of the variance to zero. Mean square convergence implies convergence in probability. Equation (20) for the case $\gamma = \rho$ then implies that the consumption share $C^1/C \to 0$ as $\theta^1 \to 0$. ■

Observe that the result in Proposition 7 does not rely on any specifications of the aggregate endowment process. Thus when all agents have the same power preferences, agents with our type of belief distortion face extinction in any continuous-time infinite-lifetime economy that fits into our general framework. This is a special case of the survival result in Kogan, Ross, Wang, and Westerfield (2009) who show that this holds for any time-separable preferences for which the relative risk aversion is bounded.

The type of economy under consideration matters for the survival results, though. Kogan, Ross, Wang, and Westerfield (2006) consider a different framework where aggregate dividend evolves according to (26) but there are no intermediate dividend payoffs, and the agents only consumes the dividend at a finite terminal date T. In their model, there exists a nontrivial region of parameter combinations for the risk aversion coefficient and the size of the belief distortion that lead to the long-run survival of the irrational agent under power utility, while the rational agent faces relative extinction (her consumption share goes to zero with probability one).

We conclude the analysis of the time-separable case with the following application of the measure transformation using the Girsanov theorem.

Corollary 8 *As long as absolute continuity conditions are satisfied, each agent survives with probability one under her own subjective beliefs.*

Proof. To provide a sketch of the proof, notice that the subjective beliefs of agent n augment the trend component of a Brownian motion by u^n. Taking agent 1 as the reference agent, expressing all subjective beliefs Q^n relative to the measure Q^1 implies augmenting the trend components of the Brownian motions under each Q^n by $-u^1$. Thus $\tilde{u}^n = u^n - u^1$ and $\tilde{u}^1 = 0$, and we can invoke the previous claim as if Q^1 was the true objective measure. Since agent 1 behaves rationally under this measure, she will survive with probability one.

■

3.2 Time-nonseparable case

For the time-nonseparable case, we have

$$\begin{aligned}\frac{q'\left(\theta^1\right)}{q\left(\theta^1\right)} &= \frac{2\left(\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)\right)}{\theta^1\left(1-\theta^1\right)\left(u^1\right)^2}-\frac{2}{\theta^1}= \\ &= 2\left(\frac{\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)}{\left(u^1\right)^2}-1\right)\frac{1}{\theta^1}+\frac{\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)}{\left(u^1\right)^2}\frac{1}{1-\theta^1} \qquad (30)\end{aligned}$$

We have already seen in (29) that for the existence of a stationary distribution, we need to assure the integrability of $q\left(\theta^1\right)$ around zero. Take a point $\underline{\theta}$ in the right neighborhood of zero and impose a hypothetical lower bound

$$\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)\geq\underline{\nu} \qquad \text{for } \theta^1\in\left(0,\underline{\theta}\right)$$

Now integrate expression (30) over $\left[\theta^1,\underline{\theta}\right]$ replacing $\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)$ on the right-hand side with $\underline{\nu}$

$$\ln q\left(\underline{\theta}\right)-\ln q\left(\theta^1\right)\geq 2\left(\frac{\underline{\nu}}{\left(u^1\right)^2}-1\right)\left(\ln\underline{\theta}-\ln\theta^1\right)-\frac{\underline{\nu}}{\left(u^1\right)^2}\left(\ln\underline{\theta}-\ln\left(1-\theta^1\right)\right)$$

Exponentiating and rearranging, we obtain

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

A sufficient condition for the density to be integrable in the neighborhood of 0 is that the exponent on $\theta^1/\underline{\theta}$ has to be larger than -1. Therefore, it is sufficient if

$$\frac{\underline{\nu}}{\left(u^1\right)^2}>\frac{1}{2}.$$

Similarly, take a point $\bar{\theta}$ in the left neighborhood of one, and impose a hypothetical upper bound

$$\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)\leq\bar{\nu} \qquad \text{for } \theta^1\in\left(\bar{\theta},1\right)$$

Again, integration leads to

$$\begin{aligned}\ln q\left(\theta^1\right)-\ln q\left(\bar{\theta}\right) &\leq 2\left(\frac{\bar{\nu}}{\left(u^1\right)^2}-1\right)\left(\ln\theta^1-\ln\bar{\theta}\right)-\frac{\bar{\nu}}{\left(u^1\right)^2}\left(\ln\left(1-\theta^1\right)-\ln\bar{\theta}\right)\\ \frac{q\left(\theta^1\right)}{q\left(\bar{\theta}\right)} &\leq \left(\frac{\theta^1}{\bar{\theta}}\right)^{2\left(\frac{\bar{\nu}}{\left(u^1\right)^2}-1\right)}\left(\frac{1-\theta^1}{\bar{\theta}}\right)^{-\frac{\bar{\nu}}{\left(u^1\right)^2}}\end{aligned}$$

and a sufficient condition for the density to be integrable in the neighborhood of 1 is that $\bar{\nu}/\left(u^1\right)^2<1$, which follows from the exponent on the last term.

These findings are summarized in the following Proposition.

Proposition 9 *Given the dynamics of the Pareto share (27), it is sufficient for the existence of a stationary density if there exist thresholds $\underline{\theta}\leq\bar{\theta}$ such that*

i) $\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)>\frac{1}{2}\left(u^1\right)^2$ *on* $\left(0,\underline{\theta}\right)$

ii) $\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)<\left(u^1\right)^2$ *on* $\left(\bar{\theta},1\right)$

iii) $\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)$ *is bounded on* $\left(\underline{\theta},\bar{\theta}\right)$.

Proof. Conditions i) and ii) follow from the above discussion. Condition iii) assures integrability on $\left(\underline{\theta},\bar{\theta}\right)$. ■

Corollary 10 *Under conditions in Proposition 9, the agent with distorted beliefs survives with probability one.*

Conditions in Proposition 9 can be checked numerically from the solution of the planner's problem. Since the volatility term in (27) provides enough mixing to the stochastic process, the stationary distribution obtained under the conditions of Proposition 9 is unique and the conditional distribution of θ^1 converges to its stationary counterpart.

Conditions i) and ii) have an intuitive interpretation. Agent 1 will survive in the long run when she becomes sufficiently patient relatively to agent 2 (i.e. when $\nu^2-\nu^1$ is above a certain threshold) as $\theta^1\to 0$. In this way, agent 1 sacrifices current consumption in favor of consuming more in the future. The increased future consumption share is linked to an increase in θ^1, thus preventing θ^1 from convergence to zero.

Proposition 9 merely states sufficient conditions for survival, it does tell us anything about whether there are parameter configurations under which the conditions are satisfied, and whether these potential configurations are empirically sensible. Our numerical calculations will now illustrate that survival indeed occurs under commonly used parameter values.

3.3 Numerical results

We take as a benchmark our economy from Section 2.6 with $\rho = \gamma = -2$, i.e. the power utility case with risk aversion 3, and study the impact of departures from time-separability.

Figure 2 displays the behavior of the quantity $\left[\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)\right] / \left(u^1\right)^2$. Proposition 9 states that one of the sufficient conditions for survival of the agent with distorted beliefs is that this quantity is larger than 1/2 in the neighborhood of zero. The solid line represents the power utility case when $\nu^2\left(\theta^1\right) = \nu^1\left(\theta^1\right) = \beta$. In this economy, the agent with distorted beliefs becomes extinct over time.

Decreasing the risk aversion ($\gamma = -1$) relative to the power utility case accelerates the speed of extinction. Figure 3 displays the evolution of the local trend $E_t\left[d\theta_t\right]$, with the solid line depicting the power utility case. Since the volatility components of $d\theta_t$ are the same across models, any model with a uniformly lower $E_t\left[d\theta_t\right]$ necessarily implies extinction of agent 1 as well. Numerical experiments show that this is the case whenever $\gamma > \rho$.

Claim 11 *In economies where the risk aversion is lower relative to the power utility case ($\gamma > \rho$), agent with distorted beliefs faces extinction.*

Going back to Figure 2, we observe that for the case with higher risk aversion ($\gamma = -3$), the quantity the quantity $\left[\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)\right] / \left(u^1\right)^2$ is sufficiently high in the neighborhood of zero for agent 1 to survive. Unfortunately, our numerical results were not able to confirm that this is the case whenever $\gamma < \rho$, i.e. for an arbitrarily small deviation from the power utility case. However, we can still make the following claim.

Claim 12 *Given a belief distortion u^1 and intertemporal elasticity of substitution ρ, agent with distorted beliefs will survive in the long run when the risk aversion in the economy is sufficiently high (i.e. γ sufficiently low).*

Numerical experiments indicate that the necessary increase in risk aversion relative to the power utility case is quantitatively small. For the example in question, we can drive γ as high as -2.05 and still confirm survival. Whether agent with distorted beliefs survives for an arbitrary $\gamma < \rho$ remains an open question.

Intuitively, the case $\gamma < \rho$ can also be interpreted as a high intertemporal elasticity of substitution, relative to the power utility case. An agent with a higher IES is willing to substitute more intertemporally, sacrificing consumption today for consumption tomorrow. This effect is amplified for the agent with distorted beliefs as $\theta^1 \to 0$, as is visible from Figure

2. The endogenously determined discount rate is much lower for agent 1 than for agent 2 as the Pareto share of agent 1 gets small, and agent 1 survives by delaying consumption.

Although we have confirmed a wide range of parameters for which the agent with distorted beliefs survives, it is not a priori clear what share of the aggregate endowment would she consume. For the just discussed parameter values, for instance, the stationary distribution is concentrated extremely close to $\theta^1 = 0$, which makes the calculation of the stationary distribution difficult. While θ^1 close to zero can still imply a nontrivial consumption share (see Figure 1), we calculate the stationary distribution for the case when $\rho = -9$. This implies a risk aversion coefficient of 10, a value that is still well within the range of common asset pricing applications.

Figure 4 displays the consumption share of the agent with distorted beliefs, and the stationary distribution of her Pareto share. While the Pareto share is quantitatively very small, the share of aggregate endowment that the agent consumes is nontrivial, between 5% and 10% of the endowment. In the next section, we will investigate the asset pricing implications of the presence of the agent with distorted beliefs.

It is also worth mentioning the slow expected rate of evolution of the Pareto share. Figure 3 indicates that for the three analyzated parametrizations, it would on average about 20 years to decrease the Pareto share of agent 1 from 0.5 to 0.3, which corresponds to a decrease in the consumption share from about 0.5 to about 0.43. These results approximately hold even for the cases $\gamma = -1$ and $\gamma = -2$ for which we know that agent 1 will ultimately become extinct[7].

4 Asset pricing

Although agents with distorted beliefs may survive in the long run when preferences are time-nonseparable, it is not a priori obvious what is the impact of their presence in the economy on prices of financial securities. In this section, we investigate the question in more detail. In particular, we are interested in the term structure of risk prices, i.e. in the pricing of the risk exposures of cash flows at different horizons to risk sources. Equation (13) reveals that belief heterogeneity contributes an additional source of risk, $u_t dW_t$, linked to the stochastic evolution of the Pareto shares of individual agents.

4.1 Stochastic discount factor

[7] Similar slow convergence has been recently noted for instance by Cogley and Sargent (2008b,a).

Duffie and Epstein (1992a) and Duffie and Skiadas (1994) show that in the case of an agent who has preferences given by the stochastic differential utility, we can write the state the stochastic discount factor between times t and $t+s$ is

$$S_{t,s} = \exp\left(\int_0^s f_v\left(C_{t+u}, V_{t+u}\right) du\right) \frac{f_c\left(C_{t+s}, V_{t+s}\right)}{f_c\left(C_t, V_t\right)} \tag{31}$$

which in the case of agent n corresponds to

$$S_{t,s}^n = \exp\left(\int_0^s -\nu_{t+u}^n du\right) \left(\frac{e_{t+s}}{e_t} \frac{\zeta_{t+s}^n}{\zeta_t^n}\right)^{\rho-1} \left(\frac{J_{\bar{\lambda}^n, t+s}}{J_{\bar{\lambda}^n, t}}\right)^{1-\frac{\rho}{\gamma}}$$

This is an expression containing quantities that we have derived as part of the planner's problem. Observe that in the case of the time-separable utility, we have $f_v \equiv -\beta$, and formula (31) reduces to the conventional ratio of discounted marginal utilities.

Given the equalization of marginal rates of substitution, we can use the stochastic discount factor (31) of any of the agents for pricing, given that we use the subjective beliefs for the evaluation of the expectation in the asset pricing formula. The price of an asset at time t that pays cash flow Z_{t+s} at time $t+s$ is given by $P_{t,s} = E_t^{Q^n}\left[S_{t,s}^n Z_{t+s}\right]$ which is valid for any point of the state space where the agent has nonzero consumption. This formula leads to a PDE for a propagation problem that can be solved backwards for a sequence of maturities starting from the terminal condition $P_{t,s} = Z_{t+s}$. Similarly to the Hamilton-Jacobi-Bellman equation for the planner's value function, the boundary conditions in the subspace given by the Pareto shares are given by solutions to analogous pricing problems for economies arising by omitting some of the agents.

4.2 Risk prices

Equilibrium returns on financial securities in general depend on the time preference of the marginal agents, their preferences toward risk, and the risk properties of the priced payoff[8]. It is common to decompose the return into the risk-free rate and the risk premium, in other words to separate agents' pure time preference from the risk components.

An equally illuminative step is the decomposition of risk premia into exposures of equilibrium payoffs to individual components of risk and the valuation of these risk exposures by agents in terms of required expected return per unit of risk exposure.

It is well-known that in the log-normal power utility model, risk premia are linear in

[8] All these components are in general an integral part of the equilibrium.

the risk exposures, and the coefficients in the linear functions therefore offer a natural interpretation as risk prices. This result does not generalize to other models — risk premia are in general nonlinear functions of the exposures.

We consider pricing of a martingale payoff

$$\frac{D_{t+s}}{D_t} = \exp\left(-\frac{1}{2}\sigma_D^2 s + \sigma_D \int_0^s dW_{t+u}\right) \tag{32}$$

exposed to the same univariate source of risk W as the aggregate consumption. We denote $P_{t,s}^D$ the time t price of a single payoff at instant $t+s$. Linearity of the pricing operator implies that the price of the whole cash flow D is

$$P_t^D = \int_0^\infty P_{t,s}^D ds.$$

Following Hansen (2008), we define the risk price to be the increase in the expected per-period return resulting from a marginal increase in the risk exposure σ_D

$$rp_{t,s}^D = \frac{d}{d\sigma_D}\left[-\frac{1}{s}\log\frac{P_{t,s}^D}{D_t}\right]$$

where we utilized the fact that $E_t\left[D_{t+s}\right] = D_t$. We focus on pricing martingale payoffs in order to correct for the fact that an increase in σ_D in the diffusion term also increases the expected growth payoff through the log-normal adjustment. Risk prices in general depend on the Markov state, and the horizon and risk exposure of the payoff.

Asset prices in our model are calculated numerically using a backward recursive method that discretizes the time axis. The method utilizes the known local evolution of the stochastic discount factor and the priced payoff, which are used to approximate the behavior over small time intervals. This is a version of the algorithm described in detail in Wachter (2005). Our specification is laid out in Appendix (**TO BE WRITTEN**).

4.3 Results

Consider again the economy from Section 2.6, this time with a pessimistic agent with distorted beliefs, $u = -0.1$. The presence of the pessimistic agent should help increase the equilibrium risk premia as in Cogley and Sargent (2008b,a). We focus on the decomposition of returns on the martingale payoff (32) locally, and across different horizons. Given that consumption growth is iid and we are using very modest parameter values in our example economy, we do not expect to provide a quantitative match to asset pricing phenomena.

Our interest lies in qualitative results that connect heterogeneous beliefs and departures from time-separability of preferences.

Since the Pareto share θ^1 is not directly observable, we prefer plotting the graphs as a function of the consumption share ζ^1. It is worth bearing in mind when making comparisons across models that ζ^1 is determined endogenously within the model. Also, the choice of ζ^1 is somewhat arbitrary — we could have chosen for instance calculated wealth levels of the individual agents as well.

Figures 5 and 6 depict the local risk-free interest rate and local risk prices for differently parametrized models. We observe that the risk-free rate is in general decreasing in the consumption share of the pessimistic agent $\zeta^1\left(\theta^1\right)$, while the risk prices are increasing. The lower risk free rate is the result of a higher equilibrium valuation of the risk-free asset vis-a-vis the risky asset which is perceived to have a lower trend component by the agent with distorted beliefs. The same effect also increases the risk prices — while risk-premia widen, the volatility of the martingale payoff stays the same, so that the increase in risk premia must be reflected in the increase in risk prices. Observe that the risk prices are exactly linear in the consumption share of agent 1.

A more interesting result is reveald by the term structure of risk-free bond yields and risk prices in Figures 7 and 8. Bottom panels in both figures depict the dependence on the consumption share of agent 1. As the cash-flow horizon increases, we observe that nonlinearity builds up in the dependence of the bond yields and the risk prices on the consumption share of agent 1. In fact, we can show the following result analytically for the logarithmic time-separable case[9]:

Claim 13 *Assume that $\gamma = \rho = 0$ (the logarithmic, time-separable case) and $0 < \zeta^1 < 1$. Then the long-run risk-free interest rate is calculated as*

$$r^b_{t,\infty} = \lim_{s\to\infty} r^b_{t,s} = \beta + \mu_e - \sigma_e^2 + \min\left(u^1, u^2\right)\sigma_e. \tag{33}$$

Numerical calculations indicate that similar conclusions hold for other parametrizations as well. Thus there is an interesting discontinuity as the maturity of the riskless claim goes to infinity — regardless of how unequal the distribution of wealth (or Pareto shares) is, the limiting long-run bond yield is always the same as in a homogeneous economy populated by the more pessimistic agent. A similar result holds for the risk prices as well.

This discontinuity is not a violation of our assumptions on the continuity of measures.

[9] The logarithmic case with the given consumption dynamics can be solved to a large extent analytically, an exercise that is left out of this paper.

We required that the measures are absolutely continuous at every finite date, but not necessarily in the limit.

The analyzed nonlinearity seems to be intriguing, but it is necessary to point out how slow the convergence is. While other models may produce faster buildup, nonlinearities in the benchmark economy hardly matter for relevant investment horizons. In general, the speed of buildup is slower when we increase risk aversion, holding the intertemporal elasticity of substitution constant, i.e. in the empirically relevant region.

5 Model uncertainty and learning

The survival and asset pricing calculations in the previous text assumed a constant belief distortion u. However, the general framework developed in Section 2 allows to incorporate more general processes as belief distortions which allows us to incorporate agents who learn about the true values of parameters in the model as they receive new information about the evolution of the economy. In order to stay in the our analytical framework, we assume that agents learn from publicly available information, and other agents are aware of this learning process as well as of the current structure of the beliefs of the learning agent.

Consider now again the economy outlined in Section 2.6. Recall that in this economy, agent's 1 belief distortion effectively makes her perceive the trend component of the aggregate endowment growth (26) to be $\tilde{\mu}_e = \mu_e + \sigma_e u^1$. One way to introduce learning is to make agent 1 update her beliefs about the trend component in the Bayesian fashion. We take a slightly different approach here that is more suitable for numerical implementation.

We depart from the assumption that each of the agents is a firm believer in a single model, and introduce model uncertainty. Consider an agent who has in mind a finite set of models indexed by $k = 1, \ldots, K$, which differ in the distorting component $u = (u_k)_{k=1}^{K}$, with the true model being ordered first, i.e. $u_0 = 0$. At time t, the agent assigns a probability distribution $\pi_t = (\pi_{k,t})_{k=1}^{K}$ to this set. Denote $\mathcal{H}_t$ the agent's information set at time t.

In the setup with time-separable utility, the problem of a Bayesian learner in this environment can be solved in two steps. Since the aggregator $f(c, u)$ in (1) is additive under time-separable utility, we can apply the law of iterated expectations, and in each step first take the expectations to account for the risk conditional on the model, and then integrate out across the models.

This method cannot be used when $f(c, u)$ is not separable. However, we will show that we can approach the problem in a similar manner as constant (or, more generally, exogenous) belief distortions, namely by constructing the appropriate distorting measure that accounts for model uncertainty. The marginal measure, integrated out across models, will again be absolutely continuous with respect to the true probability measure, and the general argument about modeling the belief distortion using the Radon-Nikodym derivative of the distorted with respect to the true measure still applies.

Consider learning from the observations of a univariate process X that evolves as

$$dX_t = \mu_x dt + \sigma_x dW_t$$

with a univariate Brownian motion W. Recall that under model k, the agent perceives the trend component to be $\tilde{\mu}_{x,k} = \mu_x + \sigma_x u_k$. It is well-known from the literature on Bayesian updating that the evolution of the probability distribution across models for a Bayesian learner follows

$$d\pi_t = \Delta(\pi_t)\left(dX_t - (\mu_x + \sigma_x u)' \pi_t dt\right)$$

where

$$\Delta(\pi_t) = |\sigma_x|^{-2}\left(\mathrm{diag}(\pi_t) - \pi_t\pi_t'\right)(\mu_x + \sigma_x u)$$

is the regresssion coefficient in the regression of the true state on the evolution of the observed variable under the agent's information set.

The agent perceives that the local trend component of the evolution of X_t is $(\mu + \sigma_x u)' \pi_t$, and thus

$$X_t - \int_0^t (\mu_x + \sigma_x u)' \pi_s ds$$

is a martingale under $\mathcal{H}_t$. This leads us to the construction of a Brownian motion under $\mathcal{H}_t$ defined as

$$d\bar{W}_t \equiv \frac{dX_t - (\mu_x + \sigma_x u)' \pi_t dt}{\sigma_x} = -u'\pi_t + dW_t$$

Thus, the Brownian motion W that is a martingale under the true measure is distorted by the trend component $u'\pi_t$ under the subjective measure. We can therefore introduce a martingale

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

instead of (11) for the learning agent, and proceed in a way analog to the case with exogenous distortions.

Observe that while we stay in the same framework, the introduction of learning comes at a cost — the distribution π_t has to become an additional vector of $K-1$ state variables. Nevertheless, the compactness of the integration region for this vector makes it suitable for numerical calculations.

5.1 Results

The numerical calculations are still to be completed.

6 Conclusion

In this work, we analyze the impact of heterogeneous beliefs under recursive preferences. We show how to obtain the set of 'optimal' allocations as a solution to a planner's problem with Pareto weights which follow endogenously determined Itô processes. In the Markov environment with the Duffie-Epstein-Zin aggregator, we derive the Hamilton-Jacobi-Bellman equation for the value function of the social planner, and describe a solution method that yields the consumption allocations and discount rate processes in the planner's problem. With the consumption allocations rules and individual recursive utility processes, we can construct the stochastic discount factors of the agents and price the securities.

Finally, it is straightforward to incorporate a learning mechanism into this framework. We can think of an agent who has a prior over a set of models which differ in the distortion of the drift term of the underlying state variable, and solves a filtering problem to update his assessment about the correct model based on the observed realizations of the state. We show how to introduce this model uncertainty in a manner that is isomorphic to an exogenous, yet time-varying belief distortion, and derive the relevant distorting martingale.

We apply the theoretical framework in a parsimonious economy with iid consumption growth. The central survival result states that sufficient conditions for the survival of the agent with distorted beliefs are satisfied when the economy is populated by agents with a sufficiently high risk aversion relative to the inverse of the intertemporal elasticity of substitution.

Belief heterogeneity leads to long-horizon risk-free bond yields and risk prices that are highly nonlinear in the consumption share of the agent with distorted beliefs. The convergence of bond yields rates and risk prices is not uniform across horizons. This result is not specific to nonseparable preferences but, in our bechmark economy, only operates over horizons which are empirically hardly relevant.

Our analysis extends the literature on market selection under belief heterogeneity into the environment with time-nonseparable preferences. This extension comes at a cost — we cannot utilize the time-separability in order to easily characterize marginal rates of substitution across horizons that can be used to make conclusions about survival. The need for numerical solutions prevents us from taking a general approach. Nevertheless, we document that for a large set of empirically relevant parameterizations of the Duffie-Epstein-Zin utilities, agents with distorted beliefs survive, and can have a nontrivial consumption share and price impact under the stationary distribution.

A Proofs of claims and propositions

This section of the appendix contains the proofs of claims and propositions omitted from the main text.

Proof of Claim 3. The first order conditions for the problem (19) with $f^n\left(C^n, J_{\bar{\lambda}^n}\right)$ defined in (2) as

$$f^n\left(C^n, J_{\bar{\lambda}^n}\right) = \frac{\beta^n}{\rho^n}\left[\left(C^n\right)^{\rho^n}\left(\gamma^n J_{\bar{\lambda}^n}\right)^{1-\frac{\rho^n}{\gamma^n}} - \left(\gamma^n J_{\bar{\lambda}^n}\right)\right]$$

are given by

$$\bar{\lambda}^n \beta^n \left(C^n\right)^{\rho^n - 1}\left(\gamma^n J_{\bar{\lambda}^n}\right)^{1-\frac{\rho^n}{\gamma^n}} = \mu \qquad n = 1, \ldots, N \tag{34}$$

where μ is the Lagrange multiplier on the feasibility constraint. Inverting the relationship implies

$$\begin{aligned}
C^n &= \mu^{\frac{1}{\rho^n - 1}}\left(\gamma^n J_{\bar{\lambda}^n}\right)^{\frac{\rho^n/\gamma^n - 1}{\rho^n - 1}}\left(\bar{\lambda}^n \beta^n\right)^{\frac{1}{1-\rho^n}} \\
e\left(X\right) &= \sum_{n=1}^{N} C^n = \sum_{n=1}^{N} \mu^{\frac{1}{\rho^n - 1}}\left(\gamma^n J_{\bar{\lambda}^n}\right)^{\frac{\rho^n/\gamma^n - 1}{\rho^n - 1}}\left(\bar{\lambda}^n \beta^n\right)^{\frac{1}{1-\rho^n}}
\end{aligned} \tag{35}$$

which can be solved for $\mu > 0$. If all agents have the same elasticity of substitution coefficient $\rho^n = \rho$, then

$$\mu = \left(e\left(X\right)\right)^{\rho - 1}\left(\sum_{n=1}^{N}\left(\gamma^n J_{\bar{\lambda}^n}\right)^{\frac{\rho/\gamma^n - 1}{\rho - 1}}\left(\bar{\lambda}^n \beta^n\right)^{\frac{1}{1-\rho}}\right)^{1-\rho}$$

Having solved for μ, equation (35) determines the optimal intratemporal allocation of consumption $C^{n*}, n = 1, \ldots, N$, which in the case of identical elasticity of substitution is

$$C^{n*} = \left(e\left(X\right)\right)\frac{\left(\gamma^n J_{\bar{\lambda}^n}\right)^{\frac{\rho/\gamma^n - 1}{\rho - 1}}\left(\bar{\lambda}^n \beta^n\right)^{\frac{1}{1-\rho}}}{\sum_{n=1}^{N}\left(\gamma^n J_{\bar{\lambda}^n}\right)^{\frac{\rho/\gamma^n - 1}{\rho - 1}}\left(\bar{\lambda}^n \beta^n\right)^{\frac{1}{1-\rho}}}$$

which establishes the result. ∎

Proof of Claim 4. The transformation of variables (22) implies the following formulas for the partial derivatives

$$
\begin{aligned}
J_t &= \tilde{J}_t \\
J_x &= \tilde{J}_x \\
J_{xx} &= \tilde{J}_{xx} \\
J_{\bar{\lambda}} &= \frac{\partial g'}{\partial \bar{\lambda}} \tilde{J}_\theta \\
J_{\bar{\lambda}x} &= \frac{\partial g'}{\partial \bar{\lambda}} \tilde{J}_{\theta x} \\
J_{\bar{\lambda}\bar{\lambda}} &= \frac{\partial g'}{\partial \bar{\lambda}} \tilde{J}_{\theta\theta} \frac{\partial g}{\partial \bar{\lambda}'} + \left[\frac{\partial}{\partial \bar{\lambda}^1} \left(\frac{\partial g'}{\partial \bar{\lambda}} \right) \tilde{J}_\theta \; \ldots \; \frac{\partial}{\partial \bar{\lambda}^N} \left(\frac{\partial g'}{\partial \bar{\lambda}} \right) \tilde{J}_\theta \right] \equiv \\
&\equiv \frac{\partial g'}{\partial \bar{\lambda}} \tilde{J}_{\theta\theta} \frac{\partial g}{\partial \bar{\lambda}'} + \mathcal{D}\left(g, \tilde{J}_\theta \right)
\end{aligned}
$$

where J and J_t are scalars, J_x is a $K \times 1$ vector, g is an $N \times 1$ vector function, $\frac{\partial g'}{\partial \bar{\lambda}}$ is an $N \times N$ matrix where the rows correspond to the partial derivatives of the vector function g with respect to each component of $\bar{\lambda}$, $J_{\bar{\lambda}x}$ is an $N \times K$ matrix, $J_{\bar{\lambda}\bar{\lambda}}$ is an $N \times N$ matrix and the matrix $\mathcal{D}\left(g, \tilde{J}_\theta \right)$ representing the expression in brackets is an $N \times N$ matrix where the individual terms form the columns. The bracketed term is a result of a tensor calculus operation.

Writing out the suggested derivatives of the transformation, recalling that

$$
\begin{aligned}
J_{zz} &= \begin{bmatrix} J_{\bar{\lambda}\bar{\lambda}} & J_{\bar{\lambda}x} \\ J_{x\bar{\lambda}} & J_{xx} \end{bmatrix} \\
\Sigma &= \begin{bmatrix} \left(\operatorname{diag}\left(\bar{\lambda} \right) u \right) \left(\operatorname{diag}\left(\bar{\lambda} \right) u \right)' & \left(\operatorname{diag}\left(\bar{\lambda} \right) u \right) \sigma_x' \\ \sigma_x \left(\operatorname{diag}\left(\bar{\lambda} \right) u \right)' & \sigma_x \sigma_x' \end{bmatrix}
\end{aligned}
$$

and observing that

$$
\operatorname{diag}\left(\bar{\lambda} \right) = \theta^N \cdot \begin{bmatrix} \theta^1 & 0 & \ldots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & \theta^{N-1} & 0 \\ 0 & \ldots & 0 & 1 - \sum_{n=1}^{N-1} \theta^n \end{bmatrix} \equiv \theta^N \Theta
$$

we obtain a transformation for the second order term in the Bellman equation (19)

$$\begin{aligned}\frac{1}{2}tr\left(J_{zz}\Sigma\right) &= \frac{1}{2}\mathrm{tr}\left(J_{\bar{\lambda}\bar{\lambda}}\left(\mathrm{diag}\left(\bar{\lambda}\right)u\right)\left(\mathrm{diag}\left(\bar{\lambda}\right)u\right)'\right)+\\ &\quad +\mathrm{tr}\left(J_{\bar{\lambda}x}\sigma_x\left(\mathrm{diag}\left(\bar{\lambda}\right)u\right)'\right)+\\ &\quad +\frac{1}{2}\mathrm{tr}\left(J_{xx}\sigma_x\sigma_x'\right)\\ &= \frac{1}{2}\mathrm{tr}\left(\left(\frac{\partial g'}{\partial\bar{\lambda}}\tilde{J}_{\theta\theta}\frac{\partial g}{\partial\bar{\lambda}'}+\mathcal{D}\left(g,\tilde{J}_{\theta}\right)\right)\left(\theta^N\Theta u\right)\left(\theta^N\Theta u\right)'\right)+\\ &\quad +\mathrm{tr}\left(\left(\frac{\partial g'}{\partial\bar{\lambda}}\tilde{J}_{\theta x}\right)\sigma_x\left(\theta^N\Theta u\right)'\right)+\\ &\quad +\frac{1}{2}\mathrm{tr}\left(\tilde{J}_{xx}\sigma_x\sigma_x'\right)\end{aligned}$$

Thus the Bellman equation (19) is transformed into

$$\begin{aligned}0 &= \theta^N\sum_{n=1}^{N}\Theta_{nn}f^n\left(C^{n*},\frac{\partial g'}{\partial\bar{\lambda}^n}\tilde{J}_{\theta}\right)+\tilde{J}_t+\tilde{J}_x\cdot\mu_x+\\ &\quad +\frac{1}{2}\mathrm{tr}\left(\left(\frac{\partial g'}{\partial\bar{\lambda}}\tilde{J}_{\theta\theta}\frac{\partial g}{\partial\bar{\lambda}'}+\mathcal{D}\left(g,\tilde{J}_{\theta}\right)\right)\left(\theta^N\Theta u\right)\left(\theta^N\Theta u\right)'\right)+\\ &\quad +\mathrm{tr}\left(\left(\frac{\partial g'}{\partial\bar{\lambda}}\tilde{J}_{\theta x}\right)\sigma_x\left(\theta^N\Theta u\right)'\right)+\frac{1}{2}\mathrm{tr}\left(\tilde{J}_{xx}\sigma_x\sigma_x'\right)\end{aligned}$$

which is the stated result. ■

B Epstein-Zin preferences in continuous time

In this section, we 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 the Epstein-Zin preferences is given by

$$\begin{aligned}\tilde{V}_t &= \left[\left(1-e^{-\beta}\right)\left(C_t\right)^{\rho}+e^{-\beta}\tilde{\mathcal{R}}_t\left(\tilde{V}_{t+1}\right)^{\rho}\right]^{\frac{1}{\rho}} \qquad (36)\\ \tilde{\mathcal{R}}_t\left(\tilde{V}_{t+1}\right) &= \left(E_t\left[\left(\tilde{V}_{t+1}\right)^{\gamma}\right]\right)^{\frac{1}{\gamma}}\end{aligned}$$

where $\rho < 1$ and $\gamma < 1$ are parameters with the same interpretation as in (2). For the sake of simplicity, we do not specifically analyze the case when $\gamma = 0$ or $\rho = 0$. Observe that when $\gamma = \rho = 0$, we could derive the time-separable logarithmic preferences using the same limiting argument.

The specification (36) will not directly yield an aggregator of the form (2) 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 the continuous-time case to a compensation using a variance multiplier that introduces an additional term to the recursive utility definition (1).

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

$$V_t = \frac{1}{\gamma}\left(\tilde{V}_t\right)^{\gamma} \tag{37}$$

which leads to the specification

$$V_t = \frac{1}{\gamma}\left[\left(1-e^{-\beta}\right)(C_t)^{\rho} + e^{-\beta}\left(\gamma E_t V_{t+1}\right)^{\frac{\rho}{\gamma}}\right]^{\frac{\gamma}{\rho}} \tag{38}$$

that reduced the certainty equivalence to an expectation[10].

Instead of using a discrete time interval of length one, we now take a time step of length ε and analyze the limit as $\varepsilon \to 0$. Express $E_t[V_{t+\varepsilon}]$ from (38) to obtain

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

and consider the limit

$$\lim_{\varepsilon \searrow 0}\frac{E_t[V_{t+\varepsilon}] - V_t}{\varepsilon}.$$

Applying the L'Hospital rule, we obtain

$$\begin{aligned}
\lim_{\varepsilon \searrow 0}\frac{E_t[V_{t+\varepsilon}] - V_t}{\varepsilon} &= \lim_{\varepsilon \searrow 0}\frac{\left[e^{\beta\varepsilon}(V_t)^{\frac{\rho}{\gamma}} - \left(e^{\beta\varepsilon}-1\right)\gamma^{-\frac{\rho}{\gamma}}(C_t)^{\rho}\right]^{\frac{\gamma}{\rho}} - V_t}{\varepsilon} = \\
&= \lim_{\varepsilon \searrow 0}\frac{\gamma}{\rho}\left[e^{\beta\varepsilon}(V_t)^{\frac{\rho}{\gamma}} - \left(e^{\beta\varepsilon}-1\right)\gamma^{-\frac{\rho}{\gamma}}(C_t)^{\rho}\right]^{\frac{\gamma}{\rho}-1} \cdot \\
&\qquad \cdot\left(\beta e^{\beta\varepsilon}(V_t)^{\frac{\rho}{\gamma}} - \beta e^{\beta\varepsilon}\gamma^{-\frac{\rho}{\gamma}}(C_t)^{\rho}\right) \\
&= \beta\frac{\gamma}{\rho}(V_t)^{1-\frac{\rho}{\gamma}}\left((V_t)^{\frac{\rho}{\gamma}} - \gamma^{-\frac{\rho}{\gamma}}(C_t)^{\rho}\right) = \\
&= -\frac{\beta}{\rho}\frac{(C_t)^{\rho} - (\gamma V_t)^{\frac{\rho}{\gamma}}}{(\gamma V_t)^{\frac{\rho}{\gamma}-1}} \equiv \\
&\equiv -f(C_t, V_t)
\end{aligned} \tag{39}$$

[10]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.

Integrating this expression over time and taking expectations, we have

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

which, assuming $V_T = 0$, implies formula (1).

C Planner's problem

Recall that in the construction of the utility process for the individual agent, we used an ordinally equivalent transformation of the utilities (37). This does not change the preference relation of the individual agent, since that is an ordinal property. However, the planner's problem is a cardinal one — by assigning weights to the utilities of individual agents, we are adding up the utility levels. Thus a weighted average calculated in (14) does not correspond to the same weighting if we used the original specification (36) instead.

In order to obtain the original weights, substitute

$$V_0^n\left(C^n\right) = \frac{1}{\gamma^n}\left(\tilde{V}_0^n\left(C^n\right)\right)^{\gamma}$$

into (14). The first-order conditions lead to equalization of the marginal utilities

$$\alpha^n\left(\tilde{V}_0^n\left(C^n\right)\right)^{\gamma-1}\nabla\tilde{V}_0^n\left(C^n, \hat{C}\right) = \alpha^{n'}\left(\tilde{V}_0^{n'}\left(C^{n'}\right)\right)^{\gamma-1}\nabla\tilde{V}_0^{n'}\left(C^{n'}, \hat{C}\right)$$

for any two agents n and n'. The derivative $\nabla\tilde{V}_0^n\left(C^n, \hat{C}\right)$ is the Gateaux derivative in the direction of the variation $\hat{C}$. An inspection of the first order conditions reveals that the weights $\tilde{\alpha}$ that would lead to the same allocation but with original utility processes $\tilde{V}^n$ in the planner's objective function need to satisfy

$$\tilde{\alpha}^n = \frac{\alpha^n\left(\tilde{V}_0^n\left(C^n\right)\right)^{\gamma-1}}{\sum_{i=1}^N \alpha^i\left(\tilde{V}_0^i\left(C^i\right)\right)^{\gamma-1}}$$

Since these weights depend on the optimal allocation, they can only be obtained after the planner's problem (14) was solved for. However, as long as the preferences satisfy certain regularity conditions, there will be a unique transformation between the weights α and $\tilde{\alpha}$, and thus solutions of the problem (14) parametrized by α will trace out the same set of Pareto optimal allocations as the solution of the planner's problem where the planner maximizes a weighted average of continuation values $\tilde{V}_0^n$.

References

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

Anderson, Evan W., 1998. "Uncertainty and the Dynamics of Pareto Optimal Allocations." PhD dissertation, University of Chicago.

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.

Cogley, Timothy, and Thomas J. Sargent, 2008a. "Diverse Beliefs, Survival, and the Market Price of Risk." University of California, Davis and New York University.

Cogley, Timothy, and Thomas J. Sargent, 2008b. "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?." *Journal of Monetary Economics.* 55 (3): 454–476.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990. "Noise Trader Risk in Financial Markets." *The Journal of Political Economy.* 98 (4): 703–738.

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 Costis Skiadas, 1994. "Continuous-time security pricing: A utility gradient approach." *Journal of Mathematical Economics.* 23 (2): 107–131.

Dumas, Bernard, 1989. "Two-Person Dynamic Equilibrium in the Capital Market." *The Review of Financial Studies*. 2 (2): 157–188.

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.

Fama, Eugene F., 1965. "The Behavior of Stock-Market Prices." *The Journal of Business*. 38 (1): 34–105.

Fama, Eugene F., and Kenneth R. French, 2007. "Disagreement, Tastes, and Asset Prices." *Journal of Financial Economics*. 83 (3): 667–689.

Friedman, Milton (ed.), 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.

Hansen, Lars Peter, 2008. "Modeling the Long Run: Valuation in Dynamic Stochastic Economies." Mimeo.

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, 2004. "The Term Structure of Interest Rate in a Pure Exchange Economy where Investors have Heterogeneous Recursive Preferences." Forthcoming in the Quarterly Review of Economics and Finance.

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.

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

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.

Wachter, Jessica A., 2005. “Solving models with external habit.” *Finance Research Letters.* 2 (4): 210–226.

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

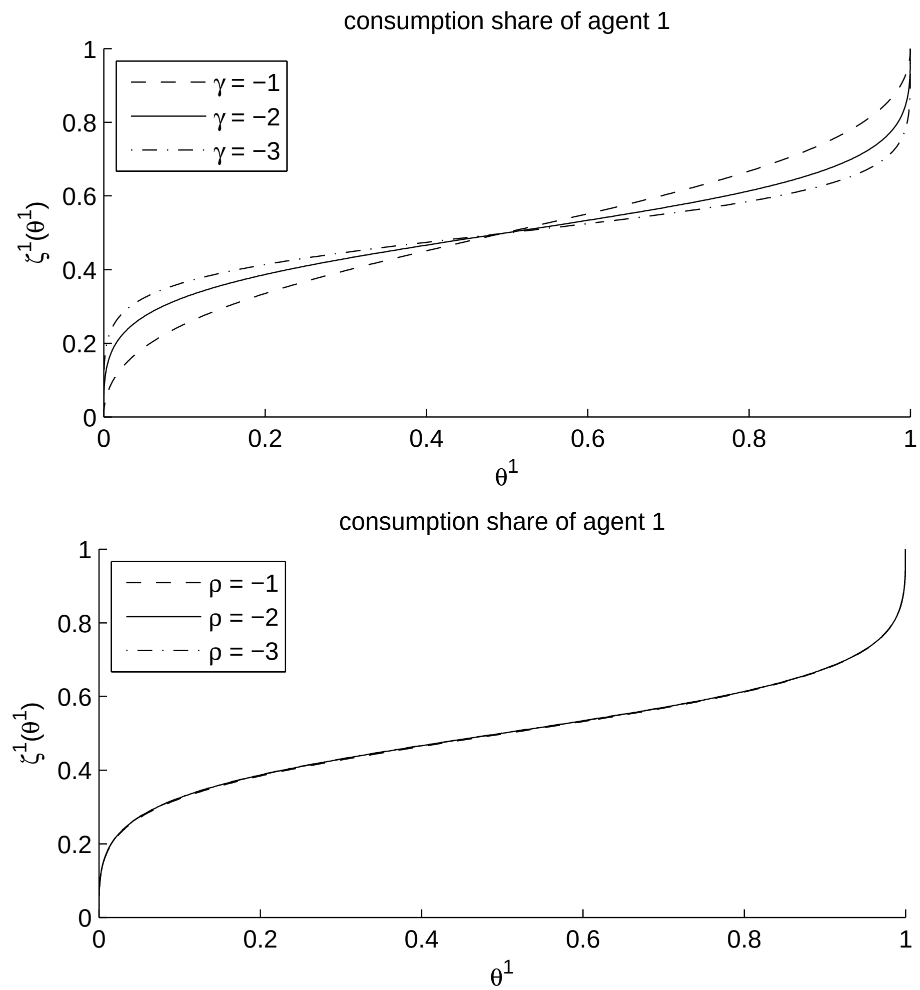

Figure 1: Consumption shares $\zeta^1\left(\theta^1\right)$ as functions of the Pareto weight of the agent with distorted beliefs. The baseline economy has $\rho = \gamma = -2$, i.e. it is the power utility case with risk aversion equal to 3. The top panel shows the result of altering risk aversion parameter while holding the intertemporal elasticity of substitution parameter fixed. The bottom panel shows the result of altering the IES, holding risk aversion fixed.

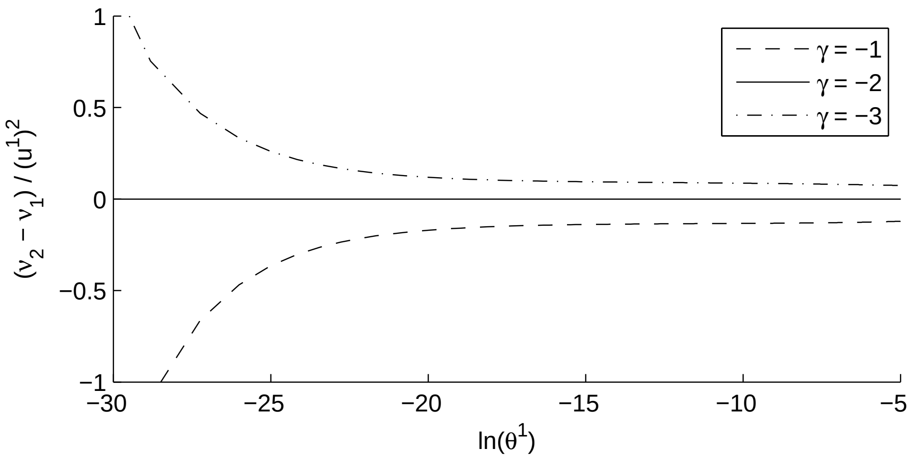

Figure 2: Behavior of the quantity $\left[\nu^2\left(\theta^1\right)-\nu^1\left(\theta^1\right)\right]/\left(u^1\right)^2$ as $\theta^1 \to 0$ for the economy in Section 2.6. $\rho=-2$, $u^1=0.1$. The case $\gamma=-1$ ($\gamma=-3$) corresponds to preferences with a lower (higher) risk aversion relative to the power utility case, holding the intertemporal elasticity of substitution fixed. Horizontal axis in logarithm.

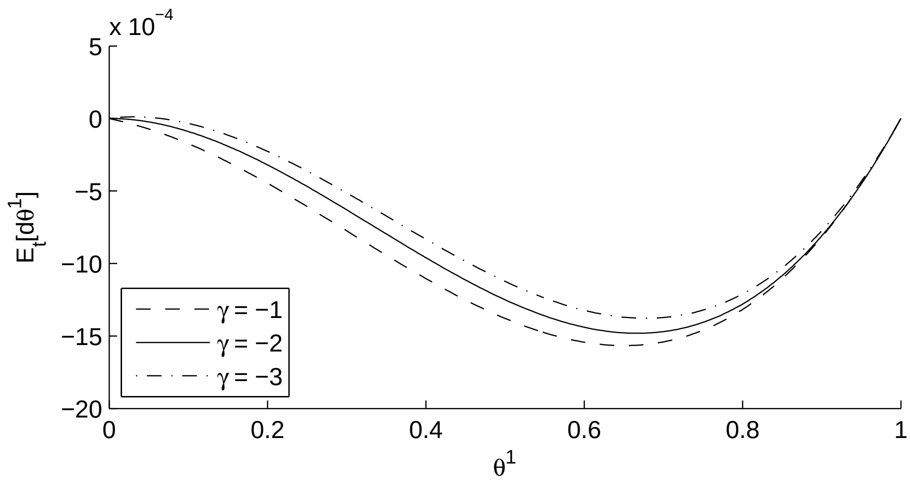

Figure 3: Behavior of the quantity local trend $E_t\left[d\theta_t^1\right]$ for the economy in Section 2.6. $\rho=-2$, $u^1=0.1$. The case $\gamma=-1$ ($\gamma=-3$) corresponds to preferences with a lower (higher) risk aversion relative to the power utility case, holding the intertemporal elasticity of substitution fixed. Negative $E_t\left[d\theta_t^1\right]$ implies that agent 1 is losing her Pareto share in expectation.

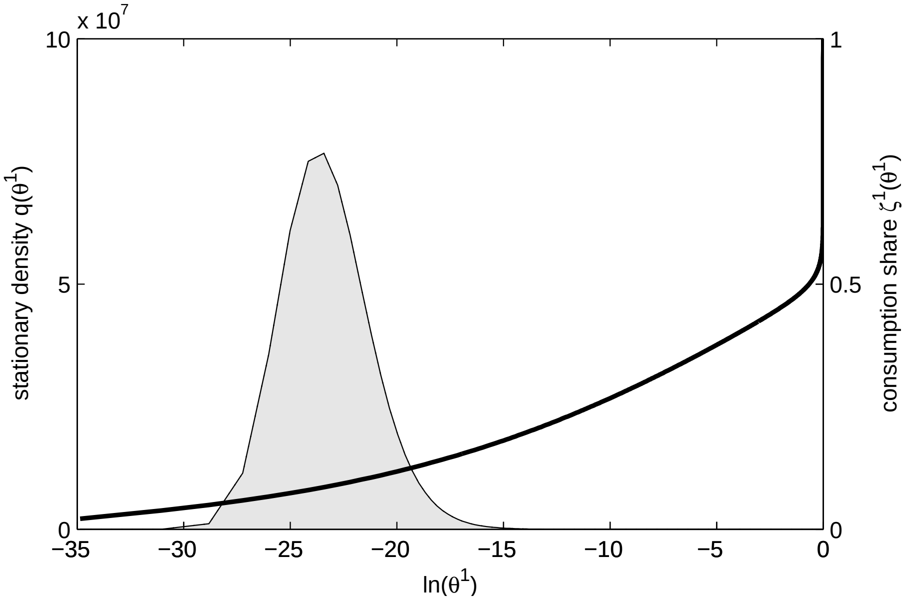

Figure 4: Consumption share of agent 1 and the stationary distribution (shaded) of the Pareto share θ^1 for the economy in Section 2.6. $\rho = -2$, $\gamma = -9$, $u^1 = 0.1$. Horizontal axis in logarithm.

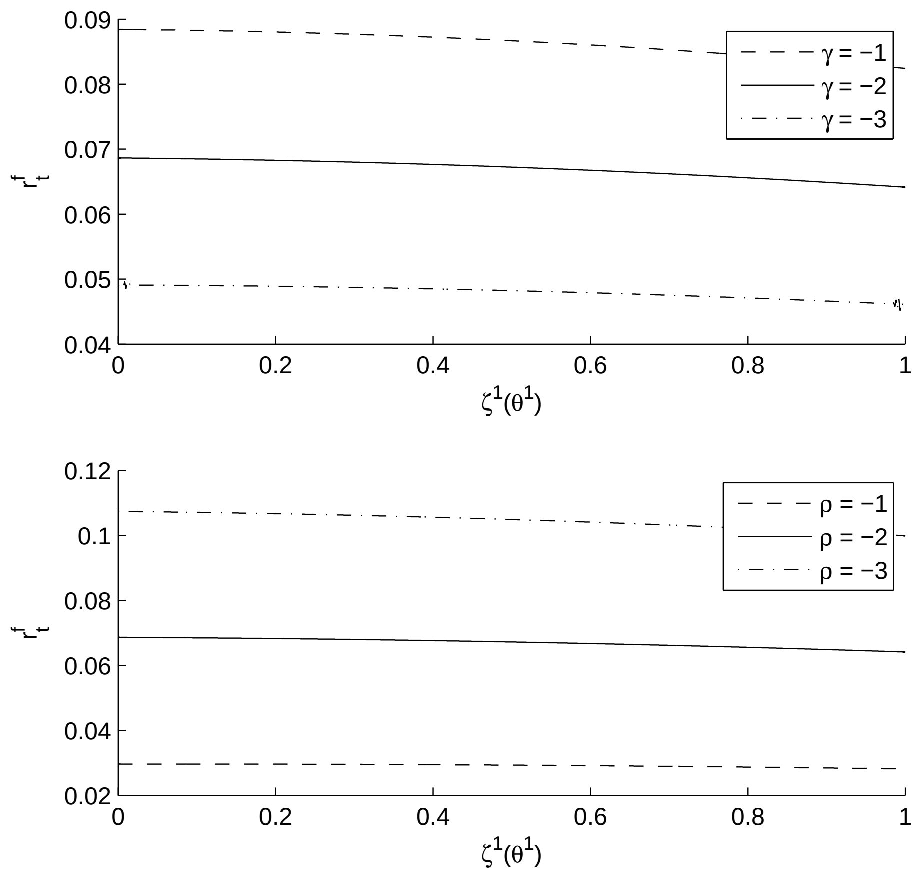

Figure 5: Local risk free rate for the economy in Section 2.6 as a function of the consumption share of agent 1. The benchmark is $\rho = \gamma = -2$, $u^1 = -0.1$, $\beta = 0.01$, the legend denotes parameter deviations from the benchmark.

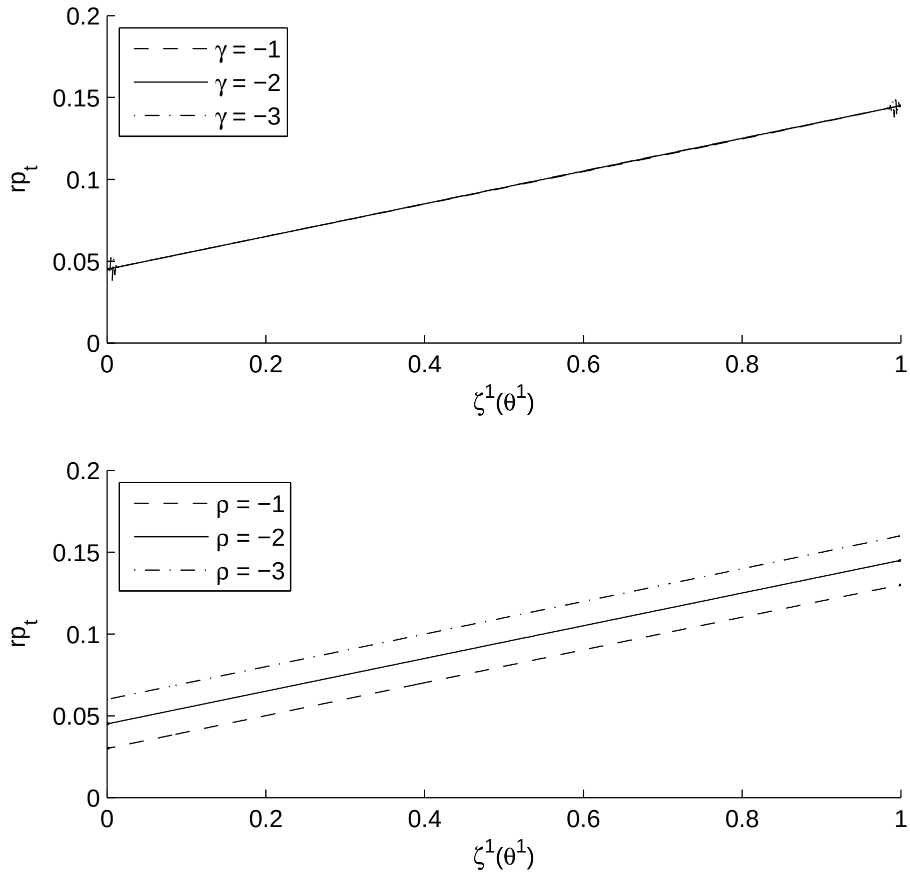

Figure 6: Local risk prices for the economy in Section 2.6 as a function of the consumption share of agent 1. The benchmark is $\rho = \gamma = -2$, $u^1 = -0.1$, $\beta = 0.01$, the legend denotes parameter deviations from the benchmark.

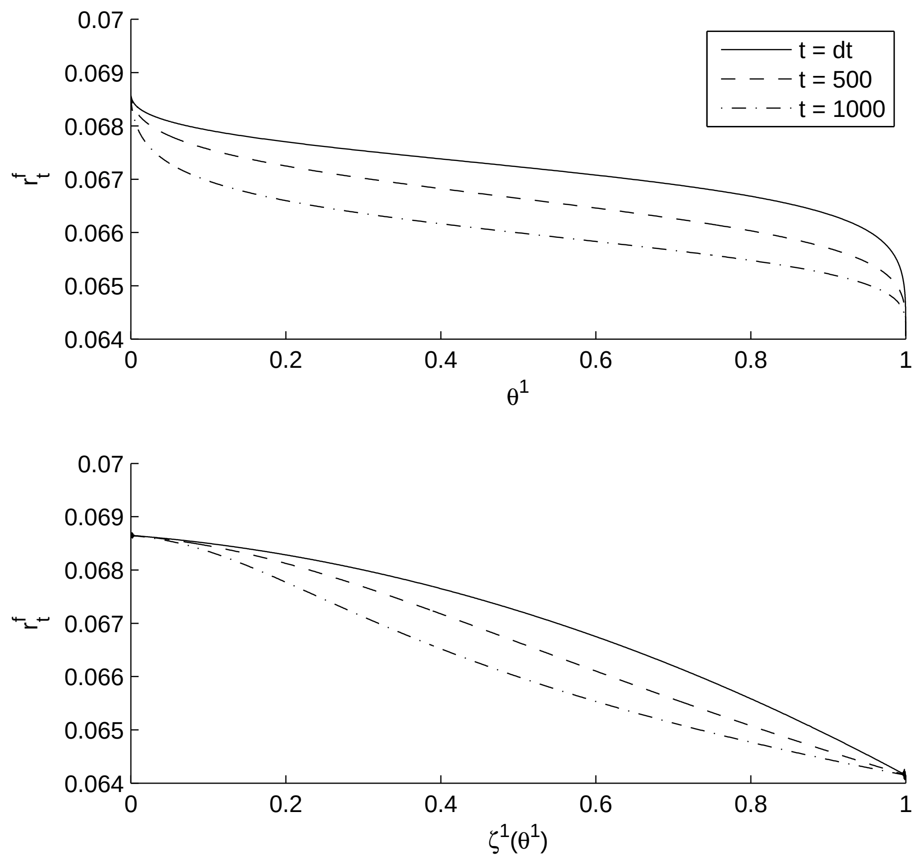

Figure 7: The term structure of interest rates for the economy in Section 2.6. The parametrization is $\rho = -2$, $\gamma = -3$, $u^1 = -0.1$, $\beta = 0.01$. Time in years.

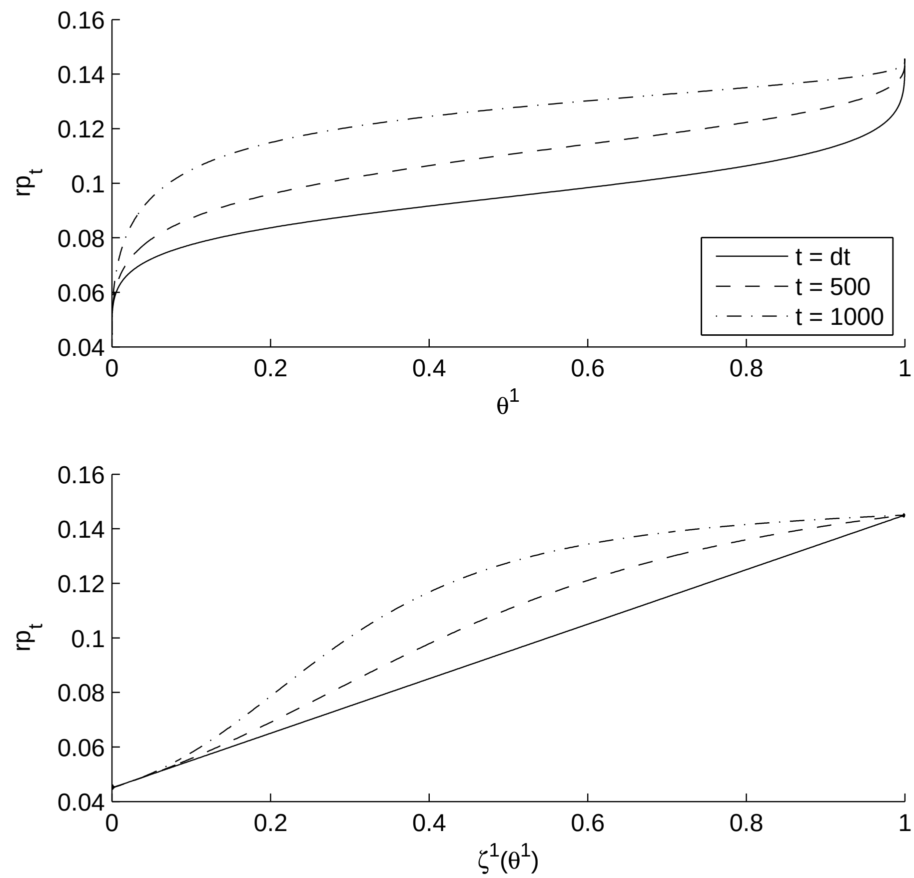

Figure 8: The term structure of risk prices for the economy in Section 2.6. The parametrization is $\rho = -2$, $\gamma = -3$, $u^1 = -0.1$, $\beta = 0.01$. Time in years.