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

Online appendix — not for publication*

Jaroslav Borovička

University of Chicago

Federal Reserve Bank of Chicago

borovicka@uchicago.edu

January 2, 2012

Abstract

This appendix provides further background and extends the results of the main paper. It is not intended for publication.

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

1 Introduction

This online appendix contains several specific results that illuminate and extend the analysis from the paper. Unless noted otherwise, the framework is the same as in the paper. The appendix is not fully self-contained, and occasionally refers to the main text.

In Section 2, I provide an extended discussion of recursive preferences that justifies the link between the discrete-time and continuous-time version of recursive preferences, and between the stochastic differential utility and variational utility approaches in continuous time. Section 3 summarizes additional survival results that are not included in the paper: the case of multiple, mutually correlated shocks, the role of change of measure under which survival is evaluated, survival regions under distortions that are symmetric around the rational case, and the exponential rate of convergence of the Pareto share to its stationary distribution. Section 4 discusses in more detail the extensions considered in the paper, including model uncertainty and learning, and robust utility.

2 Recursive preferences

The paper utilizes a continuous-time characterization of recursive preferences based on a more general variational utility approach studied by Geoffard (1996) in the deterministic case and El Karoui, Peng, and Quenez (1997) in a stochastic environment. This section provides more details on the link between the discrete-time version of recursive preferences specified in Kreps and Porteus (1978) and Epstein and Zin (1989), the continuous-time, stochastic differential utility of Duffie and Epstein (1992b), and the variational utility.

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

Kreps and Porteus (1978) relaxed the separability assumption by axiomatizing discrete-time preferences where temporal resolution of uncertainty matters and preferences are not separable over time. While intratemporal lotteries in the Kreps-Porteus axiomatization still satisfy the von Neumann-Morgenstern expected utility axioms, intertemporal lotteries cannot in general be reduced to atemporal ones. 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.

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.

2.1 Epstein-Zin preferences in continuous time

The survival analysis in the paper is conducted in a continuous-time environment, primarily for tractability reasons. The continuous-time setup leads to a straightforward characterization of the boundary conditions for survival, and an easy decentralization of the economy using only two assets and dynamic trading strategies. However, some intuition for the survival results is provided using the discrete-time version of the recursive preference specification that explicitly reveals the role of risk aversion and intertemporal elasticity of substitution. The derivation of the continuous-time, stochastic differential utility specification closely follows Duffie and Epstein (1992b).

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

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

with parameters satisfying $\gamma, \rho < 1$, and $\beta > 0$. These preferences are homothetic and exhibit a constant relative risk aversion with respect to intratemporal wealth gambles $\alpha = 1-\gamma$ and (under intratemporal certainty) a constant intertemporal elasticity of substitution $\eta = \frac{1}{1-\rho}$. Parameter β is the time preference coefficient. Assumptions provided in the paper restrict parameters to assure sufficient discounting for the continuation values to be finite. In the case when $\gamma = \rho$, the utility reduces to the separable CRRA utility with the coefficient of relative risk aversion α. Notice that the risk adjustment given by the certainty equivalence operator $\tilde{\mathcal{R}}$ operates over the next period continuation value, and the continuation value process is defined in units of current-period consumption. For the sake of simplicity, I omit the situations when $\gamma = 0$ or $\rho = 0$, but these can be treated as appropriate limiting cases.

Since the certainty equivalence $\tilde{\mathcal{R}}_t\left(\tilde{V}_{t+1}\right) = h^{-1}\left(E_t\left[h\left(\tilde{V}_{t+1}\right)\right]\right)$ is not linear in $\tilde{V}$, the continuous-time limit leads to a compensation using a variance multiplier that introduces an additional term to the continuous-time recursion. In order to avoid this issue, it is advantageous to consider an ordinally equivalent transformation of the utility process

$$V_t = \frac{1}{\gamma}\left(\tilde{V}_t\right)^{\gamma} \tag{2}$$

that implies the recursion

$$V_t = \frac{1}{\gamma}\left[\left(1-e^{-\beta}\right)\left(C_t\right)^{\rho} + e^{-\beta}\left(\gamma E_t^Q V_{t+1}\right)^{\frac{\rho}{\gamma}}\right]^{\frac{\gamma}{\rho}}. \tag{3}$$

This transformation reduces the certainty equivalence $\mathcal{R}_t\left(V_{t+1}\right) = E_t^Q V_{t+1}$ to an expectation.[1]

Instead of using a discrete time interval of length one, take a time step of length ε and analyze the limit as $\varepsilon \to 0$. Express $E_t^Q\left[V_{t+\varepsilon}\right]$ from (3) to obtain

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

Applying the L'Hospital rule leads to

$$\begin{aligned}
\lim_{\varepsilon\searrow 0}\frac{E_t^Q\left[V_{t+\varepsilon}\right]-V_t}{\varepsilon} &= \lim_{\varepsilon\searrow 0}\frac{\left[e^{\beta\varepsilon}\left(V_t\right)^{\frac{\rho}{\gamma}} - \left(e^{\beta\varepsilon}-1\right)\gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right]^{\frac{\gamma}{\rho}} - V_t}{\varepsilon} = \\
&= \lim_{\varepsilon\searrow 0}\frac{\gamma}{\rho}\left[e^{\beta\varepsilon}\left(V_t\right)^{\frac{\rho}{\gamma}} - \left(e^{\beta\varepsilon}-1\right)\gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right]^{\frac{\gamma}{\rho}-1}\cdot \\
&\qquad \cdot\left(\beta e^{\beta\varepsilon}\left(V_t\right)^{\frac{\rho}{\gamma}} - \beta e^{\beta\varepsilon}\gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right) \\
&= \beta\frac{\gamma}{\rho}\left(V_t\right)^{1-\frac{\rho}{\gamma}}\left(\left(V_t\right)^{\frac{\rho}{\gamma}} - \gamma^{-\frac{\rho}{\gamma}}\left(C_t\right)^{\rho}\right) = \\
&= -\frac{\beta}{\rho}\frac{\left(C_t\right)^{\rho} - \left(\gamma V_t\right)^{\frac{\rho}{\gamma}}}{\left(\gamma V_t\right)^{\frac{\rho}{\gamma}-1}} \doteqdot -f\left(C_t, V_t\right)
\end{aligned} \tag{4}$$

The function $f\left(C, V\right)$ is called the aggregator function. Integrating this expression over time

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

and taking expectations yields

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

which, assuming the transversality condition $\lim_{T\to\infty}E_t^Q\left[V_T\right]=0$, implies the formula for the stochastic differential utility of Duffie and Epstein (1992b):

$$V_t=E_t^Q\left[\int_t^\infty f\left(C_s,V_s\right)ds\right]\tag{5}$$

with the aggregator defined as

$$f\left(C,V\right)=\frac{\beta}{\rho}\left[\left(C\right)^\rho\left(\gamma V\right)^{1-\frac{\rho}{\gamma}}-\left(\gamma V\right)\right].\tag{6}$$

The aggregator $f\left(C_s,V_s\right)$ links together consumption C_s at time $s\in[t,\infty)$ with the continuation value V_s. Agents prefer early resolution of uncertainty when the aggregator is convex in its second argument. Separability of 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.

An important question is the existence and concavity of the stochastic differential utility $V\left(C\right)$. Duffie and Epstein (1992b) focus on the finite-horizon case and prove concavity only for a concave aggregator f. Appendix C in their paper discusses the infinite-horizon case but the sufficient conditions are too strict for this paper. However, the Markov structure of the problem allows me to utilize the infinite-horizon extensions demonstrated in Duffie and Lions (1992). Schroder and Skiadas (1999) prove that $V\left(C\right)$ is concave even when f is convex in its second argument, a case that is central to this work, and provide further technical details. Skiadas (1997) shows a representation theorem for the discrete time version of (5) with subjective beliefs.

2.2 Variational utility specification

Duffie, Geoffard, and Skiadas (1994) were the first to study optimal and equilibrium allocations with stochastic recursive utility as specified in (5). Dumas, Uppal, and Wang (2000) offer a different way of defining the recursive utility that is more convenient for the purposes of this paper. They show that the recursive utility process V can be equivalently represented as a solution to the maximization problem

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

subject to

$$\frac{d\lambda_t}{\lambda_t} = -\nu_t dt, \;\; t \geq 0; \;\; \lambda_0 = 1, \tag{8}$$

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

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

$$F\left(C, \nu\right) \equiv \inf_{V \in R} \left[f\left(C, V\right) + \nu V\right]. \tag{10}$$

The transformation (9-10) assumes that f is convex in its second argument. When f is concave, it suffices to swap the sup and inf operators in the above definitions.

The duality between the aggregator f and the felicity function F offers a transparent economic interpretation that relates the recursive and variational utility processes. The variational utility representation is an endogenous discounting problem. Given a discount rate ν_t, the concave felicity function F provides instantaneous utility $F\left(C_t, \nu_t\right) dt$, but the decision maker also pays the cost $\nu_t V_t dt$ in the form of increased discounting of the future continuation value. The continuation value V_t thus represents the price of a unit of discount rate ν_t. Problem (9) yields the maximized instantaneous discounted surplus $f\left(C_t, V_t\right) dt$ of the decision maker, and the recursive utility representation aggregates the maximized surplus.

For the case of the Duffie-Epstein-Zin preferences (6), transformation (10) implies

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

corresponding to the felicity function specification considered in the paper.

3 Specific survival results

3.1 Imperfectly correlated shocks

The economy in the paper is driven by a scalar Brownian motion shock W. A natural question arises what happens if there are multiple shocks over which the agents disagree and which are only imperfectly correlated with the innovations to the aggregate endowment. The answer is rather straightforward. Shocks to aggregate endowment can be orthogonalized and conditioned out of the problem, and the remaining problem then maps directly into the original setup.

In particular, consider a modification of the stochastic structure of the economy. The filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ with an augmented filtration defined by a family of σ-algebras $\{\mathcal{F}_t\}, \ t \geq 0$ generated by a bivariate Brownian motion $W = (W^1, W^2)$ with correlated innovations, $\operatorname{Corr}(dW_t^1, dW_t^2) = \varphi dt$. The aggregate endowment is driven by the first component of W,

$$\frac{dY_t}{Y_t} = \mu_y dt + \sigma_y dW_t^1, \qquad Y_0 > 0. \tag{11}$$

Agents $n \in \{1, 2\}$ disagree about the evolution of the second component of W. The ratio of their beliefs Q^n relative to the true probability measure P is given by the Radon-Nikodým derivative

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

The proces W^2 can be interpreted as a betting device that has no fundamental role in the economy, but its realizations are still observable to both agents and the agents can contract upon them. It is not difficult to imagine that such betting devices may exist in the real world, although it is much harder to think about appropriate calibrations (something left for future research).

The law of motion for the aggregate endowment can be rewritten as

$$\frac{dY_t}{Y_t} = \mu_y dt + \sigma_y \left(dW_t^1 - \varphi dW_t^2\right) + \varphi \sigma_y dW_t^2, \qquad Y_0 > 0,$$

where the innovation $dW_t^1 - \varphi dW_t^2$ is uncorrelated with dW_t^2.

Recall that the drift μ_y of the aggregate endowment process does not influence the survival thresholds because it is perceived symetrically by both agents, and is thus cancelled out from the formula for the relative patience (this would not be the case if we consider heterogeneity in IES). The same is true about the contribution of the random component $\sigma_y \left(dW_t^1 - \varphi dW_t^2\right)$

in the evolution of the aggregate endowment process that both agents agree upon. The derivation thus now proceeds as before, with $\varphi\sigma_y$ replacing σ_y. The resulting formulas for the limits of relative patience are

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

These formulas then enter the survival thresholds in the main proposition of the paper — recall that a sufficient condition for the existence of a nondegenerate long-run equilibrium is given by the pair of inequalities

$$\lim_{\theta^1 \searrow 0}\left[\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)\right] > \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right],$$
$$\lim_{\theta^1 \nearrow 1}\left[\nu^2\left(\theta^1\right) - \nu^1\left(\theta^1\right)\right] < \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right].$$

The irrelevance of the shock component that is orthogonal to the shock over which the agents disagree also suggests a possible decentralization. Consider the decentralization using a risk-free infinitesimal bond a two infinitesimal risky assets G and H that pay normalized cash flows

$$dG_t = \sigma_g dW_t^1, \quad dH_t = \sigma_h dW_t^2.$$

When $\theta^1 \searrow 0$, agent 2 holds the aggregate wealth and thus $\pi_g^2(0) = \sigma_y/\sigma_g$ and $\pi_h^2(0) = 0$. Equilibrium excess returns on the two risky assets G and H then are

$$\left[-\varphi u^2 + (1-\gamma)\sigma_y\right]\sigma_g \quad \text{and} \quad \left[-u^2 + \varphi(1-\gamma)\sigma_y\right]\sigma_h,$$

and agent 1 with infinitesimal wealth holds a portfolio with wealth shares

$$\pi_g^2(0) = \frac{\sigma_y}{\sigma_g} \quad \text{and} \quad \pi_h^2(0) = \frac{u^1-u^2}{(1-\gamma)\sigma_h}.$$

The amount of total risk held by both agents thus corresponds to the one-shock example. They both hold unlevered stock positions (see $\pi_g^n(0) + \pi_h^n(0)$ for the case $\sigma_g = \sigma_h = \sigma_y$), and bet on their belief differences using asset H, irrespective of its correlation φ with aggregate stock.

The problem can then be naturally extended to the case of multiple shocks.

3.2 Change of measure

The developed survival criteria are stated from the perspective of a rational observer. However, agents whose beliefs differ from the true probability measure evaluate their survival chances differently. Although both agents understand that the optimal (and equilibrium) allocations are given as a solution to the planner's problem outlined in the paper, they differ in their view about the future consumption dynamics. It is straightforward to restate the analysis from the perspective of the agent with incorrect beliefs. These results are known from earlier literature.

Lemma 1 *Agent n views the dynamics of the economy as if the belief distortions were given by $(u^n)_{(n)} = 0$ and $(u^{\sim n})_{(n)} = u^{\sim n} - u^n$, where $\sim n$ indexes the other agent in the economy and $\left(u^k\right)_n$ are the beliefs of agent k from the standpoint of agent n.*

Proof. The evolution of Brownian motion W under the beliefs of agent n is $dW_t = u^n dt + dW_t^n$. Since the evolution of θ^1 completely describes the dynamics of the economy, substituting this expression into the law of motion for θ^1 and reorganizing yields the desired result. ■

The Lemma implies in particular that the inequalities for survival and dominance developed in the paper apply for the survival and dominance considerations under a subjective probability measure Q^n, as long as u^k are replaced with $\left(u^k\right)_n$ for $k = 1, 2$.

The argument about the change of measure also applies to te planner's problem, and has implications for the local predictability of the modified discount factor processes $\bar{\lambda}^n$. The social planner can choose to maximize welfare as the weighted average of utilities evaluated as distorted relative to any subjective measure, as long as the absolute continuity assumption is satisfied and the distorting martingales M^n are properly constructed relative to the chosen measure. Then the modified discount factor process $\bar{\lambda}^n$ of the agent whose belief distortion coincides with the distortion of the social planner will be locally predictable.

3.3 Survival regions under mirror belief distortions

Under separable CRRA preferences, the case when belief distortions are symmetric around zero, $u^1 = -u^2 = u > 0$, is rather delicate. In this case, both agents survive, but the state variable θ^1 does not have a stationary distribution. What happens is that the conditional distribution gets pulled toward both boundaries, and states when one of the agents has a dominant share of wealth are more and more likely. However, the boundaries are not attracting, and thus given an arbitrary $\theta_0^1 \in (0,1)$, the process θ^1 visits every $\theta \in (0,1)$, P-a.s.

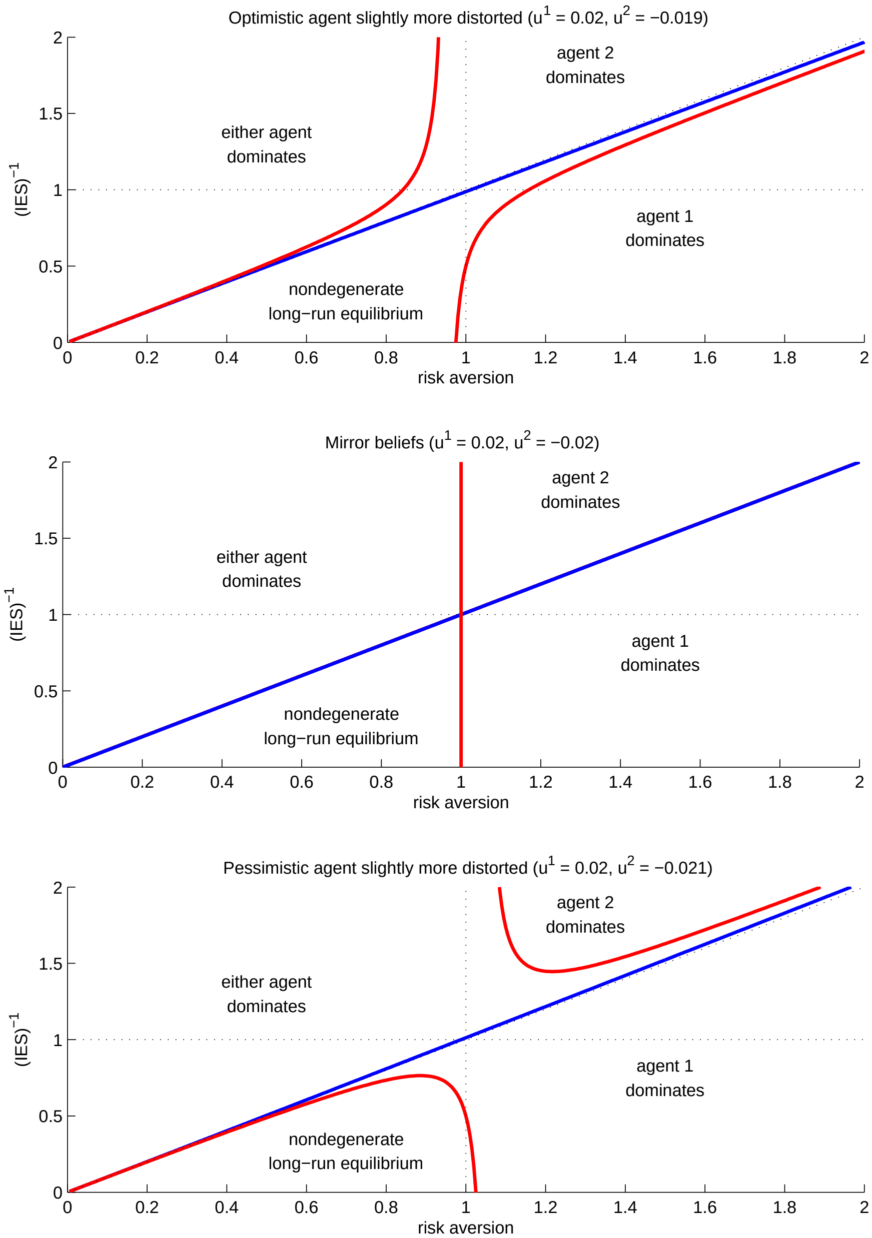

Figure 1: Survival regions for mirror belief distortions, $u^1 \approx -u^2$. All panels assume $\sigma_y = 0.02$. Belief distortion parameters u^n are shown in the titles of individual panels.

When preferences are not separable, this issue generally does not occur, and the parameter space $(1-\gamma, 1-\rho)$ is divided into four regions, and one of the survival results stated in the main survival proposition holds in each of these regions. Figure 1 shows cases when belief distortions of the two agents are close to having opposite signs, $0 < u = u^1 \approx -u^2$. The vertical line in the middle panel satisfies $(1-\gamma)\,\sigma_y = u$. The diagonal line corresponds to CRRA preference specifications which constitute the threshold cases discussed above.

The bottom panel of Figure 1 also reveals the cases when a pessimistic agent with a larger magnitude of the belief distortion can dominate the economy. This can only happen when $u^1 + u^2 + 2\sigma_y < 0$, i.e., when the sum of the two belief distortions is not large, and only when risk aversion is smaller than the inverse of IES. In this region, IES is so low that the relatively more optimistic agent has a sufficiently low saving motive vis-à-vis the high perceived returns on his portfolio, which more than compensates for the willingness of the pessimistic agent to sacrifice high expected returns in order to insure bad outcomes.

3.4 Exponential rate of convergence

When a stationary distribution for the Pareto share θ^1 exists, convergence of the process to its stationary distribution occurs at an exponential rate, so that the process θ^1 does not exhibit strong dependence properties. I state this in the paper as a fact. This result is defined and proven precisely in the following Proposition.

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

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

we have

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

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

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

These conditions are satisfied by imposing the same bounds as those for the finiteness of the speed measure

$$\lim_{\theta \searrow 0} M[\theta, \theta_h] > \infty \quad \text{and} \quad \lim_{\theta \nearrow 1} M[\theta_l, \theta] \quad \text{for some } \theta_l, \theta_h \in (0,1).$$

which in turn coincide with those for the existence of the stationary density for θ^1. ■

Despite the exponential rate of convergence, previous work in the survival literature established that in many cases, δ can be small. Yan (2008) conducts numerical experiments under separable utility when one of the agents always vanishes, and shows that the rate of extinction can indeed be very slow.

4 Extensions

In the paper, I briefly discuss further possible extensions that directly map into the analyzed model. Here, I provide a more detailed outline of the problems. The first extension introduces learning and leads to endogenously varying belief distortions u^n. The second extension incorporates robust utility models.

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

4.1 Model uncertainty and learning

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

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

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

In the setup with separable utility, the aggregator $f(C, V)$ in (5) is additive and linear in V, and the law of iterated expectations can be utilized to solve the problem of a Bayesian learner in two steps. First, calculate the continuation values in the recursive formula (5) conditional on a particular model, and then integrate out across models. This two-step solution works because posterior distributions of a Bayesian learner are martingales under the subjective probability measure of the learner.

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

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

$$dp_t^n = \Delta(p_t^n)\left(\frac{dY_t}{Y_t} - \left(\mu_y^n\right)' p_t^n dt\right), \tag{12}$$

where

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

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

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

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

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

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

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

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

is therefore the distorting martingale that replaces the original distorting martingale in the case of a learning agent. The agent acts as if there was a time-varying average distortion process $\bar{u}_t^n = (u^n)' p_t^n$. The optimization problem of the fictitious planner is extended by the filtering equation (12) and the evolution of the modified discount factor becomes

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

Conjecturing a new Markov state $Z = \left(\bar{\lambda}', Y, (p^1)', (p^2)'\right)'$, it is possible to derive a new version of the HJB equation on a multidimensional but compact set with well-defined boundary conditions that can be built up sequentially from solutions of lower-dimensional problems. The algorithm is the same as in the case of a multi-agent economy described in the paper.

In the paper, I discuss that learning under nonseparable preferences may lead to conclusions that are qualitatively difference from those in Blume and Easley (2006), who find that learning in general improves the survival chances of agents with incorrect beliefs. Under nonseparable preferences, smaller distortions may actually constitute a disadvantage for survival, and thus learning, which diminishes the distortions over time, may have an adverse impact on survival. I leave an explicit solution of this problem to future research.

4.2 Robust utility

Consider an agent who believes that the model for the aggregate endowment dynamics is misspecified and views the dynamics of the aggregate endowment

$$\frac{dY_t}{Y_t} = \mu_y dt + \sigma_y dW_t, \quad t \geq 0$$

only as a reference model that approximates the true dynamics. Anderson, Hansen, and Sargent (2003) and Skiadas (2003), among others, suggest modeling the misspecification by modifying the continuation value problem (7) as

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

subject to

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

The measure Q_u^n is specified by the Radon-Nikodým derivative

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

and the explicit subindex expresses the fact that the minimization problem also includes the choice of the appropriate subjective measure. The set of permissible processes u^n needs to satisfy some regularity conditions like square integrability.

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

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

where W^n is a Brownian motion under Q_u^n. At the same time, the agent understands that specifications that are statistically easy to discriminate from the approximate dynamics are unlikely to be correct, and thus large distortions are penalized by the penalty process $\frac{1}{2}\eta^n\left|u^n\right|^2$. Anderson, Hansen, and Sargent (2003) consider a constant η^n, while Maenhout (2004) makes η^n proportional to the continuation value V^n to retain homogeneity of the optimization problem. Epstein and Miao (2003) and Uppal and Wang (2003) construct models with ambiguity aversion where the optimal solution to the minimization problem involves a constant u^n.

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

Under separable preferences, agents who fear misspecification more (and therefore assign a lower penalty θ to deviations from the reference model) choose a more distorted worst case scenario, which worsens their survival chances.[2] However, the results for constant belief distortions u^n indicate that survival chances of the more fearful agents may well look much better for appropriate nonseparable parameterizations of preferences.

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

[2]An exact statement about survival naturally depends on the model and the choice of the process η^n for each of the agents.

to analyze long-run equilibria in heterogeneous agent economies endowed with a much wider class of preferences than the constant belief distortions that I focused on in the paper.

References

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

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

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

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

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

Duffie, Darrell, Pierre-Yves Geoffard, and Costis Skiadas, 1994. "Efficient and Equilibrium Allocations with Stochastic Differential Utility." *Journal of Mathematical Economics*. 23 (2): 133–146.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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